THE CYBER TOOLCHAIN

NO. 093

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 3, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 211 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 48 tools

▾ // THE BRIEF

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

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

Three items stand out today: SailPoint starts governing AI agents as identities, GitGuardian cuts public-secret triage by judging company

vim ? help

relevance, and ZeroPath can run AppSec scanning on-prem with container scanning and runtime validation.

> GOVERN
>
> **Put AI agents and shadow AI under identity oversight**
>
> SailPoint adds a registry and sensors for AI agents, MCP clients, and other non-human identities; Runlayer adds endpoint visibility plus blocking for risky skills and MCP servers; Zenity detonates AI agent skills before installation. Together they move AI-agent use from ad hoc trust to reviewable governance before access or execution.
>
> SailPoint · Runlayer · Zenity

> PROTECT
>
> **Run vulnerability and AppSec checks closer to the code and runtime**
>
> ZeroPath adds on-prem Kubernetes deployment with container scanning and runtime validation, while DefectDojo's Rules Engine 2.0 can branch, simulate, trace, and push vulnerability workflows into other systems. ZAP's Global Scan Policy Manager makes scan behavior centrally reusable instead of re-tuned per job.
>
> ZeroPath · defectdojo · OWASP ZAP

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

Nmap

Feedback ⓘ changes since 2026-09-11 CODE ↗

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

Nmap gains an 8-script DICOM auditing suite and a full DIMSE toolkit (C-STORE/FIND/GET/MOVE, role negotiation, DICOM-over-TLS).

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-09-11
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-11
```

└──▷ TRY IT

Fingerprint all DICOM services on a PACS subnet and surface weak TLS certificates and unauthenticated access in one pass.

```
$ nmap -p 104,11112 --script dicom-info,dicom-cfind-ls,dicom-store-
inject <target>
```

Fuzz a DICOM C-STORE endpoint to detect crashes or protocol errors in medical imaging devices during a pen test.

```
$ nmap -p 104 --script dicom-store-fuzzer <target>
```

Test whether a DICOM archive allows unauthenticated C-MOVE retrieval to an arbitrary destination AE.

```
$ nmap -p 104 --script dicom-cmove,dicom-get <target>
```

Adds `dicom-info` NSE script to fingerprint a DICOM SCP via A-ASSOCIATE, reporting Implementation Class UID, Version Name, accepted SOP classes and transfer syntaxes, max PDU, protocol version, and — over TLS — server certificate details (key type/size, signature algorithm, weak-crypto and self-signed flags) and mutual-TLS posture.

Adds `dicom-cfind-ls` NSE script to perform C-FIND queries enumerating studies, series, and patient-level metadata, flagging unauthenticated query access.

Adds `dicom-cmove` NSE script to exercise C-MOVE and test whether stored studies can be retrieved to an arbitrary destination AE.

Adds `dicom-get` NSE script to retrieve instances over C-GET with SCP/SCU role negotiation and report unauthenticated read access.

Adds `dicom-store-fuzzer` NSE script to fuzz the C-STORE service with a corpus of malformed datasets, classifying each response (accept, reject, timeout, reset, corrupt) and running periodic C-ECHO health checks to detect crashes mid-campaign.

+3 more

WAS THIS USEFUL? 👍 👎

web-check

Feedback ⓘ 2.3.0 NOTES ↗ RANK

All-in-one OSINT tool for analysing any website

> web-check 2.3.0 adds a new tracker-detection check and verified social-media stats, expands the UI with category pages, adds Render as a one-click deploy target, and ships signed provenance and SBOM attestations on its Docker image.

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

01 **Signed Docker image attestations** IMPROVED 48

Attaches provenance and SBOM attestations to the published Docker image.

– Names supply-chain mechanism but no verification command 2.3.0

02 **Render one-click deploy option** NEW 45

Adds Render as a one-click self-hosted deploy option for web-check.

– Names concrete deploy target but no steps given 2.3.0

03 **Tracker detection check** NEW 33

Adds a tracker check that detects tracking technologies present on a target website.

– Describes what it detects but no mechanism or usage detail 2.3.0

04 **Category pages and job docs in UI** IMPROVED 33

Adds category pages and improved job documentation to the UI.

– Names UI areas but no navigation path given 2.3.0

05 **Verified social media stats** NEW 28

Adds verified social stats surfacing social-media metrics for a website.

– Brief description with no detail on metrics or sources 2.3.0

2.3.0 NOTES ↗ LATEST STABLE

web-check 2.3.0 adds a tracker check, verified social stats, category pages, Render deploy, and signed Docker attestations.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2.3.0 https://github.com/lissy93/web-check.git
# already have the repo? check out this version:
$ git checkout 2.3.0
```

Adds a tracker check to detect tracking technologies present on a target website.

Adds verified social stats surfacing social-media metrics for a website.

Adds category pages and improved job documentation to the UI.

Adds Render as a one-click self-hosted deploy option.

Attaches provenance and SBOM attestations to the published Docker image.

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ changelog · Oct 3, 2026 CHANGELOG ↗ RANK

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

> Bugcrowd added a formal way for customers to request withdrawal of a submitted response.

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

01 **Request a Response Withdrawal** NEW 50

Customers can now use Request a Response Withdrawal to formally request that a submitted response be withdrawn, initiated from the relevant submission in the Bugcrowd customer portal.

When you need to retract a previously submitted response on a program, use the new withdrawal request flow in the Bugcrowd customer portal.

> 📍 In the Bugcrowd customer portal, navigate to the relevant submission, then select Request a Response Withdrawal to initiate the withdrawal process.

— Names exact UI action and flow, but no mechanism detail beyond that.

changelog-20261003-1e68171c

changelog · Oct 3, 2026 CHANGELOG ↗ LATEST STABLE

Bugcrowd adds Request a Response Withdrawal capability for customers.

└──▷ **HOW TO FIND IT**

When you need to retract a previously submitted response on a program, use the new withdrawal request flow in the Bugcrowd customer portal.

> 📍 In the Bugcrowd customer portal, navigate to the relevant submission, then select Request a Response Withdrawal to initiate the withdrawal process.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

ZeroPath

Feedback ⓘ

On-premises support, self-updating threat models, dashboard profiles, and repository inv

NOTES ↗ RANK

ZeroPath is an AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives.

ZeroPath adds on-premises Kubernetes support, self-updating threat models, dashboard profiles, a repository Inventory tab, a Scanner workspace, personal and service API keys, and new API fields for branch inventory and SCA findings control.

NEW SUPPORTED PLATFORM NEW DETECTION Adds on-premises deployment in Kubernetes with OIDC, local accounts, AI gateway, S3-compatible storage, and webhook delivery to private receivers.

Automatically updates threat models during full scans when architecture changes are detected, preserving user edits and recording section authorship and last-update timestamps.

Adds dashboard profiles to save layouts and filters across 13 widgets, kept personal or shared as organization templates, synchronized across devices.

Adds saved issue filters using CVSS 4.0 metrics, Runtime Validation results, dates, and priority scores; Fix Now rules select findings by exploitability and CVSS criteria.

Adds the Inventory tab for exploring packages, licenses, dependency risk, and AI components by branch.

+8 more

NEW INTEGRATION Gerrit and CVS repository connections

NEW INTEGRATION Multiple Bitbucket installations; Bitbucket Data Center PR scanning and instance-level tokens

UPDATED INTEGRATION AI-authorship evidence extended to GitLab, Bitbucket, and Azure DevOps

NEW COMPLIANCE FRAMEWORK ML-BOM export from Inventory tab

└──▷ BREAKING ON UPGRADE

! `issues.list`, `issues.search`, and MCP `issues_list` now include SCA findings when `detectionTypes` is omitted; callers must add `detectionTypes` to retain code-only results.

! Issue severity and score may now be `null`, meaning unavailable rather than zero.

- ! CycloneDX, SPDX, and VEX exports return `412 PRECONDITION_FAILED` for scans marked `inventoryRetiredAt`; integrations must select a current scan.
- ! When `scaPackage.dependencyPathsRetiredAt` is set, `dependencyPaths` is omitted and `transitiveDependencies` may be empty.
- ! When `agent.getRepoCloneUrl` returns `gitConfig` for Bitbucket Data Center, that config must be applied when cloning.
- ! On-premises scanner-setting updates must omit `allowDeepSeekModels`; model selection now belongs to deployment configuration.
- ! Changes to shared rules, context, and scan locations now require access to both their current and destination scopes.
- ! `trends[].meanRemediationTime` is replaced by `trends[].meanRemediationTimeInPeriod`; `trends[].resolvedIssues` provides the sample size.

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ snapshot-20261003 NOTES ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep added SCIM-based user and group provisioning via Okta and Microsoft Entra ID, plus a redesigned billing settings page for purchasing product combinations.

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

01 SCIM provisioning via Okta and Microsoft Entra ID NEW

80

Adds SCIM provisioning at `Settings > Access > Login methods` to sync users and groups from Okta and Microsoft Entra ID. Synced groups can assign deployment roles and manage Semgrep teams (Teams feature in beta); available for default-tenant deployments with SSO configured.

Provision users and groups from your identity provider so new hires automatically inherit the correct Semgrep deployment roles.

> 📍 In Semgrep Cloud, go to Settings › Access › Login methods, then enable SCIM and follow the Okta or Microsoft Entra ID setup guide to configure the SCIM endpoint and bearer token in your IdP.

— Names exact UI path and IdPs, includes setup example snapshot-20261003

THINNER COVERAGE BELOW

02 Redesigned billing settings IMPROVED

35

Redesigns `Settings > Billing` to allow selecting and purchasing any combination of products.

— Names UI path but lacks mechanism detail snapshot-20261003

snapshot-20261003 NOTES ↗ LATEST STABLE

Semgrep adds SCIM provisioning via Okta and Microsoft Entra ID, with group-based role and team management.

└──▷ **HOW TO FIND IT**

Provision users and groups from your identity provider so new hires automatically inherit the correct Semgrep deployment roles.

> 📍 In Semgrep Cloud, go to Settings › Access › Login methods, then enable SCIM and follow the Okta or Microsoft Entra ID setup guide to configure the SCIM endpoint and

bearer token in your IdP.

Adds SCIM provisioning at `Settings > Access > Login methods` to sync users and groups from Okta and Microsoft Entra ID, with groups able to assign deployment roles and manage Semgrep teams (Teams feature in beta); available for default-tenant deployments with SSO configured.

Redesigns `Settings > Billing` to allow selecting and purchasing any combination of products.

WAS THIS USEFUL? 👍 👎

Cisco skill-scanner

Feedback ⓘ 2.2.0 NOTES ↗ CODE ↗

RANK

Security Scanner for Agent Skills

> Skill Scanner's 2.2.0 release and accompanying docs refresh overhaul the LLM judge with decomposed, consensus and adjudication analysis modes, multi-cloud provider support (Bedrock, Vertex AI, Azure, Gemini, Ollama, and an experimental on-device Apple Foundation Model), and new scan policy presets with scoped suppressions. The same window adds two new static detection rules, OSV/VirusTotal/Cisco AI Defense integrations, a REST API and Python SDK, a GitHub Actions workflow and pre-commit hook, plus several breaking changes to exit codes, the default model, and verdict handling.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 **Multi-cloud LLM judge provider support** NEW 95

The LLM judge can run against Anthropic (`anthropic/claude-sonnet-5-5` , default), OpenAI, AWS Bedrock (`bedrock/<model>` or `bedrock-mantle/<model>` via `pip install 'cisco-ai-skill-scanner[bedrock]'` and IAM credentials), Google Vertex AI (`vertex_ai/<model>` , `[vertex]` extra) and Google AI Studio/Gemini (`gemini/<model>` , `[google]` extra), Azure OpenAI (`azure/<deployment-name>` with `SKILL_SCANNER_LLM_BASE_URL` and `SKILL_SCANNER_LLM_API_VERSION` , `[azure]` extra), OpenRouter, OrcaRouter, Ollama (`ollama/<model>` , defaulting `SKILL_SCANNER_LLM_BASE_URL` to `http://127.0.0.1:11434` , loopback-only), an experimental on-device Apple Foundation Model (`apple-fm/system` , macOS 26+, requires `pip install 'apple-fm-sdk>=0.2.1,<0.3'` , no API key), and any OpenAI-compatible gateway via `SKILL_SCANNER_LLM_PROVIDER=openai-compatible` .

Keep skill content fully on-prem by routing the LLM judge through a local vLLM server — no skill data leaves your machines.

```
$ export SKILL_SCANNER_LLM_PROVIDER="openai-compatible"
export SKILL_SCANNER_LLM_BASE_URL="http://127.0.0.1:8000/v1"
export SKILL_SCANNER_LLM_MODEL="gemma-4-26b-a4b"
export SKILL_SCANNER_LLM_API_KEY="unused"
skill-scanner scan ./skill --use-llm
```

Scan with AWS Bedrock as the judge backend using an IAM profile — no API key needed, and prompt budget raised for large skills.

```
$ pip install "cisco-ai-skill-scanner[bedrock]"
export AWS_REGION="us-east-1"
export AWS_PROFILE="my-bedrock-profile"
export SKILL_SCANNER_LLM_MODEL="bedrock/us.anthropic.claude-sonnet-5-5"
skill-scanner scan ./skill --use-llm
```

Run the LLM judge on a self-hosted vLLM model so skill content never leaves your machines.

```
$ export SKILL_SCANNER_LLM_PROVIDER=openai-compatible
export SKILL_SCANNER_LLM_BASE_URL=http://127.0.0.1:8000/v1
export SKILL_SCANNER_LLM_MODEL=gemma-4-26b-a4b
export SKILL_SCANNER_LLM_API_KEY=unused
skill-scanner scan ./skill --use-llm --policy low-noise --fail-on-severity high
```

Route LLM analysis through an internal OpenAI-compatible gateway instead of a first-party provider — useful in air-gapped or proxy environments.

```
$ export SKILL_SCANNER_LLM_PROVIDER=openai-compatible
export SKILL_SCANNER_LLM_API_KEY="your_gateway_key"
skill-scanner ./my-skill --use-llm
```

— Every provider, prefix and pip extra named with examples. `2.2.0`

02

Scan policy presets (default, low-noise, quiet, balanced, strict, permissive)

NEW

95

Six built-in presets are selectable via `--policy` or the `policy` key in policy YAML, and can be generated to a file with `skill-scanner generate-policy --preset <name> -o <file>`. `low-noise` demotes the 11 rules most likely to flag a harmless skill and caps low-confidence judge findings to LOW, reducing the review queue by ~3 points versus `balanced`; `quiet` demotes 19 rules and additionally caps `CONTEXTUAL_RISK` judge findings to LOW, achieving a 7.2% false-positive rate while keeping demoted findings visible at LOW. Policy YAML also exposes `llm_analysis` tuning keys `low_confidence_max_severity`, `contextual_risk_max_severity`, `max_total_prompt_chars`, `max_instruction_body_chars`, `max_code_file_chars` and `max_referenced_file_chars`.

Generate a reusable low-noise policy file to standardise scan settings across pipelines.

```
$ skill-scanner generate-policy -o my_policy.yaml --preset low-noise
```

Cap noisy low-confidence LLM findings and raise the prompt budget for large skills in your policy file.

```yaml
llm_analysis:
  low_confidence_max_severity: LOW
  contextual_risk_max_severity: LOW
  max_total_prompt_chars: 200000
  max_instruction_body_chars: 80000
  max_code_file_chars: 40000
  max_referenced_file_chars: 40000
```

Generate an editable policy from the `low-noise` preset, then customise it to suppress a noisy rule for one skill with an expiry date

```
$ skill-scanner generate-policy --preset low-noise -o .github/scan-policy.yaml
```

Start from a built-in preset to create a tuned policy with a low false-positive rate, then scan your own skills against it.

```
$ skill-scanner generate-policy --preset low-noise -o my_policy.yaml
skill-scanner scan /path/to/skill --use-llm --policy my_policy.yaml --fail-on-severity high
```

— Preset names, numbers and generate-policy command, runnable examples. 2.2.0

---

03 **GitHub Actions reusable scanning workflow** IMPROVED

95

The reusable workflow `cisco-ai-defense/skill-scanner/.github/workflows/scan-skills.yml` takes inputs `scanner_version`, `policy`, `fail_on_severity`, `use_llm`, `llm_model`, `llm_provider`, `llm_base_url` and `aws_region` (default `us-east-1`), plus a `llm_api_key` secret. It supports `bedrock/` models via a Bedrock API key (`SKILL_SCANNER_LLM_API_KEY` secret) or an OIDC role-assumed self-hosted

workflow for role-based Bedrock access, and exposes `--llm-decompose` and LLM-limit tuning options via an `allowlist` in the workflow's `extra_args` input.

Scan third-party skills with the highest-F1 preset in CI, failing on MEDIUM or above — 66.7% recall at 15.4% false-positive rate on held-out splits.

```yaml
uses: cisco-ai-defense/skill-scanner/.github/workflows/scan-skills.yml@2.2.0
with:
  scanner_version: "2.2.0"
  policy: balanced
  use_llm: true
  llm_model: anthropic/claude-sonnet-5-5
secrets:
  llm_api_key: ${{ secrets.SKILL_SCANNER_LLM_API_KEY }}
```

Use the `quiet` preset with the LLM judge to minimise false positives when gating production deploys — blocks 33.1% of malicious skills at a 5.5% false-positive rate at HIGH.

```yaml
uses: cisco-ai-defense/skill-scanner/.github/workflows/scan-skills.yml@2.2.0
with:
  scanner_version: "2.2.0"
  policy: quiet
  use_llm: true
  llm_model: anthropic/claude-sonnet-5-5
secrets:
  llm_api_key: ${{ secrets.SKILL_SCANNER_LLM_API_KEY }}
```

Gate PRs on HIGH-severity findings in your own skills, surfacing MEDIUM findings as inline annotations for reviewers.

```
$ # .github/workflows/skill-scan.yml
jobs:
  skill-scan:
    uses: cisco-ai-defense/skill-scanner/.github/workflows/scan-
skills.yml@2.2.0
    secrets:
      SKILL_SCANNER_LLM_API_KEY: ${{
secrets.SKILL_SCANNER_LLM_API_KEY }}
    with:
      scanner_version: "2.2.0"
      policy: low-noise
      fail_on_severity: high
      llm_model: anthropic/claude-sonnet-5-5
    permissions:
      actions: read
      security-events: write
      contents: read
```

Wire skill-scanner into a GitHub Actions pipeline using the reusable workflow, keeping the LLM API key in repository secrets.

yaml

```yaml
uses: cisco-ai-defense/skill-scanner/.github/workflows/scan-
skills.yml@2.2.0
with:
  scanner_version: '2.2.0'
  policy: low-noise
  fail_on_severity: high
  use_llm: true
  llm_model: anthropic/claude-sonnet-5-5
secrets:
  llm_api_key: ${{ secrets.SKILL_SCANNER_LLM_API_KEY }}
```

— All workflow inputs and secrets named with runnable YAML. 2.2.0

04

Decomposed LLM analysis via `--llm-decompose` NEW

90

`--llm-decompose` runs the LLM judge once per focus area (declared purpose, policy surface, security behaviors) and unions the findings, raising F1 from 61.2% to 68.3% on Gemma 4 26B and from 58.7% to 81.4% on Claude Haiku 4.5, at roughly 3x the model calls.

Maximize recall on a skill when model cost is acceptable — each focus area gets its own LLM pass and results are unioned.

```
$ skill-scanner scan ./skill --use-llm --llm-decompose
```

— Flag, F1 deltas and models named, runnable example given. 2.2.0

05 **Scoped rule suppressions in policy YAML** NEW 90

Policy YAML `suppressions` entries can target a specific skill via `skills`, a file path pattern via `paths`, with `severity`, `reason` and an `expires` date, so a single reviewed false positive (e.g. rule `HOMOGLYPH_ATTACK` or `ARCHIVE_FILE_DETECTED`) no longer disables a rule scan-wide.

Suppress a known false positive for a specific skill and path without disabling the rule globally — ideal for reviewed test fixtures in CI.

yaml

```yaml
suppressions:
  - rule_id: ARCHIVE_FILE_DETECTED
    skills: ["my-skill"]
    paths: ["**/fixtures/**/*.zip"]
    severity: LOW
    reason: "Reviewed test fixtures"
```

Tune a policy YAML to cap weak LLM findings, suppress a false-positive for a specific skill, and pin the judge model

yaml

```yaml
use_llm: true
llm_model: anthropic/claude-sonnet-5-5
low_confidence_max_severity: LOW
contextual_risk_max_severity: LOW
suppressions:
  - rule_id: HOMOGLYPH_ATTACK
    skills: ["docs-translator"]
    reason: "Skill legitimately contains Cyrillic prose"
    expires: 2030-12-31
```

— Exact YAML keys named with runnable suppression example. 2.2.0

06 **REST API server for scanning** NEW 90

`skill-scanner-api --port 8000` starts a REST API exposing `POST /scan-upload`, which accepts a skill ZIP along with `policy` and `use_llm` parameters.

— Exact command, endpoint and params given, directly runnable. product docs

07 **Pre-commit hook and `.skill_scannerrc` config file** NEW

90

A pre-commit hook (pinned via `rev: 2.2.0`) scans only the skills touched by a commit and blocks the commit with exit code 2 if the judge cannot be initialised. The repository-root config file `.skill_scannerrc` sets `skills_path`, `policy`, `use_llm`, `llm_model`, `severity_threshold`, `fail_fast` and (added in 2.2.0) `llm_provider`, with `SKILL_SCANNER_LLM_MODEL` and `SKILL_SCANNER_LLM_PROVIDER` env vars as fallbacks in pre-commit and self-hosted setups.

Configure the LLM judge and severity gate for local and pre-commit runs without repeating flags on every invocation.

json

```
{
  "skills_path": ".claude/skills",
  "policy": "low-noise",
  "use_llm": true,
  "llm_model": "anthropic/claude-sonnet-5-5",
  "severity_threshold": "high",
  "fail_fast": true
}
```

Wire skill-scanner into a repo so every commit only scans touched skills, with the LLM judge enabled via a repo-root config file.

json

```
{
  "skills_path": ".claude/skills",
  "policy": "low-noise",
  "use_llm": true,
  "llm_model": "anthropic/claude-sonnet-5-5",
  "severity_threshold": "high",
  "fail_fast": true
}
```

Enable the LLM judge in pre-commit so every commit to a skills repo is semantically analyzed before it lands.

yaml

```
use_llm: true
llm_model: claude-sonnet-5-5
llm_provider: anthropic
```

— Config keys and hook pin named with runnable examples. 2.2.0

08 **Breaking changes in 2.2.0: fail-closed analyzers, exit codes, defaults**

90

BREAKING

`--use-llm`, `--use-behavioral` and `--enable-meta` now fail closed: if the requested analyzer can't be built (missing API key, missing extra, unknown provider), the scan exits with code 2 and the REST API returns 400, where previously it silently ran rules-only and exited 0. Exit code 2 now specifically means a usage/configuration error (e.g. unknown policy) — previously that returned exit code 1, which still means findings at or above `--fail-on-severity`. The default LLM model is now `claude-sonnet-5-5`, retiring the previous default `claude-3-5-sonnet-20241022`. Verdict repair is now on by default (a SAFE verdict that lists findings is promoted to SUSPICIOUS; restore prior behaviour with `SKILL_SCANNER_LLM_REPAIR_INCONSISTENT_VERDICT=0`). Unpinned dependencies are now reported at LOW severity instead of MEDIUM, and the interactive wizard no longer enables the meta-analyzer by default.

— Each change states before/after behavior and the fix. 2.2.0

09 **`--use-llm` flag enables LLM judge** IMPROVED

85

The `--use-llm` flag runs an LLM-based judge alongside the deterministic rule engine, raising recall of malicious skills from roughly 8% with rules alone to 66.7% recall at the MEDIUM+ threshold. The judge model is selected via `SKILL_SCANNER_LLM_MODEL` (e.g. `export SKILL_SCANNER_LLM_MODEL=anthropic/claude-sonnet-5-5`); if the judge cannot be built the CLI now fails with a non-zero exit and the REST API returns HTTP 400.

Enable the LLM judge with a specific model before running a scan, so AI-assisted evaluation is applied to results

```
$ export SKILL_SCANNER_LLM_MODEL=anthropic/claude-sonnet-5-5
```

— Flag, env var and recall numbers given, runnable example. product docs

10 **Second-pass adjudication via `--adjudicate`** NEW

85

`--adjudicate` (off by default) runs a second-pass LLM review across all findings; it is measured at a cost of 16.4 points of recall, and its model can be overridden independently of the primary judge via `SKILL_SCANNER_ADJUDICATOR_LLM_MODEL`.

Run a second-pass adjudication review across all findings for deeper triage of high-stakes skills.

```
$ skill-scanner scan ./skill --use-llm --adjudicate
```

— Flag, override var and recall cost cited with example. product docs

11 **SARIF/JSON output and Code Scanning upload** NEW 85

`--format json` or `--format sarif` with `--output <file>` writes scan results to disk; SARIF output can be uploaded in self-hosted workflows via `github/codeql-action/upload-sarif@v4` for ingestion into GitHub Code Scanning or a SIEM.

Vet a third-party skill in CI/CD, failing the pipeline on any HIGH severity finding and emitting a SARIF report for upload.

```
$ skill-scanner scan ./my-skill --use-llm --policy low-noise --fail-on-severity high --format sarif --output results.sarif
```

Vet a third-party skill before install, blocking on HIGH and writing MEDIUM+ findings to a JSON review queue.

```
$ skill-scanner scan ./downloaded-skill --use-llm --policy balanced --fail-on-severity high --format json --output scan.json

jq -r '.findings[] | select(.severity == "MEDIUM" or .severity == "HIGH" or .severity == "CRITICAL" or .rule_id == "LLM_ANALYSIS_FAILED") | "\(.severity)\t\(.rule_id)\t\(.file_path // "-")\t\(.title)"' scan.json
```

— Format flags and upload action named with runnable examples. product docs

12 **LLM consensus runs via `--llm-consensus-runs`** NEW 80

`--llm-consensus-runs` runs the LLM analyzer multiple times independently and retains only findings with majority agreement, taking the highest severity observed across the agreeing votes. Consensus mode multiplies the number of judge runs to raise the agreement threshold, at proportional cost.

Reduce false positives in CI by requiring findings to survive across three independent LLM runs before surfacing.

```
$ skill-scanner scan ./skill --use-llm --llm-consensus-runs 3
```

— Flag and mechanism described, runnable example given. product docs

13 **Cisco AI Defense meta-analyzer via `--use-aidefense`** NEW

80

`--use-aidefense` (off by default) enables a meta-analyzer backed by Cisco AI Defense as a cloud-based content-scanning backend requiring its own API key; it is measured at a 16.4-point recall cost, and its model can be overridden independently of the primary judge via the `SKILL_SCANNER_META_LLM_*` variable family.

Run a full dependency advisory check against OSV.dev alongside LLM and behavioral analysis for a thorough audit with no API key.

```
$ skill-scanner scan ./my-skill --use-behavioral --use-llm --use-osv --use-aidefense
```

— Flag, recall cost and override var named, example shown. product docs

14 **OSV.dev dependency vulnerability checks** IMPROVED

80

`--use-osv` queries OSV.dev for known-vulnerable Python and JavaScript dependency pins, requiring network access but no API key; it was extended to also read `package.json` (npm) manifests in addition to previously supported manifest types.

Run a full dependency advisory check against OSV.dev alongside LLM and behavioral analysis for a thorough audit with no API key.

```
$ skill-scanner scan ./my-skill --use-behavioral --use-llm --use-osv --use-aidefense
```

— Flag, ecosystems and manifest extension named with example. 2.2.0

15 **New supply-chain and frontmatter detection rules** NEW

80

Adds `SUPPLY_CHAIN_REGISTRY_REDIRECT` (HIGH) to detect package-registry redirection across npm/yarn/pnpm, pip/uv, Go, Cargo, RubyGems, NuGet and Docker plus unnamed script network destinations, and `UNDECLARED_NETWORK_DESTINATION` (LOW) for undeclared network calls. Frontmatter scanning was extended to the `description` and `when_to_use` fields of `SKILL.md` — which agents load before invocation — covered by core rules, YARA and active-directive rules (optional packs still scan only the body).

— Rule IDs, severities and ecosystems named; no example. 2.2.0

16 **LLM judge gateway and output tuning variables** NEW

75

New environment variables tune the judge for compatibility and behaviour: `SKILL_SCANNER_LLM_FORCE_JSON_OBJECT=true` falls back to plain JSON mode for gateways that reject `json_schema` output; `SKILL_SCANNER_LLM_USER` sets the Chat Completions `user` field for routing or billing; `SKILL_SCANNER_LLM_TEMPERATURE` overrides the default temperature-0 behaviour (accepts a number or `none`); `SKILL_SCANNER_LLM_MAX_TOKENS` / `--llm-max-tokens N` (default 8192) controls the output token budget; and `SKILL_SCANNER_LLM_REASONING_EFFORT` / `--llm-reasoning-effort` accepts `disabled`, `minimal`, `low`, `medium`, `high`, `xhigh` or `max` to control reasoning depth.

— Every tuning variable named; no runnable example shown. product docs

---

17 **Homebrew install and prebuilt wheel coverage** NEW 75

Installable on macOS 14+ via `brew tap cisco-ai-defense/skill-scanner` and `brew install cisco-ai-defense/skill-scanner/skill-scanner`. Prebuilt wheels cover glibc Linux x86-64/ARM64, macOS x86-64/ARM64, and Windows x86-64; other platforms must build from source, requiring Go 1.27.1+ for the CEL helper.

Install skill-scanner on macOS via Homebrew for a dependency-isolated, project-independent setup.

```
$ brew tap cisco-ai-defense/skill-scanner
https://github.com/cisco-ai-defense/skill-scanner
brew install cisco-ai-defense/skill-scanner/skill-scanner
```

— Exact brew commands and platform matrix, runnable example. product docs

---

18 **Python SDK for programmatic scanning** NEW 70

A Python SDK exposes `SkillScanner`, `build_analyzers` and `ScanPolicy` (from `skill_scanner`, `skill_scanner.core.analyzer_factory` and `skill_scanner.core.scan_policy`) for programmatic scanning, including a `use_llm=True` analyzer option that pulls the model and API key from `SKILL_SCANNER_LLM_*` environment variables.

— Exact classes and modules named; no code example. product docs

---

19 **Scan command flags: recursion, severity gating, behavioral analysis**

NEW 70

`scan-all` gained `--recursive` to scan an entire registry or directory tree and `--fail-on-severity` (e.g. `high`) to gate CI on findings at or above a severity. `--use-behavioral` enables Python dataflow analysis during a scan, and an interactive wizard

mode was added to `scan` to guide selection of target, analyzers, policy and output format.

Gate CI on HIGH findings while using the LLM judge across a skills directory, with policy defined in a shared YAML file

```
$ skill-scanner scan-all ./skills --use-llm --policy
.github/scan-policy.yaml --fail-on-severity high
```

Scan a third-party skill with the LLM judge, all analyzers enabled, and fail CI on any HIGH finding — outputting a SARIF file for ingestion into a SIEM or GitHub Code Scanning.

```
$ skill-scanner scan-all ./skills --recursive --use-llm --use-
behavioral --use-osv --policy balanced --fail-on-severity high --
format sarif --output results.sarif
```

— Each flag named with purpose, demonstrated in examples. product docs

20 **System One advisory risk screening** NEW 65

`--system-one-endpoint URL` routes scans through an optional System One screening endpoint (`POST /v1/systemone`) and `--system-one-model NAME` specifies the model for it (required whenever the endpoint is set). The score is advisory only — it records a calibrated risk probability for triage and cannot change a finding, its severity, or the overall verdict.

— Flags and endpoint named; no usage example provided. product docs

THINNER COVERAGE BELOW

21 **VirusTotal cloud scanning via `--use-virustotal`** NEW 40

`--use-virustotal` opts in to VirusTotal as an optional cloud-based binary scanning backend for Tier 3 behavioral and reputation checks on third-party skills; it requires its own API key and is off by default to keep skill content on-machine.

— Only flag name and API-key requirement, no example. product docs

└──▷ BREAKING ON UPGRADE

! `--use-llm`, `--use-behavioral`, and `--enable-meta` now fail closed: if the requested analyzer cannot be built (missing API key, missing extra, unknown provider), the scan exits with code 2 and the REST API returns 400. Previously the scan ran rules only and exited 0.

! Exit code 2 now means a usage or configuration error (e.g. unknown policy); previously this condition returned exit code 1. Exit code 1 still means findings at or above `--fail-on-severity`.

! The default LLM model is now `claude-sonnet-5-5`; the previous default `claude-3-5-sonnet-20241022` is retired.

! Verdict repair is on by default: a SAFE verdict that lists findings is promoted to SUSPICIOUS. To restore the prior behavior, set `SKILL_SCANNER_LLM_REPAIR_INCONSISTENT_VERDICT=0`.

! Unpinned dependencies are now reported at LOW severity instead of MEDIUM.

! The interactive wizard no longer enables the meta-analyzer by default.

2.2.0 NOTES ↗ CODE ↗ STABLE

Skill Scanner 2.2.0 adds `low-noise` / `quiet` presets, scoped suppressions, Apple on-device LLM, OpenAI-compatible gateway support, and two new detection rules.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2.2.0 https://github.com/cisco-ai-defense/skill-scanner.git
# already have the repo? check out this version:
$ git checkout 2.2.0
```

└──▷ **USE IT**

Suppress a known false positive for a specific skill and path without disabling the rule globally — ideal for reviewed test fixtures in CI.

yaml

```yaml
suppressions:
  - rule_id: ARCHIVE_FILE_DETECTED
    skills: ["my-skill"]
    paths: ["**/fixtures/**/*.zip"]
    severity: LOW
    reason: "Reviewed test fixtures"
```

Enable the LLM judge in pre-commit so every commit to a skills repo is semantically analyzed before it lands.

.skill_scannerrc

```
use_llm: true
llm_model: claude-sonnet-5-5
llm_provider: anthropic
```

Route LLM analysis through an internal OpenAI-compatible gateway instead of a first-party provider — useful in air-gapped or proxy environments.

```
$ export SKILL_SCANNER_LLM_PROVIDER=openai-compatible
export SKILL_SCANNER_LLM_API_KEY="your_gateway_key"
skill-scanner ./my-skill --use-llm
```

Adds `low-noise` and `quiet` scan presets that demote high-false-positive rules and cap judge findings via `low_confidence_max_severity` and `contextual_risk_max_severity`.

Adds scoped rule suppressions in policy YAML: target a specific skill with `skills` or a file path pattern with `paths`, so a single false positive no longer disables a rule scan-wide.

Adds `SKILL_SCANNER_LLM_PROVIDER=openai-compatible` environment variable to route LLM analysis through any OpenAI-compatible gateway.

Adds `--llm-decompose` flag to run one focused LLM analysis pass per concern instead of a combined pass.

Adds `use_llm`, `llm_model`, and `llm_provider` keys to `.skill_scannerrc` so the pre-commit hook can invoke the LLM judge.

+6 more

BREAKING ON UPGRADE

- `--use-llm`, `--use-behavioral`, and `--enable-meta` now fail closed: if the requested analyzer cannot be built (missing API key, missing extra, unknown provider), the scan exits with code 2 and the REST API returns 400. Previously the scan ran rules only and exited 0.
- Exit code 2 now means a usage or configuration error (e.g. unknown policy); previously this condition returned exit code 1. Exit code 1 still means findings at or above `--fail-on-severity`.
- The default LLM model is now `claude-sonnet-5-5`; the previous default `claude-3-5-sonnet-20241022` is retired.
- Verdict repair is on by default: a SAFE verdict that lists findings is promoted to SUSPICIOUS. To restore the prior behavior, set `SKILL_SCANNER_LLM_REPAIR_INCONSISTENT_VERDICT=0`.
- Unpinned dependencies are now reported at LOW severity instead of MEDIUM.
- The interactive wizard no longer enables the meta-analyzer by default.

WAS THIS USEFUL?

GitHub

Feedback ⓘ

Copilot code review: API support and new default effort level (2026-10-02)

BLOG ↗ RANK

GitHub is a web-based platform for version control and collaborative software development using Git repositories.

> GitHub opened up Copilot code review via API (while quietly raising the default review effort level to Balanced), enriched the SecurityAdvisory GraphQL object with new triage fields and filters, and shipped a public-preview REST API for commenting on repository security advisories.

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

01 Copilot code review API and effort-level controls

BREAKING

75

Adds REST and GraphQL API support for requesting GitHub Copilot code reviews, with an optional review effort level per request. The default review effort level changed from `Lite` to `Balanced` for new and existing repositories and organizations as of September 28, 2026, so repos that relied on the old default now get `Balanced` reviews unless `Lite` was explicitly set. New per-level UI controls let you set the review effort level (`Lite` , `Balanced`) at the enterprise, organization, repository, or personal scope via settings pages.

Override the default effort level back to Lite for a specific repository whose reviewers prefer lighter feedback.

> 📍 In your repository settings, go to Copilot -> Code review and set the review effort level to `Lite` .

— Names API, levels, and the settings path to override the default

Copilot code review: API support and ne…

02 New fields and filters on SecurityAdvisory GraphQL API

NEW

67

Adds `severities` and `isWithdrawn` filters to the `securityAdvisories` GraphQL query for server-side narrowing of results instead of filtering client-side. The `SecurityAdvisory` GraphQL object gains four new fields: `sourceCodeLocation` (links to affected source code), `githubReviewedAt` (when GitHub reviewed the advisory), `nvdPublishedAt` (when NVD published its record), and `repositoryAdvisoryUrl` (links to the associated repository security advisory when one exists).

— Names every field and filter but gives no query example

New fields for SecurityAdvisory GraphQL…

└──▷ **BREAKING ON UPGRADE**

! The default review effort level changed from `Lite` to `Balanced` on September 28, 2026 — repositories and organizations that relied on `Lite` as the default will now receive `Balanced` reviews unless `Lite` was explicitly selected in settings.

Copilot code review: API support and new default effort level (2026-10-02)

BLOG ↗ LATEST STABLE

GitHub Copilot code review is now accessible via REST and GraphQL APIs with configurable effort levels.

└──▷ **HOW TO FIND IT**

Override the default effort level back to Lite for a specific repository whose reviewers prefer lighter feedback.

📍 In your repository settings, go to Copilot -> Code review and set the review effort level to `Lite`.

Adds REST and GraphQL API support for requesting GitHub Copilot code reviews, with an optional review effort level per request.

Changes the default review effort level from `Lite` to `Balanced` for new and existing repositories and organizations.

Adds per-level UI controls to set the review effort level (`Lite` , `Balanced`) at the enterprise, organization, repository, or personal scope via settings pages.

└──▷ **BREAKING ON UPGRADE**

! The default review effort level changed from `Lite` to `Balanced` on September 28, 2026 — repositories and organizations that relied on `Lite` as the default will now receive `Balanced` reviews unless `Lite` was explicitly selected in settings.

New fields for SecurityAdvisory GraphQL API (2026-10-02) BLOG ↗ STABLE

GitHub's SecurityAdvisory GraphQL API gains five new fields and two new server-side filters for advisory triage.

Adds `severities` and `isWithdrawn` filters to the `securityAdvisories` GraphQL query, enabling server-side narrowing of advisory results without downloading and filtering client-side.

Adds `sourceCodeLocation` field to the `SecurityAdvisory` GraphQL object linking to affected source code relevant to the advisory.

Adds `githubReviewedAt` field to the `SecurityAdvisory` GraphQL object exposing when GitHub reviewed the advisory.

Adds `nvdPublishedAt` field to the `SecurityAdvisory` GraphQL object exposing when the NVD published its record.

Adds `repositoryAdvisoryUrl` field to the `SecurityAdvisory` GraphQL object linking to the associated repository security advisory when one exists.

Repository security advisory comments API in public preview (2026-10-02)

BLOG ↗ STABLE

GitHub adds REST API endpoints to list, get, add, and edit comments on repository security advisories.

└──▷ **TRY IT**

Automatically post a triage note to a security advisory when your pipeline finishes analysis — keeping advisory discussion out of the web UI and in your toolchain.

```
$ curl -X POST \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  https://api.github.com/repos/<owner>/<repo>/security-
advisories/<ghsa_id>/comments \
  -d '{"body": "Triage complete: exploitability confirmed on v2.3.1.
Patch queued for next release."}'
```

Check which open advisories already have discussion before pulling full comment threads — useful when triaging a backlog.

```
$ curl -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  'https://api.github.com/repos/<owner>/<repo>/security-advisories'
\
  | jq '[.[] | select(.comments_count > 0) | {ghsa_id, summary,
comments_count}]'
```

Export all advisory comments updated since a given date for an audit trail or migration to another system.

```
$ curl -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  'https://api.github.com/repos/<owner>/<repo>/security-
advisories/<ghsa_id>/comments?since=2026-01-01T00:00:00Z'
```

Adds `GET` endpoint to list comments on a repository security advisory, with optional filtering by last-updated time.

Adds `GET` endpoint to retrieve a single comment on a repository security advisory.

Adds `POST` endpoint to add a comment to a repository security advisory.

Adds `PATCH` endpoint to edit a comment on a repository security advisory.

Repository security advisory responses now include a `comments_count` field so callers can identify advisories with discussion before fetching comments.

+1 more

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ changelog · Oct 3, 2026

CHANGELOG ↗ RANK

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

> SonarQube Server 2026.5 LTA ships a suite of AI agents — Remediation Agent for sandboxed, verified automated fixes, Hunter Agent for logic-flaw discovery, and Sonar Vortex for in-loop coding guidance — alongside an embedded MCP server, deeper C/C++ and Rust analysis, new compliance reports, and redesigned reporting dashboards.

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

01 Remediation Agent for automated verified code fixes NEW

94

`SonarQube Remediation Agent` is a containerized engine deployable inside Kubernetes or Docker on Enterprise/Data Center editions that connects to a cloud LLM endpoint using your own key (supporting AWS Bedrock, Azure AI Foundry, native Anthropic and OpenAI keys, custom gateways, and models including GPT-5.5 and Claude Opus 4.6), integrates with GitHub, GitLab, or Azure DevOps, and covers Java, JavaScript/TypeScript, Python and C# issues including leaked secrets and logic flaws. It re-runs Sonar analysis in a sandbox before submitting a pull request, discarding fixes that don't clear the issue or introduce new ones; an Agent activity page tracks running and finished jobs; runs can be scheduled daily or weekly per organization or project with a configurable PR cap, or triggered on-demand by filtering the Issues page for agent-fixable items; it supports air-gapped and VPC-restricted environments so code never leaves operator-controlled infrastructure.

Enable the Remediation Agent on a self-hosted SonarQube Server instance connected to Azure AI Foundry for regulated environments that cannot use external AI endpoints.

> 📍 In the SonarQube Server administration UI, go to AI Capabilities, connect your DevOps platform (GitHub, GitLab, or Azure DevOps) with the required permissions, select Azure AI Foundry as your model endpoint, and enable the Remediation Agent.

Schedule weekly automated remediation at the project level to keep technical debt low without manual triage, capping PR volume to avoid overwhelming reviewers.

> 📍 In the SonarQube Server project settings under the Remediation Agent section, set the schedule to weekly and configure the maximum number of pull requests per run, then save.

Trigger on-demand remediation before an audit by filtering the Issues page to agent-fixable items in a single project and assigning them to the agent for immediate PR generation.

> 📍 On the SonarQube Server Issues page, filter by issues the Remediation Agent can fix, select the target issues, and assign them to the agent to generate and submit pull requests grouped by rule and file type.

— Names flags, UI pages, scheduling and exact enablement steps `launch-20260929-8ac52e0b` `launch-20261001-95eff5a5` `changelog-20261003-9e5fd489`

---

02 **Rust analysis beyond Clippy** NEW 64

Adds Rust analysis beyond Clippy, layering centralized quality profiles, quality-gate enforcement, Cognitive Complexity and Cyclomatic Complexity metrics, code-coverage correlation from LCOV or Cobertura, duplication detection, and coverage of popular Rust crates.

— Names coverage formats and metrics but no setup steps `launch-20260929-8ac52e0b`

---

03 **Embedded MCP server endpoint** NEW 60

Adds an embedded `SonarQube MCP Server` endpoint letting Cursor, Claude Code, GitHub Copilot and other agentic tools reach project issues and quality gates, governed by a token-based kill-switch.

— Named endpoint and governance control, but no connect steps `launch-20260929-8ac52e0b`

THINNER COVERAGE BELOW

---

04 **Hunter Agent for logic and access-control flaws** NEW 58

`SonarQube Hunter Agent` is an on-demand background agent that uses your own Anthropic key to trace data and identity flows across a full codebase, uncovering broken access control, business logic, and authentication flaws such as IDOR, privilege escalation, and session fixation that pattern-based scanners miss, and surfaces confirmed findings with severity directly as SonarQube issues.

— Describes mechanism and flaw types but no usage path `launch-20260929-8ac52e0b` `changelog-20261003-9e5fd489`

---

05 **New MISRA, accessibility and CRA compliance reports** NEW 54

Raises MISRA C automated coverage to 38.4% for MISRA C:2012, adds a new MISRA C/C++ Compliance Report generating audit-ready, tamper-evident attestation artifacts automatically, and adds new WCAG Accessibility and CRA Compliance Reports proving conformance at the project, application, and portfolio level.

– Gives concrete coverage percentage and named report types launch-20260929-8ac52e0b

06 **New portfolio and project reporting dashboards** NEW 52

Adds customizable portfolio and project dashboards with a new overview landing page for organization-wide reporting, a project coverage dashboard showing which DevOps platforms are connected, which repositories are imported, and which projects have not yet been analyzed, and a quality gate adherence dashboard flagging risky releases that shipped despite a failing gate.

– Names three dashboards and what each surfaces launch-20260929-8ac52e0b

07 **Sonar Vortex semantic code navigation engine** NEW 48

Sonar Vortex is a semantic Code Navigation engine, self-hosted on SonarQube Server Enterprise/Data Center editions, that injects project context and constraints into coding agents and verifies changes inside the agent loop before commit, guiding AI agents to write code aligned with project standards, preventing new issues from being introduced, and reducing token consumption.

– Explains mechanism but no concrete setup steps given launch-20260929-8ac52e0b changelog-20261003-9e5fd489

08 **Expanded language and framework analysis coverage** NEW 46

Adds native support for `MuleSoft DataWeave` and `R` as new analysis languages, and expands language coverage to include Java 25, Python web frameworks (`FastAPI`, Flask, Django), Jenkins Groovy, `PowerShell`, Gosu, Ruby, and Apex.

– Lists all named languages but gives no mechanism detail launch-20260929-8ac52e0b

09 **Direct reachability analysis for Java and C#** NEW 44

Adds direct reachability analysis for Java and C# identifying which dependency vulnerabilities the application code actually calls.

– Names languages and purpose but no usage path launch-20260929-8ac52e0b

10 **Incremental taint analysis speeds up PR scans** IMPROVED 44

Adds incremental taint analysis cutting pull request scan times on large Java and C# projects by up to 90%.

– Concrete performance number but no configuration path launch-20260929-8ac52e0b

11 **Azure API Management for LLM governance** NEW 42

Adds Azure API Management support for LLM governance so Azure OpenAI customers can authenticate without static API keys.

— Named integration but mechanism and setup undescribed launch-20260929-8ac52e0b

12 **Deeper C/C++ taint and cross-file analysis** NEW 41

Adds new taint analysis for C and C++ covering command-injection and path-traversal detection, and adds cross-translation-unit awareness to C/C++ symbolic execution, making analysis consistent across file boundaries.

— Names vulnerability classes and analysis scope improvement launch-20260929-8ac52e0b

13 **Quality gate for agent-generated code** NEW 34

Adds a new quality gate purpose-built for agent-generated code that enforces stricter security, reliability, and supply-chain standards while staying permissive on coding style.

— States policy intent but no configuration detail launch-20260929-8ac52e0b

14 **Unified vulnerability model for security hotspots** IMPROVED 27

Migrates security hotspots into a unified vulnerability model preserving full history.

— Bare statement of migration with no mechanism detail launch-20260929-8ac52e0b

15 **Architecture drift detection** NEW 25

Adds architecture management that automatically catches architectural drift and surfaces it in the existing developer issue list.

— Bare description, no mechanism or scope given launch-20260929-8ac52e0b

16 **Secrets masked and redacted by default** IMPROVED 23

Detected secrets are now masked and redacted by default to reduce potential exposure.

— States default behavior change with no further detail launch-20260929-8ac52e0b

changelog · Oct 3, 2026 CHANGELOG ↗ LATEST STABLE

SonarQube Server 2026.5 LTA adds Sonar Vortex, Remediation Agent, and Hunter Agent for AI-assisted code quality and security.

Adds Sonar Vortex, a semantic Code Navigation engine that guides AI agents to write code aligned with project standards and prevents new issues from being introduced, reducing token consumption.

Adds SonarQube Remediation Agent, a containerized engine deployable inside Kubernetes or Docker that connects to a cloud LLM endpoint using your own key,

integrates with GitHub, GitLab, or Azure DevOps, and delivers SonarQube-verified automated code fixes without requiring public internet exposure.

Adds SonarQube Hunter Agent, an on-demand background agent that uses your own Anthropic key to uncover broken access control, business logic, and authentication flaws that pattern-based scanners miss, surfacing confirmed findings directly as SonarQube issues.

announced · Oct 1, 2026 BLOG ↗

SonarQube Remediation Agent arrives on SonarQube Server, auto-generating verified PRs for C#, Java, JS/TS, and Python issues inside self-hosted infrastructure.

└──▷ **HOW TO FIND IT**

Enable the Remediation Agent on a self-hosted SonarQube Server instance connected to Azure AI Foundry for regulated environments that cannot use external AI endpoints.

> 📍 In the SonarQube Server administration UI, go to AI Capabilities, connect your DevOps platform (GitHub, GitLab, or Azure DevOps) with the required permissions, select Azure AI Foundry as your model endpoint, and enable the Remediation Agent.

Schedule weekly automated remediation at the project level to keep technical debt low without manual triage, capping PR volume to avoid overwhelming reviewers.

> 📍 In the SonarQube Server project settings under the Remediation Agent section, set the schedule to weekly and configure the maximum number of pull requests per run, then save.

Trigger on-demand remediation before an audit by filtering the Issues page to agent-fixable items in a single project and assigning them to the agent for immediate PR generation.

> 📍 On the SonarQube Server Issues page, filter by issues the Remediation Agent can fix, select the target issues, and assign them to the agent to generate and submit pull requests grouped by rule and file type.

Adds the Remediation Agent to SonarQube Server Enterprise and Data Center editions, enabling automated fix generation, sandbox verification, and pull request submission without code leaving self-hosted infrastructure.

Supports AWS Bedrock, Azure AI Foundry, native Anthropic and OpenAI keys, and custom gateways as model endpoints for the Remediation Agent on SonarQube Server.

Supports C#, Java, JavaScript/TypeScript, and Python for automated remediation of maintainability, reliability, and select security issues including leaked secrets and logic flaws.

Adds an Agent activity page showing every remediation job (running and finished) with status, duration, source, and the pull requests opened.

Enables scheduling of remediation runs at the organization or project level on a daily or weekly cadence with a configurable cap on pull requests per run.

+3 more

announced · Sep 29, 2026 BLOG ↗

SonarQube Server 2026.5 LTA brings Vortex, Remediation Agent, Hunter Agent, and deep analysis upgrades to self-hosted Enterprise/DC editions.

Adds Azure API Management support for LLM governance so Azure OpenAI customers authenticate without static API keys.

Adds an embedded `SonarQube MCP Server` endpoint letting Cursor, Claude Code, GitHub Copilot and other agentic tools reach project issues and quality gates, governed by a token-based kill-switch.

Adds Sonar Vortex self-hosted deployment for SonarQube Server Enterprise and Data Center editions, injecting project context and constraints into coding agents and verifying changes inside the agent loop before commit.

Adds `SonarQube Remediation Agent` self-hosted deployment with bring-your-own-LLM support (including GPT-5.5 and Claude Opus 4.6), sandboxed fix generation and re-analysis, covering Java, JavaScript, TypeScript, Python and C# via GitHub, GitLab and Azure DevOps.

Adds `SonarQube Hunter Agent` self-hosted deployment, using your own Anthropic key to trace data and identity flows across a full codebase and detect IDOR, privilege escalation, and session fixation, delivering confirmed findings with severity into the existing SonarQube issue list.

+17 more

WAS THIS USEFUL? 👍 👎

Openlayer

Feedback ⓘ announced · Sep 23, 2026 BLOG ↗ RANK

Openlayer is the AI governance and observability platform for regulated enterprises: one place to test, monitor, and govern every model, agent, and AI system your company runs.

> Openlayer launched jevals, an open-source Python library for typed agent evals and runtime guardrails, and brought its Jev judge natively into the Openlayer platform alongside existing classifier evaluation tooling.

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

01 **jevals library with typed evaluate() API** NEW 85

Openlayer open-sourced `jevals`, installable via `pip install jevals`, a Python library for agent evals and runtime guardrails built around typed decision models that replace freeform LLM-judge text generation with structured choices, scores, and probabilities. The library exposes an evaluate() function that accepts agent messages and tool schemas and returns typed results accessible via .table().

Run security and behavior checks on a live agent trace to catch indirect injection and scope violations before acting on the result.

python

```python
from jevals import evaluate
from jevals.agent import ToolChoice, Grounded, StayedInScope
from jevals.security import IndirectInjection

result = evaluate(
    {"messages": messages, "tools": tools},
    [
        ToolChoice(),
        Grounded(),
        StayedInScope(),
        IndirectInjection(),
    ],
)
print(result.table())
```

Install jevals to start writing typed agent evals and guardrails in a new project.

```
$ pip install jevals
```

— Names install command and exact API with runnable example. launch-20260923-0abe1814

02 **37 prebuilt and custom agent evaluations** NEW 75

`jevals` ships 37 prebuilt evaluations covering tool selection, groundedness, task completion, scope adherence, prompt injection, and sensitive information, importable as `ToolChoice`, `Grounded`, `StayedInScope`, `IndirectInjection`, and more from the `jevals.agent` and `jevals.security` modules. Teams can also write custom evaluations in Python or YAML alongside these prebuilt checks.

Run security and behavior checks on a live agent trace to catch indirect injection and scope violations before acting on the result.

python

```python
from jevals import evaluate
from jevals.agent import ToolChoice, Grounded, StayedInScope
from jevals.security import IndirectInjection

result = evaluate(
    {"messages": messages, "tools": tools},
    [
        ToolChoice(),
        Grounded(),
        StayedInScope(),
        IndirectInjection(),
    ],
)
print(result.table())
```

— Names 37 checks, import paths and modules, but no custom-eval syntax shown.

launch-20260923-0abe1814

03 **Multiple backend support for evaluations** NEW 60

`jevals` supports hosted Jev, local Kev and Laya backends, as well as conventional LLM backends, with evaluation definitions remaining stable across backend switches.

— Names all backends but gives no switching mechanism or command.

launch-20260923-0abe1814

THINNER COVERAGE BELOW

04 **Gate mechanism for policy-based guardrails** NEW 55

`jevals` adds a gate mechanism that applies a policy to evaluation results and routes agent actions to proceed, block, or require human review, letting the same evaluation definition run in both development testing and production guardrails.

— Describes policy/routing mechanism but no API or config names given.

launch-20260923-0abe1814

05 **Jev as native judge inside Openlayer** NEW 45

Jev is now available as a native judge directly inside the Openlayer platform, alongside existing classifier evaluation tooling such as F1 scores, confusion matrices, drift, and explainability.

— States integration exists but no usage path or UI location given.

launch-20260923-0abe1814

announced · Sep 23, 2026 BLOG ↗

Openlayer launches jevals, an open-source Python library for typed agent evals and guardrails, plus Jev as a native judge inside Openlayer.

└──▷ **USE IT**

Run security and behavior checks on a live agent trace to catch indirect injection and scope violations before acting on the result.

python

```python
from jevals import evaluate
from jevals.agent import ToolChoice, Grounded, StayedInScope
from jevals.security import IndirectInjection

result = evaluate(
    {"messages": messages, "tools": tools},
    [
        ToolChoice(),
        Grounded(),
        StayedInScope(),
        IndirectInjection(),
    ],
)
print(result.table())
```

Install jevals to start writing typed agent evals and guardrails in a new project.

```
$ pip install jevals
```

Adds `pip install jevals` open-source Python library for agent evals and runtime guardrails built around typed decision models (Jev-style), replacing freeform LLM-judge text generation with structured choices, scores, and probabilities.

Adds evaluate() function from `jevals` accepting agent messages and tool schemas, returning typed results via .table().

Adds 37 prebuilt evaluations in `jevals` covering tool selection, groundedness, task completion, scope adherence, prompt injection, and sensitive information —

importable as `ToolChoice`, `Grounded`, `StayedInScope`, `IndirectInjection`, and more from modules `jevals.agent` and `jevals.security`.

Supports custom evaluations written in Python or YAML alongside the prebuilt checks.

Adds gate mechanism in `jevals` that applies a policy to evaluation results and routes agent actions to proceed, block, or require human review — enabling the same eval definition to run in both development testing and production guardrails.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ changelog · Oct 3, 2026

CHANGELOG ↗ RANK

GitGuardian SaaS is a cloud-based security platform that detects and prevents secrets leakage in git repositories and source code.

> GitGuardian's biggest ship this window is Agents Analysis, which automatically verdicts Public Monitoring incidents for company relevance, alongside a native Wiz secret-findings integration, a new Coverage Hub for perimeter visibility, expanded AI Hooks protection covering five more AI coding tools, and a configurable API timeout for ggshield.

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

01 **Configurable ggshield API timeout** IMPROVED 90

Adds an `api_timeout` config key and `GITGUARDIAN_API_TIMEOUT` environment variable to override the previously fixed 60-second API timeout for large scans; the key can be persisted in `~/.gitguardian.yaml` (e.g. `api_timeout: 120`).

Persist a higher API timeout in ggshield config so every scan in the environment picks it up without an env-var prefix.

yaml

```yaml
api_timeout: 120
```

— Exact config key, env var, file path and example value given

changelog-20261003-95f65a08

02 **Agents Analysis for Public Monitoring incidents** NEW 83

Public Monitoring now runs Agents Analysis automatically for all customers, labeling every analyzed public incident with a company-relevance verdict (Related, Uncertain, or Unrelated). A new Public Monitoring agents page and an Analysis tab on each incident show the verdict, the reasoning, and an Investigation details section tracing how the incident moved through triage, with counts linking to matching incidents. Analysts can submit thumbs-up/thumbs-down feedback with optional comments on the Analysis tab. The incident list gains three new saved views — Company-related, Unclear if company-related, and Not company-related — alongside the existing All and Open, plus default Company-related and Risk score columns; risk scoring now factors in company-relevance, giving Unrelated incidents a score of zero and leaving unanalyzed incidents unscored.

When triaging a flood of public exposure alerts, open the `Company-related` saved view first to see only incidents Agents Analysis has confirmed as belonging to your org, then drill into the Analysis tab on any incident to review the reasoning and Investigation details before deciding on remediation.

> 📍 In Public Monitoring, select the `Company-related` saved view › open an incident › click the `Analysis` tab to review the verdict (Related / Uncertain / Unrelated), the reasoning, and the Investigation details section.

After reviewing an incident where Agents Analysis reached the wrong company-attribution conclusion, submit a thumbs-down with a comment so GitGuardian can improve attribution accuracy for your organization over time.

> 📍 In Public Monitoring, open an incident › click the `Analysis` tab › click the thumbs-down icon › optionally type a comment explaining the misattribution › submit.

— Full mechanism and UI steps given, but no API/endpoint for the verdicts

`launch-20261001-61baee62` `changelog-20261003-95f65a08`

---

03 **AI Hooks expansion to Kiro, Junie, Claude Code, Codex, Mistral Vibe**

NEW 75

Extends AI Hooks and `ggshield` 1.55.0+ to cover Amazon Kiro (IDE and CLI) and Junie CLI, blocking secrets in commands and file reads. Also adds tool-output withholding on Claude Code, Codex, and Mistral Vibe: `ggshield` intercepts tool output containing secrets and replaces it with a safe message so the secret never reaches the model or session transcript.

— Names exact tools and ggshield version, no command shown `changelog-20261003-95f65a08`

---

04 **Coverage Hub on the perimeter page** NEW 75

Introduces the Coverage Hub on the perimeter page: four cards show installed integrations, the share of sources monitored in real time, the share scanned successfully, and the share owned by a team, and each card's action filters the sources table down to that gap.

— Mechanism and UI location described, exact path not given `changelog-20261003-95f65a08`

---

05 **Team ownership and perimeter management in sources table**

75

BREAKING

Adds a `Team` column to the sources table and a `No team assigned` filter to surface unowned sources, and replaces the `Health` filter with `Has open incidents` (saved views using the old filter migrate automatically). Adds a per-account setting to keep team access to unmonitored sources and their existing incidents instead of removing them from team perimeters, and a per-account setting to automatically remove archived sources from team perimeters. Perimeter bulk actions now gray out inapplicable actions with a tooltip explanation and show a confirmation listing skipped sources when monitoring, unmonitoring, or scanning more than 30 sources.

— Names every filter and setting but no exact UI path or API `changelog-20261003-95f65a08`

06 **Wiz secret-findings integration** NEW 65

Adds a native Wiz integration that syncs active GitGuardian secret incidents to Wiz daily as secret findings, carrying repository, file path, commit, detector, validity, severity, and a SHA-256 hash of the secret — never the raw secret value.

— Clear data flow described but no setup steps or config surface

`changelog-20261003-95f65a08`

THINNER COVERAGE BELOW

07 **Secret value and validity API extended to Public Monitoring**

IMPROVED 55

Extends the `secret value` and `validity` API endpoints to cover secrets found through Public Monitoring, not only Internal Monitoring.

— Endpoint names given, no method, path, or parameters `changelog-20261003-95f65a08`

08 **Author attribution for BYOS custom sources** NEW 45

Adds optional author attribution for leaks detected through custom sources (BYOS): an author email can be provided and is returned when leaks are found.

— Describes the field but no endpoint or schema detail `changelog-20261003-95f65a08`

09 **System hooks with non-admin tokens reach more projects**

IMPROVED 45

System hooks integrations configured with a non-admin token now monitor projects in personal namespaces and in groups the service account reaches only through a project.

— Describes new scope but no setup or config steps `changelog-20261003-95f65a08`

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `Health` filter on the perimeter sources table is replaced by `Has open incidents`; saved views using the old filter are migrated automatically.

changelog · Oct 3, 2026 CHANGELOG ↗ STABLE

GitGuardian adds Wiz secret-findings integration, Coverage Hub, AI hook expansion to Kiro/Junie, and configurable API timeout.

└──▷ **USE IT**

Persist a higher API timeout in ggshield config so every scan in the environment picks it up without an env-var prefix.

~/.gitguardian.yaml

```
api_timeout: 120
```

Adds `api_timeout` config key and `GITGUARDIAN_API_TIMEOUT` environment variable to override the previously fixed 60-second API timeout for large scans.

Adds native Wiz integration that syncs active GitGuardian secret incidents to Wiz as secret findings daily, carrying repository, file path, commit, detector, validity, severity, and a SHA-256 hash of the secret — never the raw value.

Extends AI Hooks and `ggshield` 1.55.0+ to cover Amazon Kiro (IDE and CLI) and Junie CLI, blocking secrets in commands and file reads.

Adds tool-output withholding on Claude Code, Codex, and Mistral Vibe: `ggshield` intercepts tool output containing secrets and replaces it with a safe message so the secret never reaches the model or session transcript.

Introduces the Coverage Hub on the perimeter page: four cards showing installed integrations, share of sources monitored in real time, share scanned successfully, and share owned by a team — each card action filters the sources table to the gap.

+8 more

└──▷ **BREAKING ON UPGRADE**

! The Health filter on the perimeter sources table is replaced by Has open incidents; saved views using the old filter are migrated automatically.

announced · Oct 1, 2026 BLOG ↗

GitGuardian adds Agents Analysis to Public Monitoring: automatic company-relevance verdicts, an Analysis tab with reasoning, and analyst feedback controls.

└──▷ **HOW TO FIND IT**

When triaging a flood of public exposure alerts, open the Company-related saved view first to see only incidents Agents Analysis has confirmed as belonging to your org, then drill into the Analysis tab on any incident to review the reasoning and Investigation details before deciding on remediation.

📍 In Public Monitoring, select the `Company-related` saved view › open an incident › click the `Analysis` tab to review the verdict (Related / Uncertain / Unrelated), the reasoning, and the Investigation details section.

After reviewing an incident where Agents Analysis reached the wrong company-attribution conclusion, submit a thumbs-down with a comment so GitGuardian can improve attribution accuracy for your organization over time.

📍 In Public Monitoring, open an incident › click the `Analysis` tab › click the thumbs-down icon › optionally type a comment explaining the misattribution › submit.

Adds an `Analysis` tab on public incidents showing the company-relevance verdict (Related, Uncertain, or Unrelated), the reasoning behind it, and an Investigation details section tracing how the incident moved through triage.

Adds thumbs-up/thumbs-down feedback controls with optional comments on the Analysis tab, letting analysts signal when company attribution is correct or incorrect.

Introduces Agents Analysis, which automatically labels every analyzed public incident with a Company-related verdict and is enabled automatically for all customers using Public Monitoring.

Reorganizes the Public Monitoring incident list with three new saved views — `Company-related`, `Unclear if company-related`, and `Not company-related` — alongside the existing `All` and `Open` views.

Adds `Company-related` and `Risk score` columns to the Public Monitoring incident list by default.

`+1 more`

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

DataDome

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

DataDome protects web, mobile, and API applications by detecting and blocking bot and fraud traffic.

> DataDome launched AI Key Takeaways, bringing real-time AI-generated threat and business summaries with switchable security and executive views to Enterprise plans.

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

01 AI Key Takeaways on Homepage and Trend Reports NEW

65

DataDome adds AI Key Takeaways to the Homepage and Trend Reports, generating a real-time AI summary of threat data including sophisticated bots, AI agents, and threat traffic ratio, available to Enterprise plans. A security view is the default mode, surfacing detection precision, critical escalations, unusual behavioral patterns, attack pattern analysis, and configuration recommendations, while a business view toggle reframes the same data for executive audiences with protection status, threats blocked, false positive rates, and actionable guidance for internal reporting.

Switch from the security analyst view to the executive business view inside AI Key Takeaways to prepare a monthly risk report without rebuilding data from scratch.

> 📍 In the DataDome dashboard, go to Homepage or Trend Reports › AI Key Takeaways › toggle from Security view to Business view.

— Describes mechanism and views in detail but lacks config/API specifics.

launch-20261001-12f7c8a2

> announced · Oct 1, 2026 BLOG ↗
>
> DataDome adds AI Key Takeaways to Homepage & Trend Reports, delivering real-time security and business summaries for Enterprise plans.
>
> └──▷ **HOW TO FIND IT**
>
> Switch from the security analyst view to the executive business view inside AI Key Takeaways to prepare a monthly risk report without rebuilding data from scratch.
>
> > 📍 In the DataDome dashboard, go to Homepage or Trend Reports › AI Key Takeaways › toggle from Security view to Business view.

> Adds AI Key Takeaways to the DataDome Homepage & Trend Reports, providing a real-time AI-generated summary of threat data including sophisticated bots, AI agents, and threat traffic ratio.
>
> Introduces a security view as the default mode of AI Key Takeaways, surfacing detection precision, critical escalations, unusual behavioral patterns, attack pattern analysis, and configuration recommendations.
>
> Introduces a business view toggle within AI Key Takeaways that reframes the same underlying data for executive audiences, covering protection status, threats blocked, false positive rates, and actionable guidance for internal reporting.

WAS THIS USEFUL?  

OWASP ZAP

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

The ZAP by Checkmarx Core project

> ZAP's September 2026 update introduces a unified Global Scan Policy Manager and in-progress OAuth 2.0 authentication support, while retiring the AJAX Spider and DOM XSS Active Scan Rule add-ons from Docker in favor of the Client Spider.

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

01 AJAX Spider and DOM XSS add-ons retired from Docker

BREAKING

75

The AJAX Spider and DOM XSS Active Scan Rule add-ons are removed from nightly and weekly Docker releases; the Quickstart tab now defaults to the headless Firefox Client Spider crawler, and bundled Automation Framework profiles and the Chrome AF plan have been switched to the Client Spider. Workflows depending on either removed add-on must install them manually from the Marketplace or switch to the Client Spider.

— Names exact migration path and affected components `launch-20261001-2610a6b0`

02 OAuth 2.0 authentication in Authentication Helper

NEW

72

The Authentication Helper add-on gains in-progress OAuth 2.0 support as a new authentication method, supporting `client_credentials` and `password` (Resource Owner Password Credentials) grant types, with automatic refresh token handling, configurable client authentication, and autodetection of session management and verification.

— Names grant types and mechanisms but is explicitly in-progress `launch-20261001-2610a6b0`

03 Global Scan Policy dialog (dev mode)

NEW

65

Phase 1 of the Global Scan Policy Manager adds a Global Scan Policy dialog, accessible in dev mode, that unifies configuration of HTTP passive scan rules, active scan rules, and script-based active and passive scan rules, along with migration of existing per-context scan policies into this single interface.

— Names the dialog and what it unifies, but it's dev-mode only and Phase 1

`launch-20261001-2610a6b0`

THINNER COVERAGE BELOW

04 CSP intersection analysis in passive scan rules v76

IMPROVED

52

Passive scanner rules v76 updates the Content Security Policy scan rule to analyze all active CSP headers and META policies together using browser-style intersection.

— Names version and mechanism but no usage guidance launch-20261001-2610a6b0

└──▷ BREAKING ON UPGRADE

! The AJAX Spider and DOM XSS Active Scan Rule add-ons are removed from nightly and weekly Docker releases; workflows depending on either add-on must install them manually from the Marketplace or switch to the Client Spider.

announced · Oct 1, 2026 BLOG ↗

ZAP gains a unified Global Scan Policy Manager (dev mode) and in-progress OAuth 2.0 authentication support in the Authentication Helper add-on.

Adds a Global Scan Policy dialog (accessible in dev mode) as Phase 1 of the Global Scan Policy Manager, unifying configuration of HTTP passive scan rules, active scan rules, script-based active and passive scan rules, and policy migration from per-context scan policies into a single interface.

Adds in-progress OAuth 2.0 support as a new authentication method in the Authentication Helper add-on, supporting `client_credentials` and `password` (Resource Owner Password Credentials) grant types with automatic refresh token handling, configurable client authentication, and autodetection of session management and verification.

Retires the AJAX Spider and DOM XSS Active Scan Rule add-ons from nightly and weekly Docker releases; the Quickstart tab now defaults to the headless Firefox Client Spider crawler, and bundled Automation Framework profiles and the Chrome AF plan have been switched to the Client Spider.

Updates Passive scanner rules v76 so the Content Security Policy scan rule analyzes all active CSP headers and META policies together using browser-style intersection.

└──▷ BREAKING ON UPGRADE

! The AJAX Spider and DOM XSS Active Scan Rule add-ons are removed from nightly and weekly Docker releases; workflows depending on either add-on must install them manually from the Marketplace or switch to the Client Spider.

WAS THIS USEFUL? 👍 👎

Wallarm

Feedback ⓘ announced · Sep 29, 2026 BLOG ↗ RANK

Discover the latest features, improvements, and updates in Wallarm API Security

> Wallarm launched a dedicated Incidents interface that surfaces only confirmed attacks with grouped analytics and drill-down, adds payload-level flagging and linked Security Issue status on each row, and makes filtered views shareable and exportable for incident response.

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

01 **New Incidents page with grouped analytics** NEW 70

Adds a dedicated Incidents page in the dashboard showing only confirmed incidents — attacks that hit a real vulnerability and received a confirming response — with no install or migration required. It includes a time-series chart of requests with incidents alongside breakdowns of top source IPs, endpoints, hosts, attack types/subtypes, and application response codes over a rolling month, and groups incidents by attack type, HTTP method, domain, path, and the parameter the payload landed in, with drill-down to individual requests. Links from the Threat Prevention dashboard and Security Issues now navigate directly to this new page.

⋮

— Explains mechanism and scope of the new page, UI-only so no runnable command

launch-20260929-56b05403

02 **Shareable incident filter URLs and export** NEW 70

Encodes active filters and time range in the Incidents page URL so dashboards can be bookmarked, reloaded, or shared in tickets, with the same filtered view opening for recipients; also adds export of the current filtered incident view for use in incident response or reporting workflows.

Share the exact filtered incident view with a teammate or embed it in a ticket so they land on the same scope without manual filter reconstruction.

> 📍 In the Incidents page, apply your filters (attack type, time range, endpoint, etc.), then copy the page URL from the browser address bar and paste it into your ticket or chat.

— Includes concrete steps for sharing a filtered view via URL launch-20260929-56b05403

THINNER COVERAGE BELOW

03 **Payload-level flagging within multi-payload requests** NEW

45

Marks the specific payload within a multi-payload request that triggered the incident, so analysts do not have to identify which payload caused the damage.

— Describes behavior but no UI path or example given launch-20260929-56b05403

04 **Security Issues status on incident rows** NEW 45

Shows severity and open/closed state of linked Security Issues in a Security Issues column on every incident row without requiring a click-through.

— Clear UI behavior described, no navigation steps given launch-20260929-56b05403

announced · Sep 29, 2026 BLOG ↗

Wallarm adds a dedicated Incidents interface with grouped investigation, payload flagging, shareable filter URLs, and export.

└──▷ HOW TO FIND IT

Share the exact filtered incident view with a teammate or embed it in a ticket so they land on the same scope without manual filter reconstruction.

📍 In the Incidents page, apply your filters (attack type, time range, endpoint, etc.), then copy the page URL from the browser address bar and paste it into your ticket or chat.

Adds a dedicated Incidents page in the dashboard showing only confirmed incidents — attacks that hit a real vulnerability and received a confirming response — with no install or migration required.

Adds a time-series chart of requests with incidents alongside breakdowns of top source IPs, endpoints, hosts, attack types/subtypes, and application response codes over a rolling month.

Groups incidents by attack type, HTTP method, domain, path, and the parameter the payload landed in, with drill-down to individual requests.

Marks the specific payload within a multi-payload request that triggered the incident, so analysts do not have to identify which payload caused the damage.

Shows severity and open/closed state of linked Security Issues in a Security Issues column on every incident row without requiring a click-through.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ

Challenge suspicious clients with DDoS Protection rules (2026-10-01) BLOG ↗

RANK

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

> Fastly added a non-blocking `client_challenge` action to DDoS Protection rules and launched AEDA, an autonomous AI agent in Gemini Enterprise that runs root-cause analysis and recommends edge security remediations.

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

01 Client challenge action for DDoS Protection rules NEW

80

DDoS Protection rules gain a `client_challenge` action, letting suspicious clients prove legitimacy instead of being hard-blocked during an attack. The new action is also returned alongside `default`, `block`, `log`, and `off` in list/retrieve rule responses, and can be set via `PATCH /ddos-protection/rules/{rule_id}` with `{"action": "client_challenge"}`.

During an active DDoS attack, challenge suspicious clients rather than hard-blocking them so legitimate users can still reach your service.

```
$ curl -X PATCH 'https://api.fastly.com/ddos-
protection/rules/{rule_id}' \
  -H 'Fastly-Key: <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"action": "client_challenge"}'
```

— Names exact endpoint, method, and all action values.

Challenge suspicious clients with DDoS …

02 AEDA autonomous edge defense agent in Gemini Enterprise NEW

75

AEDA (Autonomous Edge Defense Agent) runs inside Gemini Enterprise, pairing Fastly's live edge telemetry with Gemini's stateful reasoning to conduct machine-speed Root Cause Analysis (RCA) loops when an anomaly is detected. It cross-references organisation telemetry against anonymised global traffic to distinguish isolated infrastructure bugs from coordinated sector-wide attacks, generates natural-language

VCL configurations, recommends concrete remediations (rate-limiting rules, WAF signature updates, origin shielding changes) with supporting evidence and expected impact, produces evidence-backed audit trails for compliance and board reporting, and supports the Model Context Protocol (MCP) to embed its investigations into existing engineering and DevOps workflows.

Trigger AEDA to investigate a spiking 5xx error rate and get an evidence-backed RCA distinguishing a botnet from an origin issue — without manually cross-referencing dashboards.

> 📍 Contact AEDA@fastly.com or your Fastly account team to join the early access preview, then in Gemini Enterprise trigger an AEDA Root Cause Analysis (RCA) loop when an anomaly is detected.

— Rich mechanism detail but only early-access contact, no runnable surface.

launch-20260929-fa280376

Challenge suspicious clients with DDoS Protection rules (2026-10-01) BLOG ↗

LATEST STABLE

Fastly DDoS Protection rules gain a `client_challenge` action to challenge suspicious clients instead of blocking them.

└──▷ **TRY IT**

During an active DDoS attack, challenge suspicious clients rather than hard-blocking them so legitimate users can still reach your service.

```
$ curl -X PATCH 'https://api.fastly.com/ddos-
protection/rules/{rule_id}' \
  -H 'Fastly-Key: <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"action": "client_challenge"}'
```

announced · Sep 29, 2026 BLOG ↗

Fastly launches AEDA, an AI agent in Gemini Enterprise that runs autonomous root-cause analysis and recommends WAF/rate-limit remediations at edge speed.

└──▷ **HOW TO FIND IT**

Trigger AEDA to investigate a spiking 5xx error rate and get an evidence-backed RCA distinguishing a botnet from an origin issue — without manually cross-referencing dashboards.

> 📍 Contact AEDA@fastly.com or your Fastly account team to join the early access preview, then in Gemini Enterprise trigger an AEDA Root Cause Analysis (RCA) loop

when an anomaly is detected.

Adds support for the Model Context Protocol (MCP) to embed AEDA autonomous investigations into existing engineering and DevOps workflows.

Introduces AEDA (Autonomous Edge Defense Agent) inside Gemini Enterprise, pairing Fastly live edge telemetry with Gemini stateful reasoning to conduct machine-speed Root Cause Analysis (RCA) loops when an anomaly is detected.

AEDA cross-references organisation telemetry against anonymised global traffic to distinguish isolated infrastructure bugs from coordinated attacks targeting multiple organisations in the same sector.

AEDA generates natural-language VCL configurations and recommends concrete remediations — rate-limiting rules, WAF signature updates, origin shielding changes — each with supporting evidence, expected impact, and step-by-step guidance.

Produces evidence-backed audit trails intended to simplify compliance reporting and board-level risk communication for CISOs and security leaders.

WAS THIS USEFUL? 👍 👎

◆ OT / ICS Security

CrowdStrike Falcon for XIoT

Feedback ⓘ

changelog · Oct 3, 2026 API ↗ RANK

CrowdStrike Falcon for XIoT is a security platform that detects and responds to threats on extended IoT devices and operational technology networks.

> CrowdStrike Falcon for XIoT adds new network scanning capabilities for discovering and assessing XIoT devices, alongside containment allowlist controls, all exposed via the Falcon API.

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

01 **Network scanning and zone management** NEW 42

Falcon for XIoT adds network scanning to discover and assess XIoT devices across network zones, zone management to organize and segment scanned network targets, scan detection review to surface findings from network scans, and scan result reporting for network assessments.

— Describes capabilities but no endpoints, flags, or UI paths given 2026-10-03-e9a59adb

02 **Network containment allowlist management** NEW 23

Adds network containment allowlist management to control which devices are exempt from containment actions.

— Bare description with no mechanism or exact API surface named 2026-10-03-e9a59adb

changelog · Oct 3, 2026 API ↗ LATEST STABLE

Falcon for XIoT gains network scanning, zone management, scan detections, and network containment allowlists via API.

Adds network scanning capability to discover and assess XIoT devices across network zones.

Adds network zone management to organize and segment scanned network targets.

Adds scan detection review to surface findings from network scans.

Adds scan result reporting for network assessments.

Adds network containment allowlist management to control which devices are exempt from containment actions.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Weaveworks eksctl

Feedback ⓘ v0.231.0 NOTES ↗ RANK

The official CLI for Amazon EKS

> eksctl v0.231.0 adds fine-grained nodegroup connection tracking, control-plane pod garbage collection tuning, and support for placing the EKS control plane on private subnets.

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

01 **Private-subnet EKS control plane placement** NEW 70

Adds `vpc.controlPlaneOnPrivateSubnets` to cluster config to place the EKS control plane on private subnets, avoiding exposure to public network traffic.

Place the EKS control plane on private subnets to avoid exposing it to public network traffic.

yaml

```yaml
vpc:
  controlPlaneOnPrivateSubnets: true
```

— Names exact config key with example, but no deeper mechanism explained v0.231.0

THINNER COVERAGE BELOW

02 **Connection tracking config for nodegroups** NEW 40

Adds `connectionTracking` to nodegroup config to control connection tracking settings per nodegroup.

— Names config key but no values or usage shown v0.231.0

03 **Terminated pod GC threshold for control plane** NEW 40

Adds `TerminatedPodGcThreshold` to control plane component config to tune garbage collection of terminated pods.

— Names config key but lacks default value or usage example v0.231.0

v0.231.0 NOTES ↗ LATEST STABLE

> eksctl v0.231.0 adds connectionTracking for nodegroups, TerminatedPodGcThreshold for control plane, and private-subnet control plane support.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.231.0 https://github.com/eksctl-
io/eksctl.git
# already have the repo? check out this version:
$ git checkout v0.231.0
```

└──▷ **USE IT**

Place the EKS control plane on private subnets to avoid exposing it to public network traffic.

cluster.yaml

```
vpc:
  controlPlaneOnPrivateSubnets: true
```

Adds `connectionTracking` to nodegroup config to control connection tracking settings per nodegroup.

Adds `TerminatedPodGcThreshold` to control plane component config to tune garbage collection of terminated pods.

Adds `vpc.controlPlaneOnPrivateSubnets` to cluster config to place the EKS control plane on private subnets.

WAS THIS USEFUL?  

Wiz

Feedback ⓘ changelog · Sep 30, 2026 BLOG ↗ RANK

Wiz is a cloud security platform that identifies and remediates risks across cloud environments using AI-driven analysis.

> Wiz launched WizOS Helm Charts — hardened, signed Kubernetes charts built in Wiz's own pipeline with provenance and SBOMs — and populated the catalog with a dozen charts spanning delivery, observability, compliance, and automation.

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

01 WizOS Helm Charts launch with signing and SBOMs NEW

74

Wiz introduced WizOS Helm Charts — hardened, minimal Helm charts built, tested, and signed by Wiz in its own pipeline, isolating deployments from upstream supply chain risks. Every WizOS image is signed and shipped with provenance and an SBOM, enabling verification of origin and contents at install time, and the full catalog is browsable in the Secured Image Catalog in the Wiz portal.

Browse available WizOS Helm Charts and verify a chart's provenance and SBOM before installing into a cluster.

> 📍 In the Wiz portal, navigate to Secured Image Catalog, locate the desired WizOS Helm chart, and review its attached SBOM and provenance attestation before running `helm install`.

— Explains mechanism and verification flow with a concrete install-time step

2026-09-30-b842af3e

02 WizOS catalog expansion across delivery, observability, compliance and automation

NEW

63

The WizOS catalog gained Argo Rollouts (with a FIPS variant), Sealed Secrets, and KEDA for delivery and GitOps; Grafana Loki and Alloy (with anonymous usage reporting off by default) and Redis Exporter for observability; Cosign (including keyless signing) and OpenSCAP (bundled with the DISA general purpose operating system baseline) for signing and compliance; Valkey Bundle (with search, JSON, Bloom, and LDAP modules) and Redis Operator for data workloads; and OpenSSH Server and Ansible (with major cloud collections included) for access and automation.

— Names every added chart but no per-chart usage detail 2026-09-30-b842af3e

changelog · Sep 30, 2026 BLOG ↗

Wiz launches WizOS Helm Charts: hardened, signed, SBOM-shipped Kubernetes charts built and patched in Wiz's own pipeline.

└──▷ **HOW TO FIND IT**

Browse available WizOS Helm Charts and verify a chart's provenance and SBOM before installing into a cluster.

> 📍 In the Wiz portal, navigate to Secured Image Catalog, locate the desired WizOS Helm chart, and review its attached SBOM and provenance attestation before running `helm install`.

Introduces WizOS Helm Charts — hardened, minimal Helm charts built, tested, and signed by Wiz in its own pipeline, isolating deployments from upstream supply chain risks.

Signs every WizOS image and ships it with provenance and an SBOM, enabling verification of origin and contents at install time.

Adds Argo Rollouts (with a FIPS variant), Sealed Secrets, and KEDA to the WizOS catalog for delivery and GitOps use cases.

Adds Grafana Loki and Alloy (with anonymous usage reporting off by default) and Redis Exporter to the WizOS catalog for observability.

Adds Cosign (including keyless signing) and OpenSCAP (bundled with the DISA general purpose operating system baseline) to the WizOS catalog for signing and compliance.

+3 more

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - October 02, 2026

BLOG ↗ RANK

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

> Sysdig Secure introduced a Vulnerability Catalog that lets teams look up any tracked CVE — even without findings in their environment — through both the UI and a new API.

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

01 Vulnerability Catalog with CVE lookup and API NEW

75

Adds a Vulnerability Catalog to look up any tracked CVE — even with no findings in your environment — via UI and API. The CVE Catalog API lets you programmatically retrieve CVE metadata and environment-impact data, available on SaaS now and on-premises from version 7.10. Selecting a CVE in the catalog opens the CVE Details drawer, showing CVE metadata and whether the CVE affects your environment; the catalog itself is searched by CVE ID with filters for severity, fix, and exploit status.

Look up a CVE your team is investigating to see its severity, fix availability, exploit status, and whether it affects your environment — even if no findings exist yet.

> 📍 In Sysdig Secure, go to Policies › Vulnerability Catalog, enter a CVE ID in the search bar, and apply filters for severity, fix, or exploit. Select the CVE to open the CVE Details drawer.

— Names API, on-prem version, and exact UI navigation path Secure SaaS - October 02, 2026

Secure SaaS - October 02, 2026 BLOG ↗ LATEST STABLE

Sysdig adds a Vulnerability Catalog to look up any tracked CVE — even with no findings in your environment — via UI and API.

└──▷ **HOW TO FIND IT**

Look up a CVE your team is investigating to see its severity, fix availability, exploit status, and whether it affects your environment — even if no findings exist yet.

> 📍 In Sysdig Secure, go to Policies › Vulnerability Catalog, enter a CVE ID in the search bar, and apply filters for severity, fix, or exploit. Select the CVE to open the CVE Details drawer.

> Adds the CVE Catalog API to programmatically retrieve CVE metadata and environment-impact data, available on SaaS now and on-premises from version 7.10.
>
> Adds the CVE Details drawer, opened by selecting a CVE in the Vulnerability Catalog, showing CVE metadata and whether the CVE affects your environment.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Docker Compose

Feedback ⓘ v5.6.0 NOTES ↗ RANK

Define and run multi-container applications with Docker

> Docker Compose v5.6.0 adds partial support for manually-triggered jobs, expands provider-service networking with new relay capabilities, and makes the `docker compose port` command more forgiving.

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

01 Provider-service networking and relay support NEW

75

Compose adds a `get-service-config` provider request so providers can retrieve their resolved service configuration, a `publish-endpoint` capability that deploys a network relay for provider services, and `get-relay-info` so providers know where to bind for the relay. A dedicated network is created for the provider relay link, isolating relay traffic from application networks.

— Names four provider-relay surfaces but no usage walkthrough v5.6.0

02 Optional PRIVATE_PORT for docker compose port IMPROVED

65

Makes `PRIVATE_PORT` optional for `docker compose port`, so the command works without specifying an internal port.

— Exact command and flag named, directly runnable v5.6.0

THINNER COVERAGE BELOW

03 Partial jobs support via compose-go NEW

50

Adds partial support for jobs using compose-go jobs and container-spec layering; manually-triggered jobs are supported now, while scheduled jobs are not yet available.

— States scope and limitation but no command example v5.6.0

04 Warnings for unsupported compose-file attributes IMPROVED

35

Warns on unsupported Compose-file attributes so unknown keys are surfaced rather than silently ignored.

— Describes behaviour but no example output or key names v5.6.0

05 Provider services in image phase NEW

30

Includes provider services in the image phase, with image distribution on demand.

— Brief mention, no mechanism or config detail v5.6.0

v5.6.0 NOTES ↗ LATEST STABLE

Docker Compose v5.6.0 adds partial jobs support, richer provider-service networking, and a more forgiving `docker compose port` command.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v5.6.0 https://github.com/docker/compose.git
# already have the repo? check out this version:
$ git checkout v5.6.0
```

Makes `PRIVATE_PORT` optional for `docker compose port`, so the command works without specifying an internal port.

Adds partial support for jobs via compose-go jobs and container-spec layering; manually-triggered jobs are supported now, scheduled jobs are not yet available.

Adds `get-service-config` provider request so providers can retrieve their resolved service configuration.

Adds `publish-endpoint` capability that deploys a network relay for provider services.

Adds `get-relay-info` so providers know where to bind for the relay.

+3 more

WAS THIS USEFUL? 👍 👎

Anchore Grype

Feedback ⓘ v0.120.0 NOTES ↗ CODE ↗ RANK

A vulnerability scanner for container images and filesystems

> Grype's latest release adds NVD SSVC decision-point data for vulnerability prioritization and a Locations column in HTML reports, alongside documentation expanding coverage to Dart, C++, and Portage ecosystems.

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

01 **NVD SSVC decision points for prioritization** NEW 68

Grype now stores NVD SSVC decision points in a new v6 decorator table, exposing `GetSsvcs` for richer vulnerability prioritization data.

– Names table version and function but not usage mechanics v0.120.0

THINNER COVERAGE BELOW

02 **Locations column in HTML vulnerability reports** IMPROVED 40

Adds a Locations column to the HTML report template to aid vulnerability triage.

– Describes the change but not how to generate the report v0.120.0

03 **Ecosystem support docs for Dart, C++, and Portage** NEW 24

New documentation pages describe Grype's vulnerability scanning capabilities for Dart, C++, and Portage ecosystems, though the pages themselves contain no further detail in this window.

– Doc pages are empty of substantive content product docs

04 **SLES comprehensive distro coverage** IMPROVED 22

Grype v0.120.0 adds comprehensive distro coverage for SLES, per the release summary; no further detail on scope or packages was provided.

– Only mentioned in summary, no mechanism given v0.120.0

v0.120.0 NOTES ↗ CODE ↗ LATEST STABLE

Grype v0.120.0 adds NVD SSVC decision points, a Locations column in HTML reports, and SLES comprehensive distro coverage.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.120.0 https://github.com/anchore/grype.git
# already have the repo? check out this version:
$ git checkout v0.120.0
```

Stores NVD SSVC decision points in a new v6 decorator table, exposing `GetSsvcs` for richer vulnerability prioritization data.

Adds a Locations column to the HTML report template to aid vulnerability triage.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

CrowdStrike

Feedback ⓘ changelog · Oct 3, 2026 API ↗ RANK

CrowdStrike provides cloud-based endpoint security and threat intelligence services to detect, prevent, and respond to cyberattacks.

> CrowdStrike expanded its API surface with a new Seraphic enterprise browser security integration, a Scanning Orchestrator surface, Application Abuse Exclusions, and broader Audit and Prevention Policy coverage.

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

01 **Seraphic enterprise browser security integration** NEW — 45

Adds Seraphic integration with 40 API operations for managing enterprise browser security across a fleet.

— Names op count but no endpoints or examples 2026-10-03-22053f66

02 **Application Abuse Exclusions API** NEW — 40

Adds Application Abuse Exclusions API operations for defining exclusions against application abuse detections.

— Describes purpose but no endpoint names or count 2026-10-03-22053f66

03 **Audit API expanded to 86 operations** IMPROVED — 30

Expands the Audit API surface to 86 operations.

— Gives a concrete number but no detail on new ops 2026-10-03-22053f66

04 **Scanning Orchestrator API surface** NEW — 25

Adds a Scanning Orchestrator API surface for coordinating scan workflows.

— Bare name, no endpoints or mechanism given 2026-10-03-22053f66

05 **Prevention Policy API operations added** NEW — 25

Adds Prevention Policy API operations to the platform.

— No endpoint names, counts, or examples given 2026-10-03-22053f66

changelog · Oct 3, 2026 API ↗ LATEST STABLE

CrowdStrike API gains Seraphic (40 ops), Scanning Orchestrator, and expanded Audit/Prevention Policy surfaces

> Adds Seraphic integration with 40 API operations for managing enterprise browser security across a fleet.
>
> Adds Scanning Orchestrator API surface for coordinating scan workflows.
>
> Adds Application Abuse Exclusions API operations for defining exclusions against application abuse detections.
>
> Expands Audit API surface to 86 operations.
>
> Adds Prevention Policy API operations to the platform.

WAS THIS USEFUL? 👍 👎

Automox

Feedback ⓘ

Set Remote Control consent once for your whole organization (2026-10-02) BLOG ↗

RANK

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

> Automox centralized Remote Control consent management at the organization level, adding per-device overrides, a daily enforcement option, and new device table visibility into consent state.

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

01 **Org-level Remote Control consent management** NEW 75

Adds an organization-level default consent setting for Remote Control sessions, configured at `Settings > Remote Tools`, which is inherited automatically by all devices including newly added ones. Per-device consent overrides can be set individually or in bulk from the device table and persist independently of organization default changes, while an `Enforce Organization Default Daily` setting resets every device to the org default once per day and clears all per-device overrides for fleet-wide consistency. A new `Consent` column in the device table shows each device's current consent setting and whether it is inherited or an override, with filtering on both states.

Enable daily enforcement so Remote Control consent stays locked to the org default across the fleet without anyone auditing overrides by hand.

> 📍 In the Automox console, go to Settings › Remote Tools, set your organization default consent behavior, then enable `Enforce Organization Default Daily`.

Identify devices whose consent differs from the org default and decide in bulk whether to clear their overrides or keep them as exceptions.

> 📍 In the Automox console, go to the device table, open the `Consent` column, filter by `override`, select the devices you want to realign, and apply the organization default from the bulk action menu.

— Names exact settings and UI path but no API or config file.

Set Remote Control consent once for you…

> Set Remote Control consent once for your whole organization (2026-10-02)
>
> BLOG ↗ LATEST STABLE

Automox Remote Control consent is now managed at the org level, with per-device overrides, a daily enforcement option, and a new Consent column in the device table.

└──▷ HOW TO FIND IT

Enable daily enforcement so Remote Control consent stays locked to the org default across the fleet without anyone auditing overrides by hand.

> 📍 In the Automox console, go to Settings › Remote Tools, set your organization default consent behavior, then enable `Enforce Organization Default Daily`.

Identify devices whose consent differs from the org default and decide in bulk whether to clear their overrides or keep them as exceptions.

> 📍 In the Automox console, go to the device table, open the `Consent` column, filter by `override`, select the devices you want to realign, and apply the organization default from the bulk action menu.

Adds an organization-level default consent setting for Remote Control sessions, inherited automatically by all devices including newly added ones, found at Settings > Remote Tools.

Adds per-device consent overrides — individually or in bulk from the device table — that persist independently of organization default changes.

Adds an `Enforce Organization Default Daily` setting that resets every device to the organization default once per day and clears all per-device overrides, for fleet-wide consent consistency without manual checks.

Adds a `Consent` column to the device table showing each device's current consent setting and whether it is inherited or an override, with filtering on both states.

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ What's New in Iru | October 2nd NOTES ↗ RANK

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

> Iru rolled out a Privacy Library Item for polishing app permission prompts, AI-assisted Windows Update configuration using common security frameworks, and a new all-in-one Iru for Startup bundle.

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

01 **Privacy Library Item for app permissions** NEW 30

Adds a Privacy Library Item that lets admins polish the app permission prompts shown to end users.

— Names the feature but not how prompts are edited What's New in Iru | October 2nd

02 **Security frameworks for AI Windows Updates** NEW 25

Enables defining common security frameworks when using Iru AI to configure Windows Updates.

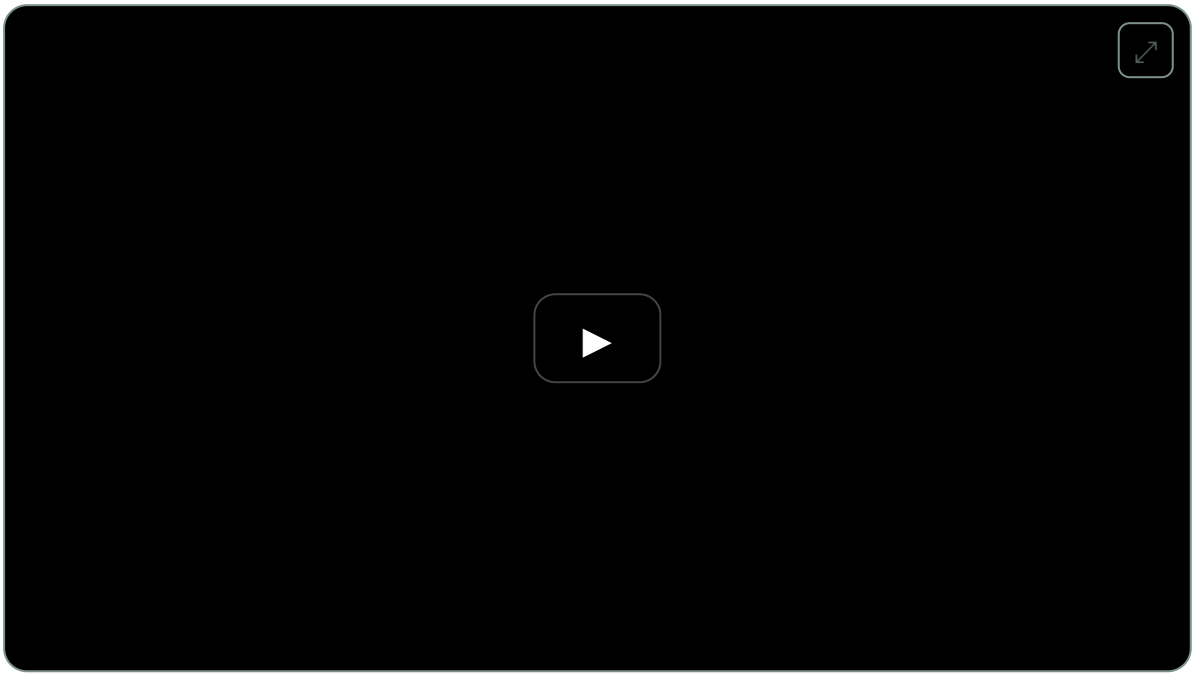
— No specific framework names or config steps given What's New in Iru | October 2nd

03 **Iru for Startup bundle** NEW 10

Introduces the Iru for Startup bundle as an all-in-one platform offering.

— Only a name, no scope or pricing details What's New in Iru | October 2nd

└──▷ **ALSO FROM THESE RELEASES**



What's New in Iru | October 2nd NOTES ↗ LATEST STABLE

Iru adds Privacy Library Item for app permissions, AI-driven Windows Update frameworks, and a new Startup bundle.

Adds Privacy Library Item to let admins polish app permission prompts shown to end users.

Enables defining common security frameworks when using Iru AI to configure Windows Updates.

Introduces the Iru for Startup bundle as an all-in-one platform offering.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Yeti

Feedback ⓘ 2.12.0 NOTES ↗ RANK

Your Everyday Threat Intelligence

> Yeti 2.12.0 adds a Google OAuth access-token exchange endpoint for API authentication, publishes ACL-change events on the event bus (a breaking change for event consumers), makes agent personas editable, redacts secrets from audit logs, and tightens the yara-python build requirement.

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

01 **Google OAuth access-token exchange endpoint** NEW 90

Yeti adds `POST /api/v2/auth/google-access-token` to exchange a Google OAuth access token for a Yeti API token, useful in headless or service-account contexts where an ID token is unavailable. The route is enabled only when `auth.google_access_token_client_ids` is set in `yeti.conf`.

Obtain a Yeti API token from a Google OAuth access token in a headless/service-account context where an ID token is unavailable.

```
$ curl -s -X POST https://<yeti-host>/api/v2/auth/google-access-token \
  -H 'Content-Type: application/json' \
  -d '{"access_token": "<google-access-token>"}'
```

Enable the Google access-token exchange route by listing the permitted OAuth client IDs in `yeti.conf`.

```ini
[auth]
google_access_token_client_ids = <oauth-client-id-1>,<oauth-client-id-2>
```

— Endpoint, config key, and runnable curl example all given. 2.12.0

02 **ACL changes now publish events on event bus** BREAKING 78

Object creation now emits three events on the event bus — the object itself, its owner ACL, and one per group in `rbac.default_acls` — and sharing/unsharing operations also publish, delivering access-control metadata to registered event tasks for the first

time. This is a breaking change: any event-bus consumer with a match-all `acts_on`, including the shipped `DatadogMetrics` plugin, will start receiving the new ACL events and report a new `yeti.object` series with `type:acl`.

— Mechanism and breaking impact named, no direct migration command. 2.12.0

THINNER COVERAGE BELOW

03 **yara-python build requirement raised** BREAKING 50

`yara-python` now requires >= 4.5.2; the image build fails if yara is compiled without OpenSSL support, rather than shipping a crypto-less binary.

— Named version and build requirement but no migration steps. 2.12.0

04 **Editable agent persona objects** NEW 30

Conversational agent instructions are now stored as editable persona objects, with persona writes publishing events to the event bus.

— Describes the change but no concrete surface or usage path. 2.12.0

05 **Audit log secret redaction** IMPROVED 18

Audit logs now redact tokens and URL-encoded secrets.

— Bare statement with no mechanism or scope. 2.12.0

└──▷ BREAKING ON UPGRADE

! Object creation now emits three events where it previously emitted one (the object, its owner ACL, and one per group in `rbac.default_acls`); any event-bus consumer with a match-all `acts_on` — including the shipped `DatadogMetrics` plugin — will start receiving the new ACL events and will report a new `yeti.object` series with `type:acl`.

! `yara-python` now requires >= 4.5.2; the image build fails if yara is compiled without OpenSSL support, rather than shipping a crypto-less binary.

2.12.0 NOTES ↗ LATEST STABLE

Yeti 2.12.0 adds a Google OAuth access-token exchange endpoint, ACL-change event publishing, and editable agent persona objects.

└──▷ GET THIS VERSION

```
$ git clone --branch 2.12.0 https://github.com/yeti-
platform/yeti.git
# already have the repo? check out this version:
$ git checkout 2.12.0
```

└──▷ TRY IT

Obtain a Yeti API token from a Google OAuth access token in a headless/service-account context where an ID token is unavailable.

```
$ curl -s -X POST https://<yeti-host>/api/v2/auth/google-access-token \
  -H 'Content-Type: application/json' \
  -d '{"access_token": "<google-access-token>"}'
```

Enable the Google access-token exchange route by listing the permitted OAuth client IDs in `yeti.conf`.

yeti.conf

```
[auth]
google_access_token_client_ids = <oauth-client-id-1>,<oauth-client-id-2>
```

Adds `POST /api/v2/auth/google-access-token` to exchange a Google OAuth access token for a Yeti API token, enabled only when `auth.google_access_token_client_ids` is set in `yeti.conf`.

ACL changes now publish events on the event bus — object creation emits three events (object, owner ACL, one per group in `rbac.default_acls`), and sharing/unsharing operations also publish, delivering access-control metadata to registered event tasks for the first time.

Conversational agent instructions are stored as editable persona objects, with persona writes publishing events to the event bus.

Audit logs now redact tokens and URL-encoded secrets.

└──▷ **BREAKING ON UPGRADE**

- ! Object creation now emits three events where it previously emitted one (the object, its owner ACL, and one per group in `rbac.default_acls`); any event-bus consumer with a match-all `acts_on` — including the shipped `DatadogMetrics` plugin — will start receiving the new ACL events and will report a new `yeti.object` series with `type:acl`.
- ! `yara-python` now requires >= 4.5.2; the image build fails if yara is compiled without OpenSSL support, rather than shipping a crypto-less binary.

WAS THIS USEFUL? 👍 👎

Meta ThreatExchange

Feedback ⓘ hma-v1.2.0 NOTES ↗

RANK

Trust & Safety tools for working together to fight digital harms.

> HMA v1.2.0 scopes exchange credentials per individual exchange and bumps the hardcoded Graph API version to v26.0.

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

01 **Per-exchange scoped credentials** IMPROVED 45

Exchange credentials are now scoped to each individual exchange rather than sharing a single credential set across all exchanges.

— Describes behaviour change but no config detail hma-v1.2.0

02 **Graph API version bump to v26.0** IMPROVED 35

Bumps the hardcoded Graph API version used by HMA to `v26.0`.

— Names exact version but no further detail hma-v1.2.0

hma-v1.2.0 NOTES ↗ LATEST STABLE

HMA v1.2.0 scopes exchange credentials per exchange and bumps Graph API to v26.0

└──▷ **GET THIS VERSION**

```
$ git clone --branch hma-v1.2.0
https://github.com/facebook/ThreatExchange.git
# already have the repo? check out this version:
$ git checkout hma-v1.2.0
```

Scopes exchange credentials to each individual exchange rather than sharing a single credential set across all exchanges.

Bumps hardcoded Graph API version to v26.0.

WAS THIS USEFUL? 👍 👎

Watcher

Feedback ⓘ v3.6.1 NOTES ↗ RANK

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

> Watcher v3.6.1 adds a site-scoped alert filter to the Website Monitoring API, pins the CertStream image to a known-good version, fixes a SearxNG deployment path, and hardens CertStream certificate matching with a new watchdog.

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

01 **CertStream image version pinning** BREAKING 85

Adds a `CERTSTREAM_VERSION` variable to Docker Compose and the Helm chart to pin the `certstream-server-go` image, currently set to `v1.8.2` (`0rickyy0/certstream-server-go:v1.8.2`). This is a breaking change: Compose files overriding the `certstream` image with `latest` will instead pull v1.9/v1.10, which cannot stream behind an HTTP proxy, and overrides to the `watcher` environment must explicitly add `certstream` to `no_proxy` or its WebSocket traffic will be misrouted through the corporate proxy.

Pin the CertStream image to v1.8.2 in a proxy-aware deployment so upgrades to the `latest` tag do not silently break certificate transparency monitoring.

ini

```
CERTSTREAM_VERSION=v1.8.2
SEARX_PATH=../Searx/searx
```

— Names env var, image tag, and exact migration pitfalls v3.6.1

02 **Site-scoped alert filtering API** NEW 80

Adds a `?site=<id>` query parameter to the `/api/site_monitoring/alert/` endpoint so alerts can be loaded scoped to a specific monitored site, instead of returning alerts across all sites.

Retrieve all alerts for a specific monitored site by ID, useful for dashboards or incident triage without wading through alerts from other sites.

```
$ curl -s 'https://<watcher-host>/api/site_monitoring/alert/?
site=42' -H 'Authorization: Token <your-api-token>'
```

— Named endpoint and query param with a runnable curl example v3.6.1

03 **SearxNG deployment path fix** BREAKING 80

Sets `SEARX_PATH=../Searx/searx` in `deployment/.env` so SearxNG correctly points at an existing `settings.yml` instead of generating a proxy-incompatible one. This changes the default from `../Searx` to `../Searx/searx` in `deployment/.env.example`; existing deployments must update `deployment/.env`, remove any auto-generated `Searx/settings.yml`, and recreate the `searxng` container.

Pin the CertStream image to v1.8.2 in a proxy-aware deployment so upgrades to the `latest` tag do not silently break certificate transparency monitoring.

```ini
CERTSTREAM_VERSION=v1.8.2
SEARX_PATH=../Searx/searx
```

— Exact config key, file path, and migration steps given v3.6.1

THINNER COVERAGE BELOW

04 **CertStream connection watchdog** NEW 45

Adds a watchdog for the CertStream connection that resets silent connections and logs a `CertStream stats` line every 5 minutes.

— Names the log line but no reader-facing action v3.6.1

05 **Case-insensitive CertStream keyword matching** IMPROVED 40

CertStream keyword matching is now case-insensitive, and certificates without a subject CN are matched against their SAN list instead of being skipped.

— Describes behavior change but gives no config or command v3.6.1

└──▷ **BREAKING ON UPGRADE**

! The `SEARX_PATH` default value has changed from `../Searx` to `../Searx/searx` in `deployment/.env.example`; deployments using the old value must update `deployment/.env` and remove any auto-generated `Searx/settings.yml`, then recreate the `searxng` container.

! The `certstream` image is now pinned to `0rickyy0/certstream-server-go:v1.8.2` via the new `CERTSTREAM_VERSION` variable; Compose files that override the `certstream` image with `latest` will pull v1.9/v1.10, which cannot stream behind an HTTP proxy.

! Compose files that override the `watcher` environment must explicitly add `certstream` to `no_proxy`, or the CertStream WebSocket will be incorrectly routed through the corporate proxy.

v3.6.1 NOTES ↗ LATEST STABLE

Watcher v3.6.1 adds a `?site=<id>` alert filter API and `/16` network memory for smarter Website Monitoring alerting.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v3.6.1 https://github.com/thalesgroup-cert/Watcher.git
# already have the repo? check out this version:
$ git checkout v3.6.1
```

└──▷ **TRY IT**

Retrieve all alerts for a specific monitored site by ID, useful for dashboards or incident triage without wading through alerts from other sites.

```
$ curl -s 'https://<watcher-host>/api/site_monitoring/alert/?site=42' -H 'Authorization: Token <your-api-token>'
```

Pin the CertStream image to v1.8.2 in a proxy-aware deployment so upgrades to the `latest` tag do not silently break certificate transparency monitoring.

deployment/.env

```
CERTSTREAM_VERSION=v1.8.2
SEARX_PATH=../Searx/searx
```

Adds `?site=<id>` query parameter to the `/api/site_monitoring/alert/` endpoint to load alerts scoped to a specific monitored site.

Adds `CERTSTREAM_VERSION` variable to Docker Compose and Helm chart to pin the `certstream-server-go` image version (currently `v1.8.2`).

Sets `SEARX_PATH=../Searx/searx` in `deployment/.env` to correctly point SearxNG at an existing `settings.yml` instead of generating a proxy-incompatible one.

CertStream keyword matching is now case-insensitive, and certificates without a subject CN are matched against their SAN list.

Adds a CertStream watchdog that resets silent connections and logs a `CertStream stats` line every 5 minutes.

└──▷ **BREAKING ON UPGRADE**

! The `SEARX_PATH` default value has changed from `../Searx` to `../Searx/searx` in `deployment/.env.example`; deployments using the old

value must update `deployment/.env` and remove any auto-generated `Searx/settings.yml`, then recreate the `searxng` container.

! The `certstream` image is now pinned to `0rickyy0/certstream-server-go:v1.8.2` via the new `CERTSTREAM_VERSION` variable; Compose files that override the `certstream` image with `latest` will pull v1.9/v1.10, which cannot stream behind an HTTP proxy.

! Compose files that override the `watcher` environment must explicitly add `certstream` to `no_proxy`, or the CertStream WebSocket will be incorrectly routed through the corporate proxy.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

JDBC Unity Catalog connections are now generally available (2026-10-02) BLOG ↗

RANK

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

> Databricks reached general availability on two governance and connectivity fronts: JDBC connections to Unity Catalog and cross-engine attribute-based access controls.

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

01 **JDBC Unity Catalog connections GA** NEW 59

JDBC Unity Catalog connections are now generally available on serverless compute, Databricks Runtime 19 and above, and SQL warehouses 2025.40 and above, with SSL certificate support included.

— Names platform versions but no connection steps

JDBC Unity Catalog connections are now …

02 **Cross-engine ABAC GA** NEW 57

Cross-engine attribute-based access controls (ABAC) are now generally available, allowing supported external engines to read Unity Catalog managed tables with ABAC policies, row filters, and column masks enforced by Databricks serverless compute.

— Explains enforcement mechanism but no setup steps

Cross-engine attribute-based access con…

> JDBC Unity Catalog connections are now generally available (2026-10-02)
>
> BLOG ↗ LATEST STABLE
>
> JDBC Unity Catalog connections reach GA on serverless compute, Databricks Runtime 19+, and SQL warehouses 2025.40+

> Cross-engine attribute-based access controls (ABAC) is now generally available (2026-1
>
> BLOG ↗ STABLE
>
> Cross-engine ABAC is now GA: external engines can read Unity Catalog tables with row filters and column masks enforced by Databricks serverless compute.

WAS THIS USEFUL?  

Exabeam

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

Exabeam is a security analytics platform that detects insider threats and advanced cyberattacks through behavioral analysis and machine learning.

October 2026 LogRhythm SIEM adds a community MCP server for local AI workflows, self-service compliance reporting engine, new rule-management APIs, AIE log ID drilldown, and Intelligence Sync proxy routing.

Adds a community Model Context Protocol (MCP) server that connects local generative AI models with LogRhythm SIEM data for on-premises queries, investigations, and alert triage.

Adds a next-generation self-service reporting engine with a library of customizable compliance graphs and reports that can be scheduled for recurring audit workflows.

Adds new APIs for automated rule creation and real-time rule statistics in the Advanced Intelligence Engine.

Adds log ID tracking to the Advanced Intelligence Engine (AIE) for precise event drilldown.

Adds proxy routing support for the LogRhythm Intelligence Sync service to meet complex network compliance requirements.

NEW INTEGRATION Generative AI Collectors for ChatGPT, Google Gemini, and GitHub Copilot

UPDATED SUPPORTED PLATFORM OpenSearch replaces Elasticsearch as the self-hosted search foundation

NEW INTEGRATION Claude Enterprise ingestion

NEW INTEGRATION Agentic SOC plugin for Claude Code and OpenAI Codex

▾ 1 more release in this issue · 2026-10-01

> announced · Oct 1, 2026 BLOG ↗
>
> October 2026 New-Scale adds a persistent Nova AI sidebar, Related Cases correlation in Threat Center, Outcomes Navigator overrides, data segregation for MSSPs, and weekly Executive Digest email reports.
>
> Adds a persistent Nova AI sidebar that follows analysts through Threat Center, search timelines, and Outcomes Navigator while retaining conversational context across workflows.
>
> Adds Exabeam Related Cases in Threat Center, correlating active and historical cases through shared artifacts and entities — including users, hosts, and IP

addresses — and surfacing earlier notes and containment playbooks.

Adds Outcomes Navigator overrides, letting security leaders adjust use case coverage scores when custom or unrecognized data sources satisfy specific security techniques.

Adds Threat Center data segregation, partitioning a single New-Scale subscription into virtual tenants with dedicated site management and role-based access controls for MSSPs and business units.

Updates Executive Digest to deliver weekly email reports that translate technical security posture improvements into business outcomes for executives.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Test v2.3.0: Structured CI test reports (2026-10-03) BLOG ↗ RANK

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

> Tenzir added a native DuckDB sink for routing events to local files, while Tenzir Test gained structured CI reporting and selective warning suppression — with the deprecated `--multi` test flag removed.

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

01 **to_duckdb operator for local database output** NEW 80

A new `to_duckdb` operator writes events into a local DuckDB database file (e.g. `"alerts.duckdb"`) using an embedded DuckDB engine, requiring no server to be running.

Route alert events into a local DuckDB file for offline analysis without standing up a database server.

```
$ to_duckdb "alerts.duckdb"
```

— Names exact operator and includes runnable command. product docs

THINNER COVERAGE BELOW

02 **Structured JSON test reports for CI** NEW 50

Tenzir Test v2.3.0 exports final test outcomes, failure diagnostics, and plain diffs as versioned JSON reports for CI integration, and adds tagged log records so CI pipelines can surface failures from sandboxed builds without parsing terminal output.

— Describes mechanism but no flag or command to invoke it.

Tenzir Test v2.3.0: Structured CI test …

03 **Removal of deprecated --multi test flag** BREAKING 50

The deprecated `--multi` option has been removed from Tenzir Test; any test configurations or scripts using it will break on upgrade.

— Names the exact removed flag with upgrade impact.

Tenzir Test v2.2.0: Selective runtime w…

04 **Selective runtime warning suppression in tests** IMPROVED

40

Tenzir Test v2.2.0 adds runtime warning suppression, allowing irrelevant warnings to be silenced while keeping errors and compilation diagnostics visible.

— Explains behavior but no named flag or config key given.

Tenzir Test v2.2.0: Selective runtime w…

└──▷ **BREAKING ON UPGRADE**

! The deprecated `--multi` option has been removed; any test configurations or scripts using it will break on upgrade.

Tenzir Test v2.3.0: Structured CI test reports (2026-10-03) BLOG ↗ LATEST

STABLE

Tenzir Test v2.3.0 adds versioned JSON test reports and tagged log records for CI-friendly failure surfacing.

Exports final test outcomes, failure diagnostics, and plain diffs as versioned JSON reports for CI integration.

Adds tagged log records so CI pipelines can surface failures from sandboxed builds without parsing terminal output.

Tenzir Test v2.2.0: Selective runtime warning suppression (2026-10-02) BLOG ↗

STABLE

Tenzir Test v2.2.0 adds runtime warning suppression so irrelevant warnings can be silenced without hiding errors.

└──▷ **BREAKING ON UPGRADE**

! The deprecated `--multi` option has been removed; any test configurations or scripts using it will break on upgrade.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Artemis

Feedback ⓘ v4.0.0 NOTES ↗ RANK

A modular vulnerability scanner with automatic report generation capabilities.

Artemis v4.0.0 adds per-scan module config, full asset export with CPEs, subdomain brute-forcing, and a login/password authorization layer.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.0.0 https://github.com/CERT-Polska/Artemis.git
# already have the repo? check out this version:
$ git checkout v4.0.0
```

Adds subdomain brute-forcing in addition to downloading from public sources.

Enables per-scan module configuration, not only global configuration.

Adds export of all found assets (domains, IPs, technologies, open ports, WordPress, WordPress plugins) with CPEs where available.

Adds an authorization layer (login + password) to the service.

NEW DETECTION 403 bypass detection

NEW DETECTION Nuclei DAST (fuzzing) templates support

NEW DETECTION Login form brute-force module

NEW DETECTION Joomla extension version check

NEW DETECTION Exposed NTLM-over-HTTP endpoint reporting

NEW DETECTION LFI detection

NEW DETECTION ORM injection detector

NEW DETECTION OS command injection detector

NEW INTEGRATION Shodan internetdb open port data download

WAS THIS USEFUL? 👍 👎

defectdojo

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo launched Rules Engine 2.0 (Beta), a major expansion of its automation engine adding branching workflows, egress actions to external systems, simulation testing, full run and delivery traceability, provenance tracking, a graph view, and new Asset Rules, alongside a migration path from the original Rules Engine.

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

01 Simulation mode for testing Rules NEW

70

Simulation mode executes Finding changes while recording egress actions without actually sending them, letting teams validate a Rule before enabling external actions. In Simulation mode, Slack messages, Jira tickets, and webhooks are logged in Deliveries as skipped rather than sent.

Validate a new Rule against real Findings without firing Slack, Jira, or webhook actions until you are confident the logic is correct.

> 📍 In the Rule editor, run the Rule in Simulation mode: Finding changes will be applied and recorded in Runs, but all Egress actions (Slack messages, Jira tickets, webhooks) are logged in Deliveries as skipped rather than sent.

— Exact before/after behaviour with a usable walkthrough launch-20261001-4c8a52cb

02 Rules Engine 2.0 (Beta) enabled via Feature Flags NEW

65

DefectDojo adds Rules Engine 2.0 (Beta), expanding automation nodes from 14 to 38 across five categories. It is enabled by toggling Rules Engine 2.0 in `Settings > Feature Flags` in the sidebar.

Enable Rules Engine 2.0 Beta so you can build branching, egress-capable triage workflows.

> 📍 In DefectDojo, go to Settings › Feature Flags in the sidebar and toggle the Rules Engine 2.0 pill to on.

— Names exact toggle location and node-count expansion launch-20261001-4c8a52cb

03 Egress actions to external systems in Rules NEW

65

Rules can now trigger actions outside DefectDojo: create Jira or other downstream tickets, send Slack or Microsoft Teams messages, deliver email, call a webhook, publish to an SNS topic, raise an in-app notification, or generate a report.

— Lists every supported egress channel launch-20261001-4c8a52cb

04 **Run traceability, delivery log, and provenance for Rules Engine 2.0**

NEW 65

Runs traceability records every Rule execution with a `node-by-node trace`, `trigger source`, `status`, `timing`, and `errors`. The Deliveries log catalogues every outbound action a Run sends, including `payload`, `destination`, `status`, `retry count`, `destination response`, and `skips`. Provenance surfaces which Rule, Run, or node caused a change directly on the Finding record, and a graph view gives a bird's-eye view of every possible outcome for a Rule in one place.

— Names every tracked field but no navigation path given launch-20261001-4c8a52cb

05 **Asset Rules for automated asset management** NEW 60

Asset Rules respond to Asset creation, update, or tagging events, or run on a schedule, and can update Asset fields, apply or remove tags, move Assets between Organizations, and reorganize the Asset hierarchy.

— Clear scope of triggers and actions, no exact steps launch-20261001-4c8a52cb

THINNER COVERAGE BELOW

06 **Branching logic in Rule workflows** NEW 55

A single Rule can evaluate a condition and route Findings down different action paths within one workflow — for example, sending Critical findings to a ticketing system while tagging lower-severity findings for weekly review.

— Mechanism explained with concrete example but no exact UI steps

launch-20261001-4c8a52cb

07 **Side-by-side Rules Engine migration** IMPROVED 40

Rules Engine and Rules Engine 2.0 can run side by side, with a converter available to migrate existing Rules Engine workflows to Rules Engine 2.0.

— States the migration path but no detail on the converter mechanism

launch-20261001-4c8a52cb

announced · Oct 1, 2026 BLOG ↗

DefectDojo launches Triage Engine with Rules Engine 2.0: branching workflows, egress actions, simulation mode, and full run traceability.

└──▷ **HOW TO FIND IT**

Enable Rules Engine 2.0 Beta so you can build branching, egress-capable triage workflows.

> 📍 In DefectDojo, go to Settings › Feature Flags in the sidebar and toggle the Rules Engine 2.0 pill to on.

Validate a new Rule against real Findings without firing Slack, Jira, or webhook actions until you are confident the logic is correct.

> 📍 In the Rule editor, run the Rule in Simulation mode: Finding changes will be applied and recorded in Runs, but all Egress actions (Slack messages, Jira tickets, webhooks) are logged in Deliveries as skipped rather than sent.

Adds Rules Engine 2.0 (Beta), enabled via the Feature Flags toggle in Settings sidebar, expanding automation nodes from 14 to 38 across five categories.

Adds branching logic so a single Rule can evaluate a condition and route Findings down different action paths (e.g. Critical to a ticketing system, lower-severity to a weekly-review tag) within one workflow.

Adds Egress capability so Rules can trigger actions outside DefectDojo: create Jira or other downstream tickets, send Slack or Microsoft Teams messages, deliver email, call a webhook, publish to an SNS topic, raise an in-app notification, or generate a report.

Adds Simulation mode, which executes Finding changes while recording egress actions without actually sending them, letting teams validate a Rule before enabling external actions.

Adds Runs traceability, recording every Rule execution with a node-by-node trace, trigger source, status, timing, and errors.

+5 more

WAS THIS USEFUL? 👍 👎

Rapid7 Insight Platform

Feedback

changelog · Oct 2, 2026 PRESS ↗ RANK

Rapid7 Insight Platform is a unified security analytics platform providing vulnerability management, incident detection, and threat intelligence across organizations.

Rapid7 Intelligence unifies threat intel, vulnerability research, and Rapid7 Labs into one engine feeding MDR and Exposure Management with behavioral, expert-validated insights.

Introduces Rapid7 Intelligence, a unified engine combining threat intelligence, vulnerability research, and Rapid7 Labs to convert real-time attacker behavior into protective action across Rapid7 products, MDR, and Exposure Management.

Delivers frontline intelligence that flows directly into existing Rapid7 products and MDR to neutralize adversary infrastructure before attacks execute, bypassing traditional research lag.

Provides behavioral correlation capability that links disparate campaigns by shared attacker behaviors rather than static indicators.

NEW DETECTION Modular Linux malware ecosystem detection — BPFDoor variants, BPF Rekoobe, droppers, AVERAT implants

WAS THIS USEFUL?

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ announced · Sep 30, 2026 BLOG ↗ RANK

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

> runZero launched AI-assisted reporting under `Reports > AI threads`, letting users query live inventory data in plain language to generate executive-ready reports, backed by a guided template gallery and multi-format export and sharing.

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

01 AI-assisted reporting via plain-language queries NEW

85

Adds `Reports > AI threads` in the console, enabling plain-language queries against live inventory data to produce executive summaries, bar charts, and risk-ranked tables in minutes. Every AI answer is grounded in live tool calls against actual inventory data, with an expandable work log showing every underlying query, tool call, and context token, and the system proactively surfaces blind spots such as mislabeled assets and governance disconnects (e.g. critical-tagged hosts with zero exploit likelihood) without explicit prompting.

Generate a prioritized critical-asset report with CISA KEV context for an executive audience, without writing any query syntax.

> 📍 In the console, go to Reports › AI threads and enter: 'find critical assets anywhere in the world to understand immediate threats and present it in an actionable report'

— Names exact UI path and a runnable prompt example launch-20260930-43497739

02 Guided report template gallery NEW

60

Includes a guided report gallery in AI threads with built-in templates: Critical systems & Blast radius, Deep outliers, End-of-Life OS analysis, Internet exposure, and Rogue & unmanaged assets.

Instantly run a pre-built End-of-Life OS assessment across all sites when preparing a migration roadmap.

> 📍 In the console, go to Reports › AI threads, open the guided report gallery, and select End-of-Life OS analysis

— Names all five templates and the UI path to select them launch-20260930-43497739

03 **Export and sharing of AI report artifacts** NEW 60

Exports AI-generated report artifacts to Markdown, HTML, PDF, CSV tables, and Starlark integration scripts; each AI-generated report, script, or visualization is preserved as a distinct private artifact, shareable across team members in the same organization or exported externally.

— Names five export formats but no command or endpoint to invoke them

launch-20260930-43497739

announced · Sep 30, 2026 BLOG ↗

runZero adds AI-assisted reporting under Reports > AI threads, with plain-language queries, built-in templates, and export to Markdown, HTML, PDF, CSV, and Starlark.

└──▷ **HOW TO FIND IT**

Generate a prioritized critical-asset report with CISA KEV context for an executive audience, without writing any query syntax.

📍 In the console, go to Reports › AI threads and enter: 'find critical assets anywhere in the world to understand immediate threats and present it in an actionable report'

Instantly run a pre-built End-of-Life OS assessment across all sites when preparing a migration roadmap.

📍 In the console, go to Reports › AI threads, open the guided report gallery, and select End-of-Life OS analysis

Adds `Reports > AI threads` in the console, enabling plain-language queries against live inventory data to produce executive summaries, bar charts, and risk-ranked tables in minutes.

Exports AI-generated report artifacts to Markdown, HTML, PDF, CSV tables, and Starlark integration scripts.

Includes a guided report gallery with built-in templates: Critical systems & Blast radius, Deep outliers, End-of-Life OS analysis, Internet exposure, and Rogue & unmanaged assets.

Grounds every AI answer in live tool calls against actual inventory data, with an expandable work log showing every underlying query, tool call, and context token.

Proactively surfaces blind spots such as mislabeled assets and governance disconnects (e.g. critical-tagged hosts with zero exploit likelihood) without explicit prompting.

+1 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

JumpCloud

Feedback ⓘ Release Notes 2022 (2026-10-02) BLOG ↗

RANK

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

> JumpCloud's biggest window items are JumpCloud Access Requests with webhook integration, JumpCloud MCP Server reaching GA, a new Custom Admin Roles Identity Management scope, and User Portal Asset Visibility/Acknowledgment in Asset Management, alongside a string of SSO/MFA hardening features, a reworked system-user reactivation API, unified Patch Management dashboards, and Password Manager desktop app updates.

└──▷ **WHAT SHIPPED · 30 FEATURES** most completely described first ? what's the number?

01 System user reactivation endpoint replaced BREAKING

90

The deprecated `GET api/systemusers/reactivate/{id}` call now returns a 404 error and is replaced by `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

Reactivate a user account using the new POST endpoint after migrating away from the deprecated GET call.

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

— Exact endpoint, method, field, runnable example Release Notes 2024 (2026-10-02)

02 Asset Visibility and Acknowledgment in Asset Management NEW

90

Asset Management adds User Portal Asset Visibility, letting employees see their assigned assets (name, serial number, status, ownership type) in the JumpCloud User Portal, and Asset Acknowledgment, letting admins send receipt-confirmation requests to an asset's current Owner from the Asset Detail view, with employees responding via `Tasks > Requests`. Every acknowledgment request and response is logged in Asset History for an audit trail, requests can be resent to the current Owner at any time for

periodic re-confirmation, and both settings are disabled by default for existing organizations but enabled by default for new ones.

Turn on Asset Visibility and Acknowledgement so end users can see their assigned hardware and admins receive confirmation receipts on every assignment.

> 📍 In the Admin Portal: Go to Device Management › Asset Management › Settings and Fields › General Settings. In the User Portal section, enable Asset Visibility and enable Asset Acknowledgement, then save.

Enable employee asset visibility and acknowledgment for an existing organization that already uses Asset Management.

> 📍 In the JumpCloud Admin Portal, go to Asset Management › Settings, then enable `User Portal Asset Visibility` and `Asset Acknowledgment`.

Send a receipt-confirmation request to an asset owner and track their response for an audit.

> 📍 In Asset Management, open the Asset Detail view for the target device, verify the correct Owner is assigned, then click `Send Acknowledgment Request`. Monitor the response status and full history under Asset History.

— Full mechanism, audit trail, and defaults given `Release Notes 2026 (2026-10-02)` `launch-20260930-76845fae`

03 Disk encryption condition in Conditional Access Policies NEW

85

Conditional Access Policies now support a disk-encryption device condition, letting admins deny access from endpoints without disk encryption enabled: in the Admin Portal go to `Security > Conditional Access Policies`, select or create a policy, add the disk encryption condition under `Conditions > Device`, and set the action to `Deny`.

Require disk encryption as a device condition in a Conditional Access Policy to deny access from unencrypted endpoints.

> 📍 In the Admin Portal, go to Security › Conditional Access Policies › (select or create a policy) › Conditions › Device, then add the disk encryption condition and set the action to `Deny` when encryption is disabled.

— Exact policy path and action given `Release Notes 2022 (2026-10-02)`

04 Dynamic Device Groups with Vendor and Public IP attributes

NEW

85

Dynamic Device Groups can auto-populate membership using new `Vendor` and `Public IP Address` attribute rules, scoping policies to devices on a given IP range or hardware vendor without manual membership management. In the Admin Portal, go to `Device Groups`, create or edit a group, enable `Dynamic`, and add attribute rules using the `Vendor` or `Public IP Address` attributes.

Auto-populate a device group with all devices on a specific public IP range or vendor to scope policies without manual membership management.

> 📍 In the Admin Portal, go to Device Groups, create or edit a group, enable `Dynamic`, and add attribute rules using the `Vendor` or `Public IP Address` attributes introduced in 2023.

— Exact attributes and UI path given `Release Notes 2023 (2026-10-02)`

---

05 **JumpCloud MCP Server reaches GA** NEW

80

JumpCloud MCP Server reached GA, connecting compatible AI tools and code editors to JumpCloud via the Model Context Protocol. Admins can enable it for end users at `Settings > JumpCloud AI > MCP Server for Users`, review or revoke active OAuth client sessions, and configure an automatic expiration window.

Enable the MCP Server for Users so staff can run User Portal tasks from AI clients, then review active OAuth sessions for audit purposes.

> 📍 In the Admin Portal: Go to Settings › JumpCloud AI and toggle on MCP Server for Users. To review or revoke active client sessions, stay on that page. To set an automatic expiration window, configure the expiration field before saving.

— GA plus admin toggle path given `Release Notes 2026 (2026-10-02)` `Release Notes 2025 (2026-10-02)`

---

06 **JumpCloud Access Requests with webhook integration** NEW

80

JumpCloud Access Requests lets admins configure approval flows, enables users to request and track access from the User Portal with events recorded in Directory Insights, and supports webhook integration to create tickets and receive real-time notifications in external systems such as JIRA. Configure via `Access Requests > Create Approval Flow` and set the webhook URL under Webhooks.

Set up an Access Request approval flow so users request access from the User Portal and a ticket is auto-created in JIRA on submission.

📍 In the JumpCloud Admin Portal, go to Access Requests › Create Approval Flow, assign approvers, then under Webhooks configure the JIRA webhook URL to receive real-time notifications and auto-create tickets when a request is submitted.

— Approval flow and webhook steps given `Release Notes 2025 (2026-10-02)`

---

07 **Unified Patch Management dashboards** NEW 80

New unified Patch Management dashboards live under `Device Management > Patch Management`, including a cross-platform Overview, a Windows KB-centric Patch List with CSV and JSON export, and an Apple OS Compliance view.

— Exact path, views and export formats named `Release Notes 2026 (2026-10-02)`

---

08 **Custom Admin Roles permission scope changes** BREAKING 80

Custom Admin Roles gain a dedicated Identity Management scope, separated from Directory and Groups Management, giving granular CRUD control over users alone (configure under `Roles > Custom Roles > Permissions > Identity Management`). The Policy Management scope is renamed Access Management, the System Management scope is renamed Devices, and the Authentication Policies scope is renamed Conditional Access Policies and IP Lists.

Assign a custom admin role with granular CRUD control over users only, now that Identity Management is a dedicated scope separate from Directory and Groups Management.

📍 In the JumpCloud Admin Console, go to Roles › Custom Roles › select or create a role › Permissions › Identity Management › set the permission level to Read, Create, Update, Delete, or No Access.

— All scope renames and new scope named `Roles Based Access Release Notes 2025 (…`

---

09 **Password Manager install/update command template** NEW 80

A new `Install and Update JumpCloud Password Manager App` command template (v2.0.0) lets admins deploy and keep the desktop app updated across Windows or macOS device groups; once the command runs once, end users can update the app without admin permissions.

Deploy and maintain the Password Manager desktop app org-wide without requiring end-users to have admin rights, using the new command template.

📍 In the JumpCloud Admin Portal, go to Commands, create a new command using the `Install and Update JumpCloud Password Manager App` command template v.2.0.0, and target your Windows or macOS device groups. After the command runs once, users can update the app without admin permissions.

— Exact command template name, version, steps `What's New: Password Manager Desktop Ap…`

---

10 **JumpCloud Go passwordless authentication** IMPROVED 75

JumpCloud Go can be set as the default MFA method for SSO conditional access policies, enabling passwordless login via hardware-backed Touch ID or Windows Hello once the browser extension is distributed to managed macOS and Windows devices (`Security > MFA Configurations`). JumpCloud Go sessions now also automatically refresh on device login or unlock, removing additional browser verification challenges.

Enable passwordless login for a managed device fleet so users authenticate to SSO apps via hardware-backed Touch ID or Windows Hello instead of passwords.

📍 In the Admin Portal, go to Security › MFA Configurations, set the default MFA method for SSO conditional access policies to `JumpCloud Go`, then distribute the JumpCloud Go browser extension to managed macOS and Windows devices.

— Mechanism and setup steps both given `Release Notes 2024 (2026-10-02)` `Release Notes 2023 (2026-10-02)`

---

11 **Step-up authentication via ForceAuthn and prompt=login** NEW 75

The IdP now honors `ForceAuthn="true"` in SAML AuthnRequests and `prompt=login` in OIDC requests, forcing fresh credential prompts without ending the existing user session; requires the Service Provider to send the flag, confirmed under `SSO Applications` in the Admin Portal.

Gate access to a sensitive app so it always prompts for fresh credentials, even with an active session, by relying on your SP sending the correct SAML flag — no JumpCloud config change required; verify your SP is sending `ForceAuthn="true"` in its SAML AuthnRequest or `prompt=login` in its OIDC request.

📍 In the JumpCloud Admin Portal, go to SSO Applications › select your app › confirm the SP is configured to send `ForceAuthn="true"` (SAML) or `prompt=login` (OIDC); JumpCloud will then prompt users for credentials on every access regardless of session state.

— Protocol flags named with verification steps `Release Notes 2025 (2026-10-02)`

---

12 **Auto Lock idle timeout for Password Manager** NEW 75

The Password Manager desktop app adds an Auto Lock setting (`Settings > Auto Lock`) with a configurable idle timeout, defaulting to 15 minutes with a maximum of 1,440 minutes, to automatically lock the vault on inactive sessions.

Enforce automatic vault locking for idle sessions to reduce credential exposure on shared or unattended workstations.

> 📍 In the Password Manager desktop app, go to Settings › Auto Lock, enable the auto-lock schedule, and set the idle timeout (default: 15 minutes; maximum: 1,440 minutes).

— Exact setting path and default/max values `What's New: Password Manager Desktop Ap…`

---

13 **Push Bombing Mitigation in MFA Configurations** NEW 70

MFA Configurations adds a Push Bombing Mitigation section with a configurable maximum concurrent push-attempt limit, letting admins toggle the feature and tighten limits; found at `Security > MFA Configurations > Push Bombing Mitigation` .

Protect admin accounts against push bombing attacks by reviewing and tightening the concurrent push attempt limit in MFA Configurations.

> 📍 In the Admin Portal, go to Security › MFA Configurations, locate the Push Bombing Mitigation section, and adjust the maximum concurrent attempts limit or toggle the feature on/off as appropriate for your environment.

— Named section and exact navigation path `Release Notes 2023 (2026-10-02)`

---

14 **Refined Dormant Application Access detection** IMPROVED 70

Access Risk's Dormant Application Access detection now flags re-access after inactivity above a configured threshold (set in `Access Risk > Configuration`), rather than flagging all first-time access.

— Before/after behavior and config path named `Release Notes 2026 (2026-10-02)`

---

15 **Dynamic Groups API filter operators renamed** BREAKING 65

Dynamic Groups `memberQuery` API filter operators `$le` and `$ge` are deprecated and replaced by `$lte` and `$gte` ; existing API scripts using `$le` or `$ge` must be updated.

— Operators named, no example of fix `Release Notes 2024 (2026-10-02)`

---

16 **`sso_auth` event for forced re-authentication** NEW 65

A new `sso_auth` event is logged for forced re-authentication events, giving audit visibility for compliance reviews; filter Directory Insights by event type `sso_auth` to see user, application, and timestamp details.

Audit which users triggered forced re-authentication events — useful for compliance reviews of high-stakes actions.

> 📍 In Directory Insights, filter events by type `sso_auth` to see all forced re-authentication events with user, application, and timestamp details.

— Event name and filter steps given Release Notes 2025 (2026-10-02)

---

17 **SaaS Management shadow IT discovery** NEW 60

SaaS Management discovers, tracks, warns on, and blocks shadow IT SaaS apps via the JumpCloud Go browser extension, with manual app addition, usage insights, and licensing tracking; app discovery now works even without JumpCloud Go enabled as a feature, though the browser extension is still required.

— Mechanism explained, no exact UI path Release Notes 2024 (2026-10-02)

---

18 **`.appimage` Linux build discontinued** DEPRECATED 60

The `.appimage` Linux build of the Password Manager desktop app is no longer supported as of version 2.2.197 (April 18, 2024); users must migrate to the RPM or Debian build.

— Exact version and migration path named What's New: Password Manager Desktop Ap…

THINNER COVERAGE BELOW

19 **Custom Banned Password List** NEW 50

Password Policies add a Custom Banned Password List, letting admins block specific words, terms, and phrases from being used in user passwords.

— Names feature, lacks exact UI path Release Notes 2026 (2026-10-02)

---

20 **Password Manager item organization and bulk actions** NEW 50

Password Manager desktop app items can now be sorted by nickname and multi-selected for bulk management; shared folders also support a Move All button to migrate all items to another folder at once instead of moving them individually.

Migrate all credentials from a legacy shared folder in bulk rather than item-by-item.

> 📍 Open the shared folder in the Password Manager desktop app, click the Move All button, and select the destination folder to move all items at once.

— Names three UI actions, no deep mechanism `What's New: Password Manager Desktop Ap…`

21 **Agent Release Notes RSS feed subscription** NEW 45

Admins can subscribe to the Agent Release Notes RSS Feed from the Release Notes and Bug Fixes page to receive agent update notifications automatically instead of checking the Help Center manually.

Stay automatically informed of JumpCloud agent updates without polling the Help Center manually.

📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click Subscribe to Agent Release Notes RSS Feed. 3. Open your RSS reader and paste or confirm the feed URL to begin receiving agent release note updates.

— Thin feature but exact navigation given `Subscribe to New Feature and Product Up…`

22 **AMR claims in SAML and OIDC** NEW 40

SAML assertions and OIDC ID tokens now include Authentication Methods References (AMR) claims, proving which factors (e.g., password + OTP) were used at login.

— Mechanism named, no usage steps `Release Notes 2025 (2026-10-02)`

23 **New Admin Portal homepage widgets** NEW 35

The Admin Portal and Multi-Tenant Portal homepage gains an MDM Upcoming Certificate Expirations widget for at-a-glance MDM, VPP, and Automated Device Enrollment certificate lifecycle visibility, plus a trial recap view summarizing feature usage during a JumpCloud feature trial.

— No navigation path given for new widgets `Release Notes 2022 (2026-10-02)`

24 **Remote Assist for Windows server login screen** NEW 35

Remote Assist now lets admins view and interact with the initial login screen for Windows servers.

— Capability named, no mechanism detail `Release Notes 2023 (2026-10-02)`

25 **ADI support on Windows Server 2022** IMPROVED 30

Active Directory Integration (ADI) now supports Windows Server 2022 domain controllers and member servers without workarounds or decreased functionality.

— States scope, no steps or limits `Release Notes 2023 (2026-10-02)`

26 **Expanded OS version support in MDM and Windows management** IMPROVED 30

JumpCloud added Android 15 support in MDM and Windows 11 25H2 support in its Windows management components.

— Names versions, no steps or scope Release Notes 2025 (2026-10-02)

Release Notes 2024 (2026-10-02)

27 **Expanded macOS Personalization policies** IMPROVED 25

macOS Personalization policies expand with updated Dock Settings and Items and Lock Screen (Screen Saver) policy options.

— Named policy types, no mechanism Release Notes 2025 (2026-10-02)

28 **M365/Azure AD security group support** NEW 20

JumpCloud's Microsoft 365 directory integration now supports M365/Azure AD security groups.

— Bare mention, no steps or scope detail Release Notes 2022 (2026-10-02)

29 **Multiple IP addresses per RADIUS server** IMPROVED 20

RADIUS server configurations now support multiple IP addresses within a single server configuration.

— One-line capability, no detail Release Notes 2024 (2026-10-02)

30 **Asset cloning in Asset Management** NEW 20

Asset Management adds asset cloning to streamline creation of assets with shared hardware configurations.

— One-line feature, no mechanism Release Notes 2025 (2026-10-02)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `GET` request for `api/systemusers/reactivate/{id}` is deprecated and now returns a 404 error; replace with `POST` `https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

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

! The Policy Management scope is renamed to Access Management.

! The System Management scope is renamed to Devices.

! The Authentication Policies scope is renamed to Conditional Access Policies and IP Lists.

! The `.appimage` Linux build is no longer supported as of version 2.2.197 (April 18, 2024); users must migrate to the RPM or Debian build.

Release Notes 2022 (2026-10-02) BLOG ↗ STABLE

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

└──▷ **HOW TO FIND IT**

Require disk encryption as a device condition in a Conditional Access Policy to deny access from unencrypted endpoints.

> 📍 In the Admin Portal, go to Security › Conditional Access Policies › (select or create a policy) › Conditions › Device, then add the disk encryption condition and set the action to `Deny` when encryption is disabled.

Adds MDM Upcoming Certificate Expirations widget to the Admin Portal and Multi-Tenant Portal homepage for at-a-glance MDM, VPP, and Automated Device Enrollment certificate lifecycle visibility.

Adds trial recap view summarising feature usage during a JumpCloud feature trial.

`Release Notes 2023 (2026-10-02)` **BLOG ↗** **STABLE**

JumpCloud: Adds Remote Assist support for admins to view and interact with the initial login screen for Windows servers.

└──▷ **HOW TO FIND IT**

Enable passwordless login for a managed device fleet so users authenticate to SSO apps via hardware-backed Touch ID or Windows Hello instead of passwords.

> 📍 In the Admin Portal, go to Security › MFA Configurations, set the default MFA method for SSO conditional access policies to `JumpCloud Go`, then distribute the JumpCloud Go browser extension to managed macOS and Windows devices.

Auto-populate a device group with all devices on a specific public IP range or vendor to scope policies without manual membership management.

> 📍 In the Admin Portal, go to Device Groups, create or edit a group, enable `Dynamic`, and add attribute rules using the `Vendor` or `Public IP Address` attributes introduced in 2023.

Protect admin accounts against push bombing attacks by reviewing and tightening the concurrent push attempt limit in MFA Configurations.

> 📍 In the Admin Portal, go to Security › MFA Configurations, locate the Push Bombing Mitigation section, and adjust the maximum concurrent attempts limit or toggle the feature on/off as appropriate for your environment.

Adds Active Directory Integration (ADI) support on 2022 domain controllers and

member servers without workarounds or decreased functionality.

Release Notes 2024 (2026-10-02) BLOG ↗ STABLE

JumpCloud: Adds `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` endpoint to replace the now-defunct `GET` variant for reactivating system users.

└──▷ **TRY IT**

Reactivate a user account using the new POST endpoint after migrating away from the deprecated GET call.

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

Adds SaaS Management capability to discover, track, warn, and block shadow IT SaaS apps via the JumpCloud Go browser extension, with manual app addition, usage insights, and licensing tracking.

Adds support for multiple IP addresses in a single RADIUS server configuration.

Adds Android 15 support in MDM.

Enables SaaS Management app discovery without requiring JumpCloud Go to be enabled as a feature (browser extension still required when JumpCloud Go is not enabled).

JumpCloud Go sessions now automatically refresh on device login or unlock, removing additional browser verification challenges.

└──▷ **BREAKING ON UPGRADE**

! The `GET` request for `api/systemusers/reactivate/{id}` is deprecated and now returns a 404 error; replace with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

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

Release Notes 2025 (2026-10-02) BLOG ↗ STABLE

JumpCloud adds Access Requests, MCP Server, Shadow AI Dashboard, Custom Admin Roles, and major SSO/MFA upgrades in 2025

└──▷ **HOW TO FIND IT**

Audit which users triggered forced re-authentication events — useful for compliance reviews of high-stakes actions.

> 📍 In Directory Insights, filter events by type `sso_auth` to see all forced re-authentication events with user, application, and timestamp details.

Gate access to a sensitive app so it always prompts for fresh credentials, even with an active session, by relying on your SP sending the correct SAML flag — no JumpCloud config change required; verify your SP is sending `ForceAuthn="true"` in its SAML AuthnRequest or `prompt=login` in its OIDC request.

> 📍 In the JumpCloud Admin Portal, go to SSO Applications › select your app › confirm the SP is configured to send `ForceAuthn="true"` (SAML) or `prompt=login` (OIDC); JumpCloud will then prompt users for credentials on every access regardless of session state.

Set up an Access Request approval flow so users request access from the User Portal and a ticket is auto-created in JIRA on submission.

> 📍 In the JumpCloud Admin Portal, go to Access Requests › Create Approval Flow, assign approvers, then under Webhooks configure the JIRA webhook URL to receive real-time notifications and auto-create tickets when a request is submitted.

Adds `sso_auth` event to logs for forced re-authentication events, providing audit visibility for compliance reviews.

IdP now honors `ForceAuthn="true"` in SAML and `prompt=login` in OIDC, enabling step-up authentication without ending the user session.

Includes Authentication Methods References (AMR) claims in SAML assertions and OIDC ID tokens to prove which factors (e.g., password + OTP) were used at login.

Introduces JumpCloud MCP Server (GA) — connects compatible AI tools and code editors to JumpCloud via the Model Context Protocol.

Introduces JumpCloud Access Requests — configure approval flows, manage access requests from the User Portal, and track events in Directory Insights.

`+4 more`

`Release Notes 2026 (2026-10-02)` BLOG ↗ STABLE

JumpCloud: Adds Custom Banned Password List to Password Policies, letting admins block specific words, terms, and phrases from user passwords.

└──▷ **HOW TO FIND IT**

Enable the MCP Server for Users so staff can run User Portal tasks from AI clients, then review active OAuth sessions for audit purposes.

📍 In the Admin Portal: Go to Settings › JumpCloud AI and toggle on MCP Server for Users. To review or revoke active client sessions, stay on that page. To set an automatic expiration window, configure the expiration field before saving.

Turn on Asset Visibility and Acknowledgement so end users can see their assigned hardware and admins receive confirmation receipts on every assignment.

📍 In the Admin Portal: Go to Device Management › Asset Management › Settings and Fields › General Settings. In the User Portal section, enable Asset Visibility and enable Asset Acknowledgement, then save.

Adds unified `Patch Management` dashboards under `Device Management > Patch Management`, including cross-platform Overview, Windows KB-centric Patch List (with CSV and JSON export), and Apple OS Compliance views.

Refines `Dormant Application Access` detection in Access Risk to flag re-access after inactivity above the configured threshold (set in `Access Risk > Configuration`) rather than flagging all first-time access.

`What's New: Password Manager Desktop App (2026-10-02)` `BLOG ↗` `STABLE`

JumpCloud: Adds the ability to sort items by nicknames.

└──▷ **HOW TO FIND IT**

Deploy and maintain the Password Manager desktop app org-wide without requiring end-users to have admin rights, using the new command template.

📍 In the JumpCloud Admin Portal, go to Commands, create a new command using the `Install and Update JumpCloud Password Manager App` command template v.2.0.0, and target your Windows or macOS device groups. After the command runs once, users can update the app without admin permissions.

Enforce automatic vault locking for idle sessions to reduce credential exposure on shared or unattended workstations.

📍 In the Password Manager desktop app, go to Settings › Auto Lock, enable the auto-lock schedule, and set the idle timeout (default: 15 minutes; maximum: 1,440 minutes).

Migrate all credentials from a legacy shared folder in bulk rather than item-by-item.

📍 Open the shared folder in the Password Manager desktop app, click the `Move All` button, and select the destination folder to move all items at once.

Enables users to multi-select items for bulk management.

└──▷ **BREAKING ON UPGRADE**

! The `.appimage` Linux build is no longer supported as of version 2.2.197 (April 18, 2024); users must migrate to the RPM or Debian build.

announced · Sep 30, 2026 **BLOG ↗**

JumpCloud Asset Management adds User Portal Asset Visibility and Asset Acknowledgment so employees can view and confirm assigned devices.

└──▷ **HOW TO FIND IT**

Enable employee asset visibility and acknowledgment for an existing organization that already uses Asset Management.

📍 In the JumpCloud Admin Portal, go to Asset Management › Settings, then enable User Portal Asset Visibility and Asset Acknowledgment.

Send a receipt-confirmation request to an asset owner and track their response for an audit.

📍 In Asset Management, open the Asset Detail view for the target device, verify the correct Owner is assigned, then click Send Acknowledgment Request. Monitor the response status and full history under Asset History.

Adds User Portal Asset Visibility setting to Asset Management, letting employees see their assigned assets (name, serial number, status, ownership type) directly in the JumpCloud User Portal.

Adds Asset Acknowledgment setting so admins can send receipt-confirmation requests to an asset's current Owner from the Asset Detail view; employees respond via `Tasks > Requests` in the User Portal.

Logs every acknowledgment request and response in Asset History, providing an audit trail of who was asked, who responded, and when.

Supports resending acknowledgment requests to the current Owner at any time, enabling periodic audits or re-confirmation after ownership changes without building recurring campaigns.

Both settings are disabled by default for existing Asset Management organizations; new organizations have both enabled by default.

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ announced · Oct 2, 2026 BLOG ↗ RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

> ConductorOne's launch week introduced a cluster of new governance services: C1 AppHub for app deployment, C1.ai as an OIDC identity provider with an OpenID AuthZEN authorization endpoint, credential vending paired with the C1 Egress proxy to remove hardcoded secrets, and C1 LLM Gateway for policy-controlled AI inference routing and spend attribution.

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

01 **C1 LLM Gateway for AI inference routing** NEW 75

C1 LLM Gateway routes model traffic by identity, data sensitivity, performance, and cost policy, with configurable routing rules based on provider, model, deployment, region, and data-handling requirements to determine eligible routes per workload. Each inference call is attributed to a caller — person, application, workload, or agent — with reporting broken down by app, department, person, or model, including cost per million tokens; model access and budgets are managed as entitlements with approval paths and end dates, and managed credential paths keep provider secrets out of application code.

.
.
.

— Rich mechanism and cost metrics given, but no API or config key named.

launch-20261002-8f5c3fad launch-20261001-016fe30e

02 **Credential vending for workload secrets** NEW 65

Credential vending mints scoped, time-limited credentials at workload runtime, restricts them to approved IP ranges, delivers them through a vault, and provides a central audit inventory of every issuance, use, and revocation. It supports per-task credential issuance for AI agents, contractors, and CI/CD pipelines governed by the same policy engine used for employee access, with instant revocation taking effect on the next call rather than requiring a new deployment.

— Mechanism and revocation behavior detailed, no API surface shown.

launch-20261002-8f5c3fad launch-20260930-835392b5

03 **C1 Egress governed proxy** NEW 65

C1 Egress is a governed egress proxy that sits between an internal app or AI agent and the network, injecting the real credential onto an outbound request only when the destination is allowed by policy. It blocks internal addresses and cloud metadata endpoints by default and logs every call's source, destination, and policy decision without ever recording the secret itself.

— Default-block behavior and logging described, no config flag named.

launch-20261002-8f5c3fad launch-20260930-835392b5

THINNER COVERAGE BELOW

04 **OpenID AuthZEN authorization endpoint** NEW 55

A new authorization endpoint built on OpenID AuthZEN lets apps query in real time whether a person can perform an action on a resource, with denials triggering an in-app access request workflow.

— Names the standard and flow, no endpoint path given. launch-20261002-8f5c3fad

05 **C1 AppHub deployment platform** NEW 55

C1 AppHub is an internal deployment platform that connects apps to identity, authorization, credentials, network access, and model routing via C1.ai, and ships with a skills file that teaches coding agents the company's governed patterns from the first commit.

— Describes connections and skills file, no concrete config shown.

launch-20261002-8f5c3fad

06 **C1.ai as an OIDC identity provider** NEW 45

C1.ai now acts as a standard OIDC provider that federates authentication to an existing identity provider, making app access a C1.ai entitlement subject to request, approval, review, certification, and revocation.

— Explains entitlement lifecycle, lacks endpoint or config detail.

launch-20261002-8f5c3fad

announced · Oct 2, 2026 BLOG ↗

ConductorOne ships four new services: AppHub, OIDC sign-in with OpenID AuthZEN authz, credential vending with C1 Egress proxy, and C1 LLM Gateway.

Adds C1 AppHub, an internal deployment platform that connects apps to identity, authorization, credentials, network access, and model routing via C1.ai, with a skills file that teaches coding agents the company's governed patterns from the first commit.

Adds C1.ai as a standard OIDC provider that federates authentication to an existing identity provider, making app access a C1.ai entitlement subject to request, approval, review, certification, and revocation.

Adds an authorization endpoint built on OpenID AuthZEN that lets apps query in real time whether a person can perform an action on a resource, with denials triggering an in-app access request workflow.

Adds credential vending that mints scoped, time-limited credentials at workload runtime, restricts them to approved IP ranges, delivers them through a vault, and provides a central audit inventory of every issuance, use, and revocation.

Adds C1 Egress, a governed egress proxy that injects real credentials into outbound requests only when policy allows the destination, blocks internal addresses and cloud metadata endpoints by default, and logs every call without recording the secret.

+1 more

announced · Oct 1, 2026 BLOG ↗

ConductorOne launches C1 LLM Gateway: a policy-controlled endpoint for routing AI inference across public, private, and customer-controlled model deployments.

Enables configurable routing rules based on provider, model, deployment, region, and data-handling requirements to determine eligible routes per workload.

Attributes inference usage and cost to the responsible caller — person, application, workload, or agent — with reporting broken down by app, department, person, or model, including cost per million tokens.

Supports managed credential paths to keep provider secrets out of application code when connecting to approved routes.

announced · Sep 30, 2026 BLOG ↗

ConductorOne launches credential vending and C1 Egress proxy to eliminate hardcoded secrets from apps and AI agents.

Adds C1 Egress, a governed proxy that sits between an internal app or AI agent and the network, injecting the real credential onto a request only when the destination is allowed by policy.

Logs every outbound call with its source, destination, and policy decision without ever recording the secret itself.

Enables per-task credential issuance for AI agents, contractors, and CI/CD pipelines governed by the same policy engine used for employee access.

Supports instant revocation that takes effect on the next call rather than requiring a new deployment.

WAS THIS USEFUL? 👍 👎

Keycloak

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

Keycloak - the open source identity and access management solution. Add single-sign-on and authentication to applications and secure services with minimum effort.

> Keycloak's 26.8 release promotes multi-cluster deployments without external Infinispan to supported, and previews token exchange delegation, OID4VCI mdoc credentials, and a new declarative Client Admin API.

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

01 Stateless multi-cluster support promoted to supported

IMPROVED 80

Multi-cluster v2 and the `stateless` feature have been promoted from preview to supported, enabling multi-cluster Keycloak deployments without needing an external Infinispan cluster. Enable it with `--features=stateless`, e.g. `bin/kc.sh start --features=stateless`.

Run a multi-cluster Keycloak deployment without an external Infinispan cluster by enabling the stateless feature.

```
$ bin/kc.sh start --features=stateless
```

— Exact flag and command given; underlying cluster mechanics not detailed

launch-20261001-cb9b31a1

02 Token exchange delegation in Fine-Grained Admin Permissions

NEW 75

Token exchange delegation moves to preview using a new `delegation:client:<client-id>` parameterized scope, which inserts an `act` claim identifying the delegating client. Fine-Grained Admin Permissions V2 gains `delegate` and `delegate-members` scopes to control authorization, alongside a new client policy executor that restricts `may_act` claims.

— Names scopes, claim and executor but gives no runnable example launch-20261001-cb9b31a1

03 Client Admin API v2 for declarative client management NEW

75

Enabling `--features=client-admin-api:v2` turns on the Client Admin API v2, which lets OIDC and SAML clients be managed declaratively with strict validation and an accurate OpenAPI spec.

Enable the Client Admin API v2 to manage OIDC and SAML clients declaratively with strict validation and an accurate OpenAPI spec.

```
$ bin/kc.sh start --features=client-admin-api:v2
```

— Exact flag and behaviour described, but no endpoint paths given

launch-20261001-cb9b31a1

04 **OID4VCI credential issuance with experimental mdoc format**

NEW 65

OID4VCI credential issuance and the experimental mdoc format can be enabled together with `--features=oid4vc-vci,oid4vc-mdoc`.

Enable OID4VCI credential issuance and the experimental mdoc format in a single start command.

```
$ bin/kc.sh start --features=oid4vc-vci,oid4vc-mdoc
```

— Exact command given but no explanation of the formats themselves

launch-20261001-cb9b31a1

announced · Oct 1, 2026 BLOG ↗

Keycloak: Promotes multi-cluster v2 and the `stateless` feature to supported, enabling multi-cluster Keycloak without an external Infinispan cluster; enabled with `--features=stateless`.

└──▷ **TRY IT**

Run a multi-cluster Keycloak deployment without an external Infinispan cluster by enabling the stateless feature.

```
$ bin/kc.sh start --features=stateless
```

Enable OID4VCI credential issuance and the experimental mdoc format in a single start command.

```
$ bin/kc.sh start --features=oid4vc-vci,oid4vc-mdoc
```

Enable the Client Admin API v2 to manage OIDC and SAML clients declaratively with strict validation and an accurate OpenAPI spec.

```
$ bin/kc.sh start --features=client-admin-api:v2
```

Adds token exchange delegation to preview using the new `delegation:client:<client-id>` parameterized scope, inserting an `act` claim identifying the delegating client, with Fine-Grained Admin Permissions V2 `delegate` and `delegate-members` scopes controlling authorization and a new client policy executor restricting `may_act` claims.

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Weekly release 2026.09.2 (2026-09-30) BLOG ↗ RANK

Secure, scalable, and highly available authentication and user management for any app.

Adds Certificate Authorities API (EA), Extensible AI Agent Profile Schema (Beta), editable MCP server resource URLs, unique MCP server name enforcement, and automatic phishing resistance for FastPass.

Adds Certificate Authorities API (EA) for managing certificate authorities programmatically.

Introduces Extensible AI Agent Profile Schema and Custom Attributes in Beta, enabling customization of AI agent profiles.

UPDATED INTEGRATION Enables editable resource URL for MCP servers.

Enforces unique MCP server names to prevent naming conflicts.

Enables automatic phishing resistance for FastPass in authentication chains.

UPDATED INTEGRATION Agent SSO support for third-party OIDC apps

UPDATED INTEGRATION Cross App Access support for AI agents and apps

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ announced · Sep 30, 2026 BLOG ↗ RANK

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

> Clerk shipped a new CLI-first onboarding flow built around `npx clerk@latest init`, letting developers and agents scaffold Clerk into a Next.js app and get working API keys without ever creating a Clerk account, alongside companion commands for linking apps, syncing environment keys, diagnosing setup issues, and deploying to production.

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

01 **`npx clerk@latest doctor` health checks and auto-fix** NEW

93

`npx clerk@latest doctor` inspects CLI session, linked application, and project health, reporting each failing check with a remedy. The `--fix` flag automatically applies remediable issues (signing in, linking the project, running env pull), and `--json` outputs structured results including the `remedy` field for agents instead of interactive prompts.

Let an agent audit the Clerk setup and parse structured remedies without triggering interactive prompts.

```
$ npx clerk@latest doctor --json
```

— Mechanism, three flags, and a runnable example given launch-20260930-aa1c330c

02 **Accountless setup via `npx clerk@latest init`** NEW

85

`npx clerk@latest init` scaffolds Clerk into a Next.js project and writes development keys to the local env file without requiring a Clerk account or Dashboard login. The `--accountless` flag forces accountless app creation even when already signed in, and `--login` forces the authenticated flow instead.

Bootstrap Clerk auth in a new Next.js project — no account, no Dashboard, no copy-pasting keys — so an agent can continue immediately.

```
$ npx clerk@latest init
```

— Exact command and flags given, runnable as shown launch-20260930-aa1c330c

03 **`npx clerk@latest env pull` for dev and prod keys** NEW

85

`npx clerk@latest env pull` writes a linked app's development keys into the local env file. The `--instance prod` flag pulls production keys instead, and `--file <path>` writes keys to a specific env file rather than the auto-detected default.

Pull production keys into a dedicated file to keep them out of the dev env file your local server reads.

```
$ npx clerk@latest env pull --instance prod --file
.env.production.local
```

— Named flags with a runnable example command launch-20260930-aa1c330c

04 **`npx clerk@latest link` to attach existing app** NEW

60

`npx clerk@latest link` attaches an existing Clerk app to a project instead of creating a new one, with an optional `--app <id>` flag to specify which app.

— Command and flag named but no usage example launch-20260930-aa1c330c

05 **`npx clerk@latest deploy` production wizard** NEW 60

`npx clerk@latest deploy` is an interactive wizard that creates a production Clerk instance and walks through DNS records and OAuth providers.

— Command and scope named but no example or steps launch-20260930-aa1c330c

THINNER COVERAGE BELOW

06 **Structured error when Clerk keys are missing** NEW 50

Starting in `@clerk/nextjs` 7.8.0, an app without Clerk keys throws a structured error instructing users or agents to run `npx clerk@latest init`.

— Version and trigger named, but no further mechanism launch-20260930-aa1c330c

07 **`npx clerk@latest auth login` to claim accountless app** NEW 40

`npx clerk@latest auth login` signs in and claims an accountless application, attaching it to a Clerk account.

— Command named but behavior and mechanism thinly described launch-20260930-aa1c330c

announced · Sep 30, 2026 BLOG ↗

Clerk adds accountless setup via `npx clerk@latest init`, letting agents and developers get API keys without signing up.

└──▷ **TRY IT**

Bootstrap Clerk auth in a new Next.js project — no account, no Dashboard, no copy-pasting keys — so an agent can continue immediately.

```
$ npx clerk@latest init
```

Pull production keys into a dedicated file to keep them out of the dev env file your local server reads.

```
$ npx clerk@latest env pull --instance prod --file
.env.production.local
```

Let an agent audit the Clerk setup and parse structured remedies without triggering interactive prompts.

```
$ npx clerk@latest doctor --json
```

Adds `npx clerk@latest init` to scaffold Clerk into a Next.js project and write development keys without a Clerk account or Dashboard login.

Adds `--accountless` flag to `npx clerk@latest init` to force accountless app creation even when signed in, and `--login` to force the authenticated flow.

> Adds `npx clerk@latest link` (with optional `--app <id>`) to attach an existing Clerk app to a project instead of creating a new one.
>
> Adds `npx clerk@latest env pull` to write a linked app's development keys into the local env file, and `--instance prod` to pull production keys instead.
>
> Adds `--file <path>` flag to `npx clerk@latest env pull` to write keys to a specific env file instead of the auto-detected default.
>
> +6 more

WAS THIS USEFUL?  

Nudge Security

Feedback ⓘ announced · Sep 30, 2026 BLOG ↗

RANK

Nudge Security is a SaaS security platform that discovers shadow IT assets, manages identity governance, and hardens security posture.

> Nudge Security introduced adaptive risk management — automated business criticality tiering, a dynamic multi-factor risk score, and AI-driven data sensitivity inference across SaaS and AI apps — alongside two new autonomous agents for OAuth grant review and browser extension risk.

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

01 **Dynamic per-app risk score** NEW 75

Calculates a dynamic risk score per app from more than 30 internal and external factors — including OAuth grants and scopes, MCP connections, AI agent permissions, authentication methods, SSO coverage, MFA enrollment, security posture findings, stale accounts, and vendor breach/supply-chain intelligence — drawn from a database of more than 250,000 SaaS and AI vendor security profiles.

— Extensive named factors and scale, but no action path launch-20260930-adfeaab0

02 **Risk dashboard heatmap and per-app Risk tab** NEW 70

Adds a risk dashboard heatmap pairing business criticality tiers with app risk scores to visualize risk concentration across the SaaS and AI estate, and a Risk tab per app surfacing a ranked breakdown of control gaps and the factors driving each risk score, updated continuously as usage and vendor conditions change.

Review the ranked control gaps driving risk for a specific app so you know what to remediate first before a vendor posture change escalates exposure.

> 📍 In the Nudge Security console, navigate to the app's detail page and select the Risk tab to see a ranked breakdown of control gaps and the factors contributing to its current dynamic risk score.

Spot where risk and business impact overlap most across your SaaS and AI estate to prioritize deeper vendor reviews.

> 📍 In the Nudge Security console, open the risk dashboard and use the heatmap to identify apps where a `critical` or `high` business criticality tier intersects with elevated app risk scores.

— Exact UI navigation paths given for both surfaces launch-20260930-adfeaab0

THINNER COVERAGE BELOW

03 **Automated business criticality tiering** NEW 55

Assigns every discovered SaaS and AI app a `critical`, `high`, `medium`, or `low` business criticality tier based on organizational spread and inferred data sensitivity.

— Named tiers and basis, but no direct action step given launch-20260930-adfeaab0

04 **AI-inferred data sensitivity mapping** NEW 55

Uses a proprietary AI model to infer up to 29 data types each app typically handles, mapping them to sensitivity tiers that customers can customize to align with internal data governance models.

— Concrete figure and customization noted but no UI steps launch-20260930-adfeaab0

05 **Adaptive policy-driven remediation controls** NEW 45

Enables adaptive, policy-driven controls that automatically expand security posture monitoring, strengthen authentication requirements, or route remediation to app owners, administrators, or users as an app's criticality or risk score shifts.

— Describes behavior in prose without specific triggers or paths launch-20260930-adfeaab0

announced · Sep 30, 2026 BLOG ↗

Nudge Security launches adaptive risk management with automated app tiering, dynamic risk scoring, and AI-driven remediation workflows for SaaS and AI.

└──▷ **HOW TO FIND IT**

Review the ranked control gaps driving risk for a specific app so you know what to remediate first before a vendor posture change escalates exposure.

📍 In the Nudge Security console, navigate to the app's detail page and select the Risk tab to see a ranked breakdown of control gaps and the factors contributing to its current dynamic risk score.

Spot where risk and business impact overlap most across your SaaS and AI estate to prioritize deeper vendor reviews.

📍 In the Nudge Security console, open the risk dashboard and use the heatmap to identify apps where a `critical` or `high` business criticality tier intersects with elevated app risk scores.

Adds automated business criticality tiering for every discovered SaaS and AI app, assigning each a `critical`, `high`, `medium`, or `low` tier based on organizational spread and inferred data sensitivity.

> Introduces a dynamic risk score per app calculated from more than 30 internal and external factors — including OAuth grants and scopes, MCP connections, AI agent permissions, authentication methods, SSO coverage, MFA enrollment, security posture findings, stale accounts, and vendor breach/supply-chain intelligence — drawn from a database of more than 250,000 SaaS and AI vendor security profiles.
>
> Uses a proprietary AI model to infer up to 29 data types each app typically handles, mapping them to sensitivity tiers that customers can customize to align with internal data governance models.
>
> Adds a risk dashboard heatmap pairing business criticality tiers with app risk scores to visualize risk concentration across the entire SaaS and AI estate.
>
> Adds a `Risk` tab per app surfacing a ranked breakdown of control gaps and the factors driving each risk score, updated continuously as internal usage and vendor conditions change.
>
> +2 more

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Zama fhEVM

Feedback ⓘ v0.14.2 NOTES ↗ RANK

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

> fhEVM v0.14.2 extends host-contract chain support to BNB Chain and refines scalar operand handling for FHE arithmetic.

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

01 **BNB Chain added to host-contracts chainIds** NEW 40

Adds BNB to the supported `chainIds` in host-contracts, enabling FHE workflows on the BNB chain.

— Names config surface but no setup detail v0.14.2

02 **Improved scalar operand handling in host-contracts** IMPROVED 25

Improves scalar operand handling in host-contracts for FHE arithmetic operations.

— Vague description with no mechanism given v0.14.2

v0.14.2 NOTES ↗ LATEST STABLE

fhEVM v0.14.2 adds BNB chain support to host-contract chainIds and improves scalar operand handling.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.14.2 https://github.com/zama-ai/fhevm.git
# already have the repo? check out this version:
$ git checkout v0.14.2
```

Adds BNB to the supported `chainIds` in host-contracts, enabling FHE workflows on the BNB chain.

Improves scalar operand handling in host-contracts for FHE arithmetic operations.

WAS THIS USEFUL? 👍 👎

Cyera

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗ RANK

Cyera helps organizations protect data, govern identities, and secure AI agents to reduce risk, stay in control, and adopt AI with confidence.

> Cyera shipped a new integration between Agent Guardian and Microsoft Agent 365, surfacing AI agent data security risk directly in Microsoft's admin tooling.

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

01 Agent Guardian risk signals in Microsoft Agent 365 NEW

77

Cyera Agent Guardian now pushes policy-driven risk events to the Agent 365 Risk Service via Microsoft's API in near real time, authenticated through Microsoft Entra ID using OAuth 2.0 client credentials. Cyera policy severity is mapped to the Agent 365 third-party risk schema fields `riskLevel` (low/medium/high), `riskEvidence` (policy name and description), and `status` (driven by Issue status), so findings appear directly on the Agent Registry page of the Microsoft 365 Admin Center. The integration delivers agent data security posture signals — including agent configurations, capabilities, connected tools, access paths to sensitive data, and drift from policy — alongside Microsoft Purview, Microsoft Defender, and Microsoft Entra risk signals in a single view.

— Names exact schema fields, auth method and UI location; only a runnable command would score higher.

launch-20261001-0563719b

> announced · Oct 1, 2026 BLOG ↗
>
> Cyera Agent Guardian now pushes data security risk signals for AI agents into Microsoft Agent 365's Agent Registry in near real time.
>
> Integrates Cyera Agent Guardian with Microsoft Agent 365, pushing policy-driven risk events to the Agent 365 Risk Service via Microsoft's API in near real time, authenticated through Microsoft Entra ID using OAuth 2.0 client credentials.
>
> Maps Cyera policy severity to the Agent 365 third-party risk schema fields `riskLevel` (low/medium/high), `riskEvidence` (policy name and description), and `status` (driven by Issue status), so findings appear directly in the Agent Registry page of the Microsoft 365 Admin Center.
>
> Delivers agent data security posture signals — including agent configurations, capabilities, connected tools, access paths to sensitive data, and drift from policy — alongside Microsoft Purview, Microsoft Defender, and Microsoft Entra risk signals in a single view.

WAS THIS USEFUL? 👍 👎

Bold Security

Feedback ⓘ announced · Sep 24, 2026 BLOG ↗

RANK

Bold Security runs a local AI agent on every endpoint that blocks insider threats, data leakage and shadow-AI use.

Bold now feeds endpoint coverage status into Axonius asset inventory, letting teams surface devices where Bold is missing, inactive, or outdated via Axonius queries.

NEW INTEGRATION Adds a Bold + Axonius integration that feeds Bold device coverage data into the Axonius asset inventory, enabling visibility into which endpoints have Bold deployed, inactive, or running out-of-date versions alongside the rest of the asset record.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

ScoutDNS

Feedback ⓘ

Insights tables keep their headers and show when they're loading NOTES ↗

RANK

DNS filtering and threat protection for businesses, MSPs, schools, and nonprofits. Block malware, phishing, and unwanted content at the DNS layer.

> ScoutDNS refined its Insights table UI with clearer loading states, error handling, and a reset control, while also fixing sorting to avoid unnecessary reloads.

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

01 Insights table loading, error, and reset UX IMPROVED

69

Insights tables now show a Reset view button in the toolbar that returns to the Domain tab with Last Hour, 50 rows, and no filters or search, with Browser Back restoring the previous view. A loading bar appears under column headers with dimmed, unclickable rows while a new Insights view loads, preventing accidental drill-downs into stale results. Empty tables show 'No results for this view.', failed queries show 'This view didn't load.' with a Refresh prompt, and a Retry button appears when a page of logs fails to load, pausing further log loading until clicked. Sorting an Insights table no longer triggers a full table reload.

Return Insights to its default state after drilling into filtered results, without losing the ability to go back via the browser.

> 📍 In Insights, click Reset view in the toolbar to return to the Domain tab with Last Hour, 50 rows, and no filters or search. To undo, click Browser Back.

— Names UI controls and messages but no API or config surface

Insights tables keep their headers and …

Insights tables keep their headers and show when they're loading NOTES ↗

LATEST STABLE

ScoutDNS Insights tables gain sticky headers, a Reset view button, inline load indicators, and clearer empty/error states.

└──▷ HOW TO FIND IT

Return Insights to its default state after drilling into filtered results, without losing the ability to go back via the browser.

📍 In Insights, click Reset view in the toolbar to return to the Domain tab with Last Hour, 50 rows, and no filters or search. To undo, click Browser Back.

Adds a `Reset view` button in the Insights toolbar that returns to the Domain tab with Last Hour, 50 rows, and no filters or search; Browser Back restores the previous view.

Shows a loading bar under column headers with dimmed, unclickable rows while a new Insights view loads, preventing accidental drill-downs into stale results.

Displays 'No results for this view.' when a table is empty and 'This view didn't load.' with a `Refresh` prompt on query failure.

Adds a `Retry` button when a page of logs fails to load, pausing further log loading until clicked.

Sorting an Insights table no longer triggers a full table reload.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.28.991 (2026-10-02) BLOG ↗ RANK

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

> Runlayer's biggest shift this window is the launch of Sessions & Shadow AI — a cross-platform AI Watch endpoint client plus proactive scanning and blocking of risky skills and MCP servers — alongside new agent delegation approval controls, a Cursor billing integration, and required Rev decision-model configuration for the Operator.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 topicCpu.rev required for Operator Procedures

BREAKING

85

The Runlayer Operator now requires `topicCpu.rev` to be configured for Procedures; creating or validating a Procedure without it returns a 503 error. Admins can enable the self-hosted Rev decision model with `topicCpu.rev.enabled: true`, `topicCpu.rev.gunicornWorkers`, and the `images.topicCpuRev` image reference so risk scoring runs locally instead of calling an external service.

Enable the self-hosted Rev decision model so the Runlayer Operator runs local risk scoring instead of calling an external service.

yaml

```yaml
images:
  topicCpuRev: REGISTRY/toolguard:runlayer-rev-guard-TAG

topicCpu:
  rev:
    enabled: true
    gunicornWorkers: 1
```

— Exact config keys and before/after 503 behavior given. product docs

02 Environment variable inheritance for local stdio connectors

NEW

78

Local stdio MCP connectors can list named environment variables such as `GH_TOKEN` and `NODE_OPTIONS` that `runlayer run` copies from each user's shell into the MCP process at runtime; the values themselves are never sent to Runlayer.

Propagate a developer's local GitHub token into an MCP without storing the value in Runlayer — ideal for local stdio connectors that need per-user credentials.

> 📍 On the local stdio connector configuration, add `GH_TOKEN` to the list of named environment variables to inherit; `runlayer run` will then copy the value from each user's shell at runtime.

— Named variables and exact inheritance mechanism described. `1.28.991 (2026-10-02)`

---

03 **AgentCore sandbox role loses AssumeRole** BREAKING 74

On AgentCore, the agent sandbox AWS role no longer has `sts:AssumeRole` permission — it is now limited to pulling the sandbox image and writing runtime logs; agent code that relied on `sts:AssumeRole` inside the sandbox will fail.

— Exact permission and before/after given, no migration step. `1.28.991 (2026-10-02)`

---

04 **AI Watch endpoint client and deployment coverage** NEW 73

AI Watch is a lightweight endpoint client that detects and logs all employee AI activity across clients, agents, MCPs, skills, and tool calls on Mac, Windows, and Linux. It ships with a browser extension covering Google Chrome, Mozilla Firefox, and Microsoft Edge for browser-based AI sessions, and supports enterprise rollout via MDM providers including Jamf, Workspace ONE, and Intune.

— Lists every OS, browser, and MDM target named in source. `launch-20261001-2ac5322f`

---

05 **Cursor billing integration** NEW 71

Workspace admins can connect or disconnect a Cursor Team or Enterprise billing API key from Settings → Integrations; Runlayer verifies the key before saving and erases it on disconnect.

Connect a Cursor Enterprise account to sync billing data into Runlayer for unified usage tracking.

> 📍 In the Runlayer console, go to Settings → Integrations, paste the Cursor Enterprise billing API key and the matching Cursor organization ID, then click Connect. Runlayer verifies the key before saving.

— Exact UI path and connect/verify/erase mechanism given. `1.28.991 (2026-10-02)`

---

06 **Proactive threat detection and blocking for skills and MCP servers**

NEW 69

Skill files and MCP servers are scanned proactively at installation time, before execution, to detect prompt injection, data exfiltration, privilege escalation, and guardrail bypass techniques. Deterministic policies can then block high-risk skills and tool calls proactively, preventing execution of malicious AI artifacts.

— Names threat types and blocking mechanism, no exact UI path. launch-20261001-2ac5322f

07 **Review agent delegation approval policy** NEW 67

A new Review agent delegation policy template requires human approval for every agent-to-agent run, including scheduled ones, in workspaces with approval policies enabled, applied from Settings → Policies in the Runlayer console.

Enforce human approval on every agent-to-agent handoff, including scheduled runs, to satisfy change-control requirements.

📍 In the Runlayer console, go to Settings → Policies and add a new policy using the Review agent delegation template, then save and apply it to the workspace.

— UI path given, enforcement scope specified. 1.28.991 (2026-10-02)

08 **OTLP observability sidecar configuration** NEW 60

Operator deployments can configure an OTLP observability sidecar via `observability.otlp.endpoints`, specifying a name, protocol, and endpoint URL to export telemetry health metrics to a reachable collector.

Configure the OTLP observability sidecar to export telemetry health metrics to a reachable collector endpoint.

yaml

```yaml
observability:
  otlp:
    endpoints:
      - name: runlayer-metrics
        protocol: http
        endpoint: http://otel-collector.monitoring.svc:4318
```

— Config key named but no further mechanism described. product docs

THINNER COVERAGE BELOW

09 **Full agent session logs with turn-by-turn visibility** NEW 56

Runlayer provides full agent session logs with turn-by-turn visibility across user inputs, AI reasoning steps, MCP tool calls, and actions, with in-line alerts on high-risk behavior.

— Describes scope of logging but no access path given. launch-20261001-2ac5322f

10 **Agent delegation toggle** NEW 46

Each agent's settings now include a Can hand work to other agents toggle that controls whether the agent can discover and delegate tasks to other agents.

— Names the toggle but no further mechanism detail. 1.28.991 (2026-10-02)

11 **Explanations for bypassed blocks** IMPROVED 46

Tool calls marked Block bypassed in a session now include plain-language explanations describing why the block did not take effect and how to stop similar calls.

— Explains behavior change but no UI path or example. 1.28.991 (2026-10-02)

12 **Skill usage analytics via connector** NEW 39

The Runlayer connector exposes skill usage totals, daily trends, and top users by asking connected agents for this data.

— States what's exposed but not how to access it. 1.28.991 (2026-10-02)

└──▷ BREAKING ON UPGRADE

! On AgentCore, the agent sandbox AWS role no longer has `sts:AssumeRole` permission — it is now limited to pulling the sandbox image and writing runtime logs; any agent code that relied on `sts:AssumeRole` inside the sandbox will fail.

1.28.991 (2026-10-02) BLOG ↗ LATEST STABLE

Runlayer 1.28.991 adds agent delegation approval policies, Cursor billing integration, and environment variable inheritance for local MCPs.

└──▷ HOW TO FIND IT

Enforce human approval on every agent-to-agent handoff, including scheduled runs, to satisfy change-control requirements.

📍 In the Runlayer console, go to Settings → Policies and add a new policy using the Review agent delegation template, then save and apply it to the workspace.

Propagate a developer's local GitHub token into an MCP without storing the value in Runlayer — ideal for local stdio connectors that need per-user credentials.

📍 On the local stdio connector configuration, add `GH_TOKEN` to the list of named environment variables to inherit; `runlayer run` will then copy the value from each user's shell at runtime.

Connect a Cursor Enterprise account to sync billing data into Runlayer for unified usage tracking.

📍 In the Runlayer console, go to Settings → Integrations, paste the Cursor Enterprise billing API key and the matching Cursor organization ID, then click Connect. Runlayer

verifies the key before saving.

Adds `GH_TOKEN`, `NODE_OPTIONS`, and other named environment variables to local stdio connectors so admins can list variables that `runlayer run` copies from each user's shell into the MCP process — values are never sent to Runlayer.

Adds a Review agent delegation policy template requiring approval for every agent-to-agent run, including scheduled ones, in workspaces with approval policies.

Enables workspace admins to connect or disconnect a Cursor Team or Enterprise billing API key from Settings → Integrations, with key verification before saving and key erasure on disconnect.

Adds a Can hand work to other agents toggle in each agent's settings to control whether it can discover and delegate tasks to other agents.

Exposes skill usage totals, daily trends, and top users through the Runlayer connector by asking connected agents.

+1 more

└─▷ **BREAKING ON UPGRADE**

! On AgentCore, the agent sandbox AWS role no longer has `sts:AssumeRole` permission — it is now limited to pulling the sandbox image and writing runtime logs; any agent code that relied on `sts:AssumeRole` inside the sandbox will fail.

announced · Oct 1, 2026 BLOG ↗

Runlayer launches Sessions & Shadow AI: proactive skill/MCP scanning, AI Watch client, and browser extension across Mac, Windows, and Linux

Introduces AI Watch, a lightweight endpoint client that detects and logs all employee AI activity across clients, agents, MCPs, skills, and tool calls on Mac, Windows, and Linux.

Adds proactive scanning of skill files and MCP servers at installation time — before execution — detecting prompt injection, data exfiltration, privilege escalation, and guardrail bypass techniques.

Ships a browser extension with AI Watch to cover browser-based AI sessions, supporting Google Chrome, Mozilla Firefox, and Microsoft Edge.

Supports deployment of AI Watch via MDM providers including Jamf, Workspace ONE, and Intune for enterprise-scale rollout.

Provides full agent session logs with turn-by-turn visibility across user inputs, AI reasoning steps, MCP tool calls, and actions, with in-line alerts on high-risk behavior.

+1 more

WAS THIS USEFUL?  

GOVERN

◆ AI/LLM Security

Zenity

Feedback ⓘ announced · Oct 3, 2026 BLOG ↗ RANK

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

> Zenity launched AI Total, a free sandbox detonation service that reveals the real runtime behavior of AI agent skills before they're installed.

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

01 **AI Total sandbox detonation service** NEW 81

AI Total is a free service at `aitotal.io` that detonates AI agent skills inside a live sandboxed agent, recording every domain contacted, package pulled, file touched, command run, and tool call to reveal runtime behavior hidden from static analysis. The sandbox is seeded with planted credentials and sensitive files as bait (e.g. `~/.aws/credentials`) so skills that hunt for secrets or phone home to unknown domains expose themselves during detonation rather than in production. It supports skills used with Claude Code, OpenClaw, and other agent platforms that load skills on demand.

Before adding a public skill from a registry to your Claude Code or OpenClaw agent, submit it to AI Total to see exactly what it does at runtime — domains it contacts, files it touches, and commands it runs.

> 📍 Visit `https://aitotal.io`, submit the skill package URL or file, and review the behavioral evidence report before installing.

— Names mechanism, bait technique, platforms, and a URL to act on

launch-20261003-59f3ed6f

announced · Oct 3, 2026 BLOG ↗

Zenity launches AI Total, a free sandbox detonation service that runs AI agent skills and reports their actual runtime behavior before deployment.

└──▷ HOW TO FIND IT

Before adding a public skill from a registry to your Claude Code or OpenClaw agent, submit it to AI Total to see exactly what it does at runtime — domains it contacts, files it touches, and commands it runs.

> 📍 Visit `https://aitotal.io`, submit the skill package URL or file, and review the behavioral evidence report before installing.

> Introduces AI Total, a free service at `aitotal.io` that detonates AI agent skills inside a live sandboxed agent — recording every domain contacted, package pulled, file touched, command run, and tool call — to reveal runtime behavior hidden from static analysis.
>
> Seeds the sandbox with planted credentials and sensitive files as bait, so skills that hunt for secrets (e.g. reading `~/.aws/credentials` or phoning home to unknown domains) expose themselves during detonation rather than in production.
>
> Supports skills used with Claude Code, OpenClaw, and other agent platforms that load skills on demand, making pre-install detonation available across the growing agent skill ecosystem.

WAS THIS USEFUL? 👍 👎

CrowdStrike Falcon AI Detection and Response

Feedback ⓘ changelog · Oct 3, 2026 API ↗ RANK

CrowdStrike Falcon uses AI to detect, investigate, and respond to cyber threats across endpoints and cloud environments in real-time.

> CrowdStrike Falcon expands into AI security, adding a dedicated AI Detection and Response capability for AI workloads alongside new surfaces for building, managing, and invoking AI agents and skills.

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

01 **AI agent and skill building platform** NEW 52

Adds a new Agents surface for building and managing AI agents, a Skills surface for defining and invoking AI skills, Agent Invocation capability to trigger agents programmatically, and knowledge base management to support agents and skills with curated content.

— Names four new surfaces but no commands or config shown 2026-10-03-c6b7e13d

02 **AI Detection and Response for AI workloads** NEW 35

Adds `AIDR` (AI Detection and Response) capability to detect and respond to threats targeting AI workloads.

— Names the capability but no mechanism or scope detail 2026-10-03-c6b7e13d

> changelog · Oct 3, 2026 API ↗ LATEST STABLE
>
> CrowdStrike Falcon adds AI agent/skill building, knowledge base management, and AI workload threat detection (AIDR).
>
> Adds AI Detection and Response (`AIDR`) capability to detect and respond to threats targeting AI workloads.
>
> Adds support for building and managing AI agents via a new Agents surface.
>
> Adds support for defining and invoking AI skills via a new Skills surface.
>
> Adds Agent Invocation capability to trigger AI agents programmatically.
>
> Adds knowledge base management to support AI agents and skills with curated content.

WAS THIS USEFUL? 👍 👎

ai-safe2-framework

Feedback ⓘ

2026-10-02_AISM_Scoring_Update NOTES ↗ RANK

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

> AI SAFE² introduces the AI Sovereign Maturity (AISM) Model, adding a Technology Contribution Profile and Evidence Assurance framework alongside a reusable Technology Card template and new validation/incident guidance for scoped AI control reviews.

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

01 **Reusable Technology Card template** NEW 65

A new template capturing `version`, `scope`, `controls`, `mechanisms`, `evidence`, `failure behavior`, `retained authority`, and `adoption conditions` for scoped AI control reviews.

— Enumerates every field the template captures, directly usable

2026-10-02_AISM_Scoring_Update

THINNER COVERAGE BELOW

02 **AISM Technology Contribution Profile v1.0** NEW 55

A new structured review method for what an artifact addresses, implements, enforces, validates, evidences, or challenges, scoped within a stated boundary.

— Names the framework and its review categories, no worked example

2026-10-02_AISM_Scoring_Update

03 **Validation and incident guidance documents** NEW 55

New documents cover clean/attacked pairs, side-effect and bypass tests, policy-integrity checks, reconstruction, and notification ownership for incident response.

— Names specific test and guidance types but not full procedures

2026-10-02_AISM_Scoring_Update

04 **AI SAFE² Evidence Assurance v1.0** NEW 50

Attributes evidence strength, origin, independence, and limitations to each individual claim made about a technology, separating claim-level evidence quality from broader scoring.

— Describes mechanism but no applied example or template shown

2026-10-02_AISM_Scoring_Update

05 **AI Sovereign Maturity (AISM) Model rebrand and navigation**

IMPROVED 40

Adopts the name AI Sovereign Maturity (AISM) Model, adding dated AISM change summaries and repository-wide navigation connecting methods to pillars, implementation guidance, research, Challenge Lab, and agent discovery.

— Names linked areas but navigation structure itself undescribed

2026-10-02_AISM_Scoring_Update

06 **Sovereignty Score scope clarification** IMPROVED 30

Clarifies the existing AISM Sovereignty Score to explicitly keep organizational maturity separate from technology contribution and claim-level evidence assurance.

— States the distinction but no mechanism for how it's enforced

2026-10-02_AISM_Scoring_Update

07 **Nine source-attributed review candidates** NEW 20

Adds nine source-attributed review candidates intended to frame scoped acceptance work without adopting unsupported scores.

— Bare mention, no candidates or criteria named 2026-10-02_AISM_Scoring_Update

2026-10-02_AISM_Scoring_Update NOTES ↗ LATEST STABLE

AI SAFE² adds Technology Contribution Profile v1.0, Evidence Assurance v1.0, and a reusable Technology Card for scoped AI control reviews.

└──▷ GET THIS VERSION

```
$ git clone --branch 2026-10-02_AISM_Scoring_Update
https://github.com/CyberStrategyInstitute/ai-safe2-framework.git
# already have the repo? check out this version:
$ git checkout 2026-10-02_AISM_Scoring_Update
```

Introduces AISM Technology Contribution Profile v1.0 to structure review of what an artifact addresses, implements, enforces, validates, evidences, or challenges within a stated boundary.

Introduces AI SAFE² Evidence Assurance v1.0 to attribute evidence strength, origin, independence, and limitations to each individual claim.

Adds a reusable Technology Card template capturing version, scope, controls, mechanisms, evidence, failure behavior, retained authority, and adoption conditions.

Adds validation and incident guidance documents covering clean/attacked pairs, side-effect and bypass tests, policy-integrity checks, reconstruction, and notification

ownership.

Adds nine source-attributed review candidates to frame scoped acceptance work without adopting unsupported scores.

+2 more

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ announced · Oct 1, 2026 BLOG ↗

RANK

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

> Docker launched Cloud Sandboxes, running agent tasks in isolated microVMs on Docker-managed cloud compute, alongside a new open Sandbox Kit specification, the `sbx` CLI, and runtime governance features like credential proxying and audit logging.

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

01 **Sandbox Kit specification and governance** NEW 70

The Docker Sandbox Kit specification is an OCI image format containing an agent, its tools, and declared access requirements, enabling teams to build, push, pull, scan, and pin agent environments by digest. It is published under Apache 2.0 with a commitment to CNCF neutral governance, with Docker Sandboxes as the first runtime implementation.

– Describes format and licensing but no file/API details launch-20261001-8151baca

02 **`sbx` CLI for sandbox lifecycle** NEW 67

The `sbx` command-line tool starts, moves, and retrieves sandboxes between local and cloud environments, allowing longer agent tasks to continue after a laptop is closed.

– Named CLI with concrete commands but no syntax shown launch-20261001-8151baca

THINNER COVERAGE BELOW

03 **`sbxenv` files for environment configuration** NEW 57

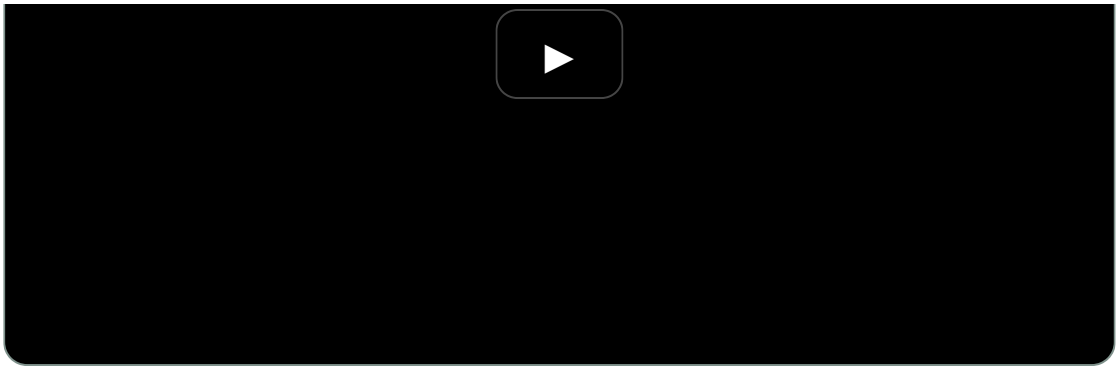
`sbxenv` files describe a set of sandbox environments as configuration, allowing agent environments to be reproduced consistently across machines.

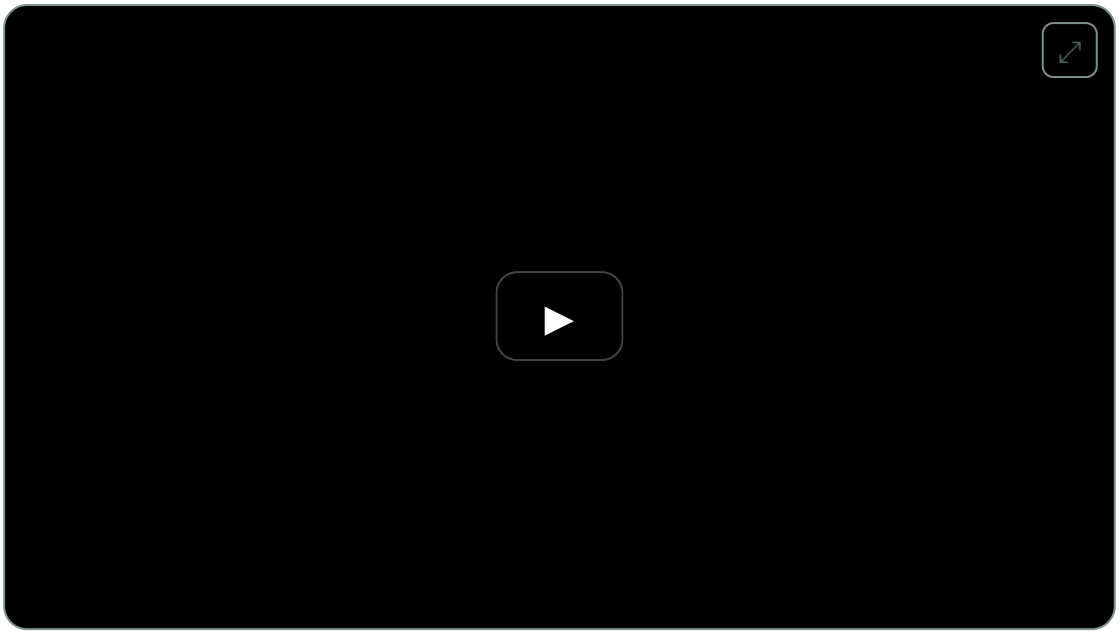
– Named config file format but no schema shown launch-20261001-8151baca

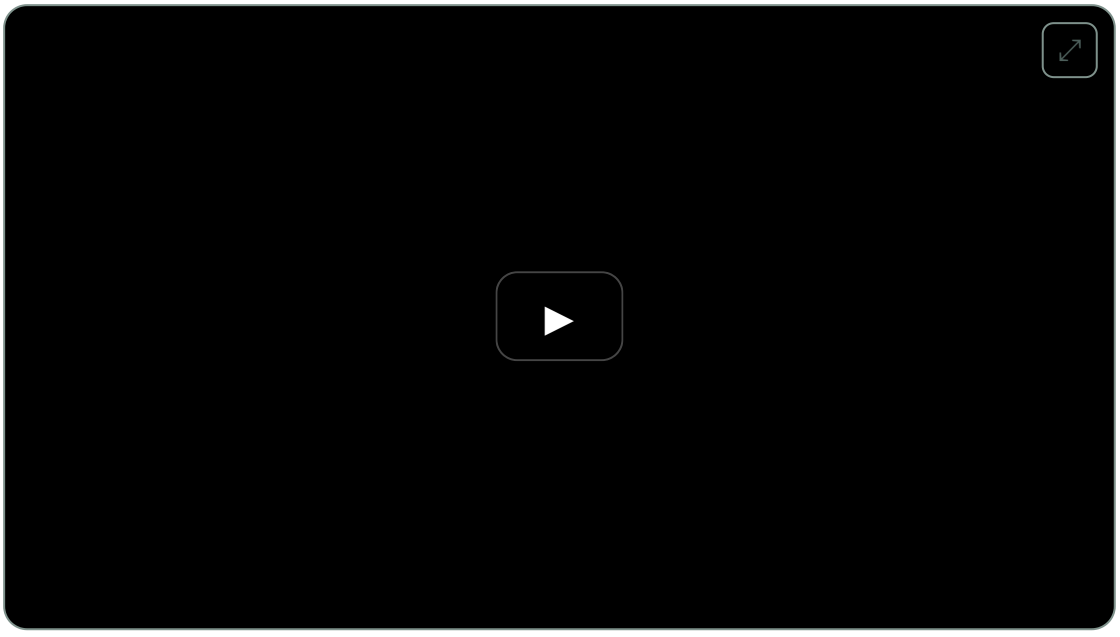
04 **Docker Cloud Sandboxes microVM isolation** NEW 55

Docker Cloud Sandboxes runs each agent task in an isolated microVM with its own kernel and Docker daemon on Docker-managed cloud compute, billed by the second.









— Names mechanism and billing but no setup steps launch-20261001-8151baca

05 **Runtime proxy for credential injection** NEW 50

A runtime proxy can supply credentials to a sandbox without exposing the raw API key inside the agent environment.

— Explains mechanism but no config surface named launch-20261001-8151baca

06 **Audit view for sandbox policy decisions** NEW 45

An audit view records policy decisions, including blocked network destinations, so developers can inspect agent activity after a task completes.

— Describes output but no UI path or API given launch-20261001-8151baca

07 **MCP support for sandboxed agents** NEW 40

Docker Sandboxes support the Model Context Protocol (MCP) for connecting tools to agents running inside sandboxes.

— Names protocol but no integration detail launch-20261001-8151baca

announced · Oct 1, 2026 BLOG ↗

Docker Cloud Sandboxes launch: run agent tasks in isolated microVMs on Docker-managed cloud compute using the `sbx` CLI

Introduces Docker Cloud Sandboxes, which run each agent task in an isolated microVM with its own kernel and Docker daemon on Docker-managed cloud compute, billed by the second.

Supports the `sbx` command-line tool for starting, moving, and retrieving sandboxes between local and cloud environments, so longer agent tasks can continue after closing a laptop.

Introduces the Docker Sandbox Kit specification — an OCI image format containing an agent, its tools, and declared access requirements — enabling teams to build, push, pull, scan, and pin agent environments by digest.

Supports `sbxenv` files, which describe a set of sandbox environments as configuration for reproducing agent environments across machines.

Publishes the Sandbox Kit specification under Apache 2.0 and commits to CNCF neutral governance, with Docker Sandboxes as the first runtime implementation.

+3 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 3, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue