THE CYBER TOOLCHAIN

NO. 143

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 18, 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

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 263 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 80 of 81 tools

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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

Three unrelated things stand out. StepSecurity's Secure Registry now hardens npm/pnpm/Bun installs for free, closing the install-time

vim ? help

window where a malicious postinstall script runs with your CI credentials. GitGuardian added 13 new detectors and 20+ analyzers, so leaked Supabase, Jamf and Grafana keys surface instead of sliding past. And Docker Engine 29.10.0 stops implicitly trusting loopback registries.

PROTECT

Stop a malicious npm postinstall before it reaches your CI runner

StepSecurity's free community Secure Registry routes npm/pnpm/Bun installs through Harden-Runner, which is the only point where a dependency's install script can be contained before it touches runner credentials; a new Critical control also audits whether fork PR workflows can run without approval, the classic pwn-request path. Chainguard separately disclosed and fixed a batch of silently-patched Java vulnerabilities, which are the ones that never get a CVE and so never show up in your scanner.

StepSecurity · Chainguard

PROTECT

Catch leaked keys for the SaaS you actually run, not just the big three clouds

13 new detectors and 20+ new analyzers extend coverage to Jamf, Supabase, Grafana and Cloudflare alongside Google and AWS. Analyzers matter more than detectors here: they tell you whether a found token is live and what it can reach, so triage stops being a manual curl against someone else's API.

GitGuardian Detection Engine

RESPOND

Find the one case from six months ago without remembering its title

Tines added full-text case search with nested boolean filters, which turns a closed-case archive into something you can actually interrogate during a repeat incident instead of scrolling. Dradis Pro v5.5 adds AIVSS-SSVC and STRIDE calculators plus concurrent-edit warnings, so two testers writing the same finding stop silently overwriting each other.

Tines · Dradis

PROTECT

Give a sandboxed coding agent internet access through one allowed host

> The Firecrawl kit lets an agent search, scrape and crawl via a tightly scoped proxy rather than opening general egress — the difference between an agent that can read the web and one that can exfiltrate to anywhere. minimal's Box Egress Proxy takes the same posture for its boxes, with certificate-based and audited network control.
>
> Docker Sandboxes

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗ RANK

Hacktron is a review, Whitebox, the API, and MCP.

> Hacktron's latest release adds compliance-ready reports for clean pentests, session-cookie login for the pentest auth agent, and per-severity SLA tracking in PR and scan views, alongside smaller fixes to integration connection handling and Slack notifications.

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

01 SLA compliance ring per severity

Not checked yet NEW

75

PR review and scan tabs now show an SLA compliance ring per severity, with MTTR and a Past SLA count that links directly to overdue findings.

Track SLA adherence per severity across all scanned pull requests without leaving the scan tab.

📍 Open the PR review or scan tab and view the SLA compliance ring for each severity; click the Past SLA count to filter to overdue findings.

— Names exact UI labels and gives navigation to use it. changelog-20261118-73a2dc19

02 Zero-findings pentest report downloads

Not checked yet NEW

65

Completed pentests with no confirmed findings now offer Full and Executive report downloads that include scope, methodology, zero finding counts, and an explicit no-findings statement, giving teams audit-trail evidence even when nothing was found.

Download a full audit-trail report after a pentest scope returns no confirmed findings — useful for compliance evidence.

📍 In the pentest report panel, click Full or Executive under report downloads to get a report that includes scope, methodology, zero finding counts, and a no-findings statement.

— Names both report types and their contents clearly. changelog-20261118-73a2dc19

THINNER COVERAGE BELOW

03 Session-cookie authentication for pentest login agent

Not checked yet NEW

55

Pentest auth instructions now accept a pasted session cookie, letting the login agent sign in to a target without a password — useful for session-based SSO targets.

Authenticate the login agent against a target that uses session-based SSO where you cannot supply a plain password.

> 📍 In the pentest auth instructions form, paste your captured session cookie into the session cookie field instead of entering a username and password.

— Clear mechanism and usage step, but no field name given. changelog-20261118-73a2dc19

---

04 **Connection form improvements for GitLab, Azure DevOps, GHES**

Not checked yet IMPROVED 50

GitLab, Azure DevOps, and GitHub Enterprise Server connection forms now show collapsible Setup instructions and retain any connection error until you edit a field or retry.

— Names the UI element and three providers, no deeper mechanism.

changelog-20261118-73a2dc19

---

05 **Correct author attribution for GitLab and Bitbucket scans**

Not checked yet IMPROVED 45

GitLab merge request and Bitbucket pull request scans now show the MR/PR author as the creator in Slack finding messages and the scan list, instead of whoever connected the integration.

— Explains before/after but offers no action step. changelog-20261118-73a2dc19

---

06 **Revoked integration flagging and reconnect** Not checked yet

IMPROVED 35

A revoked integration connection is now flagged and can be reconnected directly from the integration sheet.

— Brief description, names the integration sheet only. changelog-20261118-73a2dc19

---

07 **Slack thread completion reaction swap** Not checked yet IMPROVED

30

Slack threads now swap the eyes reaction for a check mark when Hacktron finishes a thread, and switch back to eyes for follow-ups.

— Cosmetic change with no further mechanism or action. changelog-20261118-73a2dc19

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

NetSPI OCInferno

Feedback ⓘ changes since 2026-11-17 CODE ↗

OCInferno adds a `process_download_cached` module and renames `--limit` to `--object-limit` with a 5,000-object default for object storage enumeration

└──▷ TRY IT

Recover downloads for all object storage buckets already enumerated in a previous run, without re-hitting live OCI APIs.

```
$ python -m ocinferno process_download_cached --download buckets
```

Enumerate object storage but cap at 500 objects per bucket to speed up broad tenant discovery scans.

```
$ python -m ocinferno enum_objectstorage --object-limit 500
```

Re-download all content-category artifacts (objects, vault secrets, artifacts) from DB without re-enumerating, excluding vault secrets.

```
$ python -m ocinferno process_download_cached --download content --not-downloads vault_secrets
```

Re-download enumerated content from the workspace DB without re-running live enumeration using the new `process_download_cached` module, which accepts the same `--download` and `--not-downloads` token system as `enum_all` (categories: `content`, `metadata`, `buckets`, `vault_secrets`, and others). Not checked yet

Limit object storage enumeration per bucket with `--object-limit` (default: 5,000; set to 0 for no limit), replacing the old `--limit` flag. Not checked yet

OpenGraph pathfinder now shows the granting policy name, compartment, and statement at each hop in a privilege-escalation path. Not checked yet

OpenGraph default priv-esc targets now narrow to `OCIAllResources` nodes plus `OCI_CAN_CREATE_POLICY`, `OCI_CREATE_POLICY`, and `OCI_UPDATE_POLICY` policy edges, with IDD admin-only targets; family-specific `OCI_MANAGE` targets are excluded by default but still traversed as intermediate hops. Not checked yet

Object storage enumeration now shows per-bucket progress output during scans. Not checked yet

└──▷ **BREAKING ON UPGRADE**

! The `--limit` flag in `enum_objectstorage` is renamed to `--object-limit` and its default changes from 0 (no limit) to 5,000 objects per bucket; existing scripts using `--limit` or relying on unlimited enumeration will break.

WAS THIS USEFUL?  

pwntools

Feedback ⓘ 4.15.1 NOTES ↗ RANK

Pwntools is a CTF framework and exploit development library

> pwntools 4.15.1 adds a persistent config-level I/O timeout, fixes silent truncation in libc.rip lookups, extends Python 3.14 support, and adds big-endian AArch64 emulation support.

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

01 Persistent context.timeout in pwn.conf

Not checked yet NEW

80

You can now set a persistent `context.timeout` default in `~/.pwn.conf` instead of setting it per-script, avoiding silent hangs during exploit development.

Persist a default I/O timeout for all scripts so you never hit a silent hang during exploit development.

ini

```
[context]
timeout = 10
```

— Named config key and file with runnable example. 4.15.1

THINNER COVERAGE BELOW

02 Paginated libc.rip lookups in libcdb

Not checked yet IMPROVED

55

pwntools now supports paginated libc.rip API requests, so `libcdb` lookups no longer silently truncate results against large databases.

— Explains before/after behaviour but no usage example. 4.15.1

03 Python 3.14 support via relaxed Unicorn constraint

Not checked yet IMPROVED

50

pwntools now runs on Python 3.14 by relaxing the Unicorn Engine version constraint to allow Unicorn 2.1.3 and 2.1.4; note that MIPS emulation raises at runtime under this change.

— Names versions and caveat but no action to take. 4.15.1

04 aarch64_be QEMU user-mode mapping

Not checked yet NEW

40

Adds `aarch64_be` QEMU user-mode mapping, enabling big-endian AArch64 process emulation.

— Names the mapping but gives no usage detail. 4.15.1

WAS THIS USEFUL?  

Dradis

Feedback ⓘ New in Dradis Pro v5.5 (2026-11-17) BLOG ↗ RANK

Dradis is a pentest reporting for consultancies that never put a client

> Dradis Pro v5.5 is the headline release, adding per-project Issue Library translations with a new v2 API, AIVSS-SSVC and STRIDE calculators for scoring findings and threats, concurrent-edit warnings, expanded webhook coverage, and Boards/Evidence API additions, while a docs update added colored QA-state badges and an in-app release notification.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 **Per-project Issue Library translations and v2 API** Not checked yet

NEW 90

Adds per-project Issue Library locale translation, set on create or edit, used by Rules Engine matching and report imports. Adds the new `v2` Issue Library API for accessing translations and entry reading by locale, plus `IssueLibrary` API scopes for Personal Access Tokens (PAT) to control access to the new v2 Issue Library endpoints.

— Names API version and PAT scopes but no exact endpoint paths.

New in Dradis Pro v5.5 (2026-11-17)

02 **New threat and risk calculators for Issues** Not checked yet NEW

80

Adds the AIVSS-SSVC calculator for scoring findings in agentic AI systems, letting users select a remediation outcome of Defer, Scheduled, Out-of-Cycle, or Immediate, and the STRIDE calculator for classifying threats by Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege, both available from the Issue view.

Score an agentic AI finding and pick a remediation outcome using the new AIVSS-SSVC calculator.

> 📍 In Dradis, open an Issue, select `AIVSS-SSVC` from its calculators, complete the scoring fields, and choose a remediation outcome of Defer, Scheduled, Out-of-Cycle, or Immediate.

— Usage example gives exact steps to use the calculator. Names all six categories but gives no walkthrough.

New in Dradis Pro v5.5 (2026-11-17)

03 **Remediation Tracker ticket content and linked fields**

Not checked yet IMPROVED 70

Remediation Tracker tickets now render Liquid Dynamic Content to final text when an Issue is sent, and each ticket gains Project, Issue, and Evidence fields with labeled links back to the source plus an evidence list.

— Describes rendering and field additions, no UI path.

New in Dradis Pro v5.5 (2026-11-17)

---

04 **Expanded webhook event coverage** Not checked yet IMPROVED

70

Webhooks now fire on Board, List, and Card CRUD events; Content Block, Evidence, Node, and Note CRUD events; Issue and Content Block state changes; inline thread lifecycle events; Comment and Echo Session creation; Personal Access Token lifecycle events; and Evidence QA state change events.

— Lists every event type but no webhook config instructions.

New in Dradis Pro v5.5 (2026-11-17)

---

05 **Evidence state field in API v3 Evidence endpoint** Not checked yet

NEW

70

Select Evidence state on create or update via the API v3 Evidence endpoint, and receive state in responses.

— Names exact endpoint and field but omits method/path.

New in Dradis Pro v5.5 (2026-11-17)

---

06 **Card and list repositioning via Boards API** Not checked yet NEW

70

Adds card and list repositioning within and between boards through the Boards API, along with Personal Access Token (PAT) update scopes supporting this card and list positioning.

— Names Boards API and PAT scopes but no endpoint path.

New in Dradis Pro v5.5 (2026-11-17)

---

07 **Concurrent edit warning across record types** Not checked yet NEW

60

Dradis now warns you when a teammate is already editing the record you have opened — covering Issues, Evidence, Notes, Content Blocks, Cards, Issue Library entries, and Remediation Tracker tickets.

— Lists all covered record types, no config detail. New in Dradis Pro v5.5 (2026-11-17)

THINNER COVERAGE BELOW

---

08 **Configurable language for Roslin's LanguageTool checks**

Not checked yet IMPROVED

45

Change the grammar and spell-checking language for Roslin's LanguageTool integration in Echo.

– Names Roslin and Echo but no exact setting location.

New in Dradis Pro v5.5 (2026-11-17)

09 **Colored QA-state badges for issues, evidence, content blocks**

Not checked yet NEW 40

Dradis now shows issue, evidence, and content block QA states as colored badges in tables and tabs.

– States what changed, no exact navigation given. product docs

10 **Dashboard grouping by report template** Not checked yet IMPROVED

35

Dashboard now supports grouping issues by report template.

– One-line addition, no further detail. New in Dradis Pro v5.5 (2026-11-17)

11 **Add node directly from node tree** Not checked yet NEW

35

Add a node directly from the node tree.

– Minimal description of a simple UI action. New in Dradis Pro v5.5 (2026-11-17)

12 **In-app notification for new Dradis releases** Not checked yet NEW

20

Dradis now notifies users when a new Dradis release is available.

– Single sentence, no detail on where or how. product docs

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Bugcrowd

Feedback ⓘ

Savant Vista General Availability (2.0) (2026-11-17) BLOG ↗ RANK

Bugcrowd is a crowdsourced security platform connecting organizations with ethical hackers to discover and manage vulnerabilities through bug bounty and penetration testing programs.

> Bugcrowd launched Savant Vista as a generally available, fully supported product for continuous external attack surface management, and rounded it out with bulk asset import, more granular RBAC, and improved Pathseeker target scoping.

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

01 **Savant Vista onboarding, access and scoping additions**

Not checked yet IMPROVED 41

Savant Vista now supports bulk asset import, enabling teams to onboard large asset inventories at once; adds more granular role-based access control for managing user permissions across the platform; and improves target scoping for Pathseeker testing.

— Three concrete additions named but no configuration detail given.

Savant Vista General Availability (2.0)…

02 **Savant Vista general availability** Not checked yet NEW

40

Bugcrowd adds Savant Vista as a generally available, fully supported product covering continuous external attack surface management, automated vulnerability scanning, and asset discovery.

— Names capability areas but no mechanism or setup steps.

Savant Vista General Availability (2.0)…

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Software Supply Chain Security

StepSecurity

Feedback ⓘ announced · Nov 18, 2026 BLOG ↗ RANK

StepSecurity secures GitHub Actions CI/CD pipelines by hardening workflows, detecting anomalies, and preventing supply chain attacks through automated security controls.

> StepSecurity shipped a free community tier for its Secure Registry that hardens npm/pnpm/Bun installs via Harden-Runner, a new Critical-rated control auditing fork PR workflow approval policies, and remediation tracking for Threat Center findings with new API endpoints and audit logging.

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

01 **Secure Registry community tier in Harden-Runner** Not checked yet

NEW SUPPORTED PLATFORM NEW 95

Enable Secure Registry on any job by adding `secure-registry: true` to a `step-security/harden-runner@v2` step; it automatically routes npm, pnpm, and Bun installs through `community.registry.stepsecurity.io/javascript/`. The free community tier (for public repositories on GitHub-hosted runners) applies a default two-day package cooldown and blocks known compromised packages. Customize behaviour via the `secure-registry` config block: set `npm.cooldown.days` (e.g. `days: 3`) and list exemptions such as `typescript@*` or `[email protected]`, or disable it for a job (e.g. one running `npm publish`) with `npm: mode: off`. Harden-Runner Insights also adds a Packages tab recording every package request's type, status, and evaluated controls.

Add Secure Registry protection to a CI job so all npm installs are checked for the two-day cooldown and known-compromised packages before they reach your build.

yaml

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: true
```

Extend the cooldown to three days and exempt a pinned TypeScript version so the registry holds other new releases but lets a known-good version through.

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: |
      npm:
        cooldown:
          days: 3
          exemptions:
            - "typescript@*"
            - "@typescript/*"
```

Turn off Secure Registry in a publish job so `npm publish` can reach the public registry directly without being intercepted.

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: |
      npm:
        mode: off
```

— Exact config keys, defaults, and runnable YAML given `launch-20261118-cefeb282`

02 **Remediation tracking in Threat Center** Not checked yet NEW

89

Mark individual Threat Center findings remediated from the row menu with an optional comment, or use Remediate all to clear every active finding in the current view at once. The incident API endpoint now returns each finding's remediation status, and new endpoints let you mark findings remediated or move them back to open, either for one organization or tenant-wide. Matched components and Matched IOCs tables show status, who remediated it, when, and their comment, with All, Active, and Remediated filter tabs and source-based search; incidents move to a new Remediated tab once every finding is remediated and move back automatically if a finding reappears, flagged with a 'Reappeared:' title prefix and a notification through existing threat intel channels. Every remediation action is recorded in the audit logs, a summary banner on each incident shows outstanding vs. resolved findings, and remediating requires Read & write access to the Threat Center permission.

Track remediation progress through the Threat Center UI, using the status tabs to separate active from resolved incidents.

> 📍 In Threat Center, go to the Affecting You tab and open an incident. Use the All / Active / Remediated tabs in the findings table to filter by status, then select Remediate all to close every active finding in the current view, or choose Remediate from an individual finding's row menu to add a comment before marking it done.

— Rich mechanism detail but API endpoints not named exactly

Threat Center: Remediation Tracking (20…

---

03 **Fork PR workflow approval control** Not checked yet NEW

85

A new 'Fork pull request workflows should require approval' control, rated Critical and listed under Code Injection in the Overview dashboard, checks every GitHub organization and public repository's fork PR workflow approval policy every hour and links directly to the fix at Settings > Actions > General in GitHub. Findings from the control are exposed through the existing control tools in the StepSecurity MCP (Model Context Protocol) server, and the control requires the StepSecurity GitHub App to have Administration: read permission on the checked organizations and repositories.

Review the new fork PR approval control findings for your organization via the StepSecurity Overview dashboard after granting the required permission.

> 📍 In StepSecurity, go to Overview dashboard › All Controls › Code Injection, then select 'Fork pull request workflows should require approval' to see which organizations and public repositories are on the least restrictive policy, and click Open Settings on any finding to jump directly to Settings › Actions › General in GitHub.

— Names dashboard path, permission and cadence, not an API launch-20261118-fff0dbd5

NEW SUPPORTED PLATFORM Conda package inventory and configuration audits

NEW SUPPORTED PLATFORM Java package inventory and configuration audits

NEW SUPPORTED PLATFORM Swift package inventory and configuration audits

NEW SUPPORTED PLATFORM Pi coding-agent package and extension inventory

NEW SUPPORTED PLATFORM Support for pooled and multi-session Windows virtual desktops

> announced · Nov 18, 2026 BLOG ↗

Harden-Runner agent v1.9.7 adds runtime monitoring and egress policy enforcement inside microVM-isolated ARC runner pods, covering network, DNS, process, and HTTPS events.

Set `runtimeClassName: isolated-sandbox` on runner pods to enable Harden-Runner monitoring and egress enforcement inside Amazon EKS Nitro Sandbox microVMs. Not checked yet

HTTPS visibility inside microVMs now covers requests from Go programs and HTTP/2 traffic — including calls from `gh` and other Go tools that carry their own TLS and bypass system OpenSSL — plus Docker-in-Docker image pulls. Not checked yet

Dev Machine Guard v1.19.1 (2026-11-17) BLOG ↗ STABLE

Dev Machine Guard adds stable device identity and per-session scanning for pooled virtual desktops, device tagging and retirement on uninstall, virtualization detection, and prompt-free agent configuration via `configure` flags.

Set `device_id` to `{hostname}` in agent config to keep a stable device record and history when a session host is rebuilt from an image, suiting Azure Virtual Desktop, Citrix, and Horizon host pools. Not checked yet

Exclude an imaging machine from scanning by running `configure --image-builder` or setting MSI property `IMAGEBUILDER=1` during deployment. Not checked yet

Label a device for its host pool using `device_tag`, and remove it from the dashboard and license count on teardown with `--retire` or by setting `retire_on_uninstall`. Not checked yet

Adds per-session scanning on Windows Enterprise multi-session and Remote Desktop Session Host servers, reporting each signed-in user's inventory and findings independently under the one device record for that machine. Not checked yet

Every scan now reports whether the machine is virtual or physical and which hypervisor or cloud it runs on. Not checked yet

+2 more

WAS THIS USEFUL?  

sbom-operator

Feedback ⓘ 0.50.0 NOTES ↗ RANK

Sbom-operator is a catalogue all images of a Kubernetes cluster to multiple targets with Syft

> sbom-operator's single release in this window adds per-scope Syft cataloger configuration and extends its Dependency-Track integration with workload identity federation and Sprig-based templated project names.

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

01 **Per-SBOM-scope Syft cataloger configuration** Not checked yet NEW

50

Syft scanning catalogers can now be configured per SBOM scope via the new `syft scanning catalogers` configuration, including per-sbom scope configuration.

— Names the config area but not an exact key path or file 0.50.0

02 **Sprig templating for Dependency-Track project names**

Not checked yet IMPROVED

35

Dependency-Track templated project names now support Sprig template extensions.

— Names Sprig but gives no example template syntax 0.50.0

03 **Workload identity federation for Dependency-Track** Not checked yet

NEW

30

The Dependency-Track integration now supports workload identity federation for authentication.

— States the capability but gives no setup or config detail 0.50.0

WAS THIS USEFUL? 👍 👎

sbomify

Feedback ⓘ v26.11.1 NOTES ↗ RANK

sbomify is a product security artifact hub and a trust center.

> sbomify's v26.11.1 release introduces a full release lifecycle with supersession and retirement dates, restructures Dependency-Track around release series, and adds Trust Center stale-cache resilience, alongside smaller gains in VEX review, vulnerability advisory coverage, and component metadata.

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

01 Release lifecycle with retirement and supersession

Not checked yet NEW 95

Releases can declare predecessors via a new `supersedes` list field on `POST /api/v1/releases` and `PATCH /api/v1/releases/{release_id}` (and their v2 twins), gated by the `release:manage` permission. A release can replace another on the same line (retiring it) or start a new line (keeping both monitored), with end-of-support and end-of-life dates, withdrawn and archived states, and a state badge and Lifecycle card in the UI; dates set via the API before this release had no effect but now retire the release on deploy. Only live releases are monitored for scans, VEX re-application and Dependency-Track, and retired releases show their last scan date and the reason scanning stopped (e.g. Its support ended on…) in the app and Trust Center. Deleting a release that replaces another on the same line is refused until the predecessor link is removed from the Lifecycle card.

Register a new release from CI and declare it supersedes the previous tag, so the lifecycle updates automatically without manual edits in the app.

```
$ curl -X POST https://app.sbomify.com/api/v1/releases \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"version": "1.2.5", "supersedes": ["1.2.4"]}'
```

— Named endpoints, field and permission with runnable example v26.11.1

launch-20261117-9b301f0f

02 Dependency-Track per-release-series layout Not checked yet NEW

90

Run `python manage.py dt_tracked_scope` (or with `--dry-run` to preview) to migrate Dependency-Track to a per-release-series layout: each workspace gets a collection per live release holding one project per component, plus a Latest builds collection, and a replacing release takes over the older release's projects, renaming

them and swapping in BOMs in place to preserve history. This retires old per-SBOM projects.

Dry-run the Dependency-Track series-layout migration to preview which projects will be retired before committing.

```
$ docker compose exec -T sbomify-backend python manage.py dt_tracked_scope --dry-run
```

Preview which Dependency-Track projects will be restructured before committing the release-series backfill on a self-hosted upgrade.

```
$ python manage.py dt_tracked_scope --dry-run
```

— Names command, flag and collection structure `v26.11.1` `launch-20261117-9b301f0f`

---

03 **Valkey replaces Redis in Docker Compose** Not checked yet BREAKING

90

Docker Compose replaces Redis 8 with the bundled Valkey 9, using a new `sbomify_valkey_data` volume; Valkey cannot read Redis 8 data files, so queued tasks present at upgrade are dropped. Drain workers before upgrading, then remove the old volume with `docker volume rm <project>_sbomify_redis_data`, or set the `REDIS_IMAGE` environment variable to keep the original Redis image instead.

— Exact volume names, migration command and opt-out variable `v26.11.1` `launch-20261117-9b301f0f`

---

04 **Trust Center stale-cache resilience** Not checked yet NEW

75

Tune Trust Center stale-cache behaviour with `TRUST_CENTER_STALE_CACHE_TTL` (default one day; set to `0` to turn it off) so public Trust Center pages are served from the last good copy while Postgres is unavailable.

— Config key with default and disable value given `v26.11.1`

---

05 **Self-hosted telemetry control** Not checked yet NEW 70

Self-hosted anonymous usage telemetry is now controlled by the `TELEMETRY_ENABLED` environment variable, defaulting on; set `TELEMETRY_ENABLED=False` to opt out, and it is automatically off when `CI` is set.

Opt out of anonymous usage telemetry on a self-hosted instance.

```ini
TELEMETRY_ENABLED=False
```

— Exact env var and default/opt-out behaviour given `v26.11.1`

THINNER COVERAGE BELOW

06 **Vulnerability links to sbomify Insights** Not checked yet IMPROVED

55

Vulnerability IDs in the release findings table, Trust Center posture, and VEX review now link to sbomify Insights, which also now covers GSD, HSEC, PSF, and Drupal advisory schemes.

— Names surfaces and new advisory schemes, no config to act on `launch-20261117-9b301f0f`

07 **VEX review merges duplicate agent rulings** Not checked yet IMPROVED

50

VEX review now merges agent-proposed rulings that say the same thing into one verdict per vulnerability, listing covered packages and versions as chips with a single shared explanation.

— Explains merge mechanism but no config or endpoint `launch-20261117-9b301f0f`

08 **Source Control card on component metadata** Not checked yet NEW

40

Components can now record their source repository in a new Source Control card on the component metadata page, also available through the metadata API.

— Names the surface but not fields or exact API path `v26.11.1`

09 **Sign-in on custom Trust Center domains** Not checked yet NEW

35

Readers can now sign in on a custom Trust Center domain after requesting access and being approved.

— Thin prose description, no mechanism or config named `launch-20261117-9b301f0f`

└──▷ **BREAKING ON UPGRADE**

! Dependency-Track projects must be migrated immediately after deploy by running `python manage.py dt_tracked_scope`; old per-SBOM projects are retired and a component deleted before the backfill runs leaves nothing to find its old projects by.

! The Dependency-Track API key must be granted the `PORTFOLIO_MANAGEMENT` permission (in addition to `BOM_UPLOAD`, `PROJECT_CREATION_UPLOAD`, `VIEW_PORTFOLIO` and

`VIEW_VULNERABILITY`) before the backfill, or sbomify cannot create, rename or retire projects (401/403).

! Releases past their end-of-support or end-of-life date are now retired on deploy — no more scans and their Dependency-Track projects go inactive; dates set via the API before this release had no effect but now do.

! Docker Compose: Redis 8 is replaced by Valkey 9 using a new `sbomify_valkey_data` volume; Valkey cannot read Redis 8 data files, so any queued tasks present at upgrade are dropped — drain workers before upgrading and then remove the old volume with `docker volume rm <project>_sbomify_redis_data`.

! Self-hosted installs now send anonymous usage telemetry by default; set `TELEMETRY_ENABLED=False` to opt out.

! `POST /api/v1/releases` and `PATCH /api/v1/releases/{release_id}` (and their v2 twins): sending a `supersedes` list without the `release:manage` permission returns a 403.

! Deleting a release that replaces another on the same line is refused until the predecessor link is removed from the release's Lifecycle card.

! Redis is replaced by Valkey 9 as the bundled backing service; Valkey cannot read Redis 8 data files, so the cache and task queue start empty in a new `sbomify_valkey_data` volume — drain the worker queue before upgrading or queued tasks are dropped, then remove the old `sbomify_redis_data` volume. To keep Redis, set `REDIS_IMAGE`.

! End-of-support and end-of-life dates set through the API before this release had no effect; from v26.11.1 onward they retire the release on deploy — check old support dates before upgrading.

! Dependency-Track integration now requires the `PORTFOLIO_MANAGEMENT` permission on the API key, and a one-off backfill (`python manage.py dt_tracked_scope`) must be run immediately after deploy; a component deleted before the backfill runs leaves nothing to find its old projects by.

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Akto

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗ RANK

Akto discovers APIs, MCP servers and AI agents, tests and red-teams them, and enforces runtime guardrails and posture management.

> Akto's latest release adds governance controls for blocking AI agents and vendors by team, role, device, and user, plus an AI-assisted verification step to cut false positives in red teaming.

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

01 **Scoped blocking of AI agents and vendors** Not checked yet NEW 45

Akto can now block AI agents and vendors (such as Grok Bot) for specific teams, roles, devices, and users, governing personal AI agent access across the organization.

— Names scoping dimensions and example vendor, no UI path given

changelog-20261118-070421ab

02 **Verify with AI for red teaming** Not checked yet NEW 25

A new Verify with AI capability reduces false positives generated during red teaming.

— Names the feature but not its mechanism changelog-20261118-070421ab

WAS THIS USEFUL? 👍 👎

fcaptcha

Feedback ⓘ v1.51.0 NOTES ↗ RANK

> fcaptcha's v1.51.0 release adds signed, backward-compatible protocol negotiation across its browser widget and Node, Go, and Python servers, plus a stricter caching header on challenge responses.

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

01 **Signed challenge protocol negotiation** Not checked yet NEW

55

Adds backward-compatible challenge protocol negotiation across the browser widget and the Node, Go, and Python servers, so old clients and old servers remain compatible as protocol v1. Adds signed protocol version selection, with explicit rejection when no supported version overlaps between client and server.

– Explains mechanism and scope but no config/flag names given v1.51.0

02 **No-store caching on challenge responses** Not checked yet IMPROVED

45

Adds `Cache-Control: no-store` to all challenge responses.

– Names exact header but no further mechanism detail v1.51.0

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ React Native SDK for Axios 4.0.0 (2026-11-17)

BLOG ↗ RANK

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

> DataDome shipped a major v4.0.0 rewrite of its React Native SDK for Axios, replacing manual interceptor setup with a simple attach/eject API, adding native challenge display and full fetch-adapter coverage, and raising minimum platform requirements across the board.

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

01 Axios interceptor attach/eject API

Not checked yet BREAKING

95

DataDomeInterceptor.attach(instance) attaches DataDome protection to any Axios instance — default or custom — and DataDomeInterceptor.eject(instance) removes it. This replaces the removed new DataDomeInterceptor({ axiosInstance }) constructor and the manual interceptor setup with `requestAdapter`, `responseEventHandler`, and `responseAdapter`.

Protect all requests on a custom Axios instance so every outbound call is covered by DataDome bot detection.

javascript

```
import axios from 'axios';
import { DataDome, DataDomeInterceptor } from '@datadome/react-
native-datadome-axios';

const dataDome = new DataDome({ clientKey: '<your-client-key>'
});
const instance = axios.create({ baseURL:
'https://api.example.com' });
DataDomeInterceptor.attach(instance);

// Later, when tearing down:
DataDomeInterceptor.eject(instance);
```

— Named methods with a runnable attach/eject example.

React Native SDK for Axios 4.0.0 (2026-…

02 getCookie(url) replaces stored-cookie functions

Not checked yet

80

BREAKING

dataDome.getCookie(url) retrieves the stored DataDome cookie for a given URL, replacing the removed getStoredDataDomeCookie() and storeCookie() functions.

Read the DataDome cookie for a specific URL after a challenge is solved, e.g. to forward it to a server-side session.

```javascript
const cookie = await
dataDome.getCookie('https://api.example.com/endpoint');
console.log(cookie);
```

— Named method with a runnable retrieval example.

React Native SDK for Axios 4.0.0 (2026-…

03 **Simplified SDK configuration with clientKey** Not checked yet

70

BREAKING

new DataDome({ clientKey }), with an optional `domain`, replaces the previous multi-step setup. DataDome.getInstance().setSdkKey(), DataDome.configure(), DataDome.configureWithOrigin(), and `enableVerboseLogs` are all removed in 4.0.0.

Protect all requests on a custom Axios instance so every outbound call is covered by DataDome bot detection.

```javascript
import axios from 'axios';
import { DataDome, DataDomeInterceptor } from '@datadome/react-
native-datadome-axios';

const dataDome = new DataDome({ clientKey: '<your-client-key>'
});
const instance = axios.create({ baseURL:
'https://api.example.com' });
DataDomeInterceptor.attach(instance);

// Later, when tearing down:
DataDomeInterceptor.eject(instance);
```

— Names old and new config surfaces but little mechanism beyond that.

React Native SDK for Axios 4.0.0 (2026-…

04 **Package rename and raised minimum requirements** Not checked yet

70

BREAKING

The package is renamed from `@datadome/datadome-axios` to `@datadome/react-native-datadome-axios`; `@datadome/react-native-core-datadome` (1.0.0+) is now a required peer dependency; and minimum versions are raised — React Native to 0.85+ with the New Architecture required (previously 0.72+), iOS to 15.1+ (previously 12+), Android `minSdkVersion` to 24 (previously 21), and Axios to 1.12+ (previously 1.11+).

— Concrete version numbers give a clear migration checklist.

React Native SDK for Axios 4.0.0 (2026-…

05 **Native challenge display replaces WebView modal** Not checked yet

65

BREAKING

Challenges are displayed natively without adding any component to the view hierarchy, powered by `@datadome/react-native-core-datadome`. `DataDomeModal` is removed, and the `react-native-webview` peer dependency is dropped — apps that don't otherwise use it can now uninstall it.

— Explains mechanism and removed dependency but no usage example.

React Native SDK for Axios 4.0.0 (2026-…

THINNER COVERAGE BELOW

06 **Full Axios adapter coverage including fetch** Not checked yet NEW

55

Protection now covers all Axios adapter types — default, instance-level, or per-request — including the fetch adapter via a custom `env.fetch`.

— Names adapter types but gives no example to act on.

React Native SDK for Axios 4.0.0 (2026-…

07 **Automatic retry after solved challenge** Not checked yet IMPROVED

45

A request is automatically retried after a solved challenge from the request's adapter, so interceptors and transforms run only once per request regardless of retry count.

— Describes behavior but no named surface or example.

React Native SDK for Axios 4.0.0 (2026-…

08 **Core SDK type exports** Not checked yet NEW

35

`DataDome`, `ResponseValidationStatus`, `DataDomeConfiguration`, and `DataDomeCookie` types are now exported from the `@datadome/react-native-core-datadome` core SDK.

— Just lists exported type names, no usage shown.

React Native SDK for Axios 4.0.0 (2026-…

└──▷ BREAKING ON UPGRADE

! Package renamed from `@datadome/datadome-axios` to `@datadome/react-native-datadome-axios` — update all imports and package.json references.

! `@datadome/react-native-core-datadome` (1.0.0+) is now a required peer dependency and must be installed in the application.

! Minimum React Native version raised to 0.85+ with the New Architecture required (previously 0.72+); iOS minimum raised to 15.1+ (previously 12+); Android `minSdkVersion` raised to 24 (previously 21); Axios minimum raised to 1.12+ (previously 1.11+).

! DataDome.getInstance().setSdkKey(), DataDome.configure(), and DataDome.configureWithOrigin() are removed — replace with new DataDome({ clientKey }).

! new DataDomeInterceptor({ axiosInstance }) and the manual interceptor setup with `requestAdapter`, `responseEventHandler`, and `responseAdapter` are removed — replace with new DataDomeInterceptor(dataDome).attach(instance).

! getStoredDataDomeCookie() and storeCookie() are removed — replace with dataDome.getCookie(url).

! `DataDomeModal` is removed.

! `enableVerboseLogs` is removed.

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ

Add Fastly Object Storage logger option for services (2026-11-17) BLOG ↗

RANK

Fastly is a cloud edge platform providing CDN, DDoS protection, load balancing, and real-time content delivery to accelerate and secure web applications.

> Fastly's main release this window is Fastly Object Storage logging support — rolled out simultaneously across the service API, CLI, and Terraform provider — alongside new CLI tooling for local VCL testing and bot management, plus webhook notifications for Client-Side Protection.

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

01 **Fastly Object Storage logging support across API, CLI and Terraform**

Not checked yet NEW SUPPORTED PLATFORM NEW 95

Fastly adds a Fastly Object Storage logging endpoint for services: supply only `region` and `bucket_name` and Fastly automatically provisions a bucket-scoped access key, returning its ID in `access_key` without ever exposing the secret. CLI v17.2.0 adds `service logging fos` subcommands to configure these logging endpoints, and Fastly Terraform Provider 10.3.0 adds a `logging_fos` block inside `fastly_service_vcl` or `fastly_service_compute` resources to route logs directly to Fastly Object Storage instead of a third-party destination.

Send VCL service logs to Fastly Object Storage instead of a third-party destination.

hcl

```hcl
resource "fastly_service_vcl" "example" {
  name = "my-vcl-service"

  logging_fos {
    name       = "fos-logging"
    bucket     = "my-log-bucket"
    path       = "/logs"
  }
}
```

— Named API fields, CLI subcommands, Terraform block, and a runnable config example

Add Fastly Object Storage logger option… CLI v17.2.0 (2026-11-17)

Fastly Terraform Provider 10.3.0 (2026-…

02 **VCL compile-and-test CLI commands** Not checked yet NEW

70

Fastly CLI v17.2.0 adds `vcl check`, `vcl run`, and `vcl serve` subcommands to compile and test VCL code locally.

– Names exact subcommands but gives no usage example CLI v17.2.0 (2026-11-17)

03 **Bot-management CLI commands** Not checked yet NEW

60

Fastly CLI v17.2.0 adds new `bot-management` subcommands to manage bot-management workspaces, bot categories, and bots.

– Names the subcommand group but not individual flags or examples

CLI v17.2.0 (2026-11-17)

THINNER COVERAGE BELOW

04 **Webhook notifications for Client-Side Protection** Not checked yet

NEW

35

Fastly adds signed HTTPS webhook notifications for Client-Side Protection Page events, letting customers receive notifications instead of polling.

– Describes the mechanism but names no endpoint or config to act on

Add webhook notifications for Client-Si…

> New regions for POPs, billing, and stats for Saudi Arabia (2026-11-01) BLOG ↗
>
> STABLE

WAS THIS USEFUL? 👍 👎

◆ Code Security

Mend.io Mend SAST

Feedback ⓘ changelog · Nov 18, 2026

CHANGELOG ↗

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

Mend SAST adds a guided onboarding wizard, language-filtered Global Scan Configuration, and an AI assistant (Mend Intelligent Agent) for natural-language risk exploration.

New guided onboarding wizard lets users activate SAST scanning for repositories directly from the Mend Platform, including repository selection, pilot scan execution, and email invitations to delegate onboarding to colleagues. Not checked yet

Global Scan Configuration now filters displayed languages based on engine generation, showing only languages supported by the active engine. Not checked yet

New Mend Intelligent Agent embeds an AI assistant in the Mend Platform so security teams can explore capabilities, navigate to relevant views, and get executive-level risk summaries using natural language — enabled by default for organizations with a signed AI consent addendum. Not checked yet

UPDATED DETECTION JavaScript Gen 2 engine — DOM XSS via jQuery AJAX callbacks

UPDATED DETECTION Swift Gen 2 engine — reduced false positives for CWE-117 and CWE-295

WAS THIS USEFUL? 👍 👎

StackHawk

Feedback ⓘ updated · Nov 18, 2026 NOTES ↗ RANK

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

StackHawk adds per-rule alert limits from policy, a dev-channel installer flag, ARM64 Brotli decoding, and several hawk op subcommand improvements.

Apply alert limits per scan rule from the StackHawk scan policy using `pscans.maxAlertsPerRule` for passive alerts and `scanner.maxAlertsPerRule` for active alerts in `hawk.config`. Not checked yet

Install the newest development build of HawkScan by setting `HAWK_INSTALL_CHANNEL=dev` before running the install script. Not checked yet

Keep install scripts from modifying your PATH by setting `HAWK_NO_MODIFY_PATH=1`. Not checked yet

Link an existing application to a repository by name with `hawk op repo link --app-name` without creating a new one. Not checked yet

Run `hawk op repo link` to add an application to a repository's existing list rather than replacing it. Not checked yet

+4 more

UPDATED VULNERABILITY CHECK Brotli response decoding for passive rules on ARM64

└──▷ **BREAKING ON UPGRADE**

! Install scripts now install `hawk` to `~/.local/bin` instead of `~/.hawk/bin`; if your pipeline adds `~/.hawk/bin` to PATH (e.g. via `$GITHUB_PATH` or `$BASH_ENV`), change it to `~/.local/bin`.

! `pscans.maxAlertsPerRule` in `hawk.config` now limits only passive alerts; it no longer limits active alerts as it did since 6.6.0 — use `scanner.maxAlertsPerRule` to limit active alerts.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

OpenAI's GPT-6.2 Astra and GPT-6.2 Sol now available (2026-11-17) BLOG ↗

RANK

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

> GitHub's Copilot lineup grew substantially this window with new GPT-6.2, Fusion, and MAI Code 2 models, a public preview of repeatable behavioral testing via Copilot Evals, Microsoft Entra sign-in for the CLI, and a wave of JetBrains- and Copilot-app-specific features including cloud sandboxes, remote access, browser automation, and voice dictation.

└──▷ **WHAT SHIPPED · 20 FEATURES** most completely described first ? what's the number?

01 Microsoft Entra sign-in in Copilot CLI Not checked yet

NEW INTEGRATION NEW 88

Update Copilot CLI to version `1.0.113` or later with `copilot update` to enable Microsoft Entra authentication PREVIEW. Signing in with a Microsoft Entra account silently authenticates to a linked GitHub Enterprise Managed User (EMU) account via OIDC SSO, removing the need for a separate GitHub sign-in, and requires a Copilot Business or Enterprise license with the Copilot CLI policy enabled on the organization. A Microsoft Entra identity plus a Microsoft 365 Copilot license also grants access to models through Microsoft 365 Copilot directly from the terminal with usage-based billing, with inference and usage limits managed through Microsoft 365 settings rather than GitHub; connecting the Work IQ MCP server gives Copilot CLI access to Microsoft 365 work data as context.

Update the CLI to the minimum required version to unlock Microsoft Entra authentication before attempting sign-in.

```
$ copilot update
```

— Exact command, version, licensing and SSO mechanism given.

GitHub Copilot CLI adds Microsoft Entra…

02 Copilot Evals public preview Not checked yet NEW 88

Run repeatable behavioral tests on Copilot skills, custom agents, and MCP tools using the `/eval` command or plain-language prompts, with per-run approval before execution. Install Copilot Evals as a Visual Studio Code extension from the Visual Studio Marketplace, as a Copilot CLI plugin, or as a GitHub Copilot app plugin via the

`github/evals` marketplace entry, and compare runs before and after a customization change to detect regressions in AI behavior.

▶ Watch video

Add the Copilot Evals plugin to the GitHub Copilot app so you can run evals from the chat interface without leaving the browser.

📍 In the GitHub Copilot app, go to Customize, search for `github/evals` in the marketplace, and click Install to add the GitHub Copilot Evals plugin.

— Names command, install surfaces, and marketplace entry.

GitHub Copilot Evals is now in public p…

03 **Azure DevOps integration in Copilot app** Not checked yet

UPDATED INTEGRATION NEW 86

In public preview, enable Azure DevOps support by opening `Settings > Experimental` and turning on `Azure DevOps integration`, then adding a Microsoft account under `Accounts > Add account`. Once connected you can clone Azure DevOps repos, create branches, open pull requests, and let an agent fix CI failures and merge once checks pass.

▶ Watch video

Connect an Azure DevOps account so you can clone repos and have an agent open and merge pull requests without leaving the Copilot app.

📍 1. Open Settings › Experimental and turn on Azure DevOps integration. 2. Open Accounts › Add account and sign in with your Microsoft account used for Azure DevOps. 3. Clone a repository, create a branch, and open a pull request — click Fix on review comments or enable agent merge to let Copilot address CI failures and complete the merge.

— Exact settings paths and end-to-end workflow given.

GitHub Copilot app, Microsoft Ignite re…

04 **Browser automation in Copilot app** Not checked yet NEW

86

After installing the GitHub Copilot extension for Chrome, enable browser automation in public preview via `Settings > Computer Use`, turning on both Enable Computer Use and Enable Browser Use. This lets Copilot fill forms, pull information from pages, and test web apps in an isolated tab group in Chromium-based browsers including Google Chrome and Microsoft Edge.

▶ Watch video

Let Copilot automate browser tasks — filling forms or testing a web app — in an isolated Chrome tab group.

📍 1. Install the GitHub Copilot extension for Chrome. 2. Open Settings › Computer Use and turn on both `Enable Computer Use` and `Enable Browser Use`. 3. Ask Copilot to fill out a form or pull information from a webpage; it will act in an isolated tab group in your existing Chrome or Edge session.

— Exact flags and supported browsers named. `GitHub Copilot app, Microsoft Ignite re…`

---

05 **Repository Insights page for Copilot reviews** `Not checked yet`

`NEW DETECTION` `NEW` 83

A new Repository Insights page, the first item in a repository's Insights tab, shows Copilot code review summary signals, addressed-rate trends, and comment categories for anyone with read access. The four summary signals are `Comments addressed`, `High-severity comments addressed`, 'Pull requests with no later human change request', and `Merged pull requests without Copilot review`.

Review Copilot code review outcomes for a specific repository directly in the GitHub UI.

📍 1. Navigate to your repository on GitHub. 2. Click the `Insights` tab. 3. Select `Repository insights` (the first item in the tab). 4. Review the four summary signals: `Comments addressed`, `High-severity comments addressed`, 'Pull requests with no later human change request', and `Merged pull requests without Copilot review`.

— Names exact UI path and the four tracked signals.

`Track Copilot code review outcomes in R…`

---

06 **GPT-6.2 Astra and Sol models in Copilot** `Not checked yet`

`NEW SUPPORTED PLATFORM` `NEW` 78

GitHub Copilot adds `GPT-6.2 Astra`, tuned for agentic coding and terminal work and catching more code review issues than GPT-6 Astra while using fewer tokens, and `GPT-6.2 Sol`, for complex agentic coding and in-depth code review, catching nearly twice as many real issues as GPT-6.1 Sol. Both can be selected via the model picker in Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, the Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse, with access managed for Copilot Business and Enterprise plans through the model policy in Copilot settings.

Switch to GPT-6.2 Sol for a thorough code review session in VS Code when you want maximum issue detection on a critical PR.

📍 In the GitHub Copilot model picker, select GPT-6.2 Sol before starting your code review.

— Names models, comparisons, surfaces, and policy control.

OpenAI's GPT-6.2 Astra and GPT-6.2 Sol …

07 **Copilot Fusion availability and new models** Not checked yet

IMPROVED

78

Copilot Fusion , which lets Copilot choose and coordinate models per task at a user-set effort level, is now selectable from the model picker on github.com, Visual Studio, and the GitHub Mobile beta apps for iOS and Android. Its lineup grew with `Claude Sonnet 5.5` in Fusion Max, `GPT-6.1 Terra` in Fusion High, and MAI-Code-2 in Fusion Standard; Copilot Fusion Max is now listed on the Terminal-Bench 4.1 public leaderboards hosted by Stanford, Harbor, and the Laude Institute, achieving a 71.2% resolution rate.

Switch to Copilot Fusion on github.com or Visual Studio to let Copilot automatically route your task to the best model across Max, High, or Standard tiers.

📍 In the Copilot chat model picker, select Copilot Fusion (or Copilot Fusion Max for the highest-effort tier), set your effort level, and submit your task.

— Names tiers, models, surfaces, and benchmark result.

GitHub Copilot Fusion updates: more ava…

08 **Browser use in JetBrains Copilot** Not checked yet NEW

69

Enable Browser Use via `GitHub Copilot > Chat > Enable Browser Use` in JetBrains PUBLIC PREVIEW so Copilot can open and navigate web pages, read content, and interact with links, buttons, and forms.

Let Copilot automate a web research task during a chat session without leaving your IDE.

📍 In JetBrains, go to GitHub Copilot › Chat › Enable Browser Use to turn on the browser use plugin, then ask Copilot in chat to visit a URL and summarise its content.

— Exact menu path and interaction scope given. GitHub Copilot in JetBrains, Ignite rel…

09 **Agent session steering keybindings** Not checked yet NEW

65

Press `Enter` to steer a running Copilot agent session immediately, or `Ctrl+Enter` (`Command+Enter` on macOS) to queue a follow-up that runs after the current step completes.

Quickly send a follow-up without interrupting a running agent session, or steer it in a new direction mid-run.

> 📍 While a Copilot agent session is running, press Enter to steer it immediately, or press Ctrl+Enter (Command+Enter on macOS) to queue a follow-up that runs after the current step completes.

— Exact keybindings make this directly usable. GitHub Copilot in JetBrains, Ignite rel…

---

10 **GitHub Enterprise Server account support in Copilot app**

Not checked yet NEW 60

The Copilot app reaches general availability for GitHub Enterprise Server accounts; after updating the app to version `1.2.0` or later, add a GitHub Enterprise Server account under `Accounts > Add account` to work with on-premises repositories, issues, and pull requests.

Add a GitHub Enterprise Server account to work with on-premises repositories, issues, and pull requests after upgrading to version 1.2.0.

> 📍 1. Update the GitHub Copilot app to version 1.2.0 or later. 2. Sign in with your `GitHub.com` or GHE.com account. 3. Open Settings › Accounts › Add account and choose `GitHub Enterprise Server`.

— Names version and exact menu path, little mechanism.

GitHub Copilot app, Microsoft Ignite re…

---

11 **Automatic Maven vulnerability checks** Not checked yet NEW 60

Maven vulnerability checks are now automatic: opening or importing a Maven project sends dependency coordinates to GitHub Advisories to check for vulnerabilities.

— Names the mechanism and target service, no config flag.

GitHub Copilot in JetBrains, Ignite rel…

THINNER COVERAGE BELOW

12 **Remote access to Copilot chat sessions** Not checked yet NEW 58

Opt in to remote access PUBLIC PREVIEW so Copilot chats started in JetBrains are available from other Copilot clients such as VS Code and the Copilot app, with shared and historical sessions accessible across devices.

Continue a Copilot session started in JetBrains from VS Code or the Copilot app without losing conversation history.

> 📍 In JetBrains, enable remote access via the GitHub Copilot settings opt-in, then open the Copilot app or VS Code and select the shared session from the remotely accessible computers list.

— Names target clients, lacks exact settings path.

GitHub Copilot in JetBrains, Ignite rel…

13 **JetBrains IDE 2025.2 support ended** Not checked yet BREAKING

50

Support for JetBrains IDE 2025.2 has ended; upgrade to 2025.3 or later to continue receiving GitHub Copilot updates.

— Exact versions give a clear upgrade action. GitHub Copilot in JetBrains, Ignite rel…

14 **On-device voice dictation in JetBrains** Not checked yet NEW

47

Dictate chat messages on supported native plugin packages with Editor Preview enabled PUBLIC PREVIEW; transcripts remain as editable drafts until you select Send.

— Mechanism given, no exact toggle path named. GitHub Copilot in JetBrains, Ignite rel…

15 **MAI Code 2 and MAI Code 2 Flash models** Not checked yet NEW

45

GitHub Copilot now includes MAI Code 2, a general-purpose coding model for complex tasks like deep debugging and multi-file refactoring, alongside MAI Code 2 Flash.

Switch to MAI Code 2 for a complex multi-file refactor in VS Code using the model picker.

> 📍 In VS Code, open the Copilot chat panel, click the model picker dropdown, and select MAI Code 2.

— Names models but gives no mechanism detail. MAI Code 2 and MAI Code 2 Flash in GitH…

16 **Cloud sandbox sessions in JetBrains** Not checked yet NEW

45

GitHub Copilot can now run sessions in an isolated, GitHub-hosted cloud sandbox PUBLIC PREVIEW, freeing local resources and allowing parallel sessions.

— Describes benefit but no setup steps named. GitHub Copilot in JetBrains, Ignite rel…

17 **Split mode remote development support** Not checked yet NEW

43

GitHub Copilot now supports JetBrains split mode remote development, with chat and editor UI running on the local client while Copilot operations execute on the remote backend.

– Explains mechanism but no setup steps. GitHub Copilot in JetBrains, Ignite rel…

18 **Model and token details in inline chat** Not checked yet IMPROVED

35

Inline chat now shows model names and per-response token details while retaining agent session credit totals.

– Describes display change with no further detail.

GitHub Copilot in JetBrains, Ignite rel…

19 **Agent review engine for code review** Not checked yet IMPROVED

30

Code review in JetBrains now uses the Copilot agent review engine, with results still surfaced in the existing inline review interface.

– States the change but no mechanism detail. GitHub Copilot in JetBrains, Ignite rel…

20 **Java modernization findings in Problems View** Not checked yet

NEW

30

Java modernization findings now appear in the Problems View on supported IDEs.

– Bare statement with no further mechanism. GitHub Copilot in JetBrains, Ignite rel…

└──▷ **BREAKING ON UPGRADE**

! Support for JetBrains IDE 2025.2 has ended; upgrade to 2025.3 or later to continue receiving GitHub Copilot updates.

Copilot code review: Adaptive effort, secret detection, and more (2026-11-17)

BLOG ↗ STABLE

GitHub Copilot code review adds Adaptive effort selection, cross-repository context, per-repository outcome tracking, and Azure DevOps GA.

Adaptive review effort automatically picks Lite, Balanced, or Max per PR based on complexity and risk; configurable per PR, repository, organization, or enterprise.

Not checked yet

Cross-repository context lets Copilot code review read related public repositories to find downstream dependents and relevant documentation during a review.

Not checked yet

A new `Repository insights` view under the `Insights` tab tracks how Copilot code review feedback is addressed and utilized per repository. Not checked yet

Copilot code review for Azure DevOps is now generally available, supporting branch-policy-triggered reviews, `Lite` / `Balanced` effort defaults, custom instructions, one-step suggestion apply-and-resolve, audit log events, and Azure Cost Management billing. Not checked yet

`Monitor GitHub Copilot activity through Microsoft Agent 365 (2026-11-17)`

BLOG ↗ STABLE

GitHub adds an audit log stream to Microsoft Agent 365 so enterprise admins can review Copilot prompts, responses, tool calls, and usage trends in Microsoft's agent control plane.

Stream GitHub Copilot agent activity into Microsoft Agent 365 to review Copilot prompts, responses, tool calls, and usage trends alongside other AI workloads — available for GitHub Enterprise Cloud customers with Enterprise Managed Users. Not checked yet

Access additional Copilot risk data in Microsoft Purview and Defender through the Agent 365 integration to investigate and mitigate risk. Not checked yet

---

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-16 CODE ↗

Syzkaller is an unsupervised coverage-guided kernel fuzzer developed by Google that automatically finds bugs in OS kernels by generating and executing random syscall sequences.

syzkaller adds NVMe-boot GCE machine type support, syz-hub GCP Secret Manager config resolution, and QEMU Unix socket QMP monitor

Support NVMe-boot GCE machine types (such as N4D) in `pkg/gce` and tools.
Not checked yet

Add a Dockerfile, Makefile, and Kubernetes manifests for `syz-hub` deployment.
Not checked yet

Use a Unix socket for the QEMU QMP monitor in `vm/qemu`, replacing the previous monitor connection method. Not checked yet

Skip AI cause bisection for KCSAN bugs automatically in the dashboard.
Not checked yet

WAS THIS USEFUL? 👍 👎

psalm-plugin-laravel

Feedback ⓘ v4.22.1 NOTES ↗ RANK

Psalm-plugin-laravel is a security-focused Laravel static analyser

> Psalm-plugin-laravel v4.22.1 adds typed closure parameters for eager-load constraint callbacks and raises its minimum required Psalm version.

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

01 Typed closure parameters in eager-load constraints

Not checked yet **IMPROVED** 75

The closure parameter of eager-load constraint callbacks in with(), withOnly(), load(), and loadMissing() is now typed as the constrained relation instead of untyped. This eliminates `MissingClosureParamType` and `MixedMethodCall` false positives, e.g. `Customer::query()->with(['vehicles' => fn ($query) => $query->where('year', '>', 2020)]);` now has a correctly typed `$query`.

Get a typed `$query` parameter inside an eager-load constraint closure so Psalm can check method calls on the relation model instead of reporting `MixedMethodCall`.

```php
Customer::query()->with(['vehicles' => fn ($query) => $query->where('year', '>', 2020)]);
```

— Names exact methods and error types fixed, with example v4.22.1

02 Raised minimum Psalm version requirement

Not checked yet **BREAKING** 65

The minimum required version of `vimeo/psalm` is now `^7.9.0` (was `^7.7.0`); upgrade both packages together with `composer update vimeo/psalm psalm/plugin-laravel`.

— Exact version bump and upgrade command given v4.22.1

└──▷ **BREAKING ON UPGRADE**

! The minimum required version of `vimeo/psalm` is now `^7.9.0` (was `^7.7.0`); upgrade both packages together with `composer update vimeo/psalm psalm/plugin-laravel`.

WAS THIS USEFUL? 👍 👎

OpenQodex

Feedback ⓘ openqodex@0.19.0 NOTES ↗ RANK

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

> OpenQodex's 0.19.0 release adds cross-file evidence to findings, turns convention-file rule violations into first-class findings, and bumps the bundled ruff linter to 0.17.3 with a 59-to-407 rule default expansion that includes new security rules — a breaking change for repos without an explicit rule selection.

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

01 Ruff linter upgraded to 0.17.3 with expanded security rule set

Not checked yet BREAKING 93

The bundled ruff Python linter is now version 0.17.3, whose default rule set grew from 59 rules to 407, adding security rules `S102`, `S110` and `S112` (ranked high severity by OpenQodex). Repositories with no ruff `select` now receive more findings, and `requires-python` from a `pyproject.toml` with no `[tool.ruff]` table is read as the target version. To keep the previous 59-rule set, add `select = ["E4", "E7", "E9", "F"]` to the repository's ruff config.

Pin a ruff rule set to avoid the new 407-rule default after upgrading, so only the original 59 rules run on repositories that had no prior ruff config.

toml

```
select = ["E4", "E7", "E9", "F"]
```

— Names exact versions, rule counts, and a runnable config fix openqodex@0.19.0

02 Convention-file rule violations reported as findings

Not checked yet IMPROVED 85

Violations of written rules in convention files listed under `review.instructions_files` in `.openqodex/config.yaml` are now reported as findings rather than observations. The cited rule is recorded in the `rule` field of `report.json` and displayed as `Rule: <file> > <section>: "..."` in the report; citations whose quoted text is absent from the file are dropped with an explanatory note.

— Names config key, field and exact display format openqodex@0.19.0

03 **Cross-file evidence in staged review findings** Not checked yet

IMPROVED 76

Staged reviews now pull cross-file evidence: callers of changed interface declarations (up to 8 per declaration, code-graph callers first, then name matches) are included in the finder brief, and verifiers can read callers and guards in other files. Findings that depend on other files now surface those files and lines in the `evidence` field of `report.json` and as Evidence in other files in the report.

— Names exact field and limit but no direct user action `openqodex@0.19.0`

└──▷ **BREAKING ON UPGRADE**

! The bundled ruff default rule set grew from 59 rules to 407; repositories with no ruff `select` will now receive substantially more findings, including security rules `S102`, `S110` and `S112`. To keep the previous set, add `select = ["E4", "E7", "E9", "F"]` to the repository's ruff config.

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Nov 17, 2026 BLOG ↗ RANK

The official changelog for the Snyk AI Trust Platform. Discover what's new in our security intelligence, AI-powered tools, and latest product enhancements.

> Snyk now bundles Snyk Studio installation directly into the Snyk CLI and the ADS installer, simplifying deployment and aligning version numbering across the two.

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

01 **Snyk Studio installation via Snyk CLI and ADS** Not checked yet

NEW 60

Snyk Studio can now be installed on a single machine without ADS using the Snyk CLI directly. Studio version numbers in ADS Settings now follow Snyk CLI version numbers, with the version selector moving from `1.0.19` to `1.1308.4`. Admins rolling out ADS through Jamf Pro or Microsoft Intune now get Studio bundled with the same ADS installer, with no separate Studio package required.

— Names install paths and version numbers but no exact CLI command. `2026-11-17-a2912b20`

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Acunetix 360 Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗

Acunetix 360 scans web applications and APIs to identify security vulnerabilities through automated testing.

Acunetix 360 adds KEV-aware vulnerability enrichment, an internet-exposed target flag, a website group size limit, automatic agent browser cleanup, and a scan duration cap.

Mark targets as internet-exposed when creating or editing them to surface publicly accessible assets and satisfy the asset exposure criterion in CISA BOD 26-04.
Not checked yet

Adds enriched vulnerability data and known exploitation status (KEV) to help prioritise findings and satisfy the known exploited vulnerability criterion in CISA BOD 26-04.
Not checked yet

Adds a separate report template for a missing HSTS preload directive with its own Best Practice severity, keeping it out of the existing HSTS errors and warnings report.
Not checked yet

Administrators can now set a maximum number of websites per group to keep system performance optimal. Not checked yet

Adds automatic cleanup of outdated browser files during agent upgrades, reducing false security scanner alerts and manual post-upgrade maintenance. Not checked yet

+1 more

NEW VULNERABILITY CHECK PHP multipart out-of-memory RCE and Node.js Brotli Decoder active checks

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite

Feedback ⓘ

Professional / Community 2026.11.1 (2026-11-17) BLOG ↗ RANK

Burp Suite is a we enable the world to secure the web. Burp Suite, the Web Security Academy, and world-leading research from PortSwigger â the AppSec communityâs home.

> Burp Suite 2026.11.1 adds an in-app Academy tab for Web Security Academy labs and simplifies switching between projects and restarting.

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

01 **Academy tab for Web Security Academy labs** Not checked yet NEW

50

Burp Suite adds a new Academy tab in the top navigation bar for accessing Web Security Academy labs directly from inside the tool, letting users browse and launch labs without leaving Burp.

Access Web Security Academy labs without leaving Burp — open the Academy tab to browse and launch labs in context.

> 📍 In Burp Suite, navigate to the Academy tab in the top navigation bar to browse and launch Web Security Academy labs.

— Names UI location but no further mechanism detail.

Professional / Community 2026.11.1 (202…

02 **Easier project switching and restart** Not checked yet IMPROVED

20

Burp Suite now lets you switch projects and restart more easily from within the UI.

— Bare description with no specifics on mechanism or UI path.

Professional / Community 2026.11.1 (202…

WAS THIS USEFUL? 👍 👎

ImmuniWeb

Feedback ⓘ announced · Nov 17, 2026 BLOG ↗ RANK

ImmuniWeb® AI Platform for Application Security Testing, Attack Surface Management & Dark Web Monitoring.

> ImmuniWeb expanded its free CyberScore service with industry-wide cybersecurity ratings and gave registered users deeper reporting access.

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

01 **CyberScore industry ratings and registered-user reports**

Not checked yet NEW 56

CyberScore now offers free, continually updated ratings tables across 14 industries — airlines, AI, automotive, banking, crypto, cybersecurity, e-commerce, education, financial services, healthcare, hospitality, insurance, space, and telecoms — accessible without registration. Registered free-account users can additionally view detailed technical reports, download PDF assessments, and refresh assessments after making security improvements.

— Names 14 sectors and account tiers but no concrete UI path or API.

launch-20261117-f0ec0807

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

changelog · Nov 18, 2026 | CHANGELOG ↗ | RANK

GitGuardian Detection Engine scans code repositories and developer workflows to automatically detect exposed secrets, API keys, and credentials in real time.

> GitGuardian Detection Engine expanded its secret-scanning coverage with 13 new detectors and 20+ new analyzers spanning Google, AWS, Jamf, Supabase, Grafana, Cloudflare and more, alongside a detector rename that breaks existing rule references.

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

01 **New Google credential detectors and analyzers** Not checked yet

NEW 70

Adds `base64_googlecloud_access_token` for base64-encoded Google Cloud OAuth2 access tokens, `base64_google_oauth2` for base64-encoded Google OAuth2 keys, `base64_google_oauth2_refresh_token` for base64-encoded refresh tokens, `base64_googleaiza` for base64-encoded Google API keys, `googlecloud_storage_hmac_key` for Google Cloud Storage HMAC keys, a detector for Google OAuth2 refresh tokens found without their client ID, and a detector for refresh tokens found with their client ID and secret, covering gcloud application default credentials. The Google API key analyzer now also supports base64-encoded keys.

— Names every new detector key, no usage example changelog-20261118-bef8ea9f

02 **New analyzers across Grafana, Cloudflare, DeepSeek, and more**

Not checked yet NEW 70

Adds new analyzers for Grafana service account tokens, Cloudinary API keys, DeepSeek API keys (reporting account balance with currency, availability status, and usable models), Groq API keys (reporting the models the key can use), Telegram bot tokens (reporting the bot username and name), GitLab self-managed personal access tokens, GitLab runner authentication tokens, New Relic admin keys, REST API keys, and user keys, Amp API tokens, Sourcegraph access tokens, Gemfury deploy, push, and full-access tokens, Mercado Pago access tokens and OAuth credentials, Cloudflare API tokens, R2 tokens, and global API keys, Facebook access tokens and app keys, LlamaCloud API keys, and Azure Entra app secrets.

— Names every new analyzer and some reporting output changelog-20261118-bef8ea9f

03 **Detector rename breaks existing rule references** Not checked yet

65

BREAKING

The `Google Cloud Express API Key` detector is renamed to `Google Gemini API Key`; the former `Google Gemini API Key` detector is renamed to `Google Bard Session Cookie`. Any rules, filters, or integrations referencing these detector names by their old labels will need to be updated.

— Names old/new labels but no migration tooling detailed `changelog-20261118-bef8ea9f`

---

04 **Private key validation against custom SSH hosts** Not checked yet

IMPROVED

65

The `private_key_rsa`, `base64_private_key_rsa`, `private_key_elliptic`, `base64_private_key_elliptic`, `private_key_openssh`, `base64_private_key_openssh`, `private_key_generic`, and `base64_private_key_generic` detectors can now validate keys against a custom SSH host such as a self-hosted GitLab or GitHub Enterprise instance.

— Names 8 detector keys but not how to configure the host `changelog-20261118-bef8ea9f`

---

05 **Vault detectors report instance host** Not checked yet IMPROVED

65

HashiCorp Vault `token`, `AppRole`, and `unseal key` detectors now also report the host of the Vault instance when it appears next to the credentials.

— Explains mechanism and trigger condition clearly `changelog-20261118-bef8ea9f`

THINNER COVERAGE BELOW

06 **New detectors for Jamf, Pydantic, Supabase, Food Safety Korea**

Not checked yet NEW

50

Adds new detectors for Jamf OAuth credentials, Pydantic Logfire tokens, Supabase personal access tokens (classic, scoped, and `sbp_v0_` prefix), and Food Safety Korea Open API keys.

— Names each service and token variant, no usage example `changelog-20261118-bef8ea9f`

---

07 **Linear API key checker returns INVALID status** Not checked yet

IMPROVED

50

The `linear_api_key` checker now returns `INVALID` instead of `FAILED_TO_CHECK` for revoked or invalid keys.

— Clear before/after status change, no broader scope `changelog-20261118-bef8ea9f`

---

08 **AWS Bedrock detection covers keys outside console** Not checked yet

IMPROVED

45

AWS Bedrock long-term API key detection now also covers keys created outside the console.

– States before/after scope but no mechanism detail changelog-20261118-bef8ea9f

09 **New AWS AppSync and CodeArtifact detectors** Not checked yet NEW

35

Adds new detectors for AWS AppSync API keys and AWS CodeArtifact authorization tokens.

– Bare naming of two new detectors, no further detail changelog-20261118-bef8ea9f

10 **Azure subscription key detector improved recall** Not checked yet

IMPROVED

30

The `azure_subscription_key` detector now has improved recall.

– Improvement named but no scope or measurement given changelog-20261118-bef8ea9f

└─▷ **BREAKING ON UPGRADE**

! The Google Cloud Express API Key detector is renamed to Google Gemini API Key; the former Google Gemini API Key detector is renamed to Google Bard Session Cookie — any rules, filters, or integrations referencing these detector names by their old labels will need to be updated.

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.167.12 NOTES ↗ RANK

Infisical is an open-source secret management platform that helps teams securely store, sync, and manage API keys, environment variables, and other sensitive configurations.

> Infisical's v0.167.12 release adds incident.io on-call auto-approvals and an AI judge for secret approval workflows, extends code-signing with certificate renewal/revocation and DigiCert/HSM visibility, introduces secret sync failure alerts, and rounds out Agent Vault with billing, metadata, and permission changes.

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

01 **Code signing certificate renewal and revocation** Not checked yet

NEW 65

Code signing now supports signer certificate renewal and revocation, and the UI surfaces the DigiCert order ID and HSM key label for issued certificates.

— Names DigiCert and HSM key label fields surfaced in UI. v0.167.12

THINNER COVERAGE BELOW

02 **Agent Vault billing, metadata, and permission updates**

Not checked yet IMPROVED 55

Agent Vault adds free-tier and per-request billing support, allows custom metadata to be attached to Agent Vault sessions, now checks granular permissions instead of requiring the admin role, and session log requests lead with a type icon for faster visual scanning.

— Groups four Agent Vault changes but lacks exact field or endpoint names. v0.167.12

03 **incident.io on-call auto-approvals** Not checked yet NEW

50

Infisical now auto-approves access requests via incident.io on-call schedule checks, integrating with incident.io's on-call scheduling to streamline approval workflows.

— Names the integration but not setup steps or config. v0.167.12

04 **Dedicated permission for audit log streams** Not checked yet

IMPROVED

45

Audit log streams now have a dedicated permission instead of relying on the organization profile permission, allowing more granular control over who can manage them.

— Names the permission change but no exact permission key or UI path. v0.167.12

05 **Secret sync failure alerts** Not checked yet NEW

35

Secret sync failures now trigger alerts through the alerting module, notifying users when a sync fails.

— Names the alerting module but no config or alert channel detail. v0.167.12

06 **Daily invitation email caps for free orgs** Not checked yet

IMPROVED

35

Daily organization invitation emails are now capped for free cloud organizations.

— States the limit exists but not the exact cap number. v0.167.12

07 **AI judge for secret approval workflows** Not checked yet NEW

30

Adds an AI judge capability for secret approval workflows, enabling automated evaluation of approval requests.

— Described only in prose with no mechanism or config detail. v0.167.12

WAS THIS USEFUL?  

GitGuardian Self-Hosted

Feedback ⓘ 2026.11 (2026-11-17)

BLOG ↗

GitGuardian Self-Hosted is a secrets detection platform deployed on-premises that scans source code and git repositories to prevent sensitive data leaks.

GitGuardian Self-Hosted 2026.11 adds ML risk scoring (0-100) for secret incidents, vaulted-secret validity via ggscout, per-database key rotation, password lockout policy, and new Helm global value propagation.

Turn on the machine learning risk score in Admin Area preferences to give every secret incident a score from 0 to 100; the Critical saved view lists open incidents scored 85 and above. Not checked yet

View a Validity column in the NHI inventory, with permissions, owner, and expiry in the identity detail, once `ggscout` checks each collected secret with its provider from inside your network. Not checked yet

Rotate each database's encryption keys separately from the Admin Area key rotation page. Not checked yet

Set ClickHouse encryption keys separately from the PostgreSQL ones using `miscEncryption.existingSecretKeys.clickhouseEncryptionKeys` to protect ClickHouse data if the PostgreSQL key leaks. Not checked yet

Control the password login lockout policy from Admin Area preferences with `max_login_failed_attempts` (set to `0` to disable) and `login_lockout_duration`; a password reset lifts a user's lockout. Not checked yet

+11 more

NEW DETECTION 10 new detectors in Secrets Detection Engine v2.174

UPDATED DETECTION 31 detector updates in Secrets Detection Engine v2.174

NEW DETECTION 2 new analyzers in Secrets Detection Engine v2.174

RETIRED DETECTION 1 deprecated detector in Secrets Detection Engine v2.174

NEW DETECTION 1 new detector in Secrets Detection Engine v2.175

NEW DETECTION 2 new checkers in Secrets Detection Engine v2.175

UPDATED DETECTION 8 checkers gain custom host support in v2.175

NEW DETECTION 1 new analyzer in Secrets Detection Engine v2.175

NEW INTEGRATION AWS S3 integration for secrets scanning

NEW INTEGRATION Harbor container registry integration (beta)

UPDATED INTEGRATION Google Artifact Registry integration adds keyless Workload Identity connection

UPDATED INTEGRATION [Federated JFrog Docker repositories now scanned]

└──▷ **BREAKING ON UPGRADE**

- ! Email and password authentication is now accepted only to sign in to the dashboard — scripts that called other endpoints with an email and password must use an API token instead.
- ! An upgrade to 2026.11 is refused with 'Not all required migrations have been applied' while background data migrations started in 2026.8 and 2026.9 are still running; wait for them to finish on the current version before retrying.
- ! Jira Cloud and Confluence Cloud historical scans now use the DB-less scan framework by default, which turns on the DB-less scan flag of each integration in the Admin Area, even where it had been turned off.
- ! The Google Artifact Registry migration to the new integration version must finish before upgrading to 2027.4.
- ! If `global.tolerations`, `global.nodeSelector`, or other propagated global values are set, ClickHouse, Keeper, and MinIO pods restart once during the upgrade and may be scheduled on other nodes; a volume already bound to a node excluded by a new selector keeps the pod Pending.

WAS THIS USEFUL? 👍 👎

◆ Application Security Posture

Cycode

Feedback ⓘ

Java4Shell (CVE-2026-114938): unauthenticated remote code execution in Java 17, 21, 25 a

BLOG ↗ RANK

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

> Cycode added two Threat Intelligence dashboard queries to help teams find exposure to Java4Shell (CVE-2026-114938), an unauthenticated RCE affecting Java 17, 21, 25 and 27.

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

01 **Java4Shell exposure queries in Threat Intelligence dashboard**

Not checked yet NEW 68

Cycode's Threat Intelligence dashboard gained two queries for CVE-2026-114938 (Java4Shell): a Dockerfiles query that lists Dockerfiles building on affected Java 17, 21, 25, or 27 images from Eclipse Temurin, Amazon Corretto, Azul Zulu, SapMachine, Liberica, Microsoft, or distroless; and an Image runtimes query that lists container images with a Java 17, 21, 25, or 27 runtime in the affected range, whether sourced from a JDK base image or a Debian, Ubuntu, Alpine, or RPM package.

— Names both queries and all supported vendors/distros, but no UI path given.

Java4Shell (CVE-2026-114938): unauthent…

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Corellium

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗ RANK

Advice, answers, updates, and more from the Corellium Team

> Corellium expanded virtual device OS coverage across iOS/iPadOS 27.x, 26.x, 18.x, 16.x, and 15.x, and brought on-premise deployments up to parity with the 7.14.0 cloud release.

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

01 **Virtual device support for new iOS/iPadOS builds** Not checked yet

NEW 60

Corellium added virtualization support for `iOS/iPadOS 27.2.0` (betas 1-4 for SA Falcon Premium only, RC1, RC2/GA), `27.0.1` (RC1/GA), and `27.0.0` (betas 1-8 for SA Falcon Premium only, RC1, RC2/GA); `26.7.2`, `26.7.1`, `26.7.0`, `26.6.2`, `26.6.1`, `26.6.0`, `26.5.2`, `26.5.1`, `26.5.0`, and `26.4.2` across GA, RC, and beta tiers (beta builds for SA Falcon Premium only); `iOS 18.7.11`, `18.7.10`, `18.7.9`, and `18.7.8`; and `iOS 16.7.16` and `iOS 15.8.8`.

. . . .

— Names exact builds but no mechanism beyond version listing changelog-20261118-86576f03

THINNER COVERAGE BELOW

02 **On-premise parity with 7.14.0 cloud release** Not checked yet

IMPROVED 25

On-premise deployments now have parity with the 7.14.0 cloud release.

— Bare statement with no detail on what changed changelog-20261118-86576f03

WAS THIS USEFUL? 👍 👎

Finite State Platform

Feedback ⓘ

November 17, 2026 - Manual component matching, replaced components, and the dependency v

BLOG ↗ RANK

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

> Finite State Platform's latest release enriches the component dependency API with per-node severity counts and adds deeper audit-trail and traceability support when components are replaced in the Bill of Materials.

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

01 **Severity counts on dependency nodes** Not checked yet IMPROVED

80

The `GET /api/public/v0/component-dependencies/{versionId}/{componentId}` endpoint now returns the highest vulnerability severity and finding counts per severity on each dependency node, using the same counts shown in the component list.

Retrieve dependency nodes for a firmware version with severity counts to quickly identify the most critical transitive dependencies.

```
$ curl -X GET
'https://platform.finitestate.io/api/public/v0/component-
dependencies/{versionId}/{componentId}' \
  -H 'Authorization: Bearer <token>'
```

— Names exact endpoint and includes a runnable curl example

November 17, 2026 - Manual component ma…

THINNER COVERAGE BELOW

02 **Audit trail for component replacement** Not checked yet NEW

50

Replacing a component now records who replaced it, when, and an accompanying comment in the component details, and both replacing and removing a replacement are now written to the Audit Trail.

— Describes behavior but no UI path or endpoint given

November 17, 2026 - Manual component ma…

03 **Original detection preserved after replacement** Not checked yet

IMPROVED

50

When a component is replaced, the original detection remains visible in the Bill of Materials row, including where it was detected (such as Binary SCA) and the original Match Evidence.

— Names BOM and Binary SCA but no exact navigation steps

November 17, 2026 - Manual component ma…

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-11-15 CODE ↗

CAPEv2 is an open-source malware sandbox that automatically executes and analyzes suspicious files to extract configurations, payloads, and behavioral indicators.

CAPEv2 adds a configurable junk filter in `web.conf` to control which archive files are skipped during demuxing

└──▷ **USE IT**

Filter archive children by magic only, disabling the name and path checks that can false-positive on files like `LicenseServer.exe` or a PE named `readme.txt`.

custom/conf/web.conf

```
[junk_filter]
enabled = yes
extensions =
.yar,.yara,.md,.txt,.yml,.yaml,.gitignore,.gitattributes,.gitmodules,.
pdb
names =
paths =
magics = Microsoft Roslyn C# debugging symbols,MSVC program
database,Microsoft C/C++ MSF,Microsoft C/C++ program database
```

Disable the junk filter entirely so every file extracted from an archive is submitted for analysis, without needing the per-task `ignore_junk_filter=1` option.

custom/conf/web.conf

```
[junk_filter]
enabled = no
```

Control junk-file skipping during archive demuxing with a new `[junk_filter]` section in `web.conf`, using keys `enabled`, `extensions`, `names`, `paths`, and `magics` — set any key to an empty value to disable that individual check, or set `enabled = no` to disable the filter entirely. Not checked yet

Populate `target.url` in URL analysis reports so report consumers reading that field (such as MAEC 5 and MAEC 4.1 reporters) now receive the submitted URL without relying solely on the processing-module result. Not checked yet

WAS THIS USEFUL? 👍 👎

karton

Feedback ⓘ v6.1.0 NOTES ↗ RANK

Karton is a distributed malware processing framework based on Python, Redis and S3.

> Karton v6.1.0 adds egress header filtering, letting service classes restrict which task types they are allowed to produce.

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

01 **Egress header filtering for karton services** Not checked yet NEW

50

Karton services can now restrict which task headers they are allowed to emit by declaring an `egress_filter` field and a `headers` claim on the service class.

— Names exact fields but lacks mechanism detail or example v6.1.0

WAS THIS USEFUL? 👍 👎

◆ **Digital Forensics & Incident Response**

ElcomSoft

Feedback ⓘ announced · Nov 17, 2026 BLOG ↗

ElcomSoft is a try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

iOS Forensic Toolkit 10.13 adds agent-based low-level extraction and partial file system extraction for iOS/iPadOS 26 through 26.7.3.

NEW SUPPORTED PLATFORM Adds agent-based low-level extraction support for the entire iOS/iPadOS 26 range, up to and including iOS 26.7.3. Not checked yet

Adds partial file system extraction that pulls application data across all iOS/iPadOS 26 devices. Not checked yet

WAS THIS USEFUL? 👍 👎

Obsidian Forensics Hindsight

Feedback ⓘ v2026.11

NOTES ↗ CODE ↗ RANK

Hindsight is a browser forensics tool for Google Chrome, other Chromium-based browsers, and Mozilla Firefox

> Hindsight's v2026.11 release centers on new CLI controls for filtering and logging, ETag decoding with cache/file integrity hashing, and LevelDB damage recovery, plus a handful of smaller parser improvements and three breaking changes affecting Python version, path resolution, and timestamp precision.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 Artifact filtering flags for parsing

Not checked yet NEW

85

Hindsight adds `--only` and `--skip` (comma-separated, repeatable) to filter which artifacts Hindsight parses, and `--list-artifacts` to list all valid artifact names accepted by those flags.

.

Quickly triage a Chrome profile by parsing only history and downloads, skipping the potentially large cache.

```
$ hindsight.py -i "C:\Users\Ryan\AppData\Local\Google\Chrome\User Data\Default" -o triage_case --only history,downloads
```

List all artifact names accepted by `--only` and `--skip` before writing a targeted collection script.

```
$ hindsight.py --list-artifacts
```

— Names exact flags with runnable example commands. v2026.11

02 Selectable log verbosity via --log-level

Not checked yet BREAKING

75

A new `--log-level` flag controls log verbosity (`debug` , `info` default, `warning` , or `error`). The log file no longer records DEBUG lines by default; the level is now selectable via this flag.

Run a full parse with debug-level logging to capture maximum detail when investigating a suspicious profile.

```
$ hindsight.py -i /home/user/.config/google-chrome/Default -o
full_case --log-level debug
```

— Flag, option values, and default-behavior change all given. v2026.11

03 **ETag decoding and cache integrity hashing** Not checked yet NEW

70

Hindsight now decodes server file time and size from download and cache ETags, checks hashes against sent bytes, and hashes cache bodies and File System backing files for integrity verification.

— Mechanism described but no flag or command to invoke it. v2026.11

04 **Path resolution for -l and --temp_dir** Not checked yet BREAKING

65

Relative `-l` and `--temp_dir` paths now resolve from the working directory; the default log is also written there (packaged executables still write next to themselves).

— Names flags and exact new path-resolution behavior. v2026.11

THINNER COVERAGE BELOW

05 **Session window and tab IDs in output** Not checked yet NEW

50

Session windows and tab/window IDs are now written to JSONL and SQLite output.

— Names output formats but no field-level detail. v2026.11

06 **LevelDB damage recovery and reporting** Not checked yet IMPROVED

45

Hindsight now recovers LevelDB records past damage and reports damage in every LevelDB parser.

— States the behavior but not the recovery mechanism. v2026.11

07 **Python 3.11 minimum version requirement** Not checked yet BREAKING

45

Python 3.11 or newer is now required (was 3.9).

— Exact version numbers give a clear upgrade action. v2026.11

08 **Chrome version detection from meta table** Not checked yet NEW

40

Hindsight now detects the Chrome version from each database's meta table.

— Single-line description with no further mechanism. v2026.11

09 **Structured IndexedDB database and object store separation**

Not checked yet IMPROVED

40

Hindsight now separates IndexedDB database from object store and keeps the value structured.

— Brief behavior note, lacks structural detail. v2026.11

10 **Credential reporting despite decryption failure** Not checked yet

IMPROVED

30

Hindsight now reports a saved credential even when its value cannot be decrypted.

— Simple behavior change, no mechanism given. v2026.11

11 **Firefox extension parse failure reporting** Not checked yet

IMPROVED

30

Hindsight now reports what was not parsed for Firefox extensions.

— Brief statement without further specifics. v2026.11

12 **Timestamp precision now source-dependent** Not checked yet

BREAKING

30

Timestamps are now written with the precision their source has, so precision varies between artifacts.

— Describes the change without naming affected artifacts. v2026.11

└──▷ ALSO FROM THESE RELEASES

.

.

└──▷ BREAKING ON UPGRADE

! Python 3.11 or newer is now required (was 3.9).

! Relative `-l` and `--temp_dir` paths now resolve from the working directory; the default log is also written there (packaged executables still write next to themselves).

! The log no longer records DEBUG lines by default; the level is now selectable via `--log-level`.

! Timestamps are now written with the precision their source has, so precision varies between artifacts.

WAS THIS USEFUL?  

mesh

Feedback ⓘ announced · Nov 16, 2026 BLOG ↗ RANK

A encrypted, censorship-resistant mesh VPN making remote mobile forensics and network monitoring seamless.

> mesh's only shipment this window was a Firecrawl kit for Docker Sandboxes that gives sandboxed agents live web search, scraping, and crawling behind a single allowed host, plus a macOS CLI install path for the `sbx` tool that drives it.

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

01 **Firecrawl kit for live web access in Docker Sandboxes**

Not checked yet NEW 95

Attaching the `docker.io/firecrawl/sbx-kit-firecrawl:1.0.0` kit via `--kit` gives a sandboxed agent live web search, page scraping, and site crawling, routing all web fetches through `api.firecrawl.dev` as the only new allowed host. The Firecrawl API key is stored with `sbx secret set firecrawl` for proxy-managed injection so the key never enters the sandbox image, appearing inside the sandbox only as `FIRECRAWL_API_KEY=proxy-managed`. The kit's `firecrawl/firecrawl.yaml` descriptor declares four typed capabilities: `lifecycle@1` (SDK install), `credential@1` (proxy-managed key), `network-policy@1` (allowing `pypi.org` and `api.firecrawl.dev`), and `agent-context@1` (AGENTS.md instructions).

Give a sandboxed coding agent live access to documentation sites that the balanced policy normally blocks, by attaching the Firecrawl kit alongside your workload and agent mixin.

```
$ sbx secret set firecrawl
sbx run docker/sbx-kit-shell:1.0.0 \
  --kit docker/sbx-kit-claude-mixin:2.1.285 \
  --kit docker.io/firecrawl/sbx-kit-firecrawl:1.0.0 .
```

Verify that the Firecrawl SDK, proxy-managed credential, and network rule are all working correctly from inside a running sandbox session.

```
$ !python3 -c "import firecrawl, importlib.metadata as m;
print('firecrawl-py', m.version('firecrawl-py'))"
!env | grep FIRECRAWL_API_KEY
!python3 - <<'PY'
from firecrawl import Firecrawl
fc = Firecrawl()
doc = fc.scrape("https://docs.pydantic.dev/latest/changelog/",
formats=["markdown"])
print(doc.markdown[:300])
PY
```

— Names exact kit, host, secret command, and all four descriptor capabilities with runnable commands.

launch-20261116-a959d091

02 **sbx CLI installation and policy requirements** Not checked yet

NEW

75

The `sbx` CLI can be installed on macOS with `brew trust docker/tap && brew install docker/tap/sbx`, requiring version 0.46 or later. A one-time `sbx policy init balanced` global network policy must be run before any sandbox can start.

— Exact install and policy commands given, but no further mechanism described.

launch-20261116-a959d091

WAS THIS USEFUL?

NETWORK SECURITY

◆ Network Analysis & Monitoring

Ettercap

Feedback ⓘ v0.8.5.2 NOTES ↗ CODE ↗ RANK

Ettercap is an open-source network security tool for man-in-the-middle attacks, enabling live packet sniffing, filtering, and protocol analysis on local networks.

> Ettercap v0.8.5.2 expands its filter scripting language with new payload-length tests, probabilistic/modulo rewrite functions, manual checksum control, and a startup flag to disable filters entirely.

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

01 **Filter language additions: `--no-filters`, `mod()`/`random()`, length tests, checksum override**

Not checked yet NEW 78

The etterfilter language gains `DATA.len` and `DECODED.len` fields to test packet payload length, mod() and random() functions for modulo- or probability-based rewrite logic, and manual checksum override in filter scripts. A new `--no-filters` command-line flag disables all filters at startup.

Run a MitM session without loading any filters, useful for clean traffic capture before iterating on a filter script.

```
$ ettercap -T -M arp /192.168.0.1// /192.168.0.2// --no-filters
```

Use random() in an etterfilter script to rewrite only roughly half of matching HTTP responses, simulating intermittent injection.

```
$ if (ip.proto == TCP && tcp.dst == 80 && DATA.len > 0) {
  if (random() % 2 == 0) {
    replace("</body>", "<script>alert(1)</script></body>");
  }
}
```

— Names flags/functions/fields with example commands but no deeper mechanism v0.8.5.2

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Docker CLI

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗

RANK

Docker CLI is a docker Documentation is the official Docker library of resources, manuals, and guides to help you containerize applications.

> Docker Engine 29.10.0 adds experimental daemon-native one-shot and cron-scheduled jobs and nftables-based Swarm support, moves the Engine API to an OpenAPI 3.2 spec (API 1.57), and makes a breaking change to insecure-registry defaults so loopback addresses are no longer implicitly trusted — alongside new AlmaLinux/Oracle/Rocky Linux packages and several smaller hardening and storage-driver fixes.

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

01 **Insecure registries no longer default to loopback** Not checked yet

95

BREAKING

Setting `"insecure-registries": []` in daemon configuration (`/etc/docker/daemon.json`) now disables all default insecure registries. As a breaking change, when `insecure-registries` is explicitly set (or `--insecure-registry` is used), Docker Engine no longer automatically treats loopback addresses (`127.0.0.0/8` , `::1/128`) as insecure — pulls and pushes to local HTTP registries may fail unless addresses such as `localhost:5000` are explicitly added to `insecure-registries` .

Restrict insecure registry access to only a specific local registry, disabling the loopback-wide default — important after upgrading when you run a local HTTP registry.

```json
{
  "insecure-registries": ["localhost:5000"]
}
```

— Exact config key, addresses, file and migration example given

changelog-20261118-35fe0aa7

02 **Native one-shot and scheduled jobs in daemon** Not checked yet

80

NEW

Enable experimental one-shot and cron-scheduled workloads managed natively by the daemon using `--enable-extension org.mobyproject.jobs.v0` .

Run one-shot or scheduled jobs natively in the daemon — useful for replacing external cron orchestration with daemon-managed tasks.

```
$ dockerd --enable-extension org.mobyproject.jobs.v0
```

— Flag and runnable command given, but experimental scope unclear

changelog-20261118-35fe0aa7

---

03 **Engine API 1.57 with OpenAPI 3.2 specification** Not checked yet

80

BREAKING

Updates the Engine API to version 1.57, with the specification now generated from OpenAPI 3.2 in `api/openapi.yaml`, while a Swagger 2.0 spec continues to be generated in `api/swagger.yaml` and a pinned Swagger 2.0 copy for API v1.57 is available at `api/docs/v1.57-swagger.yaml`. Future API versions will not get a per-version Swagger 2.0 copy; tooling depending on `api/docs/v1.57-swagger.yaml` `-style` paths must migrate to the OpenAPI 3.2 spec at `api/openapi.yaml`.

— Exact file paths named; migration path clear but no command changelog-20261118-35fe0aa7

---

04 **Embedded containerd configuration via drop-in files**

Not checked yet NEW

70

Customize Docker's embedded containerd by placing configuration files in `/etc/docker/containerd.d/`.

Tune Docker's embedded containerd (e.g. set a custom snapshot plugin) without rebuilding the daemon binary.

```toml
# /etc/docker/containerd.d/custom.toml
[plugins.'io.containerd.snapshotter.v1.overlayfs']
  root_path = '/var/lib/containerd/snapshots'
```

— Config path and a concrete example file given changelog-20261118-35fe0aa7

THINNER COVERAGE BELOW

---

05 **`restarting` attribute on die events** Not checked yet NEW

50

New `restarting` attribute on `die` events tells consumers when a container is about to restart.

— Names the field but gives no consumption example changelog-20261118-35fe0aa7

---

06 **overlay2 rootless mode with SELinux enabled** Not checked yet

IMPROVED

40

Allow the overlay2 storage driver in rootless mode when SELinux is enabled.

— Describes the change but gives no config example changelog-20261118-35fe0aa7

07 **Improved ARM64 image variant selection** Not checked yet IMPROVED

40

Improve image selection on ARM64 Linux hosts to prefer the highest image variant supported by the host CPU.

— Explains behavior change but no measurable detail changelog-20261118-35fe0aa7

08 **Blocked obsolete socket families in default profiles**

Not checked yet IMPROVED

35

Block obsolete socket families in the default AppArmor and seccomp profiles.

— States the change but not which families or profiles changelog-20261118-35fe0aa7

09 **Engine packages for AlmaLinux, Oracle Linux, Rocky Linux**

Not checked yet NEW

30

Docker Engine packages are now available for AlmaLinux, Oracle Linux, and Rocky Linux.

— Names distros but no install command or repo path changelog-20261118-35fe0aa7

10 **nftables firewall backend support for Swarm** Not checked yet NEW

25

Engines using the nftables firewall backend can now join a Swarm EXPERIMENTAL.

— Bare statement, no command or config shown changelog-20261118-35fe0aa7

└──▷ **BREAKING ON UPGRADE**

! When `insecure-registries` is explicitly set in daemon config (or `--insecure-registry` is used), Docker Engine no longer automatically treats loopback addresses (`127.0.0.0/8` , `::1/128`) as insecure — pulls and pushes to local HTTP registries may fail unless those addresses (e.g. `localhost:5000`) are explicitly added to `insecure-registries` .

! Future API versions will not get a per-version Swagger 2.0 copy; tooling that depends on `api/docs/v1.57-swagger.yaml` `-style` paths must migrate to the OpenAPI 3.2 spec at `api/openapi.yaml` .

WAS THIS USEFUL? 👍 👎

kata-containers

Feedback ⓘ 4.4.0 NOTES ↗

Kata Containers is an open source project and community working to build a standard implementation of lightweight Virtual Machines (VMs) that feel and perform like containers, but provide the workload isolation and security advantages of VMs. https://katacontainers.io/

Kata Containers 4.4.0 switches guest disk images to EROFS, raises the Cloud Hypervisor vCPU ceiling to 512 on x86_64, and retires kata-manager.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 4.4.0 https://github.com/kata-containers/kata-containers.git
# already have the repo? check out this version:
$ git checkout 4.4.0
```

NEW SUPPORTED PLATFORM Guest disk images now use EROFS format by default, replacing the previous rootfs image format. Not checked yet

UPDATED SUPPORTED PLATFORM Cloud Hypervisor on x86_64 now allows up to 512 vCPUs per VM, up from the previous limit. Not checked yet

Rootfs partitions are now selected by PARTUUID across QEMU, CLH, and Dragonball VMMs. Not checked yet

`kata-deploy` now allows appending to CRI configuration files rather than only overwriting them. Not checked yet

The `kata-ctl` tool is now required and shipped for devkit installations and included in the nvgpu bundle. Not checked yet

+2 more

NEW SUPPORTED PLATFORM VFIO-AP passthrough rootless VM on s390x

UPDATED SUPPORTED PLATFORM Kernel updated to v6.18.62

UPDATED SUPPORTED PLATFORM QEMU updated to v11.1.3

UPDATED SUPPORTED PLATFORM OVMF updated to edk2-stable202608

UPDATED SUPPORTED PLATFORM nydus updated to v2.4.5

UPDATED SUPPORTED PLATFORM nydus-snapshotter updated to v0.16.1

NEW INTEGRATION td-shim firmware built and shipped in release tarball

└──▷ BREAKING ON UPGRADE

! Minimum required containerd version is now v2.3; setups running an older containerd

will break on upgrade.

- ! The `kata-manager` utility is retired and no longer available; users must switch to supported deployment methods.
- ! The `RemoveContainerRequest.timeout` field has been removed from the agent and runtime-rs API.

WAS THIS USEFUL? 👍 👎

Docker Compose

Feedback ⓘ v5.8.0 NOTES ↗ RANK

Docker Compose is a define and run multi-container applications with Docker

> Docker Compose v5.8.0 introduces a top-level `jobs:` element for one-shot and scheduled workloads, while v5.7.2 streams live logs during `docker compose up --wait`.

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

01 Scheduled and one-shot jobs via `jobs:` element

Not checked yet

NEW

93

A new top-level `jobs:` element in the Compose file declares one-shot or scheduled workloads, with `triggers.schedule` accepting `cron` expressions. Scheduling requires Docker Engine 29.10+ with `dockerd --extension-enable org.mobyproject.jobs.v0`; manual jobs work on any engine. Running `docker compose up` or `docker compose create` registers scheduled jobs with the engine (which fires them on its own clock), `docker compose down` removes them, and `docker compose run <job>` triggers a job immediately, starting its `depends_on` services first, streaming logs, and returning the job's exit code. On an engine without jobs support, `docker compose up` refuses to start a project declaring scheduled jobs before creating anything.

Schedule a nightly database backup job that runs after the db service is healthy — declare it in your Compose file and register it with the engine.

yaml

```yaml
services:
  db:
    image: postgres
jobs:
  backup:
    image: backup-tool
    depends_on: [db]
    triggers:
      schedule:
        - cron: "0 3 * * *"
```

Register scheduled jobs with the engine (requires Docker Engine 29.10+ with the jobs extension enabled via `dockerd --extension-enable org.mobyproject.jobs.v0`).

```
$ docker compose up
```

Trigger a job immediately outside its schedule — useful for ad-hoc runs or testing the job in CI.

```
$ docker compose run backup
```

— Names config element, flags, engine version requirement, and exact commands v5.8.0

THINNER COVERAGE BELOW

02 **Live log streaming during `--wait`** Not checked yet IMPROVED

55

`docker compose up --wait` now streams service logs in real time instead of waiting silently, making it easier to diagnose failures in CI while Compose waits for services to become healthy.

Watch live logs while Compose waits for all services to pass their healthchecks — useful in CI where silent waits make failures hard to diagnose.

```
$ docker compose up --wait
```

— Clear flag and behaviour but thin on mechanism v5.7.2

└──▷ **BREAKING ON UPGRADE**

! On an engine without jobs support (Docker Engine older than 29.10 or jobs extension not enabled), `docker compose up` refuses to start a project that declares scheduled jobs before creating anything.

WAS THIS USEFUL? 👍 👎

Chainguard

Feedback ⓘ Week of 2026-11-16 NOTES ↗ RANK

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

> Chainguard rolled out self-serve subscription plans, launched a memory-safety rewrite initiative with its first Rust-based allocator, disclosed and remediated a batch of silent Java vulnerabilities, and shipped 25 new container images.

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

01 **Athena discloses silent Java vulnerabilities with remediated versions**

Not checked yet IMPROVED 83

Athena discloses 12 additional silent vulnerabilities across nine Java libraries, including two critical authentication bypasses. Remediated versions are available via Chainguard Repository using a `-0cgr.n` version qualifier, requiring only a lockfile change and rebuild.

— Gives exact qualifier and remediation mechanism Week of 2026-11-16

02 **Project Theseus memory-safe rewrite initiative** Not checked yet

NEW 65

Chainguard launches Project Theseus, an initiative to rewrite core open-source components in memory-safe languages. The first release is `challoc`, a Rust-based memory allocator built to replace glibc malloc, jemalloc, tcmalloc, and mimalloc.

— Names the allocator and its targets, no install steps given Week of 2026-11-16

THINNER COVERAGE BELOW

03 **25 new container images added** Not checked yet NEW 57

Chainguard adds 25 new container images this week, including `aws-route53-controller`, `grafana-pyroscope-fips`, `dify-web`, `kgateway-envoy-wrapper`, `kgateway-sds`, `logstash-with-opensearch-syslog-output-plugins`, `querybook`, `neuvector-registry-adapter`, `oc-mirror`, `gosu-fips`, and others across application, base, and FIPS tiers.

— Lists image names but no pull command or usage example Week of 2026-11-16

04 **Self-serve Starter and Team subscription plans** Not checked yet

NEW 50

Chainguard now offers self-serve Starter (free, no credit card required) and Team (paid, card billing) subscription plans, giving teams access to Chainguard Containers, Chainguard Libraries, and other products without needing a sales call.

— Names the plans but no signup flow or pricing detail Week of 2026-11-16

WAS THIS USEFUL?  

◆ Cloud Security Posture

Plexicus

Feedback ⓘ

Plexicus v1.2.79: Discard & Retry Remediations, Clearer Scan Failures, Team & Dashboard I

NOTES ↗ RANK

Plexicus validates security findings with evidence and helps your team fix what matters. AI Swarm Pentest, Deep Code Analysis, Remediation and XBOM.

> Plexicus shipped remediation retry controls, clearer scan and AI SAST failure diagnostics, finer-grained AI Pentest retest management, full pod isolation for the AI Pentest engine, and an SSO hardening pass across three releases.

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

01 **AI Pentest engine pod isolation** Not checked yet IMPROVED

60

AI Pentest engine pods are now fully isolated from platform storage and cluster credentials, preventing lateral movement during pentest runs.

— Explains mechanism and benefit but not user-facing

Plexicus v1.2.83: Hardened AI Pentest I…

THINNER COVERAGE BELOW

02 **AI Pentest retest management** Not checked yet IMPROVED

55

In-progress AI Pentest retests can now be cancelled and inspected from within the product, and retest runs respect excluded domains so they no longer re-attack them, narrowing retest scope to only in-scope targets.

— Describes behaviour and scope change but no UI/API detail

Plexicus v1.2.83: Hardened AI Pentest I… Plexicus v1.2.81: Retest Management, Sm…

03 **Clearer scan and AI SAST failure diagnostics** Not checked yet

IMPROVED

45

Scan results now show clear failure reasons, and AI SAST failure diagnostics surface clearer error details to help diagnose scan problems.

— Names AI SAST but no concrete error fields or UI location

Plexicus v1.2.81: Retest Management, Sm… Plexicus v1.2.79: Discard & Retry Remed…

04 **Discard and retry AI remediation** Not checked yet NEW

40

Users can discard a bad AI-generated remediation and retry to generate a fresh one.

— Names the action but no mechanism or UI path Plexicus v1.2.79: Discard & Retry Remed…

05 **SSO authentication hardening** Not checked yet IMPROVED

30

SSO login, configuration, and role management received a hardening pass covering authentication flows and access controls.

— Vague on what specifically changed in auth flows

Plexicus v1.2.83: Hardened AI Pentest I…

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - November 16, 2026

BLOG ↗ RANK

Sysdig Secure is a cloud-native security platform that detects threats, manages vulnerabilities, and ensures compliance across containers, Kubernetes, and cloud environments.

> Sysdig Secure added an exception-aware filter to the Identity Findings page so findings with active exceptions no longer inflate identity totals.

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

01 **Has Exception filter in Identity Findings** Not checked yet NEW

60

The Identity Findings page gains a Has Exception filter that hides findings with an active exception by default and excludes them from identity totals, letting users focus on actionable findings.

Review only actionable identity findings by hiding all entries that already have an active exception applied.

> 📍 In Sysdig Secure, go to Identity Findings and confirm the Has Exception filter is enabled (it is on by default) to see only findings without active exceptions.

— Names exact UI filter and default behavior, but limited scope detail.

Secure SaaS - November 16, 2026

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Custom checks run on continuous scanning:

NOTES ↗ RANK

Plerion is a cloud security platform that provides comprehensive risk management, vulnerability detection, and threat prioritization across cloud environments and workloads.

> Plerion's custom checks now re-evaluate automatically under continuous scanning via a new verdict field, alongside two new AWS checks for stale EKS Kubernetes versions and compromised IAM keys, and expanded public-accessibility coverage for resources behind AWS Global Accelerator.

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

01 relatedAssetIds field for custom check re-evaluation

Not checked yet NEW 83

Custom checks can return asset IDs in a new `relatedAssetIds` field of their verdict; under continuous scanning, if any of those related assets change, Plerion re-evaluates the finding within minutes, evaluating only the changed resources rather than the full check.

Re-evaluate a finding whenever a Lambda function's IAM role changes, by returning the role's asset ID from the custom check verdict.

```json
{
  "pass": false,
  "message": "Lambda has overly permissive role",
  "relatedAssetIds": ["arn:aws:iam::123456789012:role/my-lambda-role"]
}
```

— Names exact field and gives a runnable verdict example

Custom checks run on continuous scannin…

02 IAM compromised key quarantine detection Not checked yet NEW

65

New check `[PLERION-AWS-1255]` detects IAM identities with an `AWSCompromisedKeyQuarantine` policy (original, V2, V3, V4, or later) attached, flagging leaked access keys as Critical severity and noting when they are publicly exposed.

— Names check ID and policy variants but no command

Custom checks run on continuous scannin…

03 **Global Accelerator scope for public-accessibility checks**

Not checked yet IMPROVED 65

ECS tasks and EC2 instances behind an AWS Global Accelerator are now in scope for public-accessibility checks `[PLERION-AWS-755]` and `[PLERION-AWS-194]`, with the attack path view showing the accelerator, its listener, and endpoint group.

— Names check IDs and attack path detail, no navigation steps

Custom checks run on continuous scannin…

04 **EKS Kubernetes version end-of-support check** Not checked yet NEW

62

New check `[PLERION-AWS-1256]` detects EKS clusters or managed node groups running a Kubernetes version past end of AWS standard support, checking each control plane and node group independently.

— Names check ID and logic but no navigation or command

Custom checks run on continuous scannin…

WAS THIS USEFUL?  

Wiz Cloud

Feedback ⓘ changelog · Nov 16, 2026 BLOG ↗

Wiz Cloud is a cloud security platform that scans cloud environments to identify risks, misconfigurations, vulnerabilities, and compliance issues across multi-cloud infrastructure.

Wiz adds the AI Security Agent in Public Preview, connecting AI application-layer analysis with Wiz Security Graph cloud context to map end-to-end attack paths.

NEW INTEGRATION Enable the Wiz AI Security Agent in Public Preview via the Wiz Preview Hub to automatically inventory custom AI applications, models, system prompts, MCP servers, and tool definitions across code repositories, PaaS platforms, and self-hosted environments. Not checked yet

Wiz AI Security Agent analyzes AI application logic, instructions, and permissions to compare an agent's intended purpose against its actual infrastructure access, surfacing discrepancies such as internal agents being externally exposed or holding overprivileged data access. Not checked yet

Wiz AI Security Agent uses AI-powered reasoning to uncover AI-native risks and correlates findings with the Wiz Security Graph to expose end-to-end cloud attack paths including identities, permissions, and sensitive data. Not checked yet

NEW DETECTION AI-native risk detection: prompt injection and arbitrary code execution

WAS THIS USEFUL? 👍 👎

◆ SaaS Security Posture

Nudge Security

Feedback ⓘ

Add a note when you accept and close a finding · NOTES ↗ · RANK

Nudge Security is a SaaS security platform that discovers shadow IT assets, manages identity governance, and hardens security posture.

> Nudge Security added two targeted workflow improvements: risk-acceptance notes on findings and CSV export of filtered external file shares.

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

01 CSV export of filtered external file shares

Not checked yet · NEW

78

From `Inventory > External File Shares`, the external file shares table can be exported to CSV with active filters applied, including file/folder ID and name, owner, risk insights, collaborator emails and roles, last accessed, and last modified.

Isolate risky external file shares or those shared with a specific domain, then export exactly that filtered set for audit or incident response.

> 📍 In the console, go to Inventory › External File Shares, apply your filters (e.g. by risk level or external domain), then click Export to download the filtered results as a CSV.

– Names exact navigation path and all exported fields

Export filtered external file shares to…

THINNER COVERAGE BELOW

02 Notes on accepted and closed findings

Not checked yet · NEW

55

Admins can add an optional note when accepting and closing a finding via Accept & close, with the note captured in the finding timeline so the risk-acceptance rationale is preserved for future reviewers.

Record the rationale for an accepted risk so future reviewers can see why a finding was closed without remediation.

> 📍 In the Nudge Security console, open any finding, click Accept & close, type your justification note in the note field, then confirm.

– Clear UI path but no deeper mechanism detail · Add a note when you accept and close a …

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

XFA

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗ RANK

> XFA introduces three new device application checks — code signature verification, install protection, and risky app category detection — in beta for Max and Enterprise plans.

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

01 **Application checks: safety, install protection, risky categories**

Not checked yet NEW 65

XFA adds a new Application safety check that verifies the code signature of apps installed on devices, a new Install protection check that confirms Gatekeeper (macOS) or SmartScreen (Windows) is enabled, and a new Risky app categories check that detects apps such as crypto miners and hacking tools. These application checks (along with AI checks) are available in beta, gated to Max and Enterprise plans, and require employee opt-in consent before results are shared with an organization.

Walk through enabling the new application checks in your policy so devices are evaluated for risky apps and install protection.

📍 In the XFA Dashboard, go to Policies and select your target policy. Under Application checks, enable Application safety, Install protection, and Risky app categories, then save. Employees will be prompted to opt in before results are shared with your organization.

— Names each check and the plan/consent gating; UI steps given but no API/config keys.

changelog-20261118-a501a977

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ changelog · Nov 17, 2026

BLOG ↗ RANK

Microsoft Defender is a built-in security tool that protects Windows devices from malware, viruses, ransomware, and other cybersecurity threats in real time.

> Microsoft Defender launched an Integrated Security Operations Center combining Sentinel SIEM, UEBA, threat intelligence and threat disruption, alongside new agentic-security capabilities including Dynamic Agent Risk assessment, Purview DLP protection for agent workloads, and expanded Perception playbook connectivity.

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

01 **Integrated Security Operations Center with Sentinel SIEM bundle**

Not checked yet NEW 80

Microsoft Defender adds an Integrated Security Operations Center (ISOC) combining SIEM, UEBA, threat intelligence, and threat disruption, generally available for core capabilities. Eligible Defender Suite, Microsoft 365 E5, and Microsoft 365 E7 customers now get all Microsoft Sentinel SIEM features, 90-day default data retention for Defender data, more than 500 connectors, reduced ingestion rates, and a built-in data lake.

— Names components, numeric limits and eligible plans but no setup steps

2026-11-17-9f1d260b

02 **Perception playbook connectivity via MCP and cloud signals**

Not checked yet NEW 60

Defender's Perception playbooks can now connect to SIEM, SOAR, and IT service management tools via Model Context Protocol (MCP) and APIs, and can ingest signals from AWS, Okta, and Salesforce, both now in preview.

— Names MCP, APIs and three cloud sources but preview with no usage steps

2026-11-17-9f1d260b

THINNER COVERAGE BELOW

03 **Dynamic Agent Risk assessment** Not checked yet NEW 35

Microsoft Defender adds Dynamic Agent Risk to assess each agent's exposure, blast radius, and runtime behavior, available in preview.

— Describes purpose but no mechanism, scope or config detail 2026-11-17-9f1d260b

04 **Purview DLP protection for agentic workloads** Not checked yet

NEW

30

Microsoft Defender and Microsoft Purview Data Loss Prevention now protect runtime and sensitive data access for agentic workloads, available in preview.

— Brief prose description, no concrete controls or setup named `2026-11-17-9f1d260b`

WAS THIS USEFUL? 👍 👎

Microsoft Security

Feedback ⓘ changelog · Nov 17, 2026

BLOG ↗ RANK

Microsoft Security is a find official documentation, practical know-how, and expert guidance for builders working and troubleshooting in Microsoft products.

> Microsoft Perception reaches general availability as an orchestration platform for red, blue, and green team agents running closed-loop playbooks, alongside newly GA posture and detection agents, preview integrations into SIEM/SOAR/ITSM tools, an expanded Codename MDASH code scanner, and new cost-management tooling.

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

01 **Red team Recon and Web App Pen Test agents** Not checked yet NEW

80

Recon Agent is now generally available, and the Web App Pen Test Agent is now in preview, coordinated through red team playbooks to find and validate weaknesses proactively. In Perception, going to Playbooks > Protect web app and running the playbook triggers the Web App Pen Test Agent to recon the target application, surface exploitable weaknesses, and hand off remediation to a connected GitHub repository.

Trigger a full web application pen test via the Web App Pen Test Agent playbook in Perception to surface exploitable weaknesses and push fixes to GitHub before attackers find them.

> 📍 In Perception, go to Playbooks › Protect web app, select your target application, and run the playbook — the Web App Pen Test Agent will recon the app, surface weaknesses, and hand off remediation to the connected GitHub repository.

— Includes exact UI navigation and end-to-end workflow 2026-11-17-8b8e607a

02 **Codename MDASH scanner expands scale and models** Not checked yet

IMPROVED

63

Codename MDASH now analyzes repositories up to 5 GB, runs in regulated clouds, and adds `MAI Cyber 1.1 Flash` and `MAI-Cyber-2` model options in preview.

— Gives concrete size limit and named model options, no direct command

2026-11-17-8b8e607a

THINNER COVERAGE BELOW

03 **SIEM, SOAR, and ITSM integrations via MCP** Not checked yet NEW

55

Perception playbooks can now connect to SIEM, SOAR, and IT service management tools through Model Context Protocol (MCP) and APIs, and pull in signals from AWS, Okta, and Salesforce — these integrations are in preview.

— Names protocol and vendors but no endpoint or config detail 2026-11-17-8b8e607a

04 **Four blue team agents reach GA** Not checked yet NEW 50

Four blue team agents — Alert Triage Agent, Attack Investigation Agent, Detection Authoring Agent, and Threat Intelligence Agent — are now generally available inside Perception.

— Names four agents but gives no usage or workflow detail 2026-11-17-8b8e607a

05 **Microsoft Perception orchestration platform reaches GA**

Not checked yet NEW 45

Microsoft Perception is now generally available, orchestrating red, blue, and green team agents through playbooks in a closed loop for detection, investigation, and remediation.

— Describes orchestration concept broadly without a specific command or config 2026-11-17-8b8e607a

06 **Posture Prioritization and Risk Reduction agents** Not checked yet

NEW 43

Posture Prioritization Agent is now generally available and the Risk Reduction Agent is now in public preview, covering Microsoft, AWS, and Google Cloud environments.

— Names agents and cloud scope but no mechanism or usage path 2026-11-17-8b8e607a

07 **AI Readiness Initiatives planning tool** Not checked yet NEW

40

AI Readiness Initiatives in Perception — turning Secure Now guidance into actionable plans with shared progress tracking — are coming in preview in January 2027.

— Future preview described in prose with a date but no mechanism 2026-11-17-8b8e607a

08 **Cost management dashboards and estimator** Not checked yet NEW

33

New cost management dashboards and a cost estimator in Perception let teams track AI usage, forecast spending, and connect costs to security outcomes.

— Generic description of dashboards without specifics 2026-11-17-8b8e607a

WAS THIS USEFUL?  

Trend Vision One

Feedback ⓘ

Prevent SMTP recipient enumeration (2026-11-17) BLOG ↗ RANK

Trend Vision One is a unified cybersecurity platform that integrates threat detection, response, and risk management across endpoints, networks, and cloud environments.

> Trend Vision One's Cloud Email Gateway Protection gains a recipient filtering mode that blocks SMTP recipient enumeration attacks.

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

01 **Accept and validate after SMTP mode** Not checked yet NEW

75

A new Accept and validate after SMTP mode in Recipient filter settings of Cloud Email Gateway Protection accepts all recipients during the SMTP session, then silently drops mail for invalid recipients without sending non-delivery reports. This prevents attackers from enumerating valid email addresses on the protected domain.

Block SMTP recipient enumeration attacks against your domain by silently discarding mail to invalid addresses rather than issuing non-delivery reports that confirm address validity.

> 📍 In Cloud Email Gateway Protection, go to Recipient filter settings and set the mode to Accept and validate after SMTP.

— Names exact setting and UI path, mechanism explained clearly.

Prevent SMTP recipient enumeration (202…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Smart Routing now supports external models for coding agents (Beta) (2026-11-18)

BLOG ↗ RANK

Databricks is a unified data analytics platform that enables collaborative data engineering, machine learning, and analytics using Apache Spark in the cloud.

> Databricks expanded Unity Gateway with external model routing for coding-agent subagents and new usage/token-waste insights, while also previewing efficient Delta Lake column updates, an upgraded Files API 3.0, and a new Connectors section for Genie Agents.

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

01 External model routing for coding-agent subagents

Not checked yet

NEW

75

Unity Gateway's Smart Routing BETA can now dispatch coding-agent subagents to external models behind a model provider service by adding them to `models.sources` in the Unity Gateway Managed Configuration, alongside built-in `system.ai` models.

Add an external model source for coding-agent subagents so Smart Routing can dispatch to it alongside built-in system.ai models.

```yaml
models:
  sources:
    - system.ai
    - <your-external-model-provider-service>
```

— Names exact config key but no rollout or limits detail

Smart Routing now supports external mod…

02 Efficient Column Updates for Delta Lake

Not checked yet NEW

70

A Beta feature lets workspace admins add and backfill columns in Delta Lake without rewriting unchanged column data, enabled by going to Settings > Previews and turning on Efficient Column Updates.

Workspace admin enabling Efficient Column Updates so teams can add and backfill Delta Lake columns without full table rewrites

📍 In the Databricks workspace, go to Settings › Previews and enable Efficient Column Updates.

— Exact admin UI path given, no performance figures

Efficient column updates (Beta) (2026-1…

03 **Multipart uploads and conditional operations in Files API 3.0**

Not checked yet NEW 70

Files API 3.0, soon generally available, supports multipart uploads for files larger than 5 GB to Unity Catalog volumes and conditional operations that check the current file version before uploads, downloads, and deletions to guard against unintended changes.

— Names API version and size limit, no sample request shown

Files API 3.0 will soon be generally av…

04 **Connectors section in Genie Agent settings** Not checked yet NEW 65

Genie Agent settings now has a dedicated Connectors section BETA where users can configure connectors for an agent via Configure settings > Connectors > Add connector.

Add a connector to a Genie Agent using the new dedicated Connectors section in agent settings.

📍 In the Genie Agent settings, go to Configure settings › Connectors, then select Add connector to attach a data connector to the agent.

— Exact UI navigation given, mechanism otherwise undescribed

Genie Agents enhancements (2026-11-12)

THINNER COVERAGE BELOW

05 **Unity Gateway Insights reports** Not checked yet NEW 45

A new Beta reporting feature in Unity Gateway shows what users invoke agents for and where tokens are wasted, built from a sample of sessions in the unified trace table.

— Describes purpose but no endpoint or UI path given

Analyze agent usage and find wasted tok…

NEW SUPPORTED PLATFORM T-SQL to Lakeflow pipeline conversion (Beta)

└──▷ ALSO FROM THESE RELEASES

Genie Migrate can convert T-SQL to Lakeflow pipelines (Beta) (2026-11-13)

BLOG ↗ STABLE

Genie Migrate adds Beta T-SQL to Lakeflow pipeline conversion; workspace admins enable it via Migration Projects Experimental Features on the Previews page.

WAS THIS USEFUL?  

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261117.0 NOTES ↗ RANK

OpenCTI is an open-source threat intelligence platform that helps organizations manage, analyze, and share cyber threat intelligence data.

> OpenCTI's single release in this window adds time-series field breakdowns and variable presets to dashboards, feed-ingestion failure alerting, variable support in mailer notification titles, and new STIX 2.0 object exports in the Python client.

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

01 **Time-series widgets broken down by field value** Not checked yet

NEW 53

Dashboards now support breaking down time-series widgets (`bars` , `lines` , `areas` , `heatmaps`) by field value, letting analysts slice a single chart instead of building separate widgets per value.

– Names widget types but no exact UI path or config 7.261117.0

02 **Dashboard variable handling and presets** Not checked yet NEW

45

OpenCTI now supports dashboard variables in export, import, and public dashboards, preserving variable context when dashboards are shared or moved, and adds presets to save and restore sets of dashboard variable values so teams can switch between common filter configurations on shared dashboards.

– Describes behaviour but no concrete key or command 7.261117.0

03 **Variables in mailer notification titles** Not checked yet NEW

37

OpenCTI now allows variables in Platform Mailer notifier titles, so notification subjects can carry dynamic context.

– Names the notifier but no example syntax 7.261117.0

04 **Activity log and lock message for locked users** Not checked yet

NEW 30

OpenCTI adds activity log entries and a lock message for locked users as part of the user-lock workflow.

– Names the workflow but not the exact message or log fields 7.261117.0

05 **Feed ingestion alerting** Not checked yet NEW 28

OpenCTI now alerts on feed ingestion issues and failures, surfacing connector problems without requiring manual monitoring.

— No detail on alert mechanism or configuration 7.261117.0

06 **One-click enrichment background tasks** Not checked yet IMPROVED 25

Enrichment background tasks now launch without an extra confirmation click, reducing friction when triaging indicators.

— Minor UX change with little mechanism described 7.261117.0

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ announced · Nov 17, 2026 BLOG ↗ RANK

Tines is an apps, agents and automations, in the complete environment for enterprise AI.

> Tines shipped full-text case search with nested boolean filters, AI Agent model fallback chains and per-provider model settings, group-based App access controls with a new visual mapping tool, per-action change control, and a self-hosted Database Overview page for retention planning.

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

01 Automatic AI Agent model fallback chains

Not checked yet NEW

80

Add up to five fallback models to any AI Agent action so Tines automatically moves to the next model on outage, rate limit, timeout, or revoked key, with no extra story logic required.

Protect a high-volume AI Agent action from OpenAI outages by chaining Anthropic and AWS Bedrock as automatic fallbacks — no conditional story logic needed.

📍 In the AI Agent action editor, open the model selector, choose your primary OpenAI model, then add up to five fallback models (e.g. Anthropic Claude, then an AWS Bedrock model) in priority order; Tines will advance to the next on outage, rate limit, timeout, or revoked key.

— Concrete limit, trigger conditions, and a runnable setup example

launch-20261117-378a8fd3

02 **Full-text search with boolean filters in Cases** Not checked yet

NEW 74

Cases now support full-text search across every case primitive — name, description, blocks, comments, notes, field values, and metadata — using `and` / `or` logic, `contains`, `does not` (negative match), and nestable filter groups.

Find all cases mentioning a specific host in any field, then narrow to those that do not contain a resolved status, using the new nested filter groups in Cases search.

> 📍 In Cases, open the search/filter bar, add a contains filter for the hostname, add a does not filter for `resolved`, then nest both terms into a filter group with and logic to scope results precisely.

— Names exact operators and filter mechanics with a worked example

launch-20261117-378a8fd3

03 **Group-based App access controls** Not checked yet NEW 70

Restrict App access to specific identity-provider groups: an admin enables the capability, configures which IdP field holds the group name(s), and builders then limit sharing to those named groups.

Lock a sensitive App so only the SOC group in your IdP can access it, without touching any other team's sharing settings.

> 📍 As an admin, enable the group-based App access capability, configure the IdP field that carries group names, then in the App's sharing settings restrict access to the named SOC group.

— Mechanism and admin/builder roles spelled out with an example launch-20261117-378a8fd3

04 **Custom model settings in AI Agent actions** Not checked yet NEW 60

AI Agent actions can now set custom settings for OpenAI, Anthropic, and AWS Bedrock models, including thinking mode and reasoning effort level, with an in-place test before saving.

— Names providers and settings but no usage walkthrough launch-20261117-378a8fd3

THINNER COVERAGE BELOW

05 **Database Overview page for self-hosted deployments**

Not checked yet NEW 45

A new Database Overview page (self-hosted) shows record size by type and story activity to help apply targeted retention policies.

– Names the page and its data but no navigation steps launch-20261117-378a8fd3

06 **Visual entity mapping tool for group access** Not checked yet

IMPROVED 35

Replace hand-written JSON group mappings with a visual, form-based tool that lets admins add or remove entity mappings through the UI.

– Describes the change but no config or UI path detail launch-20261117-378a8fd3

07 **Per-action change-control requirements** Not checked yet NEW

35

Set change-control requirements on individual actions inside a story, not only at the story level.

– Brief statement with no mechanism or scope given launch-20261117-378a8fd3

└──▷ ALSO FROM THESE RELEASES





WAS THIS USEFUL? 👍 👎

◆ AI SOC

Exaforce

Feedback ⓘ announced · Nov 16, 2026 BLOG ↗ RANK

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

> Exaforce launched ExaU, a new generation of customizable Exabot agents that work alongside practitioners inside the console, retain persistent memory, browse the web safely, accept voice input, and take permission-scoped response actions.

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

01 **Sandboxed browser for agent web links** Not checked yet NEW

65

ExaU adds a sandboxed browser for external links opened during agent sessions, with the agent reporting where a redirect actually landed.

Let an ExaU agent browse an external threat-intel link safely during a 1:1 chat investigation, confirming where a redirect actually lands before acting on the URL.

> 📍 In your 1:1 ExaU chat, paste the suspicious URL. The agent opens it in the sandbox browser alongside the chat and reports the final redirect destination. Click Allow or Deny before any follow-on action is taken.

– Usage example gives exact steps to trigger and verify the feature.

launch-20261116-ce698d50

THINNER COVERAGE BELOW

02 **Permission-scoped response actions for ExaU agents**

Not checked yet NEW

55

ExaU agents operate within explicit permissions boundaries; response actions such as revoking sessions, resetting passwords, and containing a device are only available to agents granted those permissions.

– Names specific actions but no steps to configure permissions. launch-20261116-ce698d50

03 **ExaU agents collaborate live in console chat** Not checked yet

NEW

53

ExaU agents can work the Exaforce console alongside a practitioner in a 1:1 chat — opening pages, clicking, typing, and scrolling — with any data-changing action gated behind an explicit Allow or Deny.

— Describes mechanism and gating but no concrete UI steps given. launch-20261116-ce698d50

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface Management

CybelAngel

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗

RANK

CybelAngel is the all-in-one solution against external threats. We monitor the open, deep, and dark web to protect your assets, brand, and credentials 24/7

> CybelAngel's latest update adds four new platform views: an Infected Machines View for compromised-host telemetry, Vendor Exposure Monitoring for third-party attack-surface risk, a Vulnerability Intelligence view under Threat Intelligence, and a Custom CISO Dashboard for executive reporting.

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

01 **Vulnerability Intelligence view** Not checked yet NEW 45

CybelAngel adds a Vulnerability Intelligence view under the Threat Intelligence section to consolidate vulnerability data alongside other TI signals.

— Names the exact section location but no further mechanism. changelog-20261118-9f2dd427

02 **Infected Machines View** Not checked yet NEW 25

CybelAngel adds an Infected Machines View to surface compromised-host telemetry in the platform.

— Names the view but no mechanism or data source detail. changelog-20261118-9f2dd427

03 **Vendor Exposure Monitoring** Not checked yet NEW 25

CybelAngel adds Vendor Exposure Monitoring to track third-party attack-surface risk.

— States purpose only, no scope or mechanism given. changelog-20261118-9f2dd427

04 **Custom CISO Dashboard** Not checked yet NEW 20

CybelAngel adds a Custom CISO Dashboard for tailored executive-level reporting.

— Bare feature name with no detail on customization options. changelog-20261118-9f2dd427

WAS THIS USEFUL? 👍 👎

Rapid7 Surface Command

Feedback ⓘ

changelog · Nov 17, 2026 BLOG ↗

Rapid7 Surface Command is an attack surface management platform that continuously discovers, inventories, and monitors an organization's internet-exposed assets to reduce cyber risk.

Rapid7 Surface Command adds a Claude integration to review enterprise AI access — users, roles, groups, and org settings — via the Anthropic connector and Claude's Compliance API.

NEW INTEGRATION Integrate Claude user, role, group, and organizational settings visibility into Rapid7 Surface Command via the Anthropic connector in the Rapid7 Extension Library, using Claude's Compliance API. Not checked yet

Review Claude user records, organization roles, administrative status, and group memberships alongside other assets and identities in Rapid7 Surface Command's relationship graph. Not checked yet

WAS THIS USEFUL? 👍 👎

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ updated · Nov 18, 2026

NOTES ↗ RANK

Horizon3 NodeZero is a horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

> NodeZero introduces a Rapid Assessment runtime option designed to cut test duration by up to 75%.

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

01 **Rapid Assessment runtime option** Not checked yet NEW 25

NodeZero adds a Rapid Assessment runtime option that can reduce test duration by up to 75%.

— States a benefit number but no mechanism or setup steps snapshot-20261118

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Vulnerability & Exposure Management

oxo

Feedback ⓘ v2.25.0 NOTES ↗ RANK

> oxo added a new message type for surfacing HTTP security header fingerprint data on websites.

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

01 **Security headers fingerprint message type** Not checked yet

NEW DETECTION NEW 30

Adds a security headers fingerprint message type for websites, letting agents surface HTTP security header data.

– Names the message type but no schema or usage details v2.24.0

NEW DETECTION TLS service fingerprint messages for domain names and IPs

v2.25.0 NOTES ↗ LATEST STABLE

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.25.0 https://github.com/Ostorlab/oxo.git
# already have the repo? check out this version:
$ git checkout v2.25.0
```

WAS THIS USEFUL? 👍 👎

Qualys VMDR

Feedback ⓘ changelog · Nov 17, 2026 BLOG ↗ RANK

Qualys VMDR is a cloud-based vulnerability management platform that continuously discovers, assesses, prioritizes, and remediates security vulnerabilities across IT assets.

> Qualys launched ROCky for Gemini Enterprise, bringing plain-language queries over vulnerability, asset, and risk data from the Qualys Enterprise TruRisk Platform into Google Cloud's Gemini Enterprise.

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

01 **ROCky for Gemini Enterprise on Google Cloud Marketplace**

Not checked yet NEW 65

Qualys ROCky for Gemini Enterprise is now available on Google Cloud Marketplace, letting users ask plain-language questions about vulnerability, asset, and risk data from their Qualys Enterprise TruRisk Platform subscription directly inside Gemini Enterprise. ROCky surfaces prioritized findings ranked by Qualys Enterprise TruRisk score, highlighting vulnerabilities that are actively exploited, listed in the CISA Known Exploited Vulnerabilities (KEV) catalog, or associated with ransomware, with explanations for each.

Find your highest-risk internet-facing hosts without leaving Gemini Enterprise — no spreadsheet export needed.

> 📍 In Gemini Enterprise, open the ROCky agent and type: 'Which of my internet-facing hosts is the biggest risk?'

— Describes behavior and scope but not setup or API details 2026-11-17-33779916

WAS THIS USEFUL? 👍 👎

Tenable Nessus

Feedback ⓘ Tenable Nessus10.12.7 (2026-11-17)

NOTES ↗ RANK

Tenable Nessus is a vulnerability scanner that identifies security weaknesses, misconfigurations, and compliance issues across networks, systems, and applications.

> Tenable Nessus added configurable agent profile settings to control how often managed agents check in for plugin updates.

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

01 **Agent profile settings for plugin update check-in frequency**

Not checked yet NEW 35

Administrators can configure agent profile settings to control how often agents check in for plugin updates, for agents linked to Tenable One Vulnerability Management and Tenable Nessus Manager.

– Describes capability but no exact setting name or location

Tenable Nessus10.12.7 (2026-11-17)

WAS THIS USEFUL? 👍 👎

◆ Patch Management

Qualys Patch Management

Feedback ⓘ

changelog · Nov 17, 2026 BLOG ↗ RANK

Qualys Patch Management is a cloud-based tool that automates the discovery, prioritization, and deployment of patches across IT assets to reduce vulnerabilities.

> Qualys Patch Management introduced Qualys ROCky, a conversational AI agent embedded in Gemini Enterprise that lets analysts query risk and stage patch jobs in natural language.

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

01 **Qualys ROCky AI agent in Gemini Enterprise** Not checked yet NEW

80

Qualys ROCky is a new conversational AI agent inside Gemini Enterprise that answers on-demand risk questions against live Qualys data — including mean time to remediate, CISA KEV backlog aging, and scan data currency — and can rank vulnerabilities by risk and stage patch jobs for deployment with an explicit human confirmation step before any host changes. Every agent action runs under the signed-in user's own Qualys entitlements, with no shared account pooling and no cross-tenant data visibility, and the agent is available for deployment via Google Cloud Marketplace inside existing Gemini Enterprise environments.

Walk through a critical-patch response end-to-end in Gemini Enterprise: ask what your exposure is, review the ranked host list, and confirm the staged patch job in TruRisk Eliminate.

📍 1. Open Qualys ROCky inside Gemini Enterprise (available via Google Cloud Marketplace). 2. Ask: 'What is our exposure for the latest critical Windows patch and what should we do?' 3. Review the risk-ranked host list with CISA KEV, weaponized-exploit, and malware context. 4. Say: 'Stage it for my Windows hosts.' 5. Review the target set and patch stability details the agent surfaces. 6. Confirm deployment when prompted ('Do you want to deploy this patch now?') — no host changes until this step.

— Walkthrough and named integrations given; no API or config specifics.

2026-11-17-16cf14b0

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

ITarian

Feedback ⓘ

Release Notes: ITarian Platform Update - November 18, 2026 (2026-11-17) BLOG ↗

RANK

Streamline IT operations with ITarian's all-in-one IT management software. Automate tasks, enhance security, and boost efficiency.

> ITarian's November 2026 update overhauls Service Desk ticket tracking with configurable custom states, a time-based workflow trigger, and small ticket-list and reply-editor additions.

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

01 Status Duration workflow condition

Not checked yet NEW

93

Workflows now support a Status Duration condition under the Ticket category, letting you trigger actions when a ticket has been in a given status for more than or less than a specified number of hours (1–8760 range); rules are evaluated once per hour.

Auto-follow-up stale tickets waiting on a customer reply by firing a workflow after 48 hours in the Waiting on Customer status.

> 📍 In the workflow rule builder: 1. Add a condition: Category = Ticket, Field = Status Duration, Operator = Greater than, Status = Waiting on Customer, Value = 48 2. Add an action: Send reply asking whether the customer still needs help. 3. Save the workflow.

— Full mechanism, numeric range, evaluation cadence, and runnable example

Release Notes: ITarian Platform Update …

02 Custom ticket states in Admin Panel

Not checked yet NEW

70

Admins can add custom ticket states with a name, colour picker, and Active toggle under Admin Panel → Manage → States. Seven predefined states (New, In Progress, Waiting on Customer, Waiting on Third Party, On Hold, Resolved, Closed) ship by default and can be edited or deactivated.

Create a custom ticket state to track tickets pending an internal security review, visible across all Ticket List views.

> 📍 In Admin Panel → Manage → States: 1. Click Add State. 2. Enter a name (e.g. Pending Security Review), choose a colour, and set Active = On. 3. Save — the new

state appears in the State dropdown on ticket Edit screens and the sortable State column in Inbox, Unassigned, Overdue, and Closed views.

— Names exact admin path and default states, lacks limit on custom state count

Release Notes: ITarian Platform Update …

THINNER COVERAGE BELOW

03 **State column in Ticket List views** Not checked yet NEW

57

A sortable State column with a colour-coded badge and an Excel-style filter now appears in the Ticket List across the Inbox, Unassigned, Overdue, and Closed views.

Create a custom ticket state to track tickets pending an internal security review, visible across all Ticket List views.

📍 In Admin Panel → Manage → States: 1. Click Add State . 2. Enter a name (e.g. Pending Security Review), choose a colour, and set Active = On. 3. Save — the new state appears in the State dropdown on ticket Edit screens and the sortable State column in Inbox, Unassigned, Overdue, and Closed views.

— Lists exact views affected but no further mechanism given

Release Notes: ITarian Platform Update …

04 **Reporter column in ticket list** Not checked yet NEW 35

The ITarian Dashboard ticket list now includes a Reporter column showing the name and email of the person who submitted each ticket.

— Names the column but gives no navigation or config detail

Release Notes: ITarian Platform Update …

05 **Canned responses in reply editor** Not checked yet NEW

35

Agents can now insert canned responses from a saved-response list directly in the ITarian Dashboard ticket reply editor.

— States the capability but no setup steps or location detail

Release Notes: ITarian Platform Update …

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ Iru Agent for Mac Release 5.2.9 (5403) NOTES ↗ RANK

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

> Iru's Mac agent gained a local authentication fallback for Iru Passport to keep sign-in working during identity provider outages.

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

01 **Local authentication fallback in Iru Passport** Not checked yet

NEW 53

Iru Passport now falls back to local authentication when the identity provider does not respond within 30 seconds, keeping sign-in available during IdP outages.

— Clear mechanism and timeout value, but no config or flag given

Iru Agent for Mac Release 5.2.9 (5403)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Passbolt

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗ RANK

Manage and share passwords securely with Passbolt. Open source, audited and built for teams that need collaboration, compliance and control.

> Passbolt 5.18 adds two new UI themes, widens Quick Access support, and speeds up sharing of large resource sets.

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

01 **Sakura and Layette UI themes** Not checked yet NEW 30

Passbolt now includes two new UI themes: Sakura and Layette.



— Names themes but no detail on how to apply them changelog-20261118-12e8d08a

02 **Broader Quick Access support** Not checked yet IMPROVED 20

Passbolt now supports Quick Access across a broader set of contexts.

— Vague on which contexts were added changelog-20261118-12e8d08a

03 **Faster sharing of large resource sets** Not checked yet IMPROVED 15

Passbolt now shares large resource sets faster.

— No numbers or mechanism given for the speedup changelog-20261118-12e8d08a

WAS THIS USEFUL?  

ConductorOne

Feedback ⓘ changelog · Nov 18, 2026 CHANGELOG ↗

RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

> ConductorOne shipped a batch of UI refinements to access reviews, lifecycle action selection, request forms, and entitlement matching, alongside a major breaking overhaul of the c1i CLI's output format, argument conventions, and exit codes.

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

01 **Lifecycle action picker shows connector-specific names**

Not checked yet IMPROVED 65

The lifecycle action picker in access controls now lists each connector action by its own name (e.g. Suspend user, Deactivate user) rather than always showing Disable account.

– Concrete before/after example of action labels given changelog-20261118-2fc1adc5

02 **Submit reviews confirmation shows decision counts** Not checked yet

IMPROVED 60

The Submit reviews confirmation step now shows counts of open, certified, denied, and left-to-decide items before you hit Submit.

– Names all four counts and the exact confirmation step changelog-20261118-2fc1adc5

03 **Hide Justification field on access request forms** Not checked yet

NEW 60

Selecting Hide justification field and saving with no other fields now hides the Justification field on an access request form.

– Exact setting name and the steps to apply it changelog-20261118-2fc1adc5

THINNER COVERAGE BELOW

04 **Entitlement picker matching improvements** Not checked yet IMPROVED

50

Entitlement pickers for access profiles, requests, and provisioning now match on a full dotted name or a partial name, and campaign scoping now also matches entitlement names.

– Describes new matching behaviour but no exact UI path changelog-20261118-2fc1adc5

05 **Unified duration picker for Snooze and Accept risk**

Not checked yet · IMPROVED · 45

Snooze and Accept risk actions now share a single duration picker offering presets or a custom date, replacing what were previously separate controls.

— Names the two actions and picker options, no UI path given changelog-20261118-2fc1adc5

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! `apps add-owner` and `apps remove-owner` argument order changed: the app is now the first argument and the user is passed as `--user-id`; the old argument order exits with an error and sends nothing.

! Commands that create, update, or delete now print the API response without its wrapper key — update any `jq` paths that read the wrapper, such as `.app.id` to `.id` and `.policy.id` to `.id` for `policies create`.

! Commands that previously printed a text line now print JSON.

! `--dry-run` now prints one JSON object `{"dryRun":true,"requests":[...]}` instead of text.

! Passing `--fields` to a command that changes something is now an error.

! An unparseable `C1I_DRY_RUN` value (e.g. `C1I_DRY_RUN=yes`) now exits with an error instead of silently sending the change.

! Several exit codes changed: invalid input (blank ID, ID of `, malformed input file, unparseable C1I_DRY_RUN , C1I_DEBUG , or C1I_MAX_RETRIES`) now exits differently; stuck pagination and docs-site failures also changed exit codes.

WAS THIS USEFUL? 👍 👎

Microsoft Entra

Feedback ⓘ changelog · Nov 17, 2026 BLOG ↗

RANK

Microsoft Entra is an identity and access management platform that secures user authentication, manages permissions, and controls access to apps and resources.

> Microsoft Entra introduces agentic identity governance, assigning every discovered agent its own Entra Agent ID for Conditional Access governance, expanding agent discovery to 60 agent types, and giving Autopilot agents dedicated identities, alongside limited-preview dynamic permissions, dynamic Conditional Access controls, and new AI-risk investigation playbooks.

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

01 **Autopilot agent identities** Not checked yet NEW 60

Microsoft Autopilot agents now get their own Agent user account and ID, enabling Conditional Access, ID Protection, and access governance scoped to that identity.

— Names account type and three governance mechanisms, no config path. 2026-11-17-3ff6f9b8

THINNER COVERAGE BELOW

02 **Agent ID assignment and expanded agent discovery** Not checked yet NEW 52

Microsoft Entra assigns every discovered agent a Microsoft Entra Agent ID so it can be governed with Conditional Access; agent discovery has been expanded to cover 60 agent types.

— Names Agent ID and 60-type scope but no setup steps. 2026-11-17-3ff6f9b8

03 **Dynamic permissions in ID Governance** Not checked yet NEW 50

Microsoft Entra ID Governance dynamic permissions grant task-scoped, time-bound access; currently in limited preview.

— Names the capability and preview status but no how-to. 2026-11-17-3ff6f9b8

04 **Perception playbooks and green team agents** Not checked yet NEW 45

New Perception playbooks and green team agents let teams investigate and remediate AI posture risks and anomalous agent behavior; currently in limited preview.

— Names the feature and purpose but no usage detail. 2026-11-17-3ff6f9b8

05 **Local agent containment via Intune policies** Not checked yet NEW

40

Supported local agents can be contained on the device using Intune policies.

— Names Intune but gives no policy details. 2026-11-17-3ff6f9b8

06 **Conditional Access dynamic controls** Not checked yet NEW

35

Conditional Access dynamic controls adjust protections as risk changes; currently in limited preview.

— Bare description of a preview capability. 2026-11-17-3ff6f9b8

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-17) 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 adds Access Certification (IGA), Agent Delegated Access, user-level Windows registry and MDM policies, MCP-only app setup, OAuth 2.0 workflow connectors, and enhanced Device Lock PIN escrow.

New Access Certification feature launches one-time or recurring review campaigns for SSO applications or directory user groups, with multi-stage approval chains and automated remediation. Not checked yet

Export frozen, tamper-proof Access Certification campaign snapshots in CSV, JSON, or PDF formats with SHA-256 document hashing and administrator sign-off flows for compliance audits. Not checked yet

New Agent Delegated Access PRIVATE PREVIEW lets an AI agent authenticate with its own JumpCloud-managed identity while acting on behalf of an end user, with token exchange controlled via the agent's new Delegated access tab. Not checked yet

New User Level Windows Custom Registry Keys Policy manages specific Windows registry values under `HKEY_CURRENT_USER` for an assigned user's profile, assignable to users or user groups with no MDM enrollment or device restart required.
Not checked yet

New User Level Custom Windows MDM (OMA URI) Policy configures custom MDM settings for a user's profile on MDM-enrolled Windows devices, assignable to users or user groups. Not checked yet

+6 more

NEW INTEGRATION SIEM integration in Directory Insights

NEW SUPPORTED PLATFORM App Catalog additions for macOS and Windows

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password for Android 8.12.46 (2026-11-17)

BLOG ↗ RANK

All the latest releases from the 1Password team

> Across the 8.12.46 release wave for Android, iOS, Linux, Windows and Mac, 1Password added an item sharing-history view, recent search history in the search bar, and a new 1Password Environments summary in Developer settings, alongside smaller refinements to Secure snippets, MCP Server setup docs, and Windows UI layout.

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

01 View Sharing History on item detail pages

Not checked yet NEW

55

Android and iOS users can now view an item's sharing history directly from its detail page by selecting View Sharing History in the dropdown menu, letting users audit who an item has been shared with.

Audit who an item has been shared with directly from its detail view.

📍 Open any item › tap the dropdown menu › select View Sharing History

— Names exact UI path but no mechanism beyond the menu action

1Password for Android 8.12.46 (2026-11-… 1Password for iOS 8.12.46 (2026-11-17)

02 Recent search history in search bar

Not checked yet NEW

45

On Linux and Windows, items you recently opened using search now appear when you select the search bar, letting users re-open recently searched items without retyping the query.

Re-open a recently searched item quickly without retyping — click the search bar to see your recent search history.

📍 In the 1Password app, click the search bar without typing anything to view recently opened items from search.

Quickly re-open a recently searched item without retyping the query — click the search bar to see your recent results immediately.

📍 Steps: Click the search bar in 1Password for Windows — recently opened items from previous searches appear immediately without typing.

– Clear behaviour described but no config or limit given

1Password for Linux 8.12.46 (2026-11-17) | 1Password for Windows 8.12.46 (2026-11-…

03 **1Password Environments summary in Developer settings**

Not checked yet | NEW | 45

A new summary explaining what 1Password Environments are now appears under Developer > View Environments on Linux and Windows.

Understand your 1Password Environments setup at a glance before configuring developer workflows.

> 📍 In the 1Password app, go to Developer › View Environments to read the new summary of what 1Password Environments are.

Understand what 1Password Environments are before configuring them in Developer settings.

> 📍 Steps: In 1Password for Windows, go to Developer › View Environments to read the new summary explaining what Environments are.

– Exact navigation path given but content of summary not detailed

1Password for Linux 8.12.46 (2026-11-17) | 1Password for Windows 8.12.46 (2026-11-…

04 **Secure snippets formatting toolbar behaviour** Not checked yet

IMPROVED | 45

Secure snippets in Labs now always show the formatting toolbar when editing, and fill without formatting as plain text.

– Describes before/after behaviour but no navigation steps

1Password for Mac 8.12.46 (2026-11-17)

05 **MCP Server setup instructions updated** Not checked yet | IMPROVED

35

The 1Password MCP Server setup instructions now point to Developer settings and include troubleshooting steps.

– States the change but no specific steps or content shown

1Password for Linux 8.12.46 (2026-11-17)

06 **Responsive sidebar and search bar on Windows** Not checked yet

IMPROVED | 30

The app now collapses the sidebar when narrowing the window and centers the search bar on resize.

– Describes layout change but no steps to trigger or configure

1Password for Windows 8.12.46 (2026-11-…

07 **Watchtower share option renamed to 'Post to X'** Not checked yet

IMPROVED 25

The Watchtower share option now reads Post to X instead of Tweet on Windows.

— Cosmetic text change with minimal mechanism or action

1Password for Windows 8.12.46 (2026-11-…

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Network, Ory Hydra, Ory Kratos, Ory Polis, Ory Elements v26.4.18 released

NOTES ↗ RANK

Discover the latest features, improvements, and updates to Ory.

> Ory shipped two small Account Experience and network-routing improvements: a last used sign-in badge and a configurable root redirect for project domains.

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

01 **Root path redirect for project domains** Not checked yet NEW

35

Ory Network can now redirect visitors who open your project's slug domain or custom domain at `/` to a destination of your choice.

— Names the path `/` but lacks config key details

Ory Network, Ory Kratos, Ory Polis v26.…

02 **'Last used' badge on sign-in buttons** Not checked yet NEW

30

Ory Account Experience now shows a Last used badge on the social sign-in or SSO button a user previously authenticated with, helping users remember which method they used.

— Describes behaviour but no config or API surface

Ory Network, Ory Hydra, Ory Kratos, Ory…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Cyera

Feedback ⓘ announced · Nov 17, 2026 BLOG ↗

Cyera helps organizations protect data, govern identities, and secure AI agents to reduce risk, stay in control, and adopt AI with confidence.

Cyera Agent Guardian adds a Microsoft Sentinel integration PRIVATE PREVIEW that pushes structured runtime alerts into Microsoft Agent 365 dynamic agent risk scoring.

Adds a Microsoft Sentinel integration that pushes Agent Guardian runtime alerts — carrying `EventType`, `SignalCategory` (marked `RuntimeRisk`), `AlertName`, `DetectionId`, `RuleName`, `ThreatCategory`, `EntityType`, `EntityId`, `EntityName`, `IdentityType`, `EventOriginalSeverity`, and `AlertOriginalStatus` fields — to your Sentinel environment for correlation with Microsoft-calculated dynamic agent risk. Not checked yet

NEW INTEGRATION Agent Guardian now surfaces runtime alerts alongside Microsoft Defender alerts in the Microsoft 365 admin center (Agent 365), Microsoft Defender, Microsoft Entra, and Microsoft Purview, each linked to an investigation case. Not checked yet

Each Cyera alert sent to Sentinel includes what action occurred, whether sensitive data was involved and its severity, the identity of the acting user and agent (including Microsoft Entra Agent ID), and which policy was violated with a plain-language triage assessment from Cy. Not checked yet

WAS THIS USEFUL? 👍 👎

Microsoft Purview

Feedback ⓘ changelog · Nov 17, 2026

BLOG ↗ RANK

Microsoft Purview is a unified data governance and compliance platform that helps organisations discover, classify, protect, and manage sensitive data across their environments.

> Microsoft Purview is extending protection into the agentic era with grounding protection for Copilot, a new dynamic risk assessment for agents, and enhanced data classification, all slated to enter preview or general availability in January 2027.

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

01 **Grounding protection for Copilot** Not checked yet NEW

42

Grounding protection for Copilot, based on sensitive information types, lets customers adopt Copilot without labeling every file first. It enters preview in January 2027.

– Names mechanism and timing but no config detail 2026-11-17-7a0b8960

02 **Dynamic Agent Risk assessment** Not checked yet NEW

36

A new Dynamic Agent Risk assessment, entering preview, evaluates each agent's exposure, blast radius, and runtime behavior.

– Describes purpose but no usage path or surface 2026-11-17-7a0b8960

03 **Enhanced data classification** Not checked yet IMPROVED

26

Enhanced data classification aims for more precise protection with less manual effort, becoming generally available in January 2027.

– Vague improvement claim with no mechanism named 2026-11-17-7a0b8960

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ announced · Nov 17, 2026 BLOG ↗ RANK

Tailscale creates private WireGuard-based mesh networks with centralized identity and access controls.

> Tailscale introduced netmap caching, letting devices persist their tailnet network map to disk for faster startup and continued peer connectivity even when the coordination server is unreachable.

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

01 **Netmap caching for offline-resilient connections** Not checked yet

NEW 65

Netmap caching can be enabled via the tailnet policy file to cache a tailnet's network map to disk, allowing nodes to connect faster at startup and reach peers even when the coordination server is unavailable. With a cached netmap present, nodes share DISCO public keys over the existing encrypted WireGuard/DERP tunnel so they can establish direct connections without needing the coordination server.

— Names policy file config and mechanism, but no exact key name given

launch-20261117-a0648776

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ

11/17 November 17 – Web Application and API Update (2026-11-17) BLOG ↗ RANK

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

> DNSFilter expanded domain miscategorization reporting with two new proposed categories and made its Collections API documentation public.

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

01 **Advertising and Trackers miscategorization categories**

Not checked yet NEW 45

Users can now report miscategorized domains under the proposed categories Advertising and Trackers from the Domain Report and Domain Lookup tools.

— Names exact categories and UI entry points, no mechanism detail

11/17 November 17 – Web Application and…

02 **Public Collections API documentation** Not checked yet NEW

30

Collections API documentation is now publicly available.

— Names the API but gives no endpoint or usage detail

11/17 November 17 – Web Application and…

WAS THIS USEFUL? 👍 👎

Cato Networks Feedback ⓘ

Introducing Cato AgentSense: Stop Unsafe Agent Actions Before They Execute (2026-11-17)

BLOG ↗

Cato Networks provides a cloud-native SASE platform that converges networking and security into a single global service for enterprises.

Cato Networks adds AgentSense, a patented runtime AI agent protection engine that traces instructions to their source and blocks unsafe actions before execution, shipping as part of Cato AI Security for Users.

NEW DETECTION Cato Networks adds AgentSense, a runtime engine that continuously assesses AI agent behavior and blocks unsafe actions before they execute, starting with Untrusted Influence Detection to trace every tool call, file write, and outbound communication back to its instruction source. Not checked yet

AgentSense ships as part of Cato AI Security for Users and enforces on local agents such as coding agents and desktop assistants, with no rules to write at launch.

Not checked yet

WAS THIS USEFUL? 👍 👎

AGENT

◆ Agent Sandboxes

microsandbox

Feedback ⓘ v0.7.16 NOTES ↗ RANK

> microsandbox added Services support across its SDKs and CLI, introducing service-level sandbox management.

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

01 **Services support in SDKs and CLI** Not checked yet NEW

35

Adds Services to the SDKs and CLI, enabling service-level sandbox management alongside existing sandbox primitives.

— No specifics on commands, APIs, or mechanism given v0.7.16

WAS THIS USEFUL? 👍 👎

NVIDIA OpenShell

Feedback ⓘ v0.1.9 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell v0.1.9 adds a ServiceDiscovery API and configurable sandbox provisioning timeouts, alongside a wave of breaking changes spanning the Go SDK's module path, network policy, Postgres migrations, Helm chart configuration, and OCSF event taxonomy.

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

01 Configurable sandbox provisioning timeouts

Not checked yet NEW

90

`openshell sandbox create --timeout <SECONDS>` lets you override the 300-second default and `OPENSHELL_PROVISION_TIMEOUT` environment variable for how long the CLI waits for sandbox provisioning. Gateway-side timeouts for slower isolation backends are now tunable via `provisioning_repair_window_seconds`, `runtime_bootstrap_grace_seconds`, and `sandbox_attach_request_timeout_seconds`.

Override the default 300-second provisioning wait when creating a sandbox in a slow environment, such as a GPU-backed isolation backend.

```
$ openshell sandbox create --name gpu-agent --timeout 600
```

— Flag, env var, and three gateway config keys all named with example. v0.1.9

02 OIDC CA bundle scoped to discovery and JWKS

Not checked yet

85

BREAKING

The Helm chart no longer sets a process-wide `SSL_CERT_FILE` for a private OIDC CA; the CA from `server.oidc.caConfigMapName` now applies only to OIDC discovery and JWKS requests — outside Helm, set it with `[openshell.gateway.oidc] ca_bundle`, `--oidc-ca-bundle`, or `OPENSHELL_OIDC_CA_BUNDLE`.

— Names config key, flag, env var and Helm field together. v0.1.9

03 OCSF denial records use specific activity types

Not checked yet

80

BREAKING

Authorization denial OCSF records for write methods now report Create, Update, or Delete instead of Other (99) — update queries to select denials by `class_uid: 6003`

and `status_id: 2` rather than `type_uid: 600399`.

— Exact old and new query field values given. v0.1.9

04 **Output format flags for workspace commands** Not checked yet NEW

75

`openshell workspace create` and `openshell workspace get` now support `-o/--output table|json|yaml` to control output format.

Emit workspace details as JSON for scripted consumption or piping into jq.

```
$ openshell workspace get my-workspace -o json
```

— Exact flag and commands given with runnable example. v0.1.9

05 **Postgres migration 010 blocks writes during index build**

Not checked yet BREAKING

75

On Postgres, gateway migration `010` builds indexes on the `objects` table with a plain `CREATE INDEX`, which blocks writes to that table while it runs — large deployments should plan for downtime during the upgrade.

— Names migration, table, and command plus operational impact. v0.1.9

06 **Helm chart metrics toggle replaces port setting** Not checked yet

75

BREAKING

The Helm chart now enables the gateway's Prometheus endpoint with `metrics.enabled` instead of a port value; charts setting `service.metricsPort` to `0` or `false` now fail to render — set `metrics.enabled=false` and keep a valid port instead.

— Exact config keys and remediation given. v0.1.9

07 **ServiceDiscovery API for workspace routes** Not checked yet NEW

73

A new `openshell.v1.ServiceDiscovery` service adds `GetServiceRouteSnapshot` and `WatchServiceRoutes` RPCs so proxies and controllers can discover a workspace's service routes.

— Names service and both RPCs but no usage example. v0.1.9

08 **Go SDK moved to new module path** Not checked yet BREAKING

70

The Go SDK has moved to `github.com/ai-openshell/sdk-go`; the old `github.com/NVIDIA/OpenShell/sdk/go` path is deprecated with v0.1.8 as its last version — update module imports to the new path.

— Exact old and new import paths and final version given. v0.1.9

09 **Private IPs require explicit grants in network policy**

Not checked yet BREAKING 65

Wildcard network policy endpoints no longer allow private IP literals on their own; sandboxes that still need to reach a private address now require an exact endpoint or an `allowed_ips` grant.

— Names the config field and remediation, no full policy syntax. v0.1.9

THINNER COVERAGE BELOW

10 **Request-ID header on inspected HTTP responses** Not checked yet

NEW 50

Inspected HTTP responses now carry an `X-OpenShell-Request-ID` header that links them to OCSF records and supervisor traces.

— Names header but gives no usage steps. v0.1.9

11 **Streaming HTTP middleware hooks capability** Not checked yet NEW

50

Streaming HTTP middleware hooks can be enabled with the `openshell.supervisor-middleware.http-hooks` capability.

— Names capability flag but little mechanism detail. v0.1.9

12 **Supervisor distribution metrics** Not checked yet IMPROVED

50

New metrics report the number of supervisors per gateway and the busiest supervisor's sandbox count, and shared supervisors now spread evenly across gateway replicas when they connect.

— Explains metrics and balancing change but no metric names given. v0.1.9

13 **Grafana monitoring mixin for Kubernetes** Not checked yet NEW

45

A reusable Grafana monitoring mixin now provides dashboards for operational health, operational logs, and OCSF events in Kubernetes deployments.

— Describes dashboards but gives no setup path. v0.1.9

14 **Legacy HTTP middleware hooks deprecated** Not checked yet

DEPRECATED

45

The legacy `EvaluateHttpRequest` and `HttpResponsePreReturn` HTTP middleware hooks are deprecated and will be removed in 0.2.0.

— Names hooks and removal version but no migration path. v0.1.9

15 **Clearer gateway delete records** Not checked yet IMPROVED

40

Gateway delete records now name the workspace and resource deleted, and show when a delete found nothing to remove or is still in progress.

— Describes behavior change without a named endpoint or flag. v0.1.9

16 **Middleware selectors reject ports** Not checked yet BREAKING

40

Middleware selectors can no longer include a port and now fail validation — remove the port from any selectors.

— States the rule and fix but gives no selector syntax example. v0.1.9

└──▷ **BREAKING ON UPGRADE**

! The Go SDK has moved to `github.com/ai-openshell/sdk-go`; `github.com/NVIDIA/OpenShell/sdk/go` is deprecated and its last version is v0.1.8 — update module imports to the new path.

! Wildcard network policy endpoints no longer allow private IP literals on their own — add an exact endpoint or an `allowed_ips` grant for any private address your sandboxes still need to reach.

! On Postgres, gateway migration `010` builds indexes on the `objects` table with a plain `CREATE INDEX`, which blocks writes to that table while it runs — large deployments should plan for downtime during upgrade.

! The Helm chart now enables the gateway's Prometheus endpoint with `metrics.enabled` instead of a port value; charts that set `service.metricsPort` to `0` or `false` now fail to render — set `metrics.enabled=false` and keep a valid port instead.

! Middleware selectors can no longer include a port and now fail validation — remove the port from any selectors.

! The legacy `EvaluateHttpRequest` and `HttpResponsePreReturn` HTTP middleware hooks are deprecated and will be removed in 0.2.0.

! Authorization denial OCSF records for write methods now report Create, Update, or Delete instead of Other (99) — update queries to select denials by `class_uid: 6003` and `status_id: 2` rather than `type_uid: 600399`.

! The Helm chart no longer sets a process-wide `SSL_CERT_FILE` for a private OIDC CA; the CA from `server.oidc.caConfigMapName` now applies only to OIDC discovery and JWKS requests — outside Helm, set it with `[openshell.gateway.oidc] ca_bundle`, `--oidc-ca-bundle`, or `OPENSHELL_OIDC_CA_BUNDLE`.

WAS THIS USEFUL?  

Docker Sandboxes

Feedback ⓘ announced · Nov 16, 2026 BLOG ↗

RANK

Docker Sandboxes run AI coding agents in isolated environments on your machine or on Docker-managed cloud infrastructure

> Docker Sandboxes shipped a Firecrawl kit that gives sandboxed coding agents live web search, scraping, and crawling access through a tightly scoped proxy.

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

01 **Firecrawl kit for live web access in sandboxes** Not checked yet

NEW 95

Adding `--kit docker.io/firecrawl/sbx-kit-firecrawl:1.0.0` to `sbx run` routes all page fetches through the single allowed host `api.firecrawl.dev` without widening the sandbox's network allowlist further. The Firecrawl API key is stored on the host via `sbx secret set firecrawl` and injected by the proxy at request time, so the key never enters the sandbox image and `FIRECRAWL_API_KEY` inside the sandbox always reads `proxy-managed`. As a lifecycle step the kit installs the `firecrawl-py` SDK pinned to version 4.44.0, failing sandbox creation outright if the install or import check fails rather than erroring mid-task. The kit's `firecrawl/firecrawl.yaml` descriptor declares it as a `kind: mixin` with four typed capabilities — `lifecycle@1`, `credential@1`, `network-policy@1`, and `agent-context@1` — so it extends an existing workload rather than replacing it. It requires `sbx` CLI version 0.46 or later and a global network policy initialised with `sbx policy init balanced`; on macOS it is installed via `brew trust docker/tap && brew install docker/tap/sbx`.

Give a sandboxed shell agent live access to documentation sites blocked by the default balanced policy, using the Firecrawl kit alongside Claude Code.

```
$ sbx secret set firecrawl
sbx run docker/sbx-kit-shell:1.0.0 \
  --kit docker/sbx-kit-claude-mixin:2.1.285 \
  --kit docker.io/firecrawl/sbx-kit-firecrawl:1.0.0 .
```

Verify that the SDK, the proxy-managed credential, and the network rule are all working correctly from inside a running sandbox session.

```
$ !python3 -c "import firecrawl, importlib.metadata as m;
print('firecrawl-py', m.version('firecrawl-py'))"
!env | grep FIRECRAWL_API_KEY
```

— Names exact kit image, CLI commands, descriptor schema, and version requirements

launch-20261116-a959d091

WAS THIS USEFUL? 👍 👎

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.2.0 (2026-11-17) BLOG ↗ LATEST STABLE

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

Runlayer adds per-user Windows hook installs via `runlayer setup hooks --install`, a workspace Inbox for every user, faster in-place artifact edits with version history, and tighter agent connector-tool access controls.

Install per-user AI Watch hooks on Windows without MDM using `runlayer setup hooks --install`, which also preserves existing hook entries containing non-ASCII text. Not checked yet

UPDATED INTEGRATION Every workspace now has an Inbox with an updates bell that collects decisions, reconnect prompts, and pending access and connector requests. Not checked yet

Agents can now change text, colors, or numbers in an existing artifact without downloading and re-uploading the whole file; each change is saved as a restorable version. Not checked yet

Agent connector tools are now restricted so editors and importers can only add connector tools they can use themselves, with deny policies still applying to connector owners and admins. Not checked yet

NEW INTEGRATION Bitsight connector

NEW INTEGRATION Paper (Remote) connector

UPDATED DETECTION Skill risk scanning: reduced false High ratings, new lookalike host and interpreter pipeline detections

UPDATED DETECTION Skill risk scanning: shell expansion suffix pipe-to-shell patterns rated High

NEW INTEGRATION Quartr Pro connector

NEW INTEGRATION Affinity connector

UPDATED INTEGRATION AI Watch: OpenCode 2 session recording and policy enforcement

UPDATED INTEGRATION Teams connector: Adaptive Cards support

UPDATED INTEGRATION Gmail connector: include_signature, list_drafts search, HTML body_format

UPDATED INTEGRATION Firecrawl connector: browser sign-in instead of API key embedding

NEW INTEGRATION Claude Platform linked billing integration

└──▷ **BREAKING ON UPGRADE**

! Existing Google Calendar connectors must be redeployed to receive event color data on read and list operations.

▾ **2 more releases in this issue** · 2026-11-17

6.1.0 (2026-11-17) BLOG ↗ STABLE

Runlayer adds an Auditor role, Terraform resources for binary and device-sync settings, editable approval-system request bodies, and Lean 4 in agent sandboxes.

Manage per-package update policies and skill device sync defaults as code using the new `runlayer_binary_package_settings` and `runlayer_skill_device_sync_settings` Terraform resources. Not checked yet

New read-only Auditor role for compliance reviewers gives org-wide access to audit logs, session content, user info, metrics, Shadow AI and AI inventory, and usage export — assign it by setting `role auditor` in your identity provider. Not checked yet

Approval system request bodies now follow an editable template, letting you shape the payload Runlayer sends to Okta Workflows, ServiceNow, or a custom service to match the format your flow expects. Not checked yet

Agent sandboxes now include the Lean 4 proof checker (`lean`) with its core library, so agents can machine-check proofs without downloading the toolchain first. Not checked yet

Enforce file-read blocking now also blocks built-in file search and listing tools `grep`, `rg`, `glob`, `find`, `ls`, and `list_files` in every client, including Claude Code's `Grep` and `Glob`. Not checked yet

+7 more

└──▷ **BREAKING ON UPGRADE**

! Legacy enrollment keys no longer issue a device credential to users who can manage admins, organization settings, or MDM configuration; those users must sign in with `runlayer login` instead.

! The MDM configurations API now returns a masked enrollment key for legacy enrollment configurations instead of the full value.

! Failed resource reads no longer forward structured error data larger than 64 KiB or nested more than 100 levels deep — applies to every organization including those with masking off.

! A tool argument literally named `__object_ref` now appears as `_escaped__object_ref` in the audit log and SIEM export.

announced · Nov 17, 2026 BLOG ↗

Runlayer adds linked billing with Claude Platform: one-step approval, direct Anthropic billing, and spend-limit controls in the Claude Console.

Connect a Claude Platform organization to Runlayer without an API key by clicking Link your Claude organization in Runlayer and approving the link once — an owner, admin, or billing admin of the Claude organization approves, Anthropic bills Claude API usage directly to that organization. Not checked yet

Set a monthly spend limit for Runlayer inside the Claude Console, and unlink the Claude organization at any time. Not checked yet

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Zenity

Feedback ⓘ

Zenity Now Secures Agents Across Your Microsoft Ecosystem NOTES ↗

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

Zenity adds actor-identity resolution via Microsoft Entra ID for GitHub Enterprise, where GitHub settings hide actor names, enabling full attribution of coding-agent actions.

NEW INTEGRATION UPDATED INTEGRATION Where enterprise settings hide actor names, Zenity uses Microsoft Entra ID to resolve them, so the security team can see whose agent made a change. Not checked yet

NEW INTEGRATION GitHub Copilot coding agent — telemetry, findings, and inline prevention

NEW INTEGRATION Microsoft Copilot Cowork — telemetry and findings via Cowork transcripts

UPDATED INTEGRATION Microsoft Sentinel integration upgraded from raw-log export to high-fidelity alerts

WAS THIS USEFUL? 👍 👎

◆ Model & Agent Protection

Minimal Software Research minimal

Feedback ⓘ

release-5c8c4e87 NOTES ↗ RANK

Minimal is a secure, declarative runtimes for Claude Code, Gemini, Codex, Pi and more, so you can run multiple agents safely on your stack from a verified software supply chain.

> minimal's latest release centers on a new Box Egress Proxy for certificate-based, audited network control, alongside new `min` commands for secrets, login, audit trails, box inspection, and opt-in OpenTelemetry tracing.

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

01 **Box Egress Proxy and egress control flags** Not checked yet NEW

95

Adds a Box Egress Proxy (`bep`) that holds proxy keys in the host keychain (Secure Enclave on macOS, TPM or Secret Service on Linux), issues leaf certificates under a name-constrained root and signing CA, appends hash-chained audit records per proxy decision, rotates the signing CA, and enforces declared egress rules per box. The proxy is controlled via `--network` and `--egress` flags on `min task run` and `min box run` , and each software key is named in `min auth status` for the proxy's key store.

Run a task with explicit network and egress controls to restrict what the box can reach

```
$ min task run --network none --egress blocked my-task
```

— Full certificate/keychain mechanism plus named flags and runnable example.

release-5c8c4e87

02 **Box audit trail commands** Not checked yet NEW

70

`min box audit` reads a box's audit trail, `min box audit --parent` merges parent chains, and `min doctor` reports an edited audit chain.

— Named subcommands but no runnable example or deeper mechanism. release-5c8c4e87

03 **Box inspection as machine-readable JSON** Not checked yet NEW

65

`min box show <name> -o json` inspects a box's full spec and exec list as machine-readable JSON for scripting or auditing.

Inspect a box's full spec and exec list as machine-readable JSON for scripting or auditing

```
$ min box show my-box -o json
```

— Exact flag and runnable command given. release-5c8c4e87

04 **OpenTelemetry tracing export** Not checked yet NEW 65

Adds opt-in OpenTelemetry export with a local spool via `mlog`, one trace per request across `min` and `minimald`, and guest telemetry carried over the VM channel.

— Mechanism described but no command or config shown. release-5c8c4e87

05 **Secrets management via min secret** Not checked yet NEW 60

`min secret set` stores a secret into a box and `min secret list` lists what has been stored, keeping credentials out of environment variables.

Store a secret into a box and later list what has been stored, keeping credentials out of environment variables

```
$ min secret set MY_API_KEY
min secret list
```

— Runnable commands shown, but little on storage mechanism. release-5c8c4e87

THINNER COVERAGE BELOW

06 **Security bug reporting flag** Not checked yet NEW 45

`min bug --security` sends a suspected vulnerability report.

— Flag named but no detail on handling or response. release-5c8c4e87

07 **QUIC block helper in finalize-install** Not checked yet NEW 30

`min finalize-install` now names the QUIC block helper.

— Bare mention with no explanation of behaviour. release-5c8c4e87

WAS THIS USEFUL? 👍 👎

◆ AI Runtime Defense & Guardrails

Manifold Security

Feedback ⓘ announced · Nov 12, 2026

BLOG ↗ RANK

Manifold Security is an AI detection and response platform that monitors what autonomous AI agents do at runtime on enterprise endpoints and blocks risky actions.

> Manifold Security launched remote removal of risky skills, plugins and MCP servers directly from developer endpoints, with granular per-version and per-endpoint targeting, backed by audit logging and on-machine developer notifications.

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

01 Remote removal of risky skills, plugins and MCP servers

Not checked yet NEW 90

From an asset's Endpoints tab in Manifold, security teams can remove a risky skill, plugin or MCP server from every endpoint at once or target individual machines, across Claude Code, Cursor, Codex and Kiro. Removal can be scoped to a specific version — for example removing only the High `-severity` version 2.2.0 of a skill while leaving clean versions 2.0.0 and 2.1.0 in place — via the version picker and Remove asset action.

Remove only the high-severity version of a skill from the one endpoint running it, while leaving clean versions on other machines.

> 📍 In Manifold, open Supply chain › Skill: pdfhelper, go to the Endpoints tab, locate the row for ec-ubuntu-3 running Rev 3 · 2.2.0 (High), open Actions, select Remove asset, choose version Rev 3 · 2.2.0 in the version picker, and click Remove asset.

Permanently retire a flagged skill by combining removal with a block policy so it cannot be reinstalled after it is taken off endpoints.

> 📍 In Manifold, tag the asset as `blocked`, then write a policy in block mode that denies any asset with the `blocked` tag. Finally, open the asset's Endpoints tab, select Remove asset from all endpoints, choose All versions, and confirm removal.

— Names exact UI flow, versions, severity and four supported tools

launch-20261112-7c97c266

THINNER COVERAGE BELOW

02 Audit log entries for endpoint removals

Not checked yet NEW 45

Every removal is recorded in the audit log, capturing actor, action, change, source and endpoint, giving the security team a full trail of who removed what asset, when, and from which machine.

— Names logged fields but no path to view the log launch-20261112-7c97c266

03 **Desktop notification for removed skills** Not checked yet NEW

30

Developers receive an on-machine desktop notification when a skill is removed, displaying the asset name and the reason given by the organization.

— Describes behavior only, no reader action involved launch-20261112-7c97c266

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

GitHub CLI

Feedback ⓘ v2.105.0 NOTES ↗ RANK

GitHub CLI github's official command line tool

> GitHub CLI introduces an event-driven `gh event` command set that lets scripts and agents subscribe to GitHub Event Streams instead of polling the API.

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

01 gh event command set for event streams

Not checked yet NEW

95

A new preview command set lets you wait for specific events without polling using `gh event listen --resource <url> --event <event-type>`, with `--timeout`, `--count`, or `--follow` flags controlling how long to wait; events print as CloudEvent JSON objects with `data.api_url` pointing to the changed resource. Subscriptions can be managed directly with `gh event subscription create`, `gh event subscription list`, `gh event subscription view`, and `gh event subscription delete`. `gh run view` and `gh run watch` now suggest using `gh event listen` for check suite events.

Block a CI script until a pull request review is submitted, then extract the conclusion — avoiding repeated API polling.

```
$ gh event listen --resource
https://github.com/OWNER/REPO/pull/NUMBER --event
com.github.pull_request.review.submitted.v1 --timeout 30m
```

Wait for a check suite to finish and print only its conclusion, useful for gating a deployment step on CI passing.

```
$ gh event listen --resource
https://github.com/OWNER/REPO/pull/NUMBER --event
com.github.pull_request.check_suite.status.changed.v1 --jq
.data.conclusion
```

Inspect all active event subscriptions to audit what an agent or script is currently listening to.

```
$ gh event subscription list
```

— Names every flag and subcommand with runnable examples v2.105.0

WAS THIS USEFUL? 👍 👎

joplin

Feedback ⓘ announced · Nov 14, 2026 BLOG ↗ RANK

Joplin - the privacy-focused note taking app with sync capabilities for Windows, macOS, Linux, Android and iOS.

> Joplin adds macOS x-callback-url support enabling two-way linking with Hookmark, letting users connect notes to PDFs, web pages, emails, files, and tasks.

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

01 **x-callback-url actions for Hookmark integration on macOS**

Not checked yet NEW 70

Joplin introduces two new x-callback-url actions on macOS: `getCurrentNote`, which returns the title and link of the currently open note, and `createNote`, which creates a new note and returns its title and link. Together these enable Hookmark integration, letting users create two-way links between Joplin notes and PDFs, web pages, emails, files, and tasks, available from version 3.7.21.

— Names both actions and the integration, but no setup steps given.

launch-20261114-7fa5f616

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 18, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue