THE CYBER TOOLCHAIN

NO. 146

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 21, 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 140 changes

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 42 of 43 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. Qualys VMDR can now mitigate Java4Shell by applying a JVM flag to eligible hosts with one-click

vim ? help

rollback, buying time before patching. GitGuardian stops hooks and AI coding agents blocking on secrets it has confirmed dead. And GitHub Enterprise Server 3.23 RC brings AI secret detection on-appliance for teams that cannot send code out.

PROTECT

Neutralize Java4Shell on hosts you cannot patch this week

Qualys applies a JVM flag to eligible hosts and reports process-level coverage, so you can see which running JVMs are actually covered rather than assuming, and roll back in one click if an app misbehaves. GitHub's 3.23 RC pairs a Dependabot cooldown with wider malware advisory coverage, holding back freshly published package versions long enough for a compromised release to surface.

Qualys VMDR · GitHub

PROTECT

Stop burning developer time on secrets and findings that are already dead

GitGuardian filters confirmed-invalid secrets out of results, which matters most for pre-commit hooks and AI coding agents that halt on every match regardless of whether the credential still works. Socket extends reachability to Java/Android so an unexercised dependency CVE stops competing with ones your code actually calls, and Sysdig's Stage filter separates registry, pipeline and runtime findings so a runtime-only exposure is not buried in base-image noise.

GitGuardian SaaS · Socket · Sysdig Secure

GOVERN

Run an access certification campaign with evidence an auditor will accept

JumpCloud's multi-stage approval chains and signed exports replace the spreadsheet-and-email access review most teams still run quarterly, and the signature is what turns the result into evidence rather than an assertion. CISO Assistant adds findings-to-assessment linking so a control gap traces back to the specific finding that opened it.

JumpCloud

Does **Opus 5** read better? 👍 👎

◆ AI-Powered Offensive Security

dark-moon

Feedback ⓘ v1.7.0 NOTES ↗

Dark-moon is an open-source autonomous AI penetration testing. 50 specialist agents across web, API, cloud, Active Directory, Kubernetes, CI/CD and AI/LLM (OWASP LLM Top 10).

dark-moon adds Cyber Score v1, a Prioritized Treatment Plan, Terminal Explorer (OpenCode plugin), and attack paths with `enabled_by` agent declarations.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.7.0 https://github.com/ASCIT31/Dark-Moon.git
# already have the repo? check out this version:
$ git checkout v1.7.0
```

New Cyber Score v1 deterministic rating with coverage gates ('not rated / limited / rated') and a published methodology in `docs/cyber-score-methodology.md`. Not checked yet

New Prioritized Treatment Plan gives one action per asset × weakness class, with indicative effort, status history, retest-based verification, and risk acceptance with a named owner and expiry. Not checked yet

New Terminal Explorer (OpenCode plugin) adds a keyboard-driven workspace covering the executive block, projects, campaigns, findings, treatment plan, attack paths, frameworks, profile, and legend — readable at 80 columns. Not checked yet

Agents now declare attack-path links using `enabled_by` to show what enabled what, with every link held as a hypothesis unless evidence supports it. Not checked yet

`dashboard_*` writes are now bound to the calling run when a run key is present, scoping dashboard mutations to their originating run. Not checked yet

+2 more

NEW COMPLIANCE FRAMEWORK Framework mapping for ISO 27001, GDPR art. 32, NIS2 art. 21(2), AI Act art. 15

└──▷ BREAKING ON UPGRADE

! A terminal campaign status can no longer be re-finalised after it has been closed.

! Agent labels can no longer inflate the Cyber Score (coverage gate enforced server-side).

WAS THIS USEFUL?

Keygraph shannon

Feedback ⓘ v3.6.0 NOTES ↗ RANK

Shannon is an AI pentester for web applications and APIs. It analyzes your source code, identifies attack vectors, and executes real exploits to prove vulnerabilities before they reach production.

> Shannon's v3.6.0 release adds a documentation-backed skill plugin for extending agent capabilities and introduces configurable concurrency, timeout, and retry limits for scan workers.

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

01 **Documentation-backed Shannon skill plugin** Not checked yet NEW

35

Shannon adds a documentation-backed Shannon skill plugin to extend agent capabilities with knowledge drawn from your own documentation.

— Names the plugin but no config or usage detail given v3.6.0

02 **Concurrency, timeout and retry limits for scan workers**

Not checked yet IMPROVED

35

Shannon adds concurrency and timeout limits, plus a cap on agent retries, to the scan worker.

— States limits exist but no exact values or flags named v3.6.0

WAS THIS USEFUL?

◆ Hacking Hardware

ESP32 Marauder

Feedback ⓘ v1.21.0 NOTES ↗ RANK

ESP32 Marauder is a suite of WiFi/Bluetooth offensive and defensive tools for the ESP32.

> ESP32 Marauder adds the ability to stream files off the SD card over the serial interface.

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

01 **Framed SD file streaming over serial** Not checked yet NEW

35

Adds framed SD file streaming over serial, enabling file transfer from the SD card via the serial interface.

— Names mechanism but no commands or config details given v1.21.0

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

proxychains-ng

Feedback ⓘ v4.18 NOTES ↗ RANK

proxychains ng (new generation) - a preloader which hooks calls to sockets in dynamically linked programs and redirects it through one or more socks/http proxies. continuation of the unmaintained proxychains project.

> proxychains-ng v4.18 adds support for routing through a Unix domain socket as the first proxy in a chain and introduces a new configure option for building fat binaries on Apple M2.

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

01 Unix domain socket as first proxy

Not checked yet NEW

83

The first proxy in a chain can now be configured as a Unix domain socket instead of a TCP endpoint, e.g. `socks5 unix:/home/user/.ssh/sock/socks.sock` in `proxychains.conf`, letting proxychains reach a SOCKS proxy bound to a local socket (such as one created by `ssh -D` with `StreamLocalBindMask`) rather than an exposed TCP port.

Route a tool through an SSH SOCKS proxy bound to a local socket that only its owner can access, avoiding an exposed TCP port.

```
$ ssh -N -o StreamLocalBindMask=0177 -D ~/.ssh/sock/socks.sock
relay
```

Point proxychains at the socket-bound SOCKS proxy by adding a `socks5 unix:` entry as the first (and only) entry in the proxy list.

```ini
[ProxyList]
socks5  unix:/home/user/.ssh/sock/socks.sock
```

— Names exact config syntax and file, with runnable example v4.18

02 Fat binary build for Apple M2

Not checked yet NEW

65

A new `--fat-binary-m2` configure option builds a fat binary targeting Apple M2, usable as `./configure --prefix=/usr --sysconfdir=/etc --fat-binary-m2 && make`.

Compile a fat binary that runs natively on Apple M2 hardware.

```
$ ./configure --prefix=/usr --sysconfdir=/etc --fat-binary-m2 &&
make
```

— Exact flag and build command given, but no further mechanism detail v4.18

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

psalm-plugin-laravel

Feedback ⓘ v4.23.1 NOTES ↗ RANK

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

> Psalm-plugin-laravel significantly expanded its taint-tracking coverage this window — flowing taint through Stringable chains, open redirects, raw DB columns, cookies and session values — while also adding Laravel 13 console Signature-attribute support and bumping its minimum Psalm requirement.

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

01 **Taint propagation through Stringable chains** Not checked yet

IMPROVED 90

Taint now propagates through `Stringable` method chains — trim(), upper(), after(), limit(), append(), and toString() preserve taint, while hash(), encrypt(), toBase64(), and numbers() drop it, and slug() and stripTags() remove the HTML taint specifically.

Catch XSS from user input that travels through a `Stringable` chain before being echoed — previously missed, now reported as `TaintedHtml`.

php

```
$name = $request->str('name');
echo $name->trim()->upper();
// TaintedHtml — taint is preserved through the chain
```

— Names exact methods and behaviour with a worked code example v4.23.1

02 **Laravel 13 #[Signature] attribute support** Not checked yet NEW

90

Reads console command signatures from Laravel 13's `#[Signature]` attribute so that argument()/option() type narrowing, hasArgument()/hasOption(), and the `InvalidConsoleArgumentName` / `InvalidConsoleOptionName` checks all work for attribute-based commands, mirroring how the attribute overrides `protected $signature` in Laravel itself.

Annotate a Laravel 13 console command with `#[Signature]` so Psalm narrows argument and option types accurately without a `$signature` property.

```php
#[Signature('app:demo {name} {--tag=*}')]
final class Demo extends Command {
    public function handle(): void {
        $name = $this->argument('name'); // string
        foreach ($this->option('tag') as $tag) {} // array<int,
string>
    }
}
```

— Full mechanism, named checks, and runnable example v4.22.6

03 **Open redirect and URL component taint detection** Not checked yet

NEW 80

Detects user-controlled redirect targets as `TaintedRedirect` and injected URL parts as `TaintedUrlComponent` or `TaintedUrlPath` on redirect(), Redirector, `ResponseFactory` and its contract.

Catch an open redirect where user input is appended directly to a URL passed to redirect()

```php
redirect('https://example.com/' . (string) $request-
>input('page'));
// psalm-laravel reports: TaintedUrlComponent
```

— Names exact taint types and APIs, includes runnable example v4.23.0

04 **Raw DB column taint tracking** Not checked yet NEW 80

Tracks what raw `DB` / Connection queries write into each column and reports taint where a subsequent query reads that column back, without needing Eloquent.

Detect taint flowing from a raw INSERT into a later SELECT on the same column, without needing Eloquent

```php
DB::insert('INSERT INTO comments (author, body) VALUES (?, ?)',
['admin', $request->input('body')]);

$comment = DB::selectOne('SELECT author, body FROM comments');
echo (string) $comment->body;
// psalm-laravel reports: TaintedHtml
echo (string) $comment->author;
// no issue
```

— Clear mechanism with before/after code example v4.23.0

05 **Taint clearance via regex validation** Not checked yet NEW

75

Trusts a value that Validator::make() or validator() checked against a `regex` rule once the validator passes, clearing `TaintedHtml` and similar issues.

Clear taint on a string once it has been validated against a regex with Validator::make()

```php
$validator = Validator::make(['agent' => $agent], ['agent' =>
['required', 'regex:#^[a-zA-Z0-9 _.\-]+$#']]);
if ($validator->fails()) {
    return;
}
echo $agent;
// no issue — regex rule clears the taint
```

— Mechanism and API named with a working example v4.23.0

THINNER COVERAGE BELOW

06 **Insecure cookie flag detection** Not checked yet NEW 50

Detects cookies created with the secure flag explicitly off as `InsecureCookie` when using response()->cookie().

— Names the check and API but no example given v4.23.0

07 **Session value taint tracking** Not checked yet NEW 50

Reads session values back from what the application stored under their key, so only tainted-at-source values remain flagged — getOldInput() and previousUrl() continue to be treated as user input.

— Mechanism described but no example or command v4.23.0

08 **Minimum Psalm version raised to 7.13.1** Not checked yet BREAKING

50

Requires `vimeo/psalm` `^7.13.1` ; projects still on Psalm 7.12 must stay on v4.22.7.

— Exact version pin and migration note given v4.23.0

└──▷ **BREAKING ON UPGRADE**

! Requires `vimeo/psalm` `^7.13.1` ; projects on Psalm 7.12 must stay on v4.22.7.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Security advisory permissions are now generally available (2026-11-20) BLOG ↗

RANK

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

> GitHub's headline shipment is the GitHub Enterprise Server 3.23 release candidate, which bundles on-appliance AI secret detection, expiring OAuth tokens, wider malware advisory coverage, and a Dependabot cooldown, alongside GA security advisory permissions for custom roles, a Favorite repositories dashboard section, and multi-account support in GitHub Desktop.

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

02 Multiple accounts per host in GitHub Desktop Not checked yet

NEW 80

GitHub Desktop now supports adding more than one account per host (`GitHub.com`, GitHub Enterprise Server, or GHE.com) with no fixed limit, via Settings/Options > Accounts. Each repository can be associated with a specific account so clone, publish, and Git operations automatically use the right credentials; users choose which signed-in account to use when a repo is accessible by multiple accounts, or Desktop infers it when only one applies, and the association can be changed later from repository settings or the repository list's context menu.

Add a second `GitHub.com` or Enterprise account so each repository automatically uses the right credentials without manual switching.

> 📍 In GitHub Desktop 3.7.0, go to Settings/Options › Accounts and sign in with the additional account.

Re-assign an existing repository to a different signed-in account after adding multiple accounts.

> 📍 Right-click the repository in the repository list and choose the new account from the context menu, or open the repository's settings page and change the account association.

— Full workflow and UI paths described, no CLI surface

GitHub Desktop now supports multiple ac…

03 Granular security advisory permissions for custom roles

Not checked yet NEW 78

GitHub adds four security advisory permissions that can be assigned to custom repository roles under the Security section when creating or editing a role: View security advisories , Create, Edit, and Manage security advisories , 'Accept private vulnerability reports and close or reopen security advisories', and Publish security advisories and request CVEs . This lets orgs scope advisory access per-repository by assigning the role to users and teams only on the repositories they maintain, instead of granting full admin or org-wide security manager access.

Grant a triage team the ability to view and manage security advisories on a specific repo without making them admins or org-wide security managers

📍 In your GitHub organization, go to Settings › Roles › New role, add the Create, Edit, and Manage security advisories permission under the Security section, then assign the role to your triage team on only the target repository under that repository's Settings › Collaborators and teams.

— Names all four permissions and exact UI path, no API endpoint

Security advisory permissions are now g…

05 **Favorite repositories section on dashboard** Not checked yet NEW

65

Users can pin up to 25 repositories to the dashboard via the new Favorite repositories section, opened with Edit or Add repositories and saved with Save or discarded with Cancel . Favorites sync with pinned repositories in GitHub Mobile, with changes in either place appearing in the other after a refresh.

Pin the repositories you work in most often so they appear at the top of your GitHub dashboard.

📍 On github.com, go to your Dashboard › Favorite repositories › select Edit or Add repositories › search for and select repositories › select Save .

— Exact UI path and 25-repo limit, no API/config key

Pin your favorite repositories to your …

THINNER COVERAGE BELOW

└──▷ **BREAKING ON UPGRADE**

! Token expiration is the default for newly created OAuth apps; before requiring expiration on existing apps, all clients must be updated to request, store, and use refresh tokens or sign-in will break for clients that do not support them.

WAS THIS USEFUL? 👍 👎

syzkaller Feedback ⓘ changes since 2026-11-17 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds CSRF protection to dashboard web UI forms and includes AI bisections in the bug lifetimes graph

syzkaller now adds CSRF protection to dashboard web UI forms, requiring a valid `CSRFToken` on state-changing POST requests and reading the token from `csrfFormField`. Not checked yet

syzkaller now authenticates RPC connections before unpacking large payloads in `pkg/rpctype`. Not checked yet

syzkaller's AI bisection flow now marks bisections that point to release commits as inconclusive rather than reporting them as results. Not checked yet

syzkaller's AI bisection backend now uses `gemini-4-hydrogen` as the core model, replacing `gemini-3.6-flash`. Not checked yet

syzkaller now falls back to the next LLM model only on blocked requests, and drops thought signatures of other models on fallback. Not checked yet

> └──▷ **BREAKING ON UPGRADE**
>
> ! Dashboard admin and AI job handlers now require POST for state-changing UI actions; GET requests to those endpoints will no longer perform state changes.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.14 NOTES ↗

Infisical is the open-source platform for securing secrets, certificates, access, and AI Agents.

Infisical adds bulk secret approval endpoints, AWS IAM gateway sessions beyond the 1-hour STS cap, PKI certificate revocation approval, Slack alert channels, and grouped secret scan findings.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.167.14
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.167.14
```

Approve, reject, or merge secret approval requests in bulk with new batch review and merge API endpoints, including an approve-all and approve-without-merging option. Not checked yet

Broker AWS IAM CLI sessions through the gateway so they can outlive the 1-hour STS cap. Not checked yet

Require approval before revoking a PKI certificate. Not checked yet

Add auto-approval checks on the requester's user and groups for privileged access requests. Not checked yet

Secret scanning now groups findings that share the same value. Not checked yet

+3 more

NEW INTEGRATION Coolify secret sync

NEW INTEGRATION Neon dynamic secret provider

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS Feedback ⓘ

Invalid secrets no longer block your developers and AI coding agents (2026-11-20)

BLOG ↗ RANK

GitGuardian SaaS is a secrets detection platform that scans Git repositories and developer workflows to identify and remediate exposed API keys, credentials, and sensitive data.

> GitGuardian now filters confirmed-invalid secrets out of scan results so hooks and AI coding agents stop blocking on dead credentials, and Perimeter search gains repository URL lookup.

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

01 **Automatic filtering of confirmed-invalid secrets** Not checked yet

IMPROVED 86

GitGuardian's scan API now drops confirmed-invalid secrets from results for `pre-commit`, `pre-push`, `ggshield secret scan ci`, and direct calls to the `/v1/scan` and `/v1/multiscan` endpoints, so hooks and AI agents only block on valid or unverifiable secrets. The filter runs server-side, meaning no code change or ggshield upgrade is needed to benefit.

Verify that your CI pipeline now passes when only invalid secrets are detected — no code change or ggshield upgrade needed, since the filter runs server-side on the `/v1/scan` endpoint.

```
$ curl -s -X POST https://api.gitguardian.com/v1/scan \
  -H 'Authorization: Token <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"document": "<content-with-known-invalid-secret>",
"filename": "config.yaml"}'
```

— Names exact endpoints, hooks and command, with runnable example

Invalid secrets no longer block your de…

THINNER COVERAGE BELOW

02 **Repository URL search in Perimeter** Not checked yet NEW

42

Paste a full repository URL — including sub-page URLs for files or branches — into the Perimeter search or Source filter to locate the matching source.

— Clear UI starting point but no mechanism detail

Invalid secrets no longer block your de…

WAS THIS USEFUL?  

◆ WAF, Bot & API Protection

Bunkerity BunkerWeb

Feedback ⓘ v1.6.18 NOTES ↗ RANK

BunkerWeb is an open-source and cloud-native Web Application Firewall (WAF)

BunkerWeb adds Zstandard compression, per-service DNSBL overrides, JSON ban output, URL-list refresh intervals, a new TLS ClientHello plugin phase, and login throttling with configurable lockout settings.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/bunkerity/bunkerweb:1.6.18
```

git

```
$ git clone --branch v1.6.18
https://github.com/bunkerity/bunkerweb.git
# already have the repo? check out this version:
$ git checkout v1.6.18
```

published under the project's GitHub account · checked 21 Nov Wrong command?

Enable Zstandard response compression with `USE_ZSTD`, preferred over Brotli and gzip when the client accepts it; enabling it can also add `Vary: Accept-Encoding` to eligible uncompressed, gzip, and Brotli responses. Not checked yet

Set `DNSBL_LIST` per service so individual services can override the global DNS blocklist; services without an override keep the global list. Not checked yet

Output the current ban list as JSON for scripts and monitoring with `bwcli bans --json`. Not checked yet

Control how often URL lists are refreshed in the `blacklist`, `greylist`, and `whitelist` plugins with the new `*_URLS_REFRESH_INTERVAL` settings (default `1h`). Not checked yet

Restrict antibot challenges to specific URIs with the new `ANTIBOT_INCLUDE_URI` setting; `ANTIBOT_IGNORE_URI` still takes precedence. Not checked yet

+11 more

UPDATED DETECTION OWASP CRS v4 updated to 4.30.0

NEW SUPPORTED PLATFORM Fedora 45 support

NEW INTEGRATION Hosttech DNS-01 provider via certbot-dns-hosttech

└──▷ **BREAKING ON UPGRADE**

! BunkerWeb ModSecurity rules move to IDs 9,800,000-9,809,999; custom exclusions referencing old IDs must be updated (mapping in the docs).

! Stream upstream TLS settings (`REVERSE_PROXY_SSL_VERIFY=yes` , protocols, ciphers, cached client certificate) now enable TLS to the backend instead of being ignored; clear them for a plaintext backend.

! Setting regexes are stricter: list values joined without a separator (e.g. `404500`) and all-space values are refused; `AUTH_BASIC_USER` is now wider and accepts any value without `:` .

! The Lua API disk fallback now requires an explicit `API_TOKEN` line in `variables.env` ; a file without one is refused.

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ iOS SDK for URLSession 1.0.0 (2026-11-20)

BLOG ↗ RANK

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

> DataDome shipped a new iOS SDK that extends bot/fraud protection to apps using URLSession, with challenge handling and support for modern Swift concurrency patterns.

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

01 **iOS SDK for URLSession** Not checked yet NEW 75

A new `ProtectedURLSession` class intercepts and protects `data` , `upload` , `download` , and `bytes` URLSession task types on iOS. Challenges returned by a redirect destination are retried at that destination, matching URLSession's redirected-request behaviour, and the SDK supports async/await, completion handlers, and Combine as calling conventions for all protected request types.

– Names exact class and task types but no install/setup steps given

iOS SDK for URLSession 1.0.0 (2026-11-2…

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

humble

Feedback ⓘ 1.69 NOTES ↗ RANK

A humble and fast security-oriented HTTP headers analyzer.

> humble 1.69 adds safer non-interactive CI/CD behaviour and extends insecure cookie and header detection.

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

01 Non-interactive mode via `-e yes` Not checked yet NEW

75

humble adds a `yes` value for the `-e` option, allowing it to run non-interactively in CI/CD pipelines where stdin is closed without crashing, e.g. `python humble.py -u https://example.com -e yes`.

Run humble non-interactively in a CI/CD pipeline where stdin is closed, using `-e yes` to avoid crashes.

```
$ python humble.py -u https://example.com -e yes
```

— Names exact flag and runnable command 1.69

THINNER COVERAGE BELOW

02 Insecure `Partitioned` cookie detection Not checked yet NEW

50

humble now detects Partitioned cookies that are set without the Secure attribute and reports them as insecure.



— Names exact cookie attribute checked 1.69

03 **Three new insecure header checks** Not checked yet NEW

20

humble adds 3 new insecure header checks, though the specific headers are not named in the changelog.

— Count given but headers unnamed 1.69

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Socket

Feedback ⓘ

Expanded Java Coverage and More Accurate Reachability Results (2026-11-19)

BLOG ↗ RANK

Socket fights vulnerabilities and provides visibility, defense-in-depth, and proactive supply chain protection for JavaScript, Python, and Go dependen...

> Socket's biggest window changes are under the hood: expanded Java/Android reachability coverage, new webhook delivery-failure alerts, proxy support for reachability scans, and a reworked reachability-engine distribution mechanism and SBOM export format that both carry breaking changes for existing integrations.

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

01 Reachability engine distribution and version change

Not checked yet

80

BREAKING

Full application reachability scans now require outbound access to `api.socket.dev` and `storage.googleapis.com`, since Socket downloads the reachability engine directly instead of through npm; `SOCKET_CLI_COANA_LAUNCHER` no longer takes effect and an npm mirror alone is no longer sufficient. Older reachability engine versions released before this change cannot be used with `--reach-version`, and the engine itself is updated to Coana 15.15.0.

— Exact domains, env var and flag named but needs infra update

Socket Basics 3.5.0: Reachability and P…

02 Webhook failure notifications and auto-disable

Not checked yet

NEW

75

Socket now emails organization Owners and Admins when every outbound webhook delivery has failed for 12 hours, including the scheduled disable date, and automatically disables the webhook after four days of continuous failures with a follow-up email. The webhooks table also shows Unhealthy and Deliveries failing status badges so teams can track and re-enable affected endpoints.

Re-enable a webhook after fixing its endpoint so previous failures no longer count toward automatic disabling.

📍 In the Socket dashboard, go to Settings › Webhooks, find the entry marked Deliveries failing, fix the endpoint URL or configuration, then set the webhook

status back to Active.

— Clear thresholds and a UI path to re-enable webhooks

Get Notified When Webhook Deliveries Fa…

03 **Argument handling improvements for reachability scans**

Not checked yet IMPROVED 73

Reachability scans accept repeated `--exclude-paths` flags to supply multiple path exclusions, and `socket_additional_params` now preserves commas and quoted spaces within values so options such as `--reach-ecosystems=npm,pypi` are passed through correctly.

— Exact flags and parameter names a reader can apply directly

Socket Basics 3.5.0: Reachability and P…

04 **SPDX export relationship and supplier fields** Not checked yet

73

BREAKING

Socket's SPDX exports now use `DEPENDS_ON` relationships (from project to dependencies) instead of `DEPENDENCY_OF`, and include supplier information or `NOASSERTION` when the publisher is unknown. Integrations reading `DEPENDENCY_OF` relationships must be updated to read `DEPENDS_ON` instead.

— Exact field names given though no new endpoint Improved SPDX and CycloneDX SBOM Export…

05 **Java and Android reachability coverage expansion** Not checked yet

NEW 65

Socket's reachability engine now analyzes Android libraries packaged as AARs and traces Spring Boot auto-configuration, Spring XML bean definitions, service loading, JPA and servlet descriptors, and calls from Java into Scala objects to compute full application reachability.

— Names several specific mechanisms but no flag or command to invoke it

Expanded Java Coverage and More Accurat…

06 **Reachability timeout flag syntax change** Not checked yet BREAKING

65

The reachability scan timeout flag now uses the equals-sign form `--reach-analysis-timeout=90s` instead of the comma form `--reach-analysis-timeout,90s`, which is no longer correct.

Run a reachability scan covering both npm and PyPI ecosystems behind a corporate proxy, with a 90-second timeout and specific paths excluded.

```
$ HTTPS_PROXY=https://proxy.corp.example:8080 socket scan --reach-ecosystems=npm,pypi --reach-analysis-timeout=90s --exclude-paths=test/ --exclude-paths=vendor/
```

— Exact before and after flag syntax given Socket Basics 3.5.0: Reachability and P…

07 **Platform-specific Ruby gem handling in SBOM exports**

Not checked yet BREAKING 65

Platform-specific Ruby gems (such as `nokogiri`) locked in `Gemfile.lock` now appear once per platform, each with the registry version and a platform qualifier, so downstream tools can match them correctly. The first diff scan after this update will show affected platform-specific gems as removed and re-added once.

— Names exact file and example package but no UI path given

Improved SPDX and CycloneDX SBOM Export…

08 **Proxy support for reachability scans** Not checked yet NEW 60

Reachability scans now honour the `HTTPS_PROXY` and `NO_PROXY` environment variables, enabling scans to run behind a corporate proxy.

Run a reachability scan covering both npm and PyPI ecosystems behind a corporate proxy, with a 90-second timeout and specific paths excluded.

```
$ HTTPS_PROXY=https://proxy.corp.example:8080 socket scan --reach-ecosystems=npm,pypi --reach-analysis-timeout=90s --exclude-paths=test/ --exclude-paths=vendor/
```

— Named env vars with a runnable example command Socket Basics 3.5.0: Reachability and P…

THINNER COVERAGE BELOW

09 **Reachability accuracy and performance improvements**

Not checked yet IMPROVED 50

Vulnerabilities in libraries Socket cannot analyze now show Unknown instead of an incorrect Unreachable, falling back to precomputed results when a vulnerability pattern does not match the package. Projects using multiple versions of a crate now get call stacks specific to each version rather than a shared stack, and reachability analysis performance and precision improved for JavaScript, TypeScript and Python projects (with particular gains for PyTorch) as well as for Go.

— Describes engine behaviour changes with no actionable command

Expanded Java Coverage and More Accurat…

10 **CycloneDX export includes direct dependencies** Not checked yet

IMPROVED 35

Socket's CycloneDX exports now include the project's direct dependencies.

— Thin one-line description with no further detail

Improved SPDX and CycloneDX SBOM Export…

NEW DETECTION Malware scanning for Terraform and HCL files

NEW DETECTION CMake malware analysis as part of npm install-script behavior

└──▷ **BREAKING ON UPGRADE**

! `SOCKET_CLI_COANA_LAUNCHER` no longer takes effect; Socket downloads the reachability engine directly and an npm mirror alone is no longer sufficient.

! Older reachability engine versions released before this change cannot be used with `--reach-version`.

! Full application reachability scans now require outbound access to `api.socket.dev` and `storage.googleapis.com`; existing setups that relied on an npm mirror must be updated.

! The comma form `--reach-analysis-timeout,90s` is no longer correct; use `--reach-analysis-timeout=90s` instead.

! If your integration reads `DEPENDENCY_OF` relationships from Socket's SPDX exports, update it to read `DEPENDS_ON` instead.

! The first diff scan after this update will show affected platform-specific Ruby gems as removed and re-added once.

> New: Malware Scanning for Terraform, HCL, and CMake (2026-11-17) BLOG ↗
>
> STABLE

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

YARA-X

Feedback ⓘ v1.24.0 NOTES ↗ RANK

YARA-X is a rewrite of YARA in Rust.

> YARA-X v1.24.0 expands the DEX module with new containment-check functions, adds Windows ARM64 release builds and man pages, and improves scanner memory efficiency and rule complexity limits.

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

01 DEX module string, method and class checks

Not checked yet NEW

78

The `dex` module gains `has_string`, `has_method`, and `has_class` functions to test for the presence of strings, methods, or classes in DEX files; the previous names `contains_string`, `contains_method`, and `contains_class` are kept as aliases for YARA compatibility.

Query a DEX file for a specific class using the new canonical function name in a YARA-X rule.

yara

```
rule dex_example {
  condition:
    dex.has_class("Lcom/example/MaliciousClass;")
}
```

— Names exact functions and aliases with a runnable rule example. v1.24.0

02 Higher WASM instruction limit for rule conditions

Not checked yet

IMPROVED

68

The maximum WASM function instruction limit (`MAX_FUNC_INSTRS`) was increased from 100,000 to 300,000, allowing larger and more complex rule conditions to compile.

— Gives exact before/after limit values. v1.24.0

THINNER COVERAGE BELOW

03 Alias argument for module_export macro

Not checked yet NEW

57

The `#[module_export]` macro now supports an optional `alias` argument, letting module authors export functions under alternate names.

— Names the macro and argument but no usage example. v1.24.0

04 **Shared regex compilation across scanners** Not checked yet IMPROVED

55

Compiled regular expressions and regexp sets are now shared across `Scanner` instances within `Rules`, reducing memory use when running many scanners against the same rule set.

— Explains mechanism and benefit but no concrete numbers. v1.24.0

---

05 **Generated man pages for yr CLI** Not checked yet NEW

50

Man pages for `yr` and its subcommands are now generated and included under `cli/man/`.

— Names exact location but describes only the addition itself. v1.24.0

---

06 **Windows ARM64 release builds** Not checked yet NEW

45

YARA-X now ships Windows ARM64 release builds for the CLI, the C API, and Python wheels.

— Names the exact artifacts but no further mechanism. v1.24.0

---

07 **Full keyword highlighting in Playground editor** Not checked yet

IMPROVED

30

The Playground editor now highlights all YARA keywords and operators.

— Minor UI change described only briefly. v1.24.0

WAS THIS USEFUL? 👍 👎

CAPEv2 Feedback ⓘ changes since 2026-11-19 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 Granian ASGI server, migrates the host runtime from Poetry to uv with Python 3.14, and adds a `migrate-uv` installer command

└──▷ **TRY IT**

Migrate an existing CAPEv2 host from Poetry to uv and Python 3.14 without reinstalling the sandbox.

```
$ sudo bash installer/cape2.sh migrate-uv | tee migrate-uv.log
```

Migrate an existing Poetry-based host to uv with Python 3.14 in-place by running `sudo bash installer/cape2.sh migrate-uv` (documented in `docs/book/src/installation/uv_migration.rst`), installing a standalone CPython 3.14 interpreter in `/usr/local/share/uv/python` without touching the system `/usr/bin/python3`. Not checked yet

CAPEv2's web and ASGI stack now uses Granian, dropping the previous daphne and uWSGI servers. Not checked yet

The `--use-uv` installer flag is now a no-op; uv with Python 3.14 is the default for all new installs. Not checked yet

└──▷ **BREAKING ON UPGRADE**

! The `--use-uv` installer flag is now a deprecated no-op; uv with Python 3.14 is the default. Existing Poetry-based hosts must run `sudo bash installer/cape2.sh migrate-uv` to migrate.

! The web/ASGI stack has dropped daphne and uWSGI in favour of Granian; any systemd or proxy configuration referencing daphne or uWSGI must be updated.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ changelog · Nov 21, 2026 CHANGELOG ↗

RANK

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

> Chainguard opened self-serve access with new Starter and Team plans, launched the Rust-based Project Theseus initiative to rebuild core open-source components starting with a new memory allocator, disclosed a second wave of silent Java vulnerabilities via Athena, and added 25 new container images.

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

01 **Athena discloses 12 more silent Java vulnerabilities**

Not checked yet NEW 75

Athena's second wave discloses 12 more silent vulnerabilities in Java projects, including two critical authentication bypasses across nine libraries. Remediated versions are available from Chainguard Repository using the `-0cgr.n` version qualifier, adoptable via a lockfile change and rebuild.

— Names exact version qualifier and remediation path changelog-20261121-fc0559fe

02 **Project Theseus and the challoc allocator** Not checked yet NEW

65

Project Theseus is a new Chainguard initiative to rebuild core open-source components in Rust to eliminate classes of vulnerabilities rather than patching them repeatedly. The first component is `challoc`, a memory allocator replacing glibc malloc, jemalloc, tcmalloc, and mimalloc.

— Clear mechanism and named component, but no adoption steps changelog-20261121-fc0559fe

03 **Self-serve Starter and Team plans** Not checked yet NEW

60

Chainguard now offers self-serve sign-up with two new plans: free Chainguard Starter (no credit card required) giving access to Chainguard Containers, Chainguard Libraries, and other products within usage limits, and paid Chainguard Team which raises most Starter limits and adds ticket support.

— Names plans and scope but no exact limits or pricing changelog-20261121-fc0559fe

THINNER COVERAGE BELOW

04 **25 new container images added** Not checked yet NEW 50

Chainguard added 25 new container images this week, including `aws-route53-controller`, `grafana-pyroscope-fips`, `dify-web`, `kata-monitor`, `kgateway-envoy-wrapper`, `kgateway-sds`, `neuvector-registry-adapter`, `oc-mirror`, and `gosu-fips`, among others.

. . . .

— Lists concrete image names but no usage detail changelog-20261121-fc0559fe

WAS THIS USEFUL? 👍 👎

◆ SaaS Security Posture

Nudge Security Feedback ⓘ

Get findings for AI agents discovered in the browser NOTES ↗

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

Nudge Security now surfaces AI agent findings from the browser extension in the same Findings view as connected-app agent findings.

NEW DETECTION Nudge Security now generates findings for AI agents discovered through the browser extension, surfacing them in the Findings view alongside findings for agents discovered through connected apps. Not checked yet

Enable AI agent discovery in the browser extension settings to ensure browser-discovered agents are included in findings. Not checked yet

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ Secure SaaS - November 20, 2026

BLOG ↗ RANK

Sysdig Secure is a cloud-native security platform that provides runtime threat detection, vulnerability management, and compliance monitoring for containers, Kubernetes, and cloud environments.

> Sysdig Secure adds a Stage filter to its vulnerability views, letting teams see registry, pipeline, and runtime scan results together or separately.

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

01 **Stage filter for vulnerability findings** Not checked yet NEW

62

A new Stage filter on the Vulnerability Findings page and the Vulnerability Overview dashboard lets users filter by Pipeline, Registry, Runtime, or any combination, and both views now include registry and pipeline scan results alongside runtime findings.

Review all pre-runtime (registry and pipeline) vulnerability findings alongside runtime findings in one view, filtered to only pipeline stage results.

📍 In Sysdig Secure, go to Vulnerabilities › Findings, then open the Stage filter and select Pipeline (or Registry, Runtime, or any combination) to scope the results.

— Names exact filter values and two UI locations, but no mechanism detail

Secure SaaS - November 20, 2026

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ

Control how Huntress notifications are prioritized in HaloPSA NOTES ↗ RANK

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

> Huntress shipped two unrelated improvements this window: configurable HaloPSA ticket priority mapping and a simpler sign-in flow for Huntress MCP.

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

01 **Notification-severity-to-ticket-priority mapping for HaloPSA**

Not checked yet NEW 72

Admins can map each Huntress notification severity (e.g. Critical, High, Medium, Low) to a HaloPSA ticket priority, giving control over how Huntress-generated tickets are prioritized in HaloPSA. Configured in the Huntress portal under Settings > Integrations > HaloPSA, under Ticket Priority Mapping .

Set ticket priority in HaloPSA to match your team's response SLAs for each Huntress alert severity level.

> 📍 In the Huntress portal, go to Settings › Integrations › HaloPSA, then under Ticket Priority Mapping assign a HaloPSA ticket priority to each Huntress notification severity (e.g. Critical, High, Medium, Low).

— Names exact UI path and mapping mechanism Control how Huntress notifications are …

THINNER COVERAGE BELOW

02 **Portal credential sign-in for Huntress MCP** Not checked yet

IMPROVED 35

Users can sign in to Huntress MCP directly with their Huntress portal credentials, bypassing the need for an administrator to first create an API key.

— Describes benefit but no concrete steps given Sign in to Huntress MCP with your porta…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Elastic Security Labs

Feedback ⓘ

Mapping the AI threat landscape: MITRE ATLAS comes to Elastic Security (2026-11-20)

BLOG ↗

Elastic Security Labs is a powered by Elasticsearch. Ready to deploy your data?

Elastic Security 9.6 adds a new ATLAS coverage overview table, ATLAS support in the rule form and AI rule creation workflow, Agent Builder ATLAS awareness, and a new managed MITRE data service in Kibana.

A new ATLAS coverage overview table shows which ATLAS tactics and techniques your rules already address, scored in its own view separate from ATT&CK numbers.
Not checked yet

UPDATED DETECTION The detection rule creation and edit form now supports adding both ATT&CK and ATLAS mappings to custom rules or existing prebuilt rules.
Not checked yet

The AI-assisted rule creation workflow now supports adding both ATT&CK and ATLAS mappings to rules. Not checked yet

Agent Builder skills now understand ATLAS, letting you ask an agent which AI threats your installed rules cover or have it recommend prebuilt rules for specific ATLAS tactics and techniques. Not checked yet

A new managed MITRE data service in Kibana ingests MITRE ATLAS and ATT&CK entities as Saved Objects, enabling querying and filtering via a data client and API, and decoupling framework updates from Kibana release cycles. Not checked yet

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

MCP tools in Genie Agents are generally available (2026-11-20) 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 shipped GA warehouse-level statement timeouts with full API/Terraform/SDK support, GA Delta Sharing across regulatory domains, GA MCP tools for Genie Agents, a new Lakewatch CLI command group, GitHub app pull-request automation, and GA LTAP Direct Writes for synced tables.

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

01 **Warehouse-level statement timeouts** Not checked yet NEW

90

Set per-warehouse query time limits using the `statement_timeout` field across the SQL warehouses API, Databricks SDKs, Terraform, or Declarative Automation Bundles to automatically cancel runaway statements. This is now generally available and can be set on a Terraform `databricks_sql_warehouse` resource or applied directly via `PATCH /api/2.0/sql/warehouses/<warehouse-id>`.

Cap long-running queries on a shared warehouse so runaway jobs cannot monopolize compute — set `statement_timeout` in your Terraform warehouse resource.

hcl

```hcl
resource "databricks_sql_warehouse" "shared" {
  name = "shared-warehouse"
  statement_timeout = 3600
}
```

Apply a statement timeout directly via the SQL warehouses API to an existing warehouse without changing IaC.

```
$ curl -X PATCH https://<workspace-
host>/api/2.0/sql/warehouses/<warehouse-id> \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"statement_timeout": 3600}'
```

— Names exact field, API endpoint, and runnable Terraform/curl examples

Warehouse-level statement timeouts are …

02 **Lakewatch management via Databricks CLI** Not checked yet NEW

70

Manage Lakewatch datasources, parsers, detection rules, and per-catalog configuration from the command line using the new `lakewatch` command group in the Databricks CLI, now in beta. Requires CLI version 1.25.0 or above, with parser commands requiring version 1.27.0 or above.

— Names command group and version gates but no sample invocation

Manage Lakewatch with the Databricks CL…

03 **LTAP Direct Writes for synced tables** Not checked yet NEW

70

Enable LTAP Direct Writes when creating a synced table so loads write directly to Lakebase storage instead of through your compute endpoint, making large initial loads and full refreshes finish faster without competing with application queries. Now generally available, selectable under Sync method when creating a synced table in the Databricks UI.

When creating a synced table with a large initial data load, select LTAP Direct Writes so the ingestion goes straight to Lakebase storage and does not slow down running application queries.

📍 In the Databricks UI, go to Lakebase › Synced Tables › Create Synced Table, then under Sync method select LTAP Direct Writes before saving.

— Explains mechanism and gives UI navigation steps

LTAP Direct Writes for synced tables is…

04 **Delta Sharing across regulatory domains** Not checked yet NEW

60

Databricks-to-Databricks Delta Sharing across regulatory domains is now generally available, letting commercial workspaces share data with AWS GovCloud, AWS GovCloud DoD, and Azure China environments.

— Names target domains but no setup steps or command

Databricks-to-Databricks Delta Sharing …

THINNER COVERAGE BELOW

05 **MCP tools in Genie Agents** Not checked yet NEW

55

Genie Agents can now attach external MCP servers registered in Unity Gateway, in addition to Databricks-provided connectors, and Agent mode APIs now support MCP tools. This capability is now generally available.

— Names Unity Gateway and Agent mode APIs but no config example

MCP tools in Genie Agents are generally…

06 **GitHub app pull request automation** Not checked yet NEW

55

Genie Code and Genie ZeroOps can now create and manage pull requests on your behalf once the Databricks GitHub app is granted Read and write access to pull requests; existing branch protection rules and required reviews still apply.

.
.
.

— Names the permission and products but no setup command

The Databricks GitHub app will request …

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Endpoint Management (UEM/RMM)

Devicie

Feedback ⓘ announced · Nov 21, 2026 PRESS ↗

Devicie is a keep pace with application updates and configuration changes. Devicie automates recurring management natively through Microsoft Intune.

Devicie adds multi-tenant MSP portal support, letting MSPs view browser security events across all customer tenants in one place with per-tenant reporting keys.

NEW SUPPORTED PLATFORM Edge for Business Reporting Connector cross-platform support

NEW DETECTION Browser security event categories in Devicie Reporting Connector

WAS THIS USEFUL? 👍 👎

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ updated · Nov 21, 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's latest update focuses on expanding testing coverage and multi-tenant/enterprise administration, adding REST API-only webapp testing, an Audit Log dashboard, Kubernetes pentest configuration options, and a Tripwires webhook connector for SIEM/SOAR integration.

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

01 **Kubernetes pentest configuration options** Not checked yet NEW

45

Kubernetes pentests can now be configured with custom service accounts, by importing rules from YAML files, and by setting up a proxy.

– Names three config options but no exact flags or keys snapshot-20261121

02 **Client Management for parent accounts** Not checked yet NEW

45

Parent accounts in multi-tenant environments now have a top-level view of clients' operational health, license allocations, and API keys (Client Management).

– Describes scope of new view but no navigation detail snapshot-20261121

03 **Tripwires webhook connector** Not checked yet NEW 40

Tripwires alerts can now be sent to SIEM, SOAR, or automation platforms through a new webhook connector.

– Describes destination types but no endpoint or setup details snapshot-20261121

04 **Audit Log dashboard for Org Admins** Not checked yet NEW

40

Org Admins now have a dashboard to review, investigate, and export records of account activity (Audit Log).

– Names dashboard and export capability, no exact UI path snapshot-20261121

05 **REST API-only webapp testing** Not checked yet NEW 35

The WebApp pentest engine can now test REST APIs that don't expose a web UI.

— Clear capability but no mechanism or setup steps given snapshot-20261121

06 **Parallel WebApp pentest execution with expanded evidence**

Not checked yet IMPROVED 35

WebApp pentest flows now run in parallel, returning faster results, alongside expanded adaptive agents and verifiable evidence included in reports.

— Explains behaviour change but no metrics or specifics snapshot-20261121

07 **Updated BloodHound integration** Not checked yet IMPROVED 10

The BloodHound integration is now updated to a newer version.

— No version number or change details given snapshot-20261121

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Vulnerability & Exposure Management

OXO

Feedback ⓘ v2.25.4 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> oxo extended two message schemas this window, adding certificate history data to fingerprinting and a capture timestamp to UI call tracking.

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

01 **Certificate data in fingerprint history message** Not checked yet

IMPROVED 45

Adds x509 certificate data to the `v3.fingerprint.certificate.history` message for richer certificate history fingerprinting.

— Names exact message field but no mechanism detail. v2.25.4

02 **Timestamp field on UI call capture message** Not checked yet

IMPROVED 45

Adds a `captured_at` field to the `v3.capture.ui_call` message to record when a UI call was captured.

— Names exact message field but no mechanism detail. v2.25.3

WAS THIS USEFUL? 👍 👎

Qualys VMDR

Feedback ⓘ changelog · Nov 20, 2026 BLOG ↗ RANK

Qualys VMDR is a cloud-based platform that continuously discovers, assesses, prioritizes, and remediates vulnerabilities across an organization's IT assets.

> Qualys VMDR added a TruRisk Eliminate mitigation for Java4Shell that applies a JVM flag to eligible hosts, reports process-level coverage, and supports one-click rollback.

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

01 Java4Shell mitigation via TruRisk Eliminate

Not checked yet NEW

88

TruRisk Eliminate applies the `-XX:-UseBASE64Intrinsics` flag to `JAVA_TOOL_OPTIONS` on affected Linux and Windows hosts (linked to QIDs such as 389494), automatically checking each host for an AVX-512 VBMI CPU and an affected Java runtime before making any change. After mitigation, the console reports every running Java process with its effective JVM option setting in HotSpot resolution order, flagging unknown processes to surface coverage gaps, and the mitigation can be rolled back in one click, restoring the previous `JAVA_TOOL_OPTIONS` value.

When you need a bridge fix for Java4Shell before patched JDK builds are available, apply the JVM flag across all eligible hosts without touching containers.

> 📍 In the Qualys console, go to TruRisk Eliminate, select the Java4Shell mitigation (linked to QIDs such as 389494), choose your target Linux or Windows hosts, and click Apply. The Cloud Agent checks each host for an AVX-512 VBMI CPU and an affected Java runtime, then sets `-XX:-UseBASE64Intrinsics` in `JAVA_TOOL_OPTIONS`. Plan JVM restarts with application owners; Windows services require a host reboot. For containers, set the flag in the image or pod spec manually.

— Names exact flag, config var, QIDs, and console workflow with rollback.

2026-11-20-a6009168

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

openfga

Feedback ⓘ v1.23.0 NOTES ↗ RANK

Openfga is a high-performance and flexible authorization/permission engine built for developers and inspired by Google Zanzibar.

> OpenFGA v1.23.0 introduces a typed query AST and Querier interface for structured backend queries (with a SQLite implementation), optimizes PostgreSQL's latest-model lookup, and upgrades grpc-gateway with header and query-parameter validation changes.

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

01 Typed query AST and Querier interface for storage backends

Not checked yet

83

BREAKING

Adds a typed query AST and new Querier interface in `pkg/storage/adapter`, letting custom `RelationshipTupleReader` backends run queries as structured statements instead of hand-built SQL; OpenFGA now implements Querier for SQLite, enabling future SQL optimizations for SQLite-backed deployments. Custom `RelationshipTupleReader` implementations (library use only) must now implement a Querier method on the interface — returning `nil` is safe but skips future SQL optimizations.

— Names interface, package path and migration requirement clearly. v1.23.0

02 grpc-gateway v2.33.0 header and query parameter behavior change

Not checked yet

73

BREAKING

Upgrades to grpc-gateway v2.33.0, which no longer forwards Grpc-Metadata-* request headers to gRPC metadata, and now rejects malformed query parameters (e.g. non-decimal float spellings, `field[key]` on non-map fields) that were previously accepted.

— Names dependency version and exact behavior changes affected users must check. v1.23.0

03 PostgreSQL latest-model lookup optimization

Not checked yet

IMPROVED

65

PostgreSQL's `FindLatestAuthorizationModel` now reads only the rows of the store's latest model instead of every model version, so Check and other requests made without `authorization_model_id` no longer slow down as model history grows.

— Clear mechanism and benefit but nothing for reader to configure. v1.23.0

└──▷ BREAKING ON UPGRADE

! Custom `RelationshipTupleReader` implementations (library use only) must now implement a `Querier` method on the interface (see `pkg/storage/adapter`); returning `nil` is safe but skips future SQL optimizations.

! grpc-gateway v2.33.0 (included in this release) no longer forwards `Grpc-Metadata-*` request headers to gRPC metadata, and now rejects malformed query parameters (e.g. non-decimal float spellings, `field[key]` on non-map fields) that were previously accepted.

WAS THIS USEFUL?  

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-20) BLOG ↗

RANK

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

> JumpCloud's latest release headlines an audit-ready Access Certification workflow with multi-stage approval chains and signed exports, alongside new Agent Delegated Access for AI agents, capability-based MCP app setup, Workflow support in the MCP Server for Admins, OAuth-enabled workflow connectors with an expanded template library, and enhancements to macOS patch policy enforcement, Device Lock, and SCIM group naming.

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

01 **Access Certification review campaigns and exports** Not checked yet

NEW 85

JumpCloud adds an Access Certification feature supporting one-time or recurring review campaigns scoped to SSO applications or directory user groups, with multi-stage approval chains and automated remediation that revokes access or suspends identities when deadlines expire. Campaign snapshots can be exported in CSV, JSON, or PDF formats with SHA-256 document hashing and formal administrator sign-off for compliance audits.

Launch a recurring quarterly access review campaign for an SSO application to satisfy compliance audits, then export the signed snapshot for your auditor.

> 📍 In the Admin Portal, go to Access Certification, create a new campaign, select SSO Applications as the scope, set recurrence to quarterly, configure a multi-stage review chain (Managers then Resource Owners), enable automated remediation on deadline expiry, and after the review cycle completes export the campaign snapshot as PDF or CSV with administrator sign-off.

— Full mechanism, formats and hashing named, plus a walkthrough example

Release Notes 2026 (2026-11-20)

02 **Capability-based app setup for SSO, SCIM, MCP** Not checked yet

NEW 85

New applications only show configuration for the capabilities selected — SSO, SCIM, and/or MCP — with Manage MCP Servers as the selection for MCP-only access, letting admins create MCP-only applications, assign them to user groups for access control, and make the MCP server available through AI Gateway without configuring SSO or creating a bookmark. Admins can review a user group's applications from `Identity Management`

`> User Groups` by opening the group and selecting the `Applications` tab, where apps with validated MCP configs are marked as ready for AI Gateway.

– Named config selection and exact navigation path given `Release Notes 2026 (2026-11-20)`

---

03 **Workflow connector OAuth and expanded template library**

`Not checked yet` `IMPROVED` 80

Workflow connectors now support OAuth 2.0 authentication including Client Credentials and Authorization Code with PKCE, with credentials stored in Secrets Manager and access tokens refreshed automatically. The Workflow template library expands with ready-to-use templates: `Auto-Suspend Users on Inactivity Alert`, `Auto-Block Newly Discovered Generative AI Apps`, `Mac Enrollment - Auto-Revoke Admin Privileges`, `Sensitive Software Install Detection - Email Notification`, `User Group Added - Send Member Email`, `User Activated - Add to Pre-Hire Holding Group If Before Start Date`, and `Agent Enrollment - Add Device Custom Attribute`.

Block newly discovered generative AI apps automatically as they appear in your environment, using the new Workflow template.

> 📍 In the Admin Portal, open the Workflow Builder template library, select `Auto-Block Newly Discovered Generative AI Apps`, customize the trigger and action settings for your environment, then save and activate the workflow.

– All templates and OAuth flows named with a usage walkthrough

`Release Notes 2026 (2026-11-20)`

---

04 **Device Lock command PIN escrow and status column** `Not checked yet`

`IMPROVED` 70

The Device Lock MDM command for macOS, iOS, and iPadOS now supports automatic PIN generation and escrow for Mac (retrievable from `Highlights > Device Lock > View PIN`), a Lock Screen message and phone number, and a `macOS Lock Status` column in the Devices list.

– Exact navigation path and new list column named `Release Notes 2026 (2026-11-20)`

---

05 **Separate enforcement for major and minor macOS updates**

`Not checked yet` `IMPROVED` 60

JumpCloud's DDM-Based macOS Patch Policy now supports separate enforcement settings for major OS upgrades and minor OS updates, including independent grace periods and enforcement times for each.

– Mechanism described but no UI path or config key given `Release Notes 2026 (2026-11-20)`

THINNER COVERAGE BELOW

06 **Workflow support in MCP Server for Admins** Not checked yet NEW

55

The JumpCloud MCP Server for Admins adds Workflow support, letting admins create, modify, duplicate, troubleshoot, and manually trigger Workflows from a connected MCP-enabled AI client using natural language.

— Capability described but no concrete command or config shown

Release Notes 2026 (2026-11-20)

WAS THIS USEFUL? 👍 👎

AlertEnterprise Feedback ⓘ

Bring Guardian into the AI tools you already use (2026-11-20) BLOG ↗

AlertEnterprise is the future of work requires cyber-physical security convergence. Alert Enterprise provides identity & access SaaS solutions empowering businesses across the globe.

AlertEnterprise adds Guardian MCP, letting teams query and submit physical access and badge requests from AI assistants and developer tools under existing Guardian approvals.

NEW INTEGRATION New Guardian MCP integration connects Guardian to AI assistants and developer tools so access, badge, and identity requests can be made directly from those tools under existing Guardian approvals. Not checked yet

Setup instructions for each supported AI tool are available at `mcp.alertenterprise.io`, covering both cloud-based and on-premise Guardian deployments. Not checked yet

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Cross-App Access Requesting App now supports SAML in Early Access (2026-11-20)

BLOG ↗

Secure users, AI agents, and more with Auth0, a production-ready authentication and authorization platform built to help you ship faster and drive sign-ups.

Auth0 adds SAML support to Cross-App Access Requesting App, enabling silent third-party API token fetch for SAML-authenticated enterprise users via Token Vault and ID-JAG.

Enable the Requesting Application role on a SAML connection and configure the issuer or discovery URL and client ID via the Dashboard or Management API so apps can silently fetch third-party API tokens for SAML-authenticated users without requiring manual account linking. Not checked yet

NEW INTEGRATION Authenticate the Requesting App role with Private Key JWT or a client secret when configuring the SAML connection in Cross-App Access. Not checked yet

In multi-org tenants, connection resolution in Cross-App Access is now organization-aware, rejecting requests when the enterprise connection or target resource app is not enabled for the specified `org_id` or when org context is ambiguous. Not checked yet

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ Protect your MCP server with Clerk (2026-11-20)

BLOG ↗ RANK

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

> Clerk introduced tooling to secure Model Context Protocol (MCP) servers, covering spec-compliant authorization, pre-auth tool access, and Server Card publishing.

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

01 **MCP server authorization support** Not checked yet NEW 35

Clerk added the ability to build MCP servers with authorization that follows the current MCP authorization spec.

– Names the spec but no mechanism or API details Protect your MCP server with Clerk (202…

02 **Pre-auth tool exposure for MCP** Not checked yet NEW 30

Developers can open chosen MCP tools to users before sign-in.

– Describes capability but not configuration method

Protect your MCP server with Clerk (202…

03 **MCP Server Card publishing** Not checked yet NEW 26

Clerk added support for publishing Server Cards for MCP servers.

– Bare feature name with no further detail Protect your MCP server with Clerk (202…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

bytebase

Feedback ⓘ 3.25.0 NOTES ↗

Bytebase is a database governance built for humans and agents — controlling changes and access across every major database.

Bytebase adds AI SQL review (Enterprise), independent export permissions with a new SQL Editor Exporter role, MCP `administer_workspace` and `request_access_grant` tools, GitOps release visibility, and the EXECUTE statement type for procedure calls.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.25.0 https://github.com/bytebase/bytebase.git
# already have the repo? check out this version:
$ git checkout 3.25.0
```

AI SQL review comments on exact lines of a plan's SQL — flagging issues like an UPDATE without a WHERE — and blocks automatic rollout until all comment threads are resolved. Available on the Enterprise plan after configuring an AI provider. Not checked yet

NEW SUPPORTED PLATFORM Export is now its own permission (`bb.sql.export`), separate from query access. A new SQL Editor Exporter role and an Allow downloading exports setting (Settings > General > SQL Editor) control whether and where exports can go. Not checked yet

New MCP tools `administer_workspace` (Read-write mode, user-confirmed) and `request_access_grant` extend what agents and Project Developer users can do over MCP. Not checked yet

GitOps Releases page now shows per-database file status, source commit, branch, author and pull request; `bytebase-action` 3.25.0 signs in via workload identity and reports results in the job log. Not checked yet

Procedure calls and dynamic SQL are classified as EXECUTE with moderate default risk, enabling their own approval flow on every supporting engine. Not checked yet

NEW INTEGRATION SharePoint export destination integration

NEW INTEGRATION Neon database integration for PostgreSQL instance import

└──▷ BREAKING ON UPGRADE

! The release detail page is retired; links now open the associated plan. Releases can no longer be archived or restored in the console.

- ! The Releases page now requires `bb.plans.list` instead of `bb.releases.list`.
- ! Listing a project's databases now requires `bb.databases.list` instead of `bb.projects.get` in custom roles.
- ! The Allow data export without approval setting is removed; custom roles with only `bb.sql.select` lose export until `bb.sql.export` is added.
- ! New SQL Editor User assignments no longer include export; SQL Editor Exporter must be assigned separately.
- ! Upgrading a workspace that required export approval turns Allow downloading exports off automatically.
- ! Approval rules matching procedure calls as unclassified statements must be updated to include EXECUTE.
- ! MCP tool inputs and outputs changed; dependent scripts and agent configurations must be reviewed.

WAS THIS USEFUL? 👍 👎

Zama fhEVM

Feedback ⓘ v0.13.9 NOTES ↗ RANK

FHEVM, a full-stack framework for integrating Fully Homomorphic Encryption (FHE) with blockchain applications

> fhEVM's v0.13.9 patch backports key and parameter integrity checks to the SDK.

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

01 **KMS generation key and CRS integrity checks** Not checked yet

IMPROVED 38

The SDK now adds KMSGeneration key and CRS integrity checks, backported to the v0.13.x line.

— Names the check but no usage detail or mechanism given. v0.13.9

WAS THIS USEFUL? 👍 👎

MIND

Feedback ⓘ announced · Nov 19, 2026 PRESS ↗ RANK

MIND data security platform puts data loss prevention (DLP) and insider risk management (IRM) on autopilot - automatically identify, detect & prevent data leaks

> MIND launched an Email DLP module for Microsoft Exchange that classifies outbound messages and attachments and blocks inappropriate sends without installing an agent on Exchange.

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

01 **Email DLP module for Microsoft Exchange** Not checked yet NEW

55

MIND adds an Email DLP module for Microsoft Exchange that classifies outbound messages and attachments and blocks inappropriate sends, with no agent installed on Exchange. The module shares the same policies as every other environment MIND protects (GenAI, Agentic AI, SaaS, on-premise file shares, and endpoints), extending outbound coverage to mobile devices.

— Explains mechanism and scope but no config/API specifics. launch-20261119-c0822155

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ changelog · Nov 21, 2026 CHANGELOG ↗ RANK

RegScale is a continuous compliance management platform that automates security documentation, assessments, and regulatory compliance workflows for organizations.

> RegScale adds a `client_runs` CLI command group to manage and unstick open sync runs, raises the default vulnerability-to-asset link cap to 50,000 with per-plugin Tenable matching, and surfaces Microsoft Defender recommendations as vulnerabilities, alongside a batch of JCAM CLI fixes, a Wiz inventory reclassification, and cleanup of unused scanner config blocks.

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

01 Vulnerability-to-asset linking and Tenable plugin matching

Not checked yet BREAKING 90

Scanner syncs now link a vulnerability to up to 50,000 affected assets by default (up from 10,000), configurable via `maxAssetIdentifiersPerVulnerability` in `init.yaml`; Nessus file import and Tenable.sc sync now record one vulnerability per Tenable plugin mapped to every affected asset instead of one per CVE, matching Tenable's own finding definition. On the first import after upgrading, earlier per-CVE vulnerabilities and their paired issues close and new per-plugin vulnerabilities open in their place.

Cap the per-vulnerability asset link count and restore per-finding grouping for a Trivy import that must stay compatible with an existing pipeline.

```yaml
groupScanFindings: false
maxAssetIdentifiersPerVulnerability: 25000
```

— Named config key and numbers, with migration behaviour described.

changelog-20261121-8737756e

02 client_runs command group for stuck syncs

Not checked yet NEW 85

New `regscale client_runs` command group adds `list`, `status`, `find`, and `close` subcommands to inspect and unstick open sync runs directly from the CLI, so a stuck run can be found and closed before starting a new one.

Inspect and unstick a sync run that another process left open before starting a new one.

```
$ regscale client_runs list
regscale client_runs find --plan <plan-id> --source <source-name>
regscale client_runs close <run-id>
```

— Named subcommands plus a runnable example sequence. changelog-20261121-8737756e

---

03 **JCAM CLI fixes across connection, mapping and export**

Not checked yet

85

BREAKING

`regscale jcam test_connection` now exits 1 with a configuration message when `jcamToken` or `jcamURL` is unset or a placeholder instead of silently reporting success; `import_parameters --learn-from` now requires `--crosswalk` so a learned mapping is never written into the installed package where an upgrade would discard it; `regscale jcam check_config` now reports an unset `jcamFrameworkCatalog` catalogue ID for 800-53r5, which `import_controls` requires; `regscale jcam export_ssp` now sends each authorization and front-matter value as plain text from the mapped field, so drift detection no longer flags identical values and `--force` no longer fails with a serialization error; and `export_poams` now sends each POA&M's stored severity as the matching JCAM criticality instead of No Selection.

— Every JCAM subcommand and flag named with its exact behaviour change.

changelog-20261121-8737756e

---

04 **Wiz inventory Hardware-to-Software reclassification**

Not checked yet

80

BREAKING

Wiz inventory sync now moves images stored as Hardware to Software and keeps their stored OS version; list `VIRTUAL_MACHINE_IMAGE` in `wizHardwareAssetTypes` with `useWizHardwareAssetTypes` on to retain them as Hardware. This takes effect on the next sync run after upgrading.

— Named config keys give a clear opt-out path. changelog-20261121-8737756e

---

05 **Qualys VM import placeholder host ID keying** Not checked yet

70

BREAKING

Since RegScale 7.1.0, the Qualys VM scan import keys a host whose QG Host ID is a placeholder (`none` or `unknown`) on its IP address; a host imported earlier under a placeholder ID gets a new asset keyed on its IP, and its old asset stops receiving findings.

— Explains mechanism and consequence but no remediation step. changelog-20261121-8737756e

THINNER COVERAGE BELOW

06 **Defender recommendations as vulnerabilities** Not checked yet NEW

55

Microsoft Defender and Defender for Cloud recommendations now also appear as vulnerabilities on the Vulnerability Status Board and in the FedRAMP VDR report, linked to their existing issues; requires RegScale platform 6.36.2.0 or later.

— Clear scope and version requirement but no action step. changelog-20261121-8737756e

07 **Removed no-op scanner config blocks** Not checked yet DEPRECATED

50

Removes `no longer carries` blocks for `aqua`, `defenderFile`, `ecr`, `nexpose`, `snyk`, `sonarcloud`, `veracode`, and `xray` from configuration, as the CLI never read them; existing blocks are now silently ignored.

— Names all affected keys but requires no reader action. changelog-20261121-8737756e

08 **FedRAMP DRF import CVE handling fix** Not checked yet IMPROVED

50

`regscale fedramp import-drf` no longer sets a deviation's CVE from the Vulnerability Name column when the workbook has no CVE column.

— Named command and column but minimal context. changelog-20261121-8737756e

09 **Synqly filtered asset sync status fix** Not checked yet IMPROVED

50

Synqly asset syncs run with `--filter` no longer mark assets excluded by the filter as Inactive.

— Named flag but no further mechanism given. changelog-20261121-8737756e

10 **Tenable.io MAC address formatting** Not checked yet IMPROVED

45

`regscale tenable io sync_assets` now sends each MAC address in upper case split by colons.

— Named command, but a single-line formatting tweak. changelog-20261121-8737756e

└──▷ **BREAKING ON UPGRADE**

! Nessus and Tenable.sc now record one vulnerability per Tenable plugin instead of one per CVE; on the first import after upgrading, all earlier per-CVE vulnerabilities and their paired issues close, and new per-plugin vulnerabilities open in their place.

! JCAM `import_parameters --learn-from` now requires `--crosswalk`; calls without `--crosswalk` will fail.

! Wiz inventory sync will move images currently stored as Hardware to Software on the next run; set `useWizHardwareAssetTypes` and list `VIRTUAL_MACHINE_IMAGE` in `wizHardwareAssetTypes` to retain them as Hardware.

! Config blocks for `aqua`, `defenderFile`, `ecr`, `nexpose`, `snyk`, `sonarcloud`, `veracode`, and `xray` are no longer read by the CLI; existing blocks are silently ignored.

! Since 7.1.0, the Qualys VM scan import keys a host whose QG Host ID is a placeholder (`none` or `unknown`) on its IP address, so a host imported earlier under a placeholder ID gets a new asset keyed on its IP and its old asset stops receiving findings.

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v4.3.4 NOTES ↗ RANK

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 200+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

> CISO Assistant v4.3.4 adds a GLPI integration, an experimental pre-audit preparation module, new findings-to-assessment linking, TPRM task assignment, several new compliance frameworks and mappings, and warns that `curl` will be removed from the next release's Docker image.

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

01 New compliance frameworks and cross-framework mappings

Not checked yet NEW 65

Adds the SAMA IT Governance Framework, the NDMO Data Management and Personal Data Protection Standards (SDAIA), the `CADA` and `SEAL` frameworks for the European market, and Essential Eight reference controls, threats, and mappings to the library. Also adds `CMMC` to NIST 800-171 rev2, NIST 800-171 rev2 to `rev3`, and `CMMC` to `CSF 2.0` mappings to the library.

— Names every framework and mapping added, clear library location v4.3.4

THINNER COVERAGE BELOW

02 Opt-in auto-lock for overdue audits

Not checked yet IMPROVED 50

Auto-lock of overdue audits is now controlled by a new setting and is off by default, whereas it was previously automatic.

— Explains before/after behaviour but not exact setting key v4.3.4

03 curl removed from Docker health check

Not checked yet BREAKING 50

`curl` will be removed in the next version — users should update their Docker Compose health check before upgrading.

— Gives a clear pre-upgrade action but limited mechanism detail v4.3.4

04 GLPI integration for tickets and assets

Not checked yet NEW 40

Adds a GLPI provider integration for ticket and asset management.

— Names the integration but no setup or endpoint detail v4.3.4

05 **Task assignment to third-party audit representatives**

Not checked yet NEW 40

In the TPRM module, tasks raised on a third-party audit can now be assigned to its representatives.

— Names module and action but no UI path v4.3.4

06 **Audit reports show implementation, documentation, maturity scores**

Not checked yet IMPROVED 40

Audit reports now display implementation, documentation, and maturity scores.

— Names the three new score types shown v4.3.4

07 **Attach findings to requirement assessments** Not checked yet NEW 35

Findings can now be attached directly to a requirement assessment.

— States the capability without workflow detail v4.3.4

08 **German and Hungarian localization additions** Not checked yet NEW 35

Adds German translations for CRA, metrics, workflow, and threats to the library, and adds Hungarian email templates and policy templates.

— Names languages and content areas but no usage detail v4.3.4

09 **Experimental pre-audit preparation module** Not checked yet NEW 25

Adds an experimental pre-audit preparation module.

— Bare name, no mechanism or scope given v4.3.4

└──▷ **BREAKING ON UPGRADE**

! `curl` will be removed in the next version — update your Docker Compose health check before upgrading.

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

raviaxo RiskShard

Feedback ⓘ v0.15.0 NOTES ↗ RANK

RiskShard is an evidence-governed quantitative cyber risk — a trustworthy CLI and scenario engine where every number traces to a reviewed public source.

> RiskShard added evidence re-verification tracking and expanded its measurement basis vocabulary with a geometric-mean cost option.

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

01 **Evidence re-verification fields for artifacts** Not checked yet

NEW 70

New `artifact_verified_at` and `artifact_verified_sha256` fields track when evidence was re-read against its source file and detect publisher reissues automatically; the two fields must appear together or not at all.

— Names exact fields and their pairing rule. v0.15.0

02 **Geometric-mean cost measurement basis** Not checked yet NEW

60

A new `geometric_mean_total_event_cost` value models geometric-mean cost figures from loss reports, added as the 19th entry in the `measurement_basis` vocabulary.

— Names the exact value and its position in vocabulary. v0.15.0

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 19th, 2026 NOTES ↗

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

Compyl adds custom Questionnaire fields queryable in Query Builder, bulk blueprint scheduling, and a Trust Center Admin path change to `/trust-center-admin`.

Add custom fields to Questionnaires: admins can create an extension form whose fields appear on each questionnaire, save per-response data, and are queryable in Query Builder. Not checked yet

Set a bulk run schedule for blueprints to move many blueprints to a new cadence (for example, monthly to daily) in one step. Not checked yet

List fields such as the VirusTotal watch list now include search, a clear-all action, and an action to delete all entries that failed validation. Not checked yet

The Trust Center admin module now lives at `/trust-center-admin` and is renamed Trust Center Admin in navigation, freeing `/trust-center` for the public Trust Center page. Not checked yet

NEW INTEGRATION Google Workspace integration

UPDATED INTEGRATION Azure integration — SQL Server firewall rules, Sentinel alert rules, threat intelligence indicators, storage account settings

UPDATED INTEGRATION Duo managed endpoint sync

└──▷ BREAKING ON UPGRADE

! The Trust Center admin module has moved from its previous path to `/trust-center-admin`; any bookmarks or hardcoded links to the old admin URL must be updated.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.5.0 (2026-11-21) BLOG ↗ RANK

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

> Runlayer added PGP private key detection and audit-payload credential masking, Cursor blocking notices, and faster PII/token masking at scale, alongside a new access-request workflow, reply-to-quote chat, and a unified Procedures setting that auto-benchmarks and auto-promotes self-improving agents' procedure drafts.

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

01 **Access request workflow for artifacts and agents** Not checked yet

NEW 75

Users can request access to any artifact, agent, skill, or plugin they lack permission to open by clicking Request access ; the owner receives the request in Approvals and in Slack, and approving it opens the page for the requester.

Grant a teammate access to an agent they cannot open, without leaving the product, by using the built-in request flow.

> 📍 In the Runlayer web app, open the agent, artifact, skill, or plugin page you cannot access, then click Request access . The owner sees the request under Approvals (and in Slack if connected) and can share it directly from there.

— Includes exact UI steps and notification channels 6.4.0 (2026-11-20)

02 **Procedures setting for self-improving agents** Not checked yet

NEW 75

A single Procedures setting lets self-improving agents save their own procedures, automatically triggers benchmarking of drafts saved without named runs, and promotes a passing draft to live without waiting for approval. Stale procedure draft versions with no review decision after 30 days are automatically sent back with a reason and their review requests closed, while any live version keeps running; failed procedure benchmarks now retry daily automatically instead of waiting for a new version, within each version's test-run limit.

— Full lifecycle described but no config key named 6.4.0 (2026-11-20)

03 **PGP private key detection in tool-call arguments** Not checked yet

NEW 70

Runlayer now detects complete armored PGP private keys (`-----BEGIN PGP PRIVATE KEY BLOCK-----` , `PGP SECRET KEY BLOCK`) and private keys with dash-containing headers (`GPGTools Comment:` line, encrypted PEM `DEK-Info:`) in tool-call arguments, so the Mask, Block and Alert actions now apply to them.

— Names exact detected patterns but no config surface given 6.5.0 (2026-11-21)

---

04 **MCP output blocking notices in Cursor** Not checked yet IMPROVED 70

On Cursor 2.6.22 or later, Runlayer now shows a notice to the agent when an MCP tool's output is blocked in Protect mode (replacing the previous silent omission) and also shows a notice when Runlayer cannot be reached to scan an MCP tool's output. Sessions and the supported clients list now report Block bypassed for built-in tool output and Blocked for replaced MCP output.

— Clear before/after and version gate, no config key 6.5.0 (2026-11-21)

---

05 **Reply-to-quote in agent chat** Not checked yet NEW 70

Highlighting words in an agent's answer and clicking Reply quotes them into your message box; available in agent chat, the Runlayer Assistant, and artifact chat.

Quote a specific sentence from an agent response to ask a follow-up question without retyping the context.

> 📍 In agent chat, highlight the relevant words in the agent's answer, then click Reply — the quoted text appears in your message box and is sent with your next message.

— Exact UI interaction given, easy to try 6.4.0 (2026-11-20)

---

06 **Token masking improvements** Not checked yet IMPROVED 60

Hidden-character, PII and token-masking scanner time now scales proportionally to tool-call or tool-result size instead of slowing down sharply on dense inputs, so a multi-megabyte response with many personal-data matches that previously took minutes to mask now takes about a second; token masking also now fully hides a private key written right after `password=` with no `END` marker, where previously only the first few characters were masked.

— Quantifies speedup but names no config to tune it Specific before/after but no config surface

6.5.0 (2026-11-21) 6.4.0 (2026-11-20)

---

07 **Skill risk policy blocks extended to @ mentions and distillation**

Not checked yet IMPROVED 60

Skill risk policy blocks now also apply when skills are reached via `@` mentions in agent chat or through skill distillation in self-improving agents, matching the behaviour already in place for agent runs.

— Names exact mechanism but no config surface 6.4.0 (2026-11-20)

THINNER COVERAGE BELOW

08 **Session privacy in Analytics detail view** Not checked yet IMPROVED

55

When Session privacy is on, the timeline in an Analytics output's detail view now hides other people's sessions, showing a count of hidden sessions instead of their title, owner, cost or session link.

— Describes behaviour but no setting path shown 6.5.0 (2026-11-21)

09 **Credential masking in audit payloads** Not checked yet NEW

55

Recognised credentials are now masked in audit payloads before they reach the audit database, payload archives and downstream exports; security approvals retain matching without copying detected credential values into approval evidence.

— Explains pipeline stages but no API or flag named 6.5.0 (2026-11-21)

WAS THIS USEFUL? 👍 👎

mcp-security

Feedback ⓘ secops-soar-v0.2.0 NOTES ↗ RANK

> mcp-security shipped new case-management tools for its SOAR server, added proxy support across both the SOAR and GTI (VirusTotal) clients, and introduced a breaking configuration gate on GTI's file-analysis tool.

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

01 GTI_ALLOWED_UPLOAD_DIR required for file analysis

Not checked yet

95

BREAKING

The `analyse_file` tool now requires the `GTI_ALLOWED_UPLOAD_DIR` environment variable to be configured before it is enabled, restricting which directories may be uploaded; existing deployments relying on `analyse_file` without this setting will have file analysis disabled until the variable is set, e.g. `GTI_ALLOWED_UPLOAD_DIR=/tmp/gti-uploads`.

Allow the `analyse_file` tool to upload files from a specific directory by setting `GTI_ALLOWED_UPLOAD_DIR` in your environment before starting the server.

env

```
GTI_ALLOWED_UPLOAD_DIR=/tmp/gti-uploads
```

— Exact env var, breaking impact, and example command given. gti-v0.1.4

02 Proxy support via trust_env in SOAR and GTI clients

Not checked yet NEW

65

Corporate proxy support was added by setting `trust_env` in the `HttpClient` session used by the SOAR MCP server, and by setting `trust_env` during `vt.Client` initialization in the GTI VirusTotal client, so HTTP proxy environment variables are honoured automatically in both.

— Mechanism and client names given, but no exact env var shown. gti-v0.1.4

secops-soar-v0.2.0

THINNER COVERAGE BELOW

03 Case management tools in SOAR MCP server

Not checked yet NEW

55

The SOAR MCP server adds `create_case` and `close_case` tools so cases can be created and closed programmatically.

— Names new tools but gives no usage example. secops-soar-v0.2.0

04 **Identifier validation and path normalization in HttpClient**

Not checked yet IMPROVED 40

`HttpClient` now validates case and alert identifier formats and normalizes request paths automatically, reducing malformed requests to the SOAR API.

— Internal behavior change with no exposed config or command. secops-soar-v0.2.0

└──▷ **BREAKING ON UPGRADE**

! The `analyse_file` tool now requires `GTI_ALLOWED_UPLOAD_DIR` to be configured; existing deployments that relied on `analyse_file` without this setting will have file analysis disabled until the variable is set.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

ai-safe2-framework

Feedback ⓘ v3.5 NOTES ↗ RANK

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

> AI SAFE² v3.5 adds three CLI routes for querying, auditing, and assessing readiness against framework controls, an offline Rosetta Stone explorer consolidating over a thousand source-to-control relationships, and twelve new governance controls alongside expanded MCP, deployment, UAS, and threat profiles.

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

01 **New CLI routes for framework query, audit, readiness**

Not checked yet NEW 88

Adds three new CLI routes: `safe2 framework query --control CP.12` to query a specific control, `safe2 framework audit` to audit against framework controls, and `safe2 framework readiness --help` to assess readiness.

— Exact runnable commands with flags given. v3.5

02 **Offline Rosetta Stone explorer and consolidated source data**

Not checked yet NEW 85

Open the offline Rosetta Stone explorer by loading `dashboard/rosetta.html` from a local checkout to filter records by source, control, review status, and keyword, and export them as JSON. It now consolidates 1,202 recovered source-to-control relationships across 33 external source families, 64 discovery entries, and 16 audit-rule classes.

— Exact file path and filter/export mechanics given. v3.5

03 **12 new controls across trust and evidence domains** Not checked yet

NEW 75

Adds 12 new controls — `CP.11` through `CP.15`, `S1.8`, `F3.6`, `F3.7`, `A2.7`, `A2.8`, `A2.9`, and `E5.5` — covering trust lifecycle governance, final-effect authorization, control evidence independence, organizational AI management governance, design-time coordination contracts, pre-trust initialization boundaries, aggregate consumable-authority ceilings, environment binding, idempotency and side-effect ledgers, verified outcome contracts, evidence-based failure attribution, and declared success criteria gates.

— Each control ID named with its specific governance scope. v3.5

THINNER COVERAGE BELOW

04 **UAS v2 revises acceptance interpretation** Not checked yet IMPROVED

55

Adds UAS v2, retaining 27 IDs while revising acceptance interpretation for human understanding, appropriate reliance, intervention, fairness, and evidence quality, with reciprocal Cognitive Sovereignty Framework references.

— Scope and ID count given, no lookup path. v3.5

05 **157 child clauses and research notes deepen existing controls**

Not checked yet IMPROVED

50

Adds 157 explicit child clauses across 53 reviewed decisions, including two research-disposition clauses, plus 53 new research notes connecting source material, placement decisions, alternatives, trade-offs, and acceptance cases to framework controls — deepening requirement precision across existing controls.

— Counts given but no named controls or access point. v3.5

06 **Six deployment profiles for execution surfaces** Not checked yet

NEW

50

Adds six deployment profiles covering harnesses, gateways, evaluation, browser and computer use, finance, and physical systems to make applicability explicit across execution surfaces.

— Lists surfaces but no profile IDs or access path. v3.5

07 **Three new threat profiles sharpen existing threat families**

Not checked yet NEW

40

Adds three threat profiles sharpening attention to assurance compromise, unsafe authorized execution, and human-agent trust exploitation within the existing eleven threat families.

— Named threat areas but no control mapping shown. v3.5

08 **MCP profile expanded with 12 child requirements** Not checked yet

IMPROVED

35

Expands the MCP profile with 12 additional child requirements across the existing 19 stable control IDs.

— Numbers given but no specific requirement names. v3.5

09 **Machine-readable page map and agent discovery indexes**

Not checked yet NEW

35

Adds a machine-readable complete page map, root and NEXUS agent indexes, and shared module navigation to support both human and agent repository discovery.

— Describes purpose but no file paths or formats named. v3.5

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal

Feedback ⓘ

release-14b57239 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 release-14b57239 deepens its `min box` CLI with audit, inspection and egress controls, and lays groundwork for secrets management, hardware-backed key storage and OpenTelemetry tracing, though most of the latter are named without further detail.

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

01 Box inspection and introspection commands

Not checked yet NEW

80

`min box show <box> --network` shows a box's effective network policy; `min box show -o json` and `min box show -o yaml` output box details including its exec list; `min box` list/show subcommands list and show box types; `min box spec` lists each store reference with its authorities and each lifecycle hook with its declarer; and box names are shell-completable for every `min box` verb that accepts one.

Inspect the effective network policy of a running box to confirm egress rules before an agent run.

```
$ min box show mybox --network
```

— Multiple named subcommands and flags enumerated with a usage example. release-14b57239

02 Box audit trail with parent history merge

Not checked yet NEW

75

Read a box's audit trail with `min box audit`, and merge parent history into the trail with `min box audit --parent` to verify no decisions were edited.

Review the full audit trail for a box, merging parent history, to verify no decisions were edited.

```
$ min box audit mybox --parent
```

— Named flag with a runnable verification command. release-14b57239

THINNER COVERAGE BELOW

03 **Box Egress Proxy and egress flags on run commands** Not checked yet

NEW

50

minimal adds a Box Egress Proxy for controlling network egress from boxes, alongside new network and egress flags on `min task run` and `min box run`.

– Named flags and a proxy feature but no mechanism or example given. release-14b57239

04 **Vulnerability reporting via min diag --security** Not checked yet

NEW

50

`min diag --security` reports a suspected vulnerability.

– Named flag with clear purpose but no detail on vulnerability scope. release-14b57239

05 **Tampered audit chain detection in min doctor** Not checked yet

NEW

40

`min doctor` now reports an edited audit chain, flagging boxes whose audit history has been altered.

– Names the command and detection purpose, no example shown. release-14b57239

06 **Software key names shown in min auth status** Not checked yet

IMPROVED

40

`min auth status` now shows software key names for each key.

– Named command enhancement without a usage example. release-14b57239

07 **Secrets management via min secret subcommands** Not checked yet

NEW

30

minimal introduces secrets management with new `min secret` subcommands.

– Subcommand prefix named but no specific options described. release-14b57239

08 **TPM/Secure Enclave key storage** Not checked yet NEW

20

minimal adds support for storing keys in a TPM or Secure Enclave.

– Only named in the release summary, no usage detail. release-14b57239

09 **OpenTelemetry tracing support** Not checked yet NEW

20

minimal adds OpenTelemetry tracing.

– Named only, no configuration or scope detail provided. release-14b57239

WAS THIS USEFUL? 👍 👎

ToolHive

Feedback ⓘ v0.67.0 NOTES ↗ RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.67.0 adds a client_credentials grant for SPIFFE-authenticated workloads with per-association grant-type selection, structured severity findings for skill and plugin validation, and a configurable DCR upstream client name — alongside several breaking changes to SPIFFE token exchange, default transport selection, and MCPRemoteProxy resource limits.

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

01 SPIFFE token exchange breaking changes

Not checked yet BREAKING

90

SPIFFE associations (`spiffeClientAuth`) now require a non-empty `resources` list; associations without `resources` are rejected and the auth server refuses to start, and the `audiences` field is now ignored. Token-exchange requests must send exactly one `resource` parameter — requests sending `audience=` or omitting a target now receive `invalid_target` . The `act.sub` claim in a SPIFFE token exchange is now the workload's SPIFFE ID instead of the client ID, so Cedar policies matching `claim_act.sub` against a SPIFFE client ID must be updated to match `claim_client_id` . The Go API's token-exchange HTTP middleware and runner factory must now be imported from `pkg/auth/tokenexchange` , and `authserver.SPIFFEGrantTypeTokenExchange` must be replaced with `oauthproto.GrantTypeTokenExchange` .

— Enumerates every migration step needed v0.67.0

02 SPIFFE client_credentials grant and grant-type selection

Not checked yet NEW

85

`MCPExternalAuthConfig` and `VirtualMCPServer` SPIFFE associations now have a `grantTypes` field to select `client_credentials` , token exchange (`urn:ietf:params:oauth:grant-type:token-exchange`), or both, defaulting to token exchange. SPIFFE-authenticated workloads can now obtain access tokens via the `client_credentials` grant bound to their SPIFFE ID, with token lifetime capped at the SVID's remaining validity.

Allow a SPIFFE workload to authenticate with both client_credentials and token exchange on a VirtualMCPServer association.

```yaml
grantTypes:
  - client_credentials
  - urn:ietf:params:oauth:grant-type:token-exchange
```

— Names config field, grant types, and validity rule v0.67.0

03 **Default resource limits for MCPRemoteProxy** Not checked yet

75

BREAKING

`MCPRemoteProxy` now applies default CPU (`50m` request / `200m` limit) and memory (`64Mi` request / `256Mi` limit) resource constraints; proxies whose working set exceeds 256Mi will be OOMKilled on upgrade unless explicit limits are set beforehand.

— States exact default limits and upgrade risk v0.67.0

04 **Default transport change for uvx packages** Not checked yet

70

BREAKING

`thv run uvx://…` now defaults to stdio transport instead of auto-detecting; packages that serve HTTP or SSE must now explicitly pass `--transport streamable-http` or `--transport sse`.

— Gives exact flags needed to migrate v0.67.0

05 **DCR upstream client name configuration** Not checked yet NEW

65

Operators can set `clientName` in the DCR upstream configuration to control the name shown on the upstream IdP's authorization screen; it defaults to ToolHive MCP Client.

Override the client name shown on the upstream IdP consent screen when using DCR upstream configuration.

```yaml
clientName: "My Team MCP Client"
```

— Names config key and default, has example v0.67.0

06 **Structured findings in skill and plugin validation**

Not checked yet NEW

65

`thv skill validate`, `thv ai-plugin validate` and their corresponding API endpoints now return a structured `findings` array with severity levels instead of plain text output.

— Names commands and response field but no example `v0.67.0`

└──▷ BREAKING ON UPGRADE

- ! SPIFFE associations (`spiffeClientAuth`) now require a non-empty `resources` list; associations without `resources` are rejected and the auth server refuses to start. The `audiences` field is now ignored.
- ! SPIFFE token-exchange requests must send exactly one `resource` parameter; requests that send `audience=` or omit a target now receive `invalid_target`.
- ! `thv run uvx://…` now defaults to stdio transport; packages that serve HTTP or SSE must now explicitly pass `--transport streamable-http` or `--transport sse`.
- ! The `act.sub` claim in a SPIFFE token exchange is now the workload's SPIFFE ID instead of the client ID; Cedar policies matching `claim_act.sub` against a SPIFFE client ID must be updated to match `claim_client_id`.
- ! `MCPRemoteProxy` now applies default CPU (`50m` request / `200m` limit) and memory (`64Mi` request / `256Mi` limit) resource constraints; proxies whose working set exceeds 256Mi will be OOMKilled on upgrade unless explicit limits are set beforehand.
- ! Go API: the token-exchange HTTP middleware and runner factory must now be imported from `pkg/auth/tokenexchange`; `authserver.SPIFFEGrantTypeTokenExchange` must be replaced with `oauthproto.GrantTypeTokenExchange`.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 21, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue