THE CYBER TOOLCHAIN

NO. 137

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 12, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 144 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 40 of 41 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.

PROTECT

vim ? help

Find the agent config files your developers already committed — and block the next one

GitLab pairs search APIs and scheduled historic secret scanning to surface `.claude/settings.local.json` files already in history with prohibited-filename push rules and group-wide `.gitignore` templates so they stop arriving; Duo Agentic Chat drives the cleanup. Semgrep's observe-only mode lets Guardian run against a repo set and show what it would have blocked before anyone's push starts failing — the usual reason rollouts stall.

GitLab · Semgrep

GOVERN

Decide which agents and MCP servers are allowed to act on a user's behalf

Varonis's MCP Gateway gives one place to allow or deny tools, APIs, agents and MCP servers acting under a user's identity — today most teams have no inventory of what their developers wired up, let alone a switch. WitnessAI's guardrail REST endpoints wrap model and agent calls with runtime policy checks and add a discovery catalog of AI apps in use, with an agentless network-level option for estates where you cannot instrument every caller.

Varonis · WitnessAI

DETECT

Query Sentinel's cheap lake tier interactively instead of waiting on a job

Interactive KQL over lake tables means hunting across long-retention data no longer requires exporting it or standing up a separate analytics path, and Fabric federation pulls external datasets into the same query. Hunts that previously stopped at the hot-tier retention boundary can now run against the full window.

Microsoft Sentinel

PROTECT

Give agent processes a network exit you can audit and revoke

minimal's Box Egress Proxy issues certificates and keeps a hash-chained audit trail of what an agent box reached, with per-run network and egress flags — so a coding agent's outbound traffic becomes a reviewable record rather than an unknown. Docker Sandboxes adds Docker Hardened Image

> selection and tunable stalled-connection timeouts so a dropped VPN recovers instead of hanging the run.
>
> minimal

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

reverse_ssh

Feedback ⓘ v2.9.0 NOTES ↗ RANK

SSH based reverse shell

> reverse_ssh v2.9.0 adds TOML-based server configuration and improves proxy compatibility by handling Negotiate-only authentication challenges.

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

01 TOML configuration file support NEW

75

The RSSH server can now be configured via a `TOML` file by passing `--config <path>` instead of relying solely on CLI flags, enabling repeatable, version-controlled server deployments.

Start the RSSH server from a TOML config file instead of passing all options on the command line — useful for repeatable, version-controlled server deployments.

```
$ ./server --config /etc/rssh/server.toml
```

— Names exact flag and file format with runnable example v2.9.0

THINNER COVERAGE BELOW

02 Negotiate-only proxy authentication handling IMPROVED

40

reverse_ssh now handles proxies that advertise `Proxy-Authenticate: Negotiate` without also advertising NTLM.

— Names exact header but no usage steps given v2.9.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

julius

Feedback ⓘ v1.7.0 NOTES ↗

└──▷ GET THIS VERSION ▾

Go

```
$ go install github.com/praetorian-inc/julius/cmd/julius@v1.7.0
```

git

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

Go module path matches the repo · checked 12 Nov Wrong command?

NEW DETECTION Hermes Agent API server detection

WAS THIS USEFUL? 👍 👎

sippts

Feedback ⓘ v4.2.2 NOTES ↗ RANK

Sippts is a set of tools to audit SIP based VoIP Systems

> Sippts 4.2.2 brings full SIP-over-WebSocket capability to its core commands and removes the standalone wssend module in favor of integrated WebSocket sending.

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

01 SIP-over-WebSocket support across scan, enumerate and send

95

BREAKING

Sippts extends SIP-over-WebSocket support across the `scan`, `enumerate`, and `send` commands: `scan` gains `-p ws`, `-p wss`, and an expanded `-p all` that now includes ports 5061/TLS, 8088/WS, and 8089/WSS, plus support for comma-separated transport lists like `-p udp,tls,wss`; a new `-path` flag enables automatic WebSocket path discovery in `scan`, `enumerate`, and `send`; `enumerate` adds `-p ws|wss` and `-path` for method enumeration over WebSocket; and the `wssend` module has been removed, with its functionality replaced by `send -p ws` and `send -p wss`. All WebSocket traffic is handled by a new shared client, `lib/wsclient.py`, implementing RFC 7118.

Enumerate supported SIP methods on a WebSocket endpoint when a standard UDP/TCP scan yields no results.

```
$ sippts enumerate -i mysipserver.com -p wss -path
```

Send a custom SIP OPTIONS message over a WebSocket transport, replacing a previously separate `wssend` command.

```
$ sippts send -i mysipserver.com -p ws -path -m options
```

— Names every flag, port, module and the RFC it implements. v4.2.2

THINNER COVERAGE BELOW

02 **Scan results surface HTTP and SIP-reflecting ports** IMPROVED

30

Ports answering HTTP or reflecting SIP requests are now shown as rows in the `scan` results table instead of being discarded.

— States the change but gives no flag or command detail. v4.2.2

└──▷ **BREAKING ON UPGRADE**

! The `wssend` module has been removed; replace `sippts wssend` invocations with `sippts send -p ws` or `sippts send -p wss`.

WAS THIS USEFUL?  

◆ Hacking Hardware

DarkFlippers Unleashed Firmware

Feedback ⓘ

unlshd-095 NOTES ↗ RANK

Flipper Zero Unleashed Firmware, an unlocked alternative firmware.

> Unleashed Firmware's unlshd-095 release adds a Remote Analyzer app for detecting undecoded garage/gate remotes, a T5577 app with password-sniffing, new SubGHz protocol decoding for several remote families, and expands Nice O-Code and Seed Capturer with RAW-file decoding, alongside a breaking renumbering of KeeLoq CounterMode values.

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

01 **New SubGHz remote protocols and custom buttons** NEW 75

Adds decoding and replay support for KeeLoq Phantom/Hiland (now reads and sends properly, available as `KL: Phantom/Hiland 433M` in Add Manually, with `CounterMode: 10` sending the resync pair `0xFFFE` then `0xFFFF` in one press as a clone remote), FAAC TM 12 dip-switch TM 433/868 DS remotes (available as FAAC TM 433MHz and FAAC TM 868MHz in Add Manually), Nero Radio rolling remotes, Freezer remotes, and Proteco MOVER remotes with custom buttons; Dooya remotes also gain custom button support.

— Enumerates Add Manually entries and counter mode values across protocols unlshd-095

02 **CLI infrared signal transmission** NEW 70

Send infrared signals over the CLI with `ir universal <remote> <signal>` without rebooting the device.

Send a specific infrared signal from a universal remote database over the CLI without triggering a device reboot.

```
$ ir universal <remote> <signal>
```

— Exact runnable command with example, minimal mechanism detail unlshd-095

03 **Remote Analyzer app** NEW 70

Adds the Remote Analyzer app, which detects garage and gate remotes that SubGHz does not yet decode and shows their type and details, detects Life Fido remotes and Crypt / Crypt Lite encrypted remotes, and includes `Mod. sweep`, which walks

modulations one second at a time while a remote transmits and locks onto the one that works.

— Explains detection mechanism and sub-features, no exact command unlshd-095

04 **T5577 app with password sniffing** NEW 70

Adds a T5577 app that reads a T5577 tag, saves it to `/ext/lfrfid`, and writes it back, with Sniff Password to recover a tag password by listening to a reader or cloner write a tag and keep found passwords in a personal dictionary.

— Names file path and feature, lacks exact command invocation unlshd-095

05 **T5577 developer API exports** NEW 70

Exports the T5577 block read/write API for apps via `t5577.h` and `t5577_file.h`, plus `faac_slh_reset_prog_mode`, `subghz_protocol_decoder_base_get_type`, `ibutton_protocols_is_writable`, and `bit_lib_reverse_32_fast`.

— Names exact headers and functions but no usage example unlshd-095

06 **Nice O-Code decoding and cloning upgrades** IMPROVED 70

Nice O-Code app can now decode a saved RAW `.sub` file and shows the certificate beside the installer code, supports the longer 72-bit frame as Nice O-Code One 433MHz and 112-bit certificate frames as Nice O-Code+Crt 433MHz in Add Manually, and its New remote option mints a serial the gate is certain to enrol.

— Names frame types and Add Manually entries, no exact command unlshd-095

THINNER COVERAGE BELOW

07 **Seed Capturer RAW decoding and BFT Mitto support** NEW 55

Seed Capturer can decode a saved RAW `.sub` file and supports BFT Mitto remotes on 433.92 MHz using an external CC1101 module.

— Names frequency and hardware requirement, limited further detail unlshd-095

08 **KeeLoq CounterMode renumbering** BREAKING 55

KeeLoq `CounterMode: 7` now sends what `CounterMode: 8` previously sent; the old `CounterMode: 7` resync pack is now mode 8, and the previous mode 8 is now mode 9.

— States exact mode remapping but no migration steps unlshd-095

09 **Half-rate KeeLoq decoding** IMPROVED 45

KeeLoq remotes that send at half the usual rate (Te 200us) are now decoded and sent back at the rate they were received, shown as Crypt Key instead of an unknown

manufacturer.

— Describes behavior change but no interaction steps `unlshd-095`

└──▷ **BREAKING ON UPGRADE**

! KeeLoq `CounterMode: 7` now sends what `CounterMode: 8` previously sent; the old `CounterMode: 7` resync pack is now mode 8, and the previous mode 8 is now mode 9.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

Semgrep

Feedback ⓘ updated · Nov 12, 2026 NOTES ↗ RANK

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

> Semgrep shipped a hosted MCP server for editor and terminal access to findings, an observe-only mode for Guardian rollouts, and support for bringing your own AI provider keys to Agentic Workflows, alongside Supply Chain, language-parsing, and advisory-matching improvements — while dropping Intel Mac binaries after 1.182.

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

01 **semgrep ci exits with code 2 on dependency resolution failure**

IMPROVED 90

With `--allow-local-builds`, running `semgrep ci --no-suppress-errors` now exits with code `2` when dependency resolution fails for a subproject, listing the affected dependency sources in scan output.

— Exact command, flag, and exit code given snapshot-20261112

02 **Hosted MCP server and workspace scanning** NEW 88

Semgrep adds a hosted MCP server in public beta, letting Claude Code, Codex, Cursor, VS Code, or any MCP client search and triage findings, run scans, and manage policies from the editor or terminal with permissions following your AppSec Platform role. The `semgrep mcp` server can now scan files outside the MCP client's workspace root by setting `SEMGREP_MCP_ALLOW_SCAN_OUTSIDE_WORKSPACE=true`.

Let the Semgrep MCP server scan files that live outside the MCP client's workspace root, such as shared library directories.

```
$ export SEMGREP_MCP_ALLOW_SCAN_OUTSIDE_WORKSPACE=true && semgrep mcp
```

— Named server, clients, exact env var and command snapshot-20261112

03 **Guardian observe-only mode** NEW 85

Semgrep Guardian adds an observe-only mode, enabled by setting `SEMGREP_GUARDIAN_OBSERVE_ONLY=1`, which lets Guardian scan as usual while hiding results from developers and coding agents until you're ready to roll it out.

Shadow-deploy Guardian to see what it finds before surfacing results to developers or coding agents.

```
$ export SEMGREP_GUARDIAN_OBSERVE_ONLY=1
```

— Exact env var and behavior explained snapshot-20261112

04 **Bring-your-own AI provider keys for Agentic Workflows** NEW

78

Agentic Workflows can now run on your own AI provider keys for OpenAI, Anthropic, Google Gemini, xAI, and Microsoft Foundry, so configured model requests and AI-powered detection use your providers instead of Semgrep-managed keys. Microsoft Foundry (formerly Azure OpenAI) now has its own card in AI provider settings, accepting a resource endpoint, an API key, and one or more deployments each with its model.

— Names providers and config fields, no exact command snapshot-20261112

05 **Dropped Intel Mac CLI binaries** BREAKING

75

The Semgrep CLI no longer ships binaries for Intel Macs; version 1.182 is the last release supporting them. Use the `semgrep/semgrep` Docker image on Intel Mac or native arm64 Python on Apple silicon.

— Exact version cutoff and migration path given snapshot-20261112

06 **Expanded language syntax parsing** IMPROVED

68

Semgrep Code now parses newer Bash arithmetic expressions such as $(($X * 2)), Clojure files using reader-evaluated metadata (^#=(...)) or the legacy `#^` prefix, and Swift 5.9 syntax including macros, parameter packs, and `if` and `switch` expressions.

— Concrete syntax named, no runnable example snapshot-20261112

07 **Supply Chain dependency path for Go transitive findings** NEW

65

Supply Chain now shows the dependency path for transitive findings in Go projects, requiring `--allow-local-builds` and dependency paths turned on.

— Names flag and scope, lacks deeper mechanism snapshot-20261112

08 **Advisory version-range matches and Slack alerts** NEW

62

An advisory's matching projects now include a 'Potentially matching: declared in a version range' section listing each project, branch, file, and declared version. Incident Slack notifications also report how many such projects exist and which packages they declare.

— Describes new UI section and Slack reporting snapshot-20261112

THINNER COVERAGE BELOW

09 **Autotriage cost reduction** IMPROVED 55

Autotriage cost drops to 0.1 AI credits per finding (down from 1 credit) starting October 31, 2026, for paid plans, trials, and the free plan.

— Exact numbers and date, no action required snapshot-20261112

10 **Third-party licenses in CLI help output** NEW 55

Adds a `THIRD-PARTY LICENSES` section to the `--help` output of most `semgrep` commands, linking to `semgrep.dev/cli-credits`.

— Exact flag and URL named, minor feature snapshot-20261112

11 **Notes on Agentic Workflows issues** NEW 50

Adds notes to Agentic Workflows issues via the New note button at the top of an issue's activity.

— Names exact UI control but little else snapshot-20261112

12 **Guardian for GitHub Copilot CLI public beta** NEW 30

Semgrep Guardian for GitHub Copilot CLI is now in public beta on macOS and Linux.

— Thin one-line announcement, no mechanism given snapshot-20261112

└──▷ **BREAKING ON UPGRADE**

! The Semgrep CLI no longer ships binaries for Intel Macs; version 1.182 is the last release supporting them. Use the `semgrep/semgrep` Docker image on Intel Mac or native arm64 Python on Apple silicon.

WAS THIS USEFUL?  

syzkaller

Feedback ⓘ changes since 2026-11-10 CODE ↗

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

syzkaller adds vhost/vsock executor helpers and support for disabling the #syz test command per reporting stage

Adds `syz_vhost_vsock_setup` and `syz_vhost_vsock_send` executor helpers for fuzzing vhost vsock interfaces.

Supports disabling the `#syz test` command per reporting stage via `DisableTestCommandConfig` with a custom message in dashboard and `syz-cluster`.

Adds ext4, fscrypt, and fsverity scenario seeds to `sys/linux/test` for improved kernel fuzzing coverage.

Improves ext4, quota, fscrypt, and fsverity syscall descriptions in `sys/linux`.

Updates `pkg/aflow` backend to use the `gemini-4-carbon` model for deep reasoning.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Issue types and fields are available for personal accounts (2026-11-11) BLOG ↗

RANK

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

> GitHub brought issue types and issue fields to personal-account repositories, complete with REST/GraphQL and webhook support, and let users create labels and milestones directly from an issue.

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

01 **Issue types and fields for personal accounts** NEW 88

GitHub now supports issue types and issue fields on personal-account repositories, manageable via the REST and GraphQL APIs and webhooks. Up to 25 issue types can be customized from the Issue types page under Planning in account settings (defaults: Bug, Task, Feature), and up to 25 issue fields from the Issue fields page under Planning (defaults: Priority, Effort, Start date, Target date), supporting text, number, date, single-select, and multi-select field types. Issue fields can be pinned to specific issue types or left available on issues without a type, and both types and fields can be used in projects to create custom views and perform bulk updates.

Create a new issue type for a personal-account repository to standardize triage across your open-source projects.

> 📍 In your personal account settings, go to Planning › Issue types, then click New issue type and define a name, description, and color.

Create a custom single-select Priority field and pin it to the Bug issue type so it always appears when filing bugs.

> 📍 In your personal account settings, go to Planning › Issue fields, click New issue field, choose Single-select, name it Priority, add options (Critical, High, Low), then pin it to the Bug issue type.

— Names exact settings pages, limits, defaults, field types, and API surfaces.

Issue types and fields are available fo…

THINNER COVERAGE BELOW

02 **Create labels and milestones from an issue** NEW 45

Users can create labels and milestones directly from an issue, without navigating to the repository's labels or milestones pages.

.

— Describes the action but gives no steps or mechanism detail.

Issue types and fields are available fo…

WAS THIS USEFUL? 👍 👎

HashiCorp terraform

Feedback ⓘ v1.17.2 NOTES ↗ RANK

Terraform enables you to safely and predictably create, change, and improve infrastructure. It is a source-available tool that codifies APIs into declarative configuration files that can be shared amongst team members, treated as code, edited, reviewed, and versioned.

> Terraform v1.17.2 improves performance of the `query` command for large-scale configuration generation from HCP Terraform.

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

01 **Faster, lower-memory query configuration generation** IMPROVED

45

The `query` command now uses less memory and runs faster when generating configuration from large HCP Terraform queries.

— Names the command but no metrics or mechanism. v1.17.2

WAS THIS USEFUL? 👍 👎

OpenQodex

Feedback ⓘ openqodex@0.18.0 NOTES ↗ RANK

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

> OpenQodex 0.18.0 adds support for Goose as a coding agent, automating its setup via a new init flag.

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

01 **Goose agent support via init flag** NEW 88

Running `openqodex init --agent goose` sets up Goose as a coding agent: it adds a review hint to Goose's hints file and an openqodex skill to Goose's folders. The command locates Goose via `~/.config/goose` or via `XDG_CONFIG_HOME` when that environment variable is set.

Register OpenQodex with Goose so it adds a review hint and skill automatically, using either the default config location or a custom XDG path.

```
$ npx openqodex init --agent goose
```

— Names exact flag, config paths and env var; runnable command given. openqodex@0.18.0

WAS THIS USEFUL? 👍 👎

GitLab Feedback ⓘ

How to prevent leaks from Claude Code's 'Don't ask again' file (2026-11-11)

BLOG ↗ RANK

Teams and agents ship trusted software faster with GitLab, the intelligent orchestration platform with context and governance for the entire software lifecycle.

> GitLab shipped a full toolkit for finding and eliminating leaked Claude Code `.claude/settings.local.json` files across an instance, combining new search APIs, Secret Push Protection rollout guidance, scheduled historic secret scanning, group-wide `.gitignore` templates, prohibited-filename push rules, and Duo Agentic Chat-driven cleanup.

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

01 API endpoints to discover leaked Claude Code settings files

NEW 90

New API calls for locating `.claude/settings.local.json` across an instance: `GET /api/v4/groups/<id>/search?scope=blobs&search=filename:settings.local.json%20path:.claude&exclude_forks=false&per_page=100` finds all current copies across a group, `GET /api/v4/search?scope=blobs&search=filename:settings.local.json%20path:.claude` searches an entire self-managed instance without looping over groups, `GET /api/v4/groups?owned=true&top_level_only=true&per_page=100` lists every top-level group you own for bulk scanning, and `GET /api/v4/projects/<id>/forks` lists all forks of an affected project so their independent histories are also checked.

Check whether `.claude/settings.local.json` ever appeared in a project's full commit history — catches deleted files and all branches, not just the current default branch.

```
$ curl --header 'PRIVATE-TOKEN: <token>' \
  'https://gitlab.com/api/v4/projects/<id>/repository/commits?path=.claude/settings.local.json&all=true&follow=false&per_page=100'
```

— Four exact endpoints with query params given, runnable directly.

How to prevent leaks from Claude Code's…

02 Scheduled historic secret detection via scan execution policy

NEW 85

Run GitLab Secret Detection across every project in a group, including those with no `.gitlab-ci.yml`, by creating a scan execution policy under `Group > Secure > Policies > New policy > Scan execution policy` with a weekly-schedule rule and the variable `SECRET_DETECTION_HISTORIC_SCAN: 'true'` to cover full commit history.

— Names policy path and exact variable needed to enable full-history scans.

How to prevent leaks from Claude Code's…

---

03 **Prohibited filename push rule for Claude Code settings file**

NEW

75

Enforce a prohibited filename rule across a group via `Settings > Repository > Pre-defined push rules` with the pattern `(^|/)\.claude/settings\.local\.json$` under Prohibited filenames, blocking the file from ever being pushed.

Block `.claude/settings.local.json` from landing in any project — apply the push rule to every project in a group by looping over project IDs.

```
$ # Get all project IDs in the group
curl --header 'PRIVATE-TOKEN: <token>' \
  'https://gitlab.com/api/v4/groups/<group_id>/projects?
include_subgroups=true&per_page=100' \
  | jq '.[].id'

# Apply the push rule to each project (use PUT if the project
already has push rules)
curl --request POST \
  --header 'PRIVATE-TOKEN: <token>' \
  --header 'Content-Type: application/json' \
  --data '{"file_name_regex": "
(^|/)\\.claude/settings\\.local\\.json$"}' \
  'https://gitlab.com/api/v4/projects/<id>/push_rule'
```

— Exact regex pattern plus a runnable API example to apply it.

How to prevent leaks from Claude Code's…

---

04 **Secret Push Protection rollout for Self-Managed** NEW

60

Enable Secret Push Protection on GitLab Self-Managed by turning it on under `Admin > Settings > Security and compliance`, then activating it per project under `Secure > Security configuration > Secret push protection` or via the group security settings API.

— Exact settings paths given but no API or command example.

How to prevent leaks from Claude Code's…

05 **Group-wide .gitignore template inheritance** NEW 60

Inherit a `.gitignore` entry for `.claude/settings.local.json` in every new project by adding it to the group's file templates project under `Settings > General > Templates`.

Make every new project in a group inherit a `.gitignore` rule that excludes `.claude/settings.local.json`, so the gap Claude Code leaves is closed by default.

> 📍 In the group, go to Settings › General › Templates, select your file templates project, and add a .gitignore that includes the line: .claude/`settings.local.json`

— Exact UI path and concrete gitignore line to add.

How to prevent leaks from Claude Code's…

THINNER COVERAGE BELOW

06 **Duo Agentic Chat group-wide audit and cleanup** NEW 50

Use GitLab Duo Agentic Chat in a top-level group to audit all projects for the Claude Code settings file — including deleted copies and non-default branches — and to open merge requests that remove the file and add it to `.gitignore` across the whole group in two approval clicks.

— Describes capability and workflow but no concrete command shown.

How to prevent leaks from Claude Code's…

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

fcaptcha

Feedback ⓘ v1.50.0 NOTES ↗ RANK

> fcaptcha shipped an official React integration package and a new paired-observation analyzer for detecting automated agents in browser captures.

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

01 Official React package for fcaptcha NEW

75

The `@webdecoy/fcaptcha-react` npm package integrates fcaptcha into React 18 and 19 applications, including SSR and StrictMode environments. It rejects conflicting page-wide `serverUrl` configuration instead of silently redirecting widgets to the wrong server, and ships TypeScript declarations along with explicit backend token-verification guidance.

Add fcaptcha to a React 18 or 19 project, including SSR setups where a bare script tag would fail.

```
$ npm install @webdecoy/fcaptcha-react
```

— Named package, install command, and specific safeguards described. v1.49.0

THINNER COVERAGE BELOW

02 Paired-observation analyzer for agent detection NEW

43

A new paired-observation analyzer screens candidates against every captured human using consistent-direction and leave-one-pair-out checks, aimed at browser-internal agent/control captures.

— Mechanism named but no interface or usage detail given. v1.48.2

UPDATED DETECTION identity-coherence-v1 JavaScript-engine axis

> v1.50.0 NOTES ↗ LATEST STABLE
>
> The new engine axis is monitoring-only; no production scoring, token issuance, rate-limit, or blocking-policy behavior changes in this release.

└──▷ **GET THIS VERSION** ▾

Docker

```
$ docker pull ghcr.io/webdecoy/fcaptcha:1.50.0
```

git

```
$ git clone --branch v1.50.0
https://github.com/WebDecoy/FCaptcha.git
# already have the repo? check out this version:
$ git checkout v1.50.0
```

published under the project's GitHub account · checked 12 Nov Wrong command?

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Cloudsmith Feedback ⓘ

Hugging Face downloads now work on huggingface_hub 1.22 and later (2026-11-11)

BLOG ↗

Cloudsmith is a fully-managed artifact management, software supply chain security, and software distribution platform.

UPDATED INTEGRATION Hugging Face integration — Tree API support

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

kingfisher

Feedback ⓘ v2.18.1 NOTES ↗

Kingfisher is a detect secrets. Validate them live.

Kingfisher adds live finding counts during CLI scans, a Node SDK showcase (Kingfisher Workbench), and new Node SDK APIs for dependency-aware rule loading and scoped validation.

└──▷ GET THIS VERSION ▾

Homebrew

```
$ brew install kingfisher
# already installed? upgrade to this version:
$ brew upgrade kingfisher
```

Docker

```
$ docker pull ghcr.io/mongodb/kingfisher:2.18.1
```

pipx

```
$ pipx install kingfisher-bin==2.18.1
# already installed? upgrade to this version:
$ pipx install --force kingfisher-bin==2.18.1
```

git

```
$ git clone --branch v2.18.1 https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.18.1
```

Homebrew lists this version · published under the project's GitHub account · PyPI attestation · checked 12 Nov

Wrong command?

Show live reportable finding counts on stderr during CLI scanning and validation without affecting final report deduplication; progress stays on interactive stderr only.

Adds Kingfisher Workbench, a Node SDK showcase at `javascript/examples/kingfisher-workbench/` with working-tree and Git-history scans, live progress, redacted HTML reports, synthetic screenshots, and a loopback validation/revocation demo.

NEW VULNERABILITY CHECK Adds opt-in Rules.load({ includeDependencies: true }) on the Node SDK to retain transitive helper rules below the requested confidence before compilation.

Adds Validator.validate(findings, { candidates }) on the Node SDK to scope provider checks without discarding the original helper batch; candidates are preflighted like RevocationPlan.validate() and incomplete inputs are skipped before any provider request.

Adds non-copying `ScanInput.byteLength` to the Node SDK.

└──▷ **BREAKING ON UPGRADE**

- ! Dynamic confidence upgrades and demotions are removed across the CLI and SDKs; each rule now uses its configured confidence. Legacy custom YAML/TOML confidence helpers are inert.
- ! An explicitly selected rule whose confidence is below the active threshold now errors before compilation, including when used alongside `--rule all`; previously it would run silently.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ changelog · Nov 12, 2026

CHANGELOG ↗ RANK

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

> Finite State Platform shipped an estimated-count header for Audit Trail pagination, a new endpoint to poll for pending finding recalculations after component status changes, time-windowed CVE update queries, Brotli component detection with version reading in binary scans, and corrected vendor/OS naming for Debian and Ubuntu packages in vulnerability results.

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

01 Pending finding-update polling endpoint NEW

85

A `GET /public/v0/components/finding-updates/{projectVersionId}` endpoint reports which components still have findings being recalculated, surfaced via the `X-Finding-Updates-Pending` header, so automation can poll after a `PUT /public/v0/components/{projectVersionId}/{componentId}/status` status change (e.g. marking a component `false_positive`) before acting on results.

Poll whether a large component's findings have finished updating after marking it False Positive, so downstream automation does not act on stale data.

```
$ curl -s -X PUT
'https://platform.finitestate.io/public/v0/components/{projectVer
sionId}/{componentId}/status' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"status": "false_positive"}' \
  -D - | grep -i 'X-Finding-Updates-Pending'

# If header is present, check pending components:
curl -s
'https://platform.finitestate.io/public/v0/components/finding-
updates/{projectVersionId}' \
  -H 'Authorization: Bearer <token>'
```

— Two endpoints and header named with a runnable curl example.

changelog-20261112-cdb5d3ff

02 Audit Trail API total count estimation IMPROVED

78

Audit Trail API endpoints now set the `X-Total-Count-Estimated: true` response header when the total exceeds 10,000 events, letting clients detect when the reported count is approximate before relying on it for a full paginated export.

Check whether the Audit Trail total is an estimate before deciding to export a CSV for the full record, avoiding pagination gaps above 10,000 events.

```
$ curl -s -D - 'https://platform.finitestate.io/public/v0/audit-events' \
  -H 'Authorization: Bearer <token>' | grep -E 'X-Total-Count|X-Total-Count-Estimated'
```

— Exact header and threshold named with a verifying curl command.

changelog-20261112-cdb5d3ff

03 **CVE update events endpoint with time filters** IMPROVED

70

The `GET /public/v0/cves/updates` endpoint accepts `start` and `end` timestamp parameters to pull CVE update events for a given window, and returns a validated error response when an unsupported filter field is supplied.

Pull CVE update events for a wide time window and handle the validated error response when an unsupported filter field is supplied.

```
$ curl -s 'https://platform.finitestate.io/public/v0/cves/updates?start=2025-01-01T00:00:00Z&end=2026-11-12T00:00:00Z' \
  -H 'Authorization: Bearer <token>'
```

— Endpoint and parameters named but validation behaviour underspecified.

changelog-20261112-cdb5d3ff

THINNER COVERAGE BELOW

04 **Brotli component identification in binary scans** NEW

45

Binary scans now identify Brotli as a component, including reading its exported version string on supported builds.

— Names the component and version detail but no usage surface.

changelog-20261112-cdb5d3ff

05 **Debian and Ubuntu package naming in vulnerability results**

IMPROVED

41

Vulnerability results now include Debian and Ubuntu packages that previously appeared with no vendor and no operating-system component named.

— Describes the fix but gives no concrete scope or metric. changelog-20261112-cdb5d3ff

---

WAS THIS USEFUL? 👍 👎

ipatool

Feedback ⓘ v2.8.1 NOTES ↗ RANK

Ipatool is a command-line tool that allows you to search for iOS, iPadOS, tvOS, watchOS, visionOS, and macOS apps on the App Store, and download .ipa or macOS .pkg app packages.

> ipatool restores broken App Store authentication with local Anisette generation, speeds up downloads with parallel streams, and expands compatibility to cover delisted apps, legacy iOS packages, and more App Store regions.

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

01 **Local Anisette generation for App Store authentication**

IMPROVED 50

App Store authentication now works again via local Anisette generation and provisioning, removing the dependency on external Anisette servers. Store operations and signing now use the modern Apple protocol, improving compatibility with current App Store requirements.

– Explains mechanism and cause but no config/flag names v2.8.0

02 **Dynamic App Store storefront discovery** IMPROVED 40

ipatool now discovers App Store storefronts dynamically, supporting regions not in the built-in list.

– Clear behaviour change but no concrete region list v2.8.1

03 **Actionable auth error dialogs** IMPROVED 40

ipatool now shows Apple's authentication dialogs and account-action links when sign-in fails, making auth errors actionable.

– Describes user-facing behaviour, no exact UI/flag detail v2.8.1

04 **Parallel download streams for large apps** IMPROVED 30

ipatool now downloads large app packages using parallel download streams for faster transfers.

– Describes benefit, no mechanism or flag detail v2.8.1

05 **Download support for delisted and legacy apps** NEW 30

ipatool now recovers previously acquired, delisted iOS apps, and supports downloads of legacy iOS apps that lack platform declarations.

– Two thin compatibility fixes grouped, no mechanism given v2.8.0 v2.8.1

WAS THIS USEFUL?  

edb-debugger

Feedback ⓘ 1.6.0 NOTES ↗ RANK

edb is a cross-platform AArch32/x86/x86-64 debugger.

> edb-debugger 1.6.0 ships a hardware breakpoint dialog, breakpoint-aware and library-load-aware bookmark handling, cancelable dialog searches, automatic function block comments, Qt6 build support, and a UI layout reset option.

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

01 **Hardware breakpoint configuration dialog** NEW 55

edb adds a hardware breakpoint dialog for configuring hardware breakpoints through the UI, accessed from the debugger menu to select address, condition (execute, read, write), and width.

When you need to trigger on memory access or execution without consuming a software breakpoint slot, open the hardware breakpoint dialog from the debug menu to configure a hardware breakpoint.

> 📍 In the debugger menu, open the hardware breakpoint dialog, select the address, condition (execute, read, write), and width, then apply.

— Names the dialog and configurable options with UI steps. 1.6.0

02 **Bookmark save/restore with library load awareness** NEW 52

Bookmarks now persist correctly across sessions that involve dynamic library loading, so analysis positions set via the Bookmarks panel survive restarts even when library base addresses shift.

When debugging a binary that uses dynamic libraries, use bookmark save/restore to keep your analysis positions intact across restarts where library base addresses may shift.

> 📍 In the Bookmarks panel, set your bookmarks as usual; they are automatically saved and restored with library load awareness on the next session.

— Explains persistence mechanism and when it applies, with usage steps. 1.6.0

03 **Qt6 build support** IMPROVED 35

edb-debugger now supports building and running with Qt6, including Qt 6.11.

— Names specific Qt version but no build steps given. 1.6.0

04 **Automatic function block comments** NEW 30

edb adds automatic function block comments before detected functions in the disassembly view.

— Names the feature and view, but no mechanism detail. 1.6.0

05 **Cancelable search dialogs** IMPROVED 25

edb adds cancelable search in most or all dialogs, allowing long-running searches to be interrupted.

— States behaviour but no specific dialogs or commands named. 1.6.0

06 **Reset UI layout from GUI** NEW 25

edb adds the ability to reset the UI layout from within the GUI.

— Minimal description; no menu path or config detail given. 1.6.0

07 **Bookmark breakpoint indicator** NEW 20

edb now shows which bookmarks have breakpoints set, making it easier to distinguish breakpoint-annotated bookmarks at a glance.

— Bare description of a visual indicator, no mechanism detail. 1.6.0

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.16.5 NOTES ↗ CODE ↗ RANK

Netcap is a framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap's v0.16.5 release reworked the analyst web UI's chart handling and the `/api/chart/data` endpoint, redesigned the offline rule evaluation flow into three distinct steps, and overhauled the connections and services views.

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

01 **Chart data endpoint response codes** IMPROVED 90

The `/api/chart/data` endpoint now answers `204 No Content` when no plottable values exist in scope, and `200` with an empty series for `format=json`; read and configuration failures remain `500`.

– Exact endpoint path, query param and status codes given v0.16.5

02 **Dashboard chart display and layout improvements** IMPROVED 65

Categorical charts now show the 20 most frequent values, with ties broken by name for a deterministic cut-off. Dashboard charts with no data in the selected scope are now hidden and the grid closes the gap automatically, and the dashboard grid now follows its container width, keeps rows aligned, stacks below 768 px, and scrolls pie legends while truncating axis labels.

– Named thresholds but UI-only behavior, no API or command v0.16.5

03 **Three-step offline rule evaluation flow** IMPROVED 65

Offline rule evaluation now presents pack, validation and evaluation as three distinct steps with status, readable limit units, and a results summary of alerts, records, matched rules and elapsed time.

– Describes flow structure but no exact UI path given v0.16.5

THINNER COVERAGE BELOW

04 **Connections and services view redesign** IMPROVED 50

Connections and services views now collapse behavioral context by default and display observations as a table, with directional byte bars and an aligned packet table replacing stacked statistics.

– Describes layout change but no navigation path or metric names v0.16.5

WAS THIS USEFUL?  

DETECTION & RESPONSE

◆ SIEM

Microsoft Sentinel

Feedback ⓘ changelog · Nov 12, 2026

CHANGELOG ↗ RANK

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

> Microsoft Sentinel added interactive KQL lake queries over Sentinel tables and federated external data via Microsoft Fabric.

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

01 Interactive KQL lake queries via Microsoft Fabric NEW

63

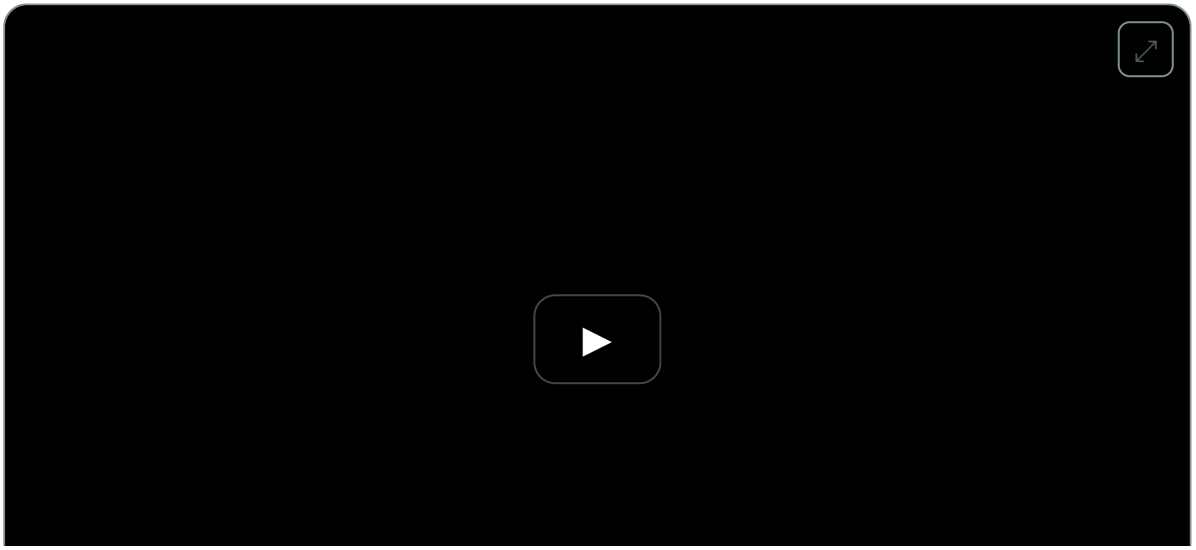
Users can run interactive KQL lake queries on Microsoft Sentinel tables and federated external data sources using Microsoft Fabric, enabling custom analytics on data lake tables and external data tables. In practice, this means opening a KQL queryset in Microsoft Fabric, connecting to a Microsoft Sentinel workspace, and running a lake query across Sentinel tables or external federated data sources.

Query Sentinel tables and federated external data sources interactively via Microsoft Fabric when you need ad-hoc analytics beyond standard Sentinel workbooks.

> 📍 In Microsoft Fabric, open a KQL queryset, connect to your Microsoft Sentinel workspace, and run a lake query across Sentinel tables or external federated data sources.

— Clear steps given but no limits or config details. changelog-20261112-cd14823c

└──▷ **ALSO FROM THESE RELEASES**





WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Bring your own ChatGPT subscription to Codex with Unity Gateway (2026-11-11)

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 a mix of AI-assistant integrations and platform extensions this window, headlined by bring-your-own ChatGPT billing through Unity Gateway, Python-based custom service policies in Unity Catalog, and new Genie One client surfaces for Windows and Chrome.

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

01 **Text and binary inputs for protobuf model serving requests**

NEW 80

Custom model serving endpoints (Public Preview) now accept text and binary data — titles, images, audio, and documents — as KServe v2 BYTES tensors alongside numeric tensors in the same protobuf request, with binary model outputs also returned as BYTES tensors.

— Names exact tensor format and request mechanism, no sample call

Text and binary inputs for protobuf req…

02 **Genie One Chrome extension (Beta)** NEW 75

Install the Genie One Chrome extension from the Chrome Web Store to open your Databricks workspace in every new Chrome tab, giving one-click access to asking questions, searching across tools, and opening dashboards.

Set Genie One as your Chrome new tab page so every tab opens directly into your Databricks workspace for instant querying and dashboard access.

📍 1. Go to the Chrome Web Store and search for Genie One. 2. Click Add to Chrome to install the extension. 3. Open a new Chrome tab — it now loads Genie One in your Databricks workspace.

— Includes exact install steps a reader can follow

The Genie One Chrome extension is in Be…

03 **Custom service policies in Python for Unity Catalog (Beta)**

NEW 70

Write custom service policies as Python functions in Unity Catalog to express rules that SQL cannot handle, such as regular expressions, Python packages, or calls to an external guardrail through a service credential.

— Explains mechanism and use cases, no exact syntax shown

Write custom service policies in Python…

04 **ChatGPT and Claude as partner platforms in ingress policies**

NEW INTEGRATION NEW 65

Context-based ingress policies now support ChatGPT and Claude as partner platforms, letting hosted MCP clients — such as the Databricks Genie plugins for ChatGPT and Claude — reach your workspace without manually tracking their IP ranges.

— Names platforms and plugins, no config path given

Allow ChatGPT and Claude as partner pla…

05 **Bring your own ChatGPT subscription to Codex via Unity Gateway**

NEW 60

Connect your own ChatGPT Go, Plus, Pro, or Enterprise subscription to Codex and the ChatGPT desktop app through Unity Gateway, with requests billed to your subscription and access governed by Unity Gateway.

— Names subscription tiers and gateway but no setup steps

Bring your own ChatGPT subscription to …

THINNER COVERAGE BELOW

06 **Genie One desktop app for Windows (Beta)** NEW 45

The Genie One native desktop app now supports Windows (Beta), letting users chat with data in natural language, manage threads and documents, and use Genie Agents outside a browser.

— Describes capability but no install or usage path

The Genie One desktop app for Windows i…

UPDATED COMPLIANCE FRAMEWORK Feature Views default availability for compliance security profile workspaces

└──▷ ALSO FROM THESE RELEASES

Connect ChatGPT to Genie One with the Databricks Genie plugin (2026-11-11)

BLOG ↗ STABLE

Databricks adds a Genie plugin that connects ChatGPT to the Genie One MCP for natural-language data queries.

> Feature Views are now available by default for workspaces with the compliance security

BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Nov 12, 2026 CHANGELOG ↗ RANK

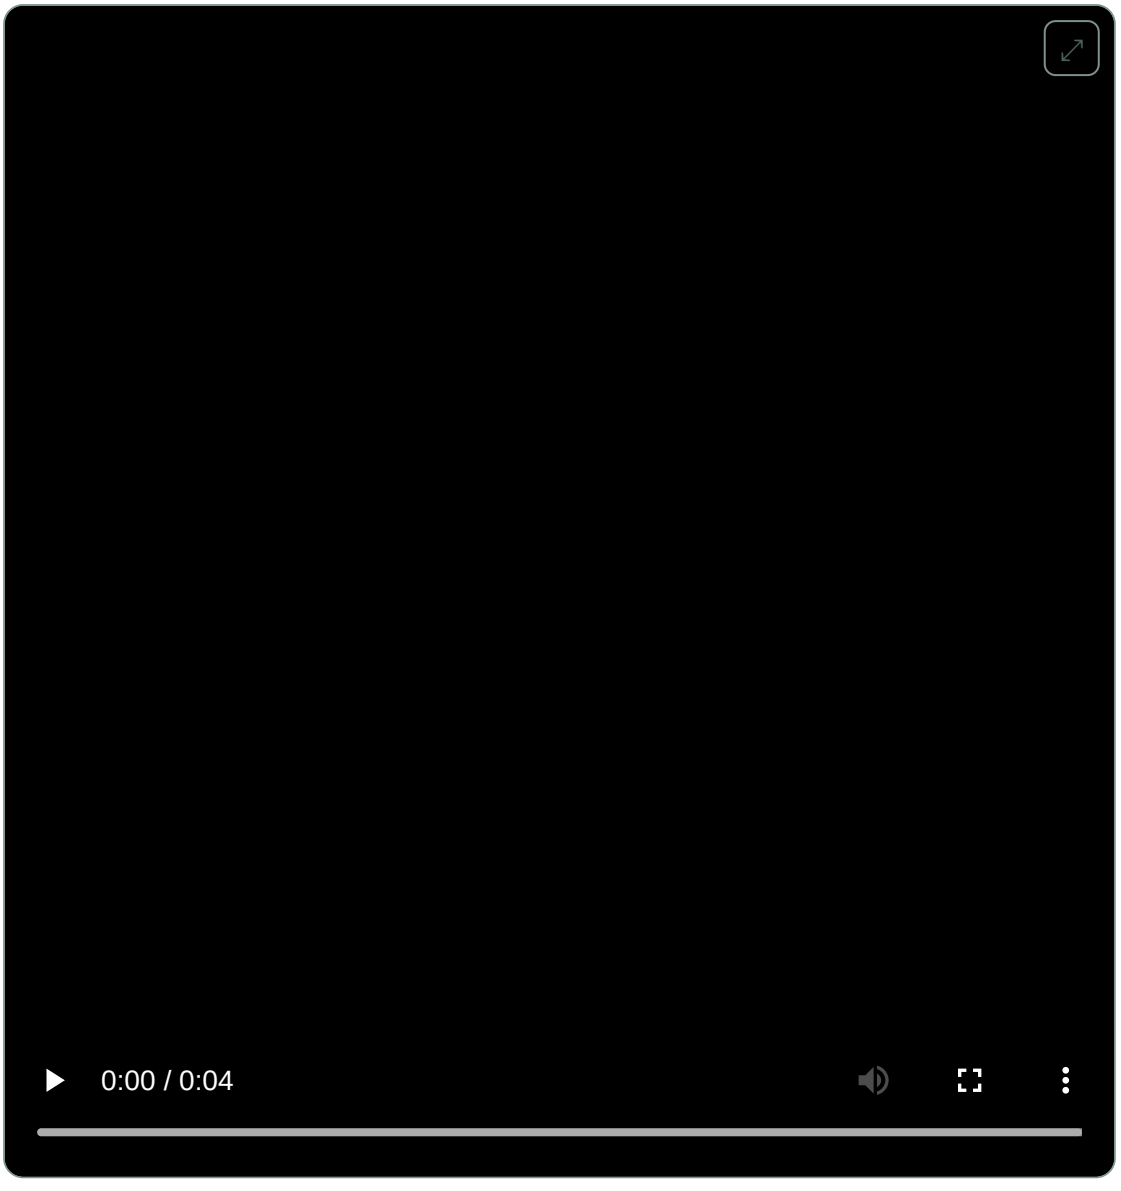
Tines is an apps, agents and automations, in the complete environment for enterprise AI.

> Tines 44.3.0 improves visibility across cases, scheduled actions and credentials, adding webhook case tracking, credential-refresh alerts, and new notification and dashboard controls.

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

01 **Custom date filter on dashboards** NEW 37

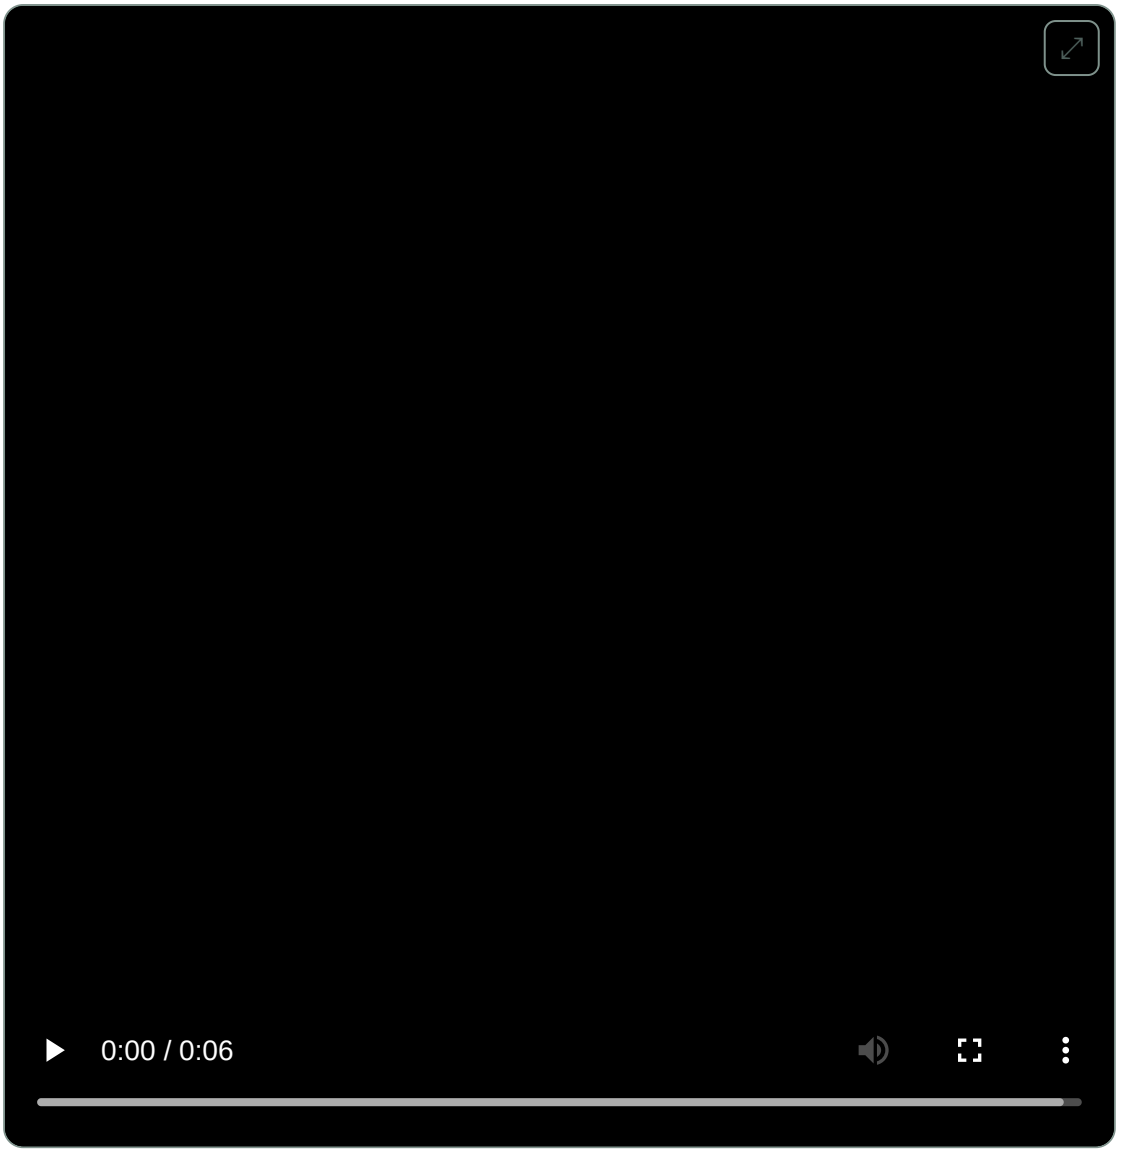
Tines now lets users apply their own date filter on dashboards for custom time-range exploration.



— Clear capability but no exact UI control named changelog-20261112-faa6528b

02 **Webhook case actions in activity feed** NEW 35

Tines now shows webhook case actions in the activity feed so teams can track case changes triggered by webhooks.



— Describes what changed but no config detail `changelog-20261112-faa6528b`

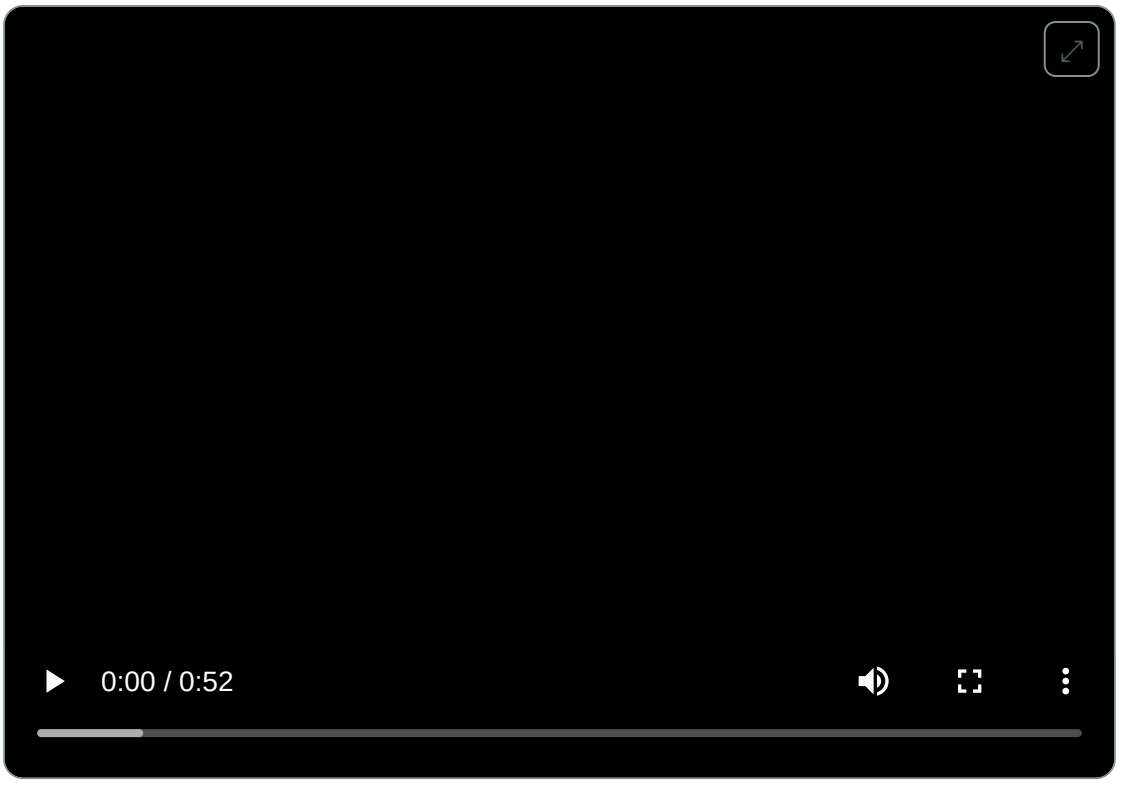
03 **Credentials in story monitoring recipients** NEW 35

Tines now supports using credentials in story monitoring recipients, allowing authenticated destinations for monitoring alerts.

— Names the feature area but no field or setup steps `changelog-20261112-faa6528b`

04 **Credential refresh failure notifications** NEW 33

Tines now sends notifications when credentials fail to refresh, alerting teams before automated workflows break.



— Explains purpose but lacks config or channel detail changelog-20261112-faa6528b

05 **Email notification control for case assignments** IMPROVED

28

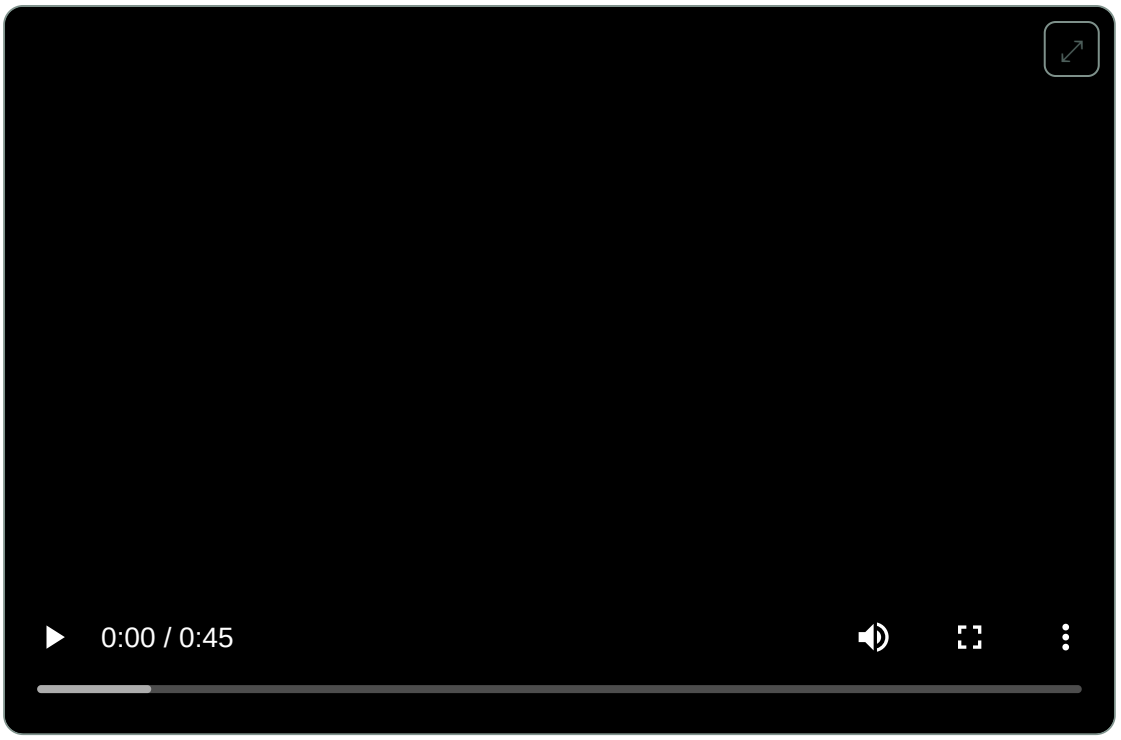
Tines now lets users control email notifications for case assignments to reduce alert noise.

— Names a control but not where to set it changelog-20261112-faa6528b

06 **Scheduled actions visibility** NEW

26

Tines now surfaces scheduled actions visibility so practitioners can see upcoming automated actions at a glance.



— Thin description, no UI path or mechanism changelog-20261112-faa6528b

07 **Open groups in story version preview** IMPROVED 26

Tines now shows open groups when previewing story versions, giving reviewers full context during version review.

— Describes behaviour, no exact navigation given changelog-20261112-faa6528b

WAS THIS USEFUL?

◆ Endpoint Security (EDR/XDR)

clamui

Feedback ⓘ v0.6.0 NOTES ↗ RANK

Clamui is a clamav Desktop application (GUI)

> ClamUI v0.6.0 preserves full scan evidence past the 64 MiB output limit, adds a protected stdin transport for very long daemon paths, adds Vietnamese localization, and improves history/quarantine browsing plus VirusTotal upload security checks.

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

01 **Protected stdin transport for long daemon paths** IMPROVED

70

Daemon paths of `1,023` bytes or more now use the protected stdin transport instead of being split, retaining exclusions, progress, and cancellation.

— Names exact byte threshold and the before/after transport change. v0.6.0

02 **Full scan evidence retention beyond 64 MiB limit** IMPROVED

60

ClamUI now retains findings, skipped files, warnings, and summary counts after the 64 MiB displayed-output limit for both ClamAV backends, so large-scan history and final results preserve evidence and exact filename identities.

— Explains mechanism and scope but no exact UI path or command. v0.6.0

THINNER COVERAGE BELOW

03 **History and quarantine browsing improvements** IMPROVED

50

History loads through a recoverable metadata index, quarantine browsing uses bounded pages with Unicode-aware search, and progress updates are coalesced without dropping threat changes or terminal events.

— Describes mechanism but no concrete config or UI path. v0.6.0

04 **VirusTotal upload pathname substitution check** IMPROVED

50

VirusTotal uploads are now rejected if the original pathname is substituted after lookup, alongside the existing no-follow descriptor identity checks.

— Describes security check but no exact flag or config exposed. v0.6.0

05 **Vietnamese localization added** NEW

45

ClamUI adds Vietnamese (`vi`) to its desktop and CLI translation catalogs.

— Names the locale code but gives no further mechanism. v0.6.0

WAS THIS USEFUL? 👍 👎

Avast Business Hub

Feedback ⓘ November 2026 NOTES ↗

RANK

Avast Business Hub is a layered IT protection platform that is easy to implement and manage, designed for SMBs that want the highest level of protection for their business.

> Avast Business Hub made a small usability improvement to its device management reporting, adding reboot history to CSV exports.

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

01 **Last reboot time in device CSV export** IMPROVED 50

Device list CSV exports now include a Last reboot time column, surfacing reboot history alongside existing device data.

— Names exact column added but no export path or limits November 2026

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Patch Management

Patch My PC

Feedback ⓘ

Patch My PC Publisher Preview Version 3.26.11013.0 Released (2026-11-11) BLOG ↗

RANK

Simplify IT with Patch My PC's endpoint management software. Automate patching, app updates, reporting, and analytics to save time and boost security.

> Patch My PC's Publisher preview release adds remote visibility into sync logs and changes default installer return code handling across major deployment targets.

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

01 **Remote access to wsyncmgr.log in console sessions** NEW

55

Publisher now lets admins view `wsyncmgr.log` directly within remote console sessions, without needing a local Publisher session on the machine performing the sync.

— Names the log file but gives no further usage steps.

Patch My PC Publisher Preview Version 3…

02 **Catalog-defined installer return codes as default** IMPROVED

35

Publisher now treats catalog-defined installer return codes as the default behavior for WSUS, Configuration Manager, and Intune, per the release summary, though the exact mechanism of the change is not detailed in the available text.

— Names three deployment targets but mechanism and impact are unclear from garbled source text.

Patch My PC Publisher Preview Version 3…

WAS THIS USEFUL? 👍 👎

◆ Vulnerability & Exposure Management

defectdojo

Feedback ⓘ 3.5.101 NOTES ↗ RANK

Defectdojo is an open-source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo improved TruffleHog secret deduplication across commits and added automatic Jira ticket comment notifications via the Jira connector.

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

01 **Stable TruffleHog secret IDs across commits** IMPROVED 55

DefectDojo now assigns a stable `unique_id_from_tool` to each TruffleHog secret so that the same secret appearing in multiple commits is not duplicated as separate findings.

— Names the field but no usage steps 3.5.101

02 **Automatic Jira ticket comments via connector** NEW 35

A new notification hook triggers automatic Jira ticket comments through the Jira connector.

— Brief description, no config details given 3.5.101

WAS THIS USEFUL? 👍 👎

oxo

Feedback ⓘ v2.22.13 NOTES ↗ RANK

> oxo expanded its domain-name fingerprinting message vocabulary with new selectors covering web-archive, HTTP service, and DOM-based signals.

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

01 New fingerprinting message types for domain names NEW

55

oxo added new message/selector types for domain-name asset fingerprinting: `v3.fingerprint.domain_name.web_archive` for web-archive-based domain fingerprinting signals between agents, and `v3.fingerprint.domain_name.service.http` plus `http.dom` selectors for HTTP service and DOM fingerprinting on domain name assets.

— Names exact message types but not how agents use them v2.22.14 v2.22.12

UPDATED SUPPORTED PLATFORM HarmonyOS file metadata message enriched with size, hashes, version, and label

> v2.22.13 NOTES ↗ STABLE
>
> oxo now includes file size, hashes, version, and label in the HarmonyOS file metadata message for richer scan context.
>
> └──▷ **GET THIS VERSION** ▾
>
> ```
> $ git clone --branch v2.22.13 https://github.com/Ostorlab/oxo.git
> # already have the repo? check out this version:
> $ git checkout v2.22.13
> ```

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ Iru Agent for Windows Endpoint Agent 1.26.15 NOTES ↗

RANK

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

> Iru's Windows Endpoint Agent extended its Auto Apps deployment feature to support ZIP-packaged applications.

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

01 **ZIP support for Auto Apps on Windows** IMPROVED 40

Iru Agent for Windows now installs Auto Apps packaged as ZIP files, extending the Auto Apps deployment mechanism beyond its previous format support.

— Names the capability and format but no limits or steps

Iru Agent for Windows Endpoint Agent 1.…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Non-Human & Agent Identity

Okta for AI Agents

Feedback ⓘ changelog · Nov 12, 2026

CHANGELOG ↗

Okta for AI Agents governs AI agent identities, permissions, and access across autonomous workflows.

Okta for AI Agents adds import of staged and preview AI agents, and Salesforce agent imports now show the latest active or published version.

Import AI agents that are in a staged or preview state, expanding beyond production-only imports.

Salesforce AI agent imports now display the latest active version, or the latest published inactive version when no active version is available.

NEW INTEGRATION ANSI Standards Connect SAML integration

NEW INTEGRATION Datatrac Web OIDC integration

NEW INTEGRATION Datatrac Web Universal Logout integration

UPDATED INTEGRATION Dropbox Business XAA support

UPDATED INTEGRATION Miro XAA support

NEW INTEGRATION Modal SCIM integration

NEW INTEGRATION Monday.com XAA integration

NEW INTEGRATION Rogue Security API Service integration

UPDATED INTEGRATION Rubrik Security Cloud additional OAuth scopes

NEW INTEGRATION StatusDrift.com SAML integration

NEW INTEGRATION StatusDrift.com Universal Logout integration

NEW INTEGRATION VeriKnox OIDC integration

UPDATED INTEGRATION Zoom XAA support

UPDATED SUPPORTED PLATFORM Android 14-17 patch level 2026-10-01

UPDATED SUPPORTED PLATFORM Windows 10 26H2 support

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ updated · Nov 12, 2026 NOTES ↗ RANK

Varonis connects data, AI, and threat detection into one complete security platform that allows companies to keep innovating, without increasing risk.

> Varonis introduced the MCP Gateway within its Atlas AI Gateway, giving organizations a single control point to govern which tools, APIs, agents, and MCP servers can act on behalf of users.

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

01 **MCP Gateway for agent tool access control** NEW 60

The MCP Gateway is a single control point within the Atlas AI Gateway that enforces identity-aware rules over which tools, APIs, agents, and MCP servers can be reached and which actions they may perform. Atlas validates agent identity through existing identity providers using delegated credentials so agents operate under real user permissions rather than broad standing access, and every tool call is tied to a real user identity and recorded alongside prompts in an investigation timeline for accountability.

.
.
.

— Explains mechanism and identity model but no config/API names or steps.

snapshot-20261112

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

amneziawg-installer

Feedback ⓘ v6.3.1 NOTES ↗ RANK

AmneziaWG-installer is a script that installs AmneziaWG, a modified WireGuard VPN with obfuscation features to bypass deep packet inspection and censorship.

> AmneziaWG-installer v6.3.1 adds a repair path for kernel module issues after OS upgrades, source-address control for cascade tunnels, and tightens how obfuscation parameters and exit-address validation are generated.

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

01 **Repair-module flag for kernel helper** NEW 85

Running `install_amneziawg_en.sh --repair-module` deploys the new module helper version 4 and rebuilds the kernel module without triggering a full install, intended for use after a 24.04 HWE-to-26.04 upgrade where DKMS had cached the old build.

After upgrading a server from Ubuntu 24.04 HWE to 26.04, redeploy the module helper and rebuild the kernel module without re-running the full install.

```
$ sudo bash ./install_amneziawg_en.sh --repair-module
```

— Names exact flag, version, and upgrade scenario it fixes v6.3.1

02 **AWG1_SRC for cascade-tunnel source address** NEW 80

The new `AWG1_SRC` setting gives cascade-tunnel source-address control, letting traffic to the exit server be routed from a specific address on an entry server that has multiple public IPv4 addresses, rather than sharing the main interface IP.

In a cascade setup where the entry server has multiple public IPv4 addresses, route tunnel traffic to the exit server from a specific third address so it does not share the main interface IP.

```bash
AWG1_SRC=203.0.113.3
```

— Names exact variable with a working config example v6.3.1

03 **S4=16 excluded from parameter generator** IMPROVED 73

The parameter generator now excludes `S4` = 16 from the allowed range (4–27 in the 2.0 profile, 12–27 in 3.1), so new installs no longer produce packet lengths that are

multiples of 16 at every payload size, matching a pattern plain WireGuard packets exhibit and that could aid detection.

— Names exact parameter and ranges, no direct command given v6.3.1

THINNER COVERAGE BELOW

04 **Exit-address validation against DNS-resolved names** IMPROVED

20

Exit-address validation is tightened against DNS-resolved entry names, per the release summary; no further mechanism or example is provided.

— Only a one-line summary mention, no detail or example v6.3.1

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 7.1.0 (2026-11-11) BLOG ↗

RegScale is a continuous compliance automation platform that helps organizations manage security controls, audit evidence, and regulatory requirements in real time.

Nessus imports now show an overall file progress bar and log one line per file.

NEW COMPLIANCE FRAMEWORK Tanium DISA STIG results import via sync_stig_benchmarks

NEW COMPLIANCE FRAMEWORK Tanium CIS benchmark and STIG results import via sync_benchmarks

NEW INTEGRATION Tanium integration for STIG and CIS benchmark sync

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v4.2.3 NOTES ↗

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 adds evidence review mode, scenario selection across risk workflows, memcached caching, a new `view_user_full` IAM permission, maturity scores in Framework Insights, and a dedicated third-party app shell.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.2.3 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v4.2.3
```

Gate user sign-in and administration detail visibility with the new `view_user_full` IAM permission.

Filter risk assessments, action plans, and exports by selecting specific scenarios with the new scenario selection feature across the risk assessment workflow.

Add memcached for consistent caching and rate-limiting across the application.

Review submitted evidence in a new dedicated evidence review mode.

Show documentation and maturity scores inside Framework Insights.

+1 more

NEW COMPLIANCE FRAMEWORK LYSS switch security framework with ISO/IEC 27001:2022 mapping

NEW COMPLIANCE FRAMEWORK ITSP.10.033 framework (EN+FR) with NIST SP 800-53 rev5 mapping

WAS THIS USEFUL? 👍 👎

OneTrust

Feedback ⓘ

November 18, 2026 - cmpapi-202611.1.0 (2026-11-11) BLOG ↗ RANK

The OneTrust Developer Portal allows you to leverage APIs and SDKs to integrate external systems seamlessly and streamline the data flow with the OneTrust platform.

> OneTrust's November 2026 CMP API update extends Web SDK (`OT-Device-Type: web`) support across the Log Consent, Get Banner Data, Get Preference Center Data, and Get IAB and Google Vendors APIs, while adding new OS-detection and error-code headers, A/B test variant visibility, and granular location-based reconsent.

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

01 Web consent logging and purpose notes in Log Consent API NEW

85

The Log Consent API now accepts `OT-Device-Type: web` and web-specific `interactionType` values `API_ALLOW_ALL`, `API_REJECT_ALL`, `API_UPDATE_CONSENT`, `GPC_VALUE_CHANGED`, `DNT_VALUE_CHANGED`, `ANONYMOUS_LOGGED_OUT_CONSENT`, `CROSS_DEVICE_SYNC_LOCAL`, and `NO_INTERACTION` for logging consent from web implementations. It also returns `canGenerateNotGivenReceipts` and `authenticatedLoggedOutConsent` fields under `appConfig` to indicate receipt support, and now includes `noteType` and `noteText` alongside `noteId` in `PurposeNote` responses.

— Names exact header, eight interactionType values, and response fields

November 18, 2026 - cmpapi-202611.1.0 (…

02 OS-detection and error-code headers on CMP APIs NEW

80

CMP APIs add an optional `OT-OS` request header (`Apple` or `Android`) to identify the underlying OS in cross-platform SDKs such as .NET MAUI — it takes precedence over `OT-Platform` and enables OS-specific behavior for ATT and Google advertising identifiers — and a new `OT-Error-Code` response header that returns comma-separated unique error codes on CMP API error responses, letting callers read failures without parsing the JSON body.

— Names both headers, their values, and precedence behavior

November 18, 2026 - cmpapi-202611.1.0 (…

03 Web banner data via Get Banner Data API NEW

73

The Get Banner Data API, called with `OT-Device-Type: web`, now returns `bannerUIData` covering layouts, positioning, accessibility, focus management,

branding, and custom styling, plus web-specific `appConfig` settings for geolocation handling, Google Consent Mode, Additional Consent Mode, Multi-Cookie Management, and global opt-out configuration.

— Names endpoint, header, and specific appConfig settings

November 18, 2026 - cmpapi-202611.1.0 (…

04 **Web preference center data via Get Preference Center Data API**

NEW

73

The Get Preference Center Data API, called with `OT-Device-Type: web`, now returns `pcUIData` covering layouts, styling, branding, accordion behavior, accessibility, purpose hierarchies, GPC/DNT signals, tracking technology disclosures (cookies, local storage, session storage), and search and filtering for cookies, vendors, and tracking technologies.

— Names endpoint, header, and the data fields it returns

November 18, 2026 - cmpapi-202611.1.0 (…

05 **Web vendor data via Get IAB and Google Vendors API** NEW

73

The Get IAB and Google Vendors API, called with `OT-Device-Type: web`, retrieves IAB TCF v2.4, Google, and general vendor data for web consent experiences, including localized labels via `vendorListDescription`, `filterAriaLabel`, `searchPlaceholderText`, and `searchAriaLabel` fields.

— Names endpoint, header, TCF version, and localization fields

November 18, 2026 - cmpapi-202611.1.0 (…

THINNER COVERAGE BELOW

06 **A/B test variant details in appConfig** NEW

55

CMP API responses now include a `variant` object under `appConfig` returning the variant `id` and `name`, letting callers confirm which banner variant was delivered and validate audience- and rule-based variant selection.

— Names the field and its two sub-fields November 18, 2026 - cmpapi-202611.1.0 (…

07 **Location-specific reconsent dates** NEW

40

Location-specific reconsent dates can now be applied at the state/region, country, or geolocation-rule level across CMP APIs and the Web SDK when granular geo-level reconsent is enabled.

— Describes scope in prose without a named field or endpoint

November 18, 2026 - cmpapi-202611.1.0 (…

WAS THIS USEFUL? 👍 👎

AGENT

◆ Agent Sandboxes

Docker Sandboxes

Feedback ⓘ v0.53.0 NOTES ↗ RANK

Docker Sandboxes run AI coding agents in isolated environments on your machine or on Docker-managed cloud infrastructure

> Docker Sandboxes v0.53.0 adds local disk usage reporting, Docker Hardened Image selection, Claude in Chrome support, and tunable stalled-connection timeouts for VPN recovery, alongside a cluster of CLI usability fixes and several breaking changes to agent handling, secret providers, and HTTP policy matching.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 **Local disk usage reporting with sbx df** NEW 80

`sbx df` reports local image, container, and managed-volume disk usage and estimated reclaimable space, with `--verbose` giving a per-resource breakdown.

Check how much local disk space sandboxes are consuming and what could be reclaimed before running a prune.

```
$ sbx df --verbose
```

— Runnable command with exact flag and scope named v0.53.0

02 **Docker Hardened Image selection via settings** BREAKING 70

Docker Hardened Image templates are now selected with `sbx settings set platform.images.useDHI true`, replacing the retired `DOCKER_SANDBOXES_USE_DHI` environment variable.

— Names exact setting key and the retired env var v0.53.0

03 **Stalled outbound connection timeout** NEW 70

Stalled outbound connection timeouts can be tuned with `egress.stalledConnectionTimeoutSeconds`, or set to `0` to disable; the daemon must be restarted for the change to apply to existing sandboxes, aiding VPN recovery.

— Config key, disable value, and restart requirement all named v0.53.0

04 **Claude in Chrome access for sandboxed Claude Code** NEW 70

Sandboxed Claude Code can now access Claude in Chrome via the `claude.remoteControl` and `claude.chrome` settings, started with `--chrome` on Claude Code.

— Both settings and the launch flag are named v0.53.0

05 **sbx run entrypoint argument handling change** BREAKING

65

`sbx run <agent> -- <args>` now preserves the image's entrypoint arguments when a prompt or subcommand is passed; if the entrypoint already includes a subcommand, it should be omitted from the arguments to avoid running it twice.

— Exact command form and a concrete migration caveat v0.53.0

06 **Untagged and digest-only image listing** NEW

60

`sbx template ls --all --digests --no-trunc` lists untagged and digest-only local images, with full image IDs preserved in JSON and YAML output.

Find and remove dangling or digest-only local sandbox images that do not appear in a plain listing.

```
$ sbx template ls --all --digests --no-trunc
```

— Runnable command with all flags named v0.53.0

07 **Secret reference provider execution changes** BREAKING

60

Built-in secret reference providers now run host tools from a temporary directory and ignore empty or relative `PATH` entries; provider tools should be placed in absolute `PATH` directories and asdf or mise versions configured globally.

— Mechanism and remediation steps named, no exact command v0.53.0

THINNER COVERAGE BELOW

08 **Named sandbox selection for MCP registrations** NEW

55

When MCP registration names are shared across sandboxes, a sandbox-owned registration can be selected by name with `--sandbox` on `sbx mcp inspect` or `sbx mcp rm --registration-only`.

— Exact flags and subcommands named, mechanism thin v0.53.0

09 **Flags removed from sbx exec and env exec** BREAKING

55

The `--detach` and `--privileged` flags have been removed from `sbx exec` and `sbx env exec`; local exec commands also now reject `--detach-keys` instead of silently ignoring it.

— Exact flags and commands named, no replacement given v0.53.0

10 **sbx logout requires confirmation** BREAKING 50

`sbx logout` now requires confirmation even without a running daemon; scripts should pass `--yes` or `-y` to proceed non-interactively.

— Exact flags given to restore scripted behaviour v0.53.0

11 **HTTP policy matching now case-insensitive** BREAKING 50

HTTP policy matching now treats ASCII letters in URL paths and query parameter names and values case-insensitively; rules relying on case to distinguish allowed and denied requests should be reviewed.

— Behaviour change named, no example of affected rule v0.53.0

12 **Directory copy improvements in sbx cp** IMPROVED 45

`sbx cp` now supports copying a directory to a new path ending in `/` and reports a missing destination directory before uploading, rather than failing partway through.

— Names the command and the two behaviour changes v0.53.0

13 **sbx inspect discoverability and output fields** IMPROVED 45

`sbx inspect` is now listed in help and shell completion, and local JSON and YAML inspection output now includes the sandbox ID and creation time.

— Names command and new output fields, thin on mechanism v0.53.0

14 **Built-in gemini agent retired** BREAKING 45

The built-in `gemini` agent name is retired; existing Gemini sandboxes keep running but can no longer be moved with `sbx move`.

— Clear before/after but no migration path given v0.53.0

15 **Shell completion and command alias expansion** IMPROVED 40

Shell completion now covers pull policies, timeout actions, capture modes, shared skills modes, and move destinations. MCP listings accept `list` as an alias for `ls`, SSH setup

removal accepts `rm` as an alias for `remove`, and policy checks and MCP authorization accept `-v` as shorthand for `--verbose`.

— Several named aliases and completions, no single runnable example `v0.53.0`

└──▷ **BREAKING ON UPGRADE**

! The built-in `gemini` agent name is retired; existing Gemini sandboxes keep running but cannot be moved with `sbx move`.

! `sbx run <agent> -- <args>` now preserves the image's entrypoint arguments when you pass a prompt or subcommand; if the entrypoint already includes a subcommand, omit it from your arguments to avoid running it twice.

! Built-in secret reference providers now run host tools from a temporary directory and ignore empty or relative `PATH` entries; put provider tools in absolute `PATH` directories and configure asdf or mise versions globally.

! `DOCKER_SANDBOXES_USE_DHI` is no longer read; use `sbx settings set platform.images.useDHI true` to select Docker Hardened Image templates.

! `sbx logout` now requires confirmation even without a running daemon; use `--yes` or `-y` in scripts.

! The `--detach` and `--privileged` flags have been removed from `sbx exec` and `sbx env exec`; local exec commands also reject `--detach-keys` instead of ignoring it.

! HTTP policy matching now treats ASCII letters in URL paths and query parameter names and values case-insensitively; review rules that rely on case to distinguish allowed and denied requests.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

h5i

Feedback ⓘ v0.5.5 NOTES ↗ RANK

H5i is an Axum-based Rust web framework with Lean 4-verifiable application kernels that run on servers, bare-metal embedded devices, and GPUs, plus integrated security testing.

> h5i's v0.5.5 release adds a built-in proxy for capturing and auditing browser sessions, installable websec and recon plugins for request replay and target reconnaissance, a command for running attack-flow security regression tests, and a local activity dashboard.

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

01 Browser session capture via built-in proxy NEW 80

Run browser sessions through a built-in proxy with `h5i browser open <url>`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` — every request is policy-checked and logged before it moves.

Inspect every HTTP request a browser session made — and what was refused — after an automated audit run.

```
$ h5i browser open https://staging.example.com
h5i browser audit
h5i browser requests
```

— Names four commands and the policy-check/log mechanism. v0.5.5

02 Websec HTTP workbench plugin NEW 80

Install the websec plugin with `h5i plugin install websec --from ./h5i-websec` to get an HTTP workbench that replays and modifies captured requests, e.g. `h5i websec replay req_42 --set query.id=456`.

Replay a captured request with a modified parameter to test how the target handles tampered input.

```
$ h5i websec replay req_42 --set query.id=456
```

— Install and replay commands with a flag example are fully runnable. v0.5.5

03 Attack-flow security regression testing NEW 72

Run `h5i test <path> --target <url>` to execute attack-flow tests against a target, with every request allowlist-checked and logged before any bytes move.

— Runnable command with allowlist/logging mechanism described. v0.5.5

THINNER COVERAGE BELOW

04 **Recon plugin for target enumeration** NEW 55

Install the recon plugin with `h5i plugin install recon --from ./h5i-recon` to record what a target exposes, tracking each candidate URL to the stored message that confirmed it.

— Install command given but no further subcommands named. v0.5.5

05 **Local activity dashboard** NEW 35

Launch a local dashboard with `h5i ui` to view browser and session activity.

— Command named but scope of what's shown is vague. v0.5.5

WAS THIS USEFUL? 👍 👎

tenuo Feedback ⓘ v0.3.7 NOTES ↗ RANK

> Tenuo's v0.3.7 release adds optional argument support with presence-aware constraints, a Python preflight attenuation checker, multi-party approval builders, and OCI-based Helm chart distribution, alongside breaking changes to type coercion in constraint matching.

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

01 **Authorizer Helm chart via GHCR OCI artifact** NEW 90

The authorizer Helm chart can now be installed directly from GHCR as an OCI artifact using `helm upgrade --install tenuo-authorizer oci://ghcr.io/tenuo-ai/charts/tenuo-authorizer --version 0.3.7`; an immutable digest `sha256:c6dcd5a07178d2161f1397b947796e54b025e69eb43520514a7cfe57e2a68cba` is attached and recorded in `chart-reference.txt`.

Deploy the authorizer from the immutable OCI Helm chart on GHCR without a separate chart repository.

```
$ helm upgrade --install tenuo-authorizer oci://ghcr.io/tenuo-ai/charts/tenuo-authorizer --version 0.3.7
```

— Exact runnable command, version, and immutable digest given v0.3.7

02 **Optional arguments in constraint engine** NEW 87

Arguments can now be omitted entirely while still enforcing their value constraint when supplied, via `ConstraintSet::insert_optional` in the Rust constraint engine; diagnostics, signed CBOR warrants, and attenuation all preserve this presence rule, and a required argument cannot become optional during delegation. Warrants using this optional argument metadata require compatible verifiers — upgrade receivers before issuing them, since strict older decoders reject the metadata and older permissive decoders treat the arguments as required.

Mark a tool argument as optional so callers may omit it entirely, but enforce the value constraint whenever they do supply it.

```rust
use tenuo::ConstraintSet;

let mut cs = ConstraintSet::new();
cs.insert_optional("replicas", constraint);
```

— Names exact method, behaviour, and upgrade/compatibility caveat v0.3.7

03 **Breaking change to Exact/OneOf type coercion** BREAKING

85

`Annotated` constraint checks and the A2A server no longer convert an argument to a string before comparing with `Exact` or `OneOf`; a policy using a number as a string such as Exact("8080") or OneOf(["1", "2"]) no longer matches an integer argument and the call is denied — write the typed literal instead, e.g. Exact(8080), OneOf([1, 2]). A2A grants now preserve the JSON types of `OneOf` and `NotOneOf` instead of coercing to strings, and `Exact`, `OneOf`, and `NotOneOf` now reject `bytes` and `bytearray` literals.

— Gives exact before/after examples and migration fix, no runnable snippet v0.3.7

04 **Python validate_attenuation() preflight check** NEW

75

The new Python validate_attenuation(parent, child) function checks cross-type constraint narrowing before issuing a warrant, raising typed exceptions on mismatch, and uses the same Rust core as grants and authorization so results match actual enforcement.

Preflight-check that a child warrant's constraints are a valid narrowing of a parent's before issuing, to catch policy errors early without a full round-trip.

```python
from tenuo import validate_attenuation

validate_attenuation(parent_warrant, child_warrant)
```

— Has runnable example but limited mechanism detail v0.3.7

THINNER COVERAGE BELOW

05 **Multi-party approval requirements in builders** NEW

59

Python grant and native attenuation builders gain add_approvers() and raise_min_approvals() to add multi-party approval requirements, backed by the existing monotone Rust operations.

— Names the two methods but no usage example given v0.3.7

└──▷ **BREAKING ON UPGRADE**

! `Annotated` constraint checks and the A2A server no longer convert an argument to a string before comparing with `Exact` or `OneOf`; a policy using a number as a string such as Exact("8080") or OneOf(["1", "2"]) no longer matches an integer argument and the call is denied — write the typed literal instead: Exact(8080), OneOf([1, 2]).

! A2A grants now keep the JSON types of `OneOf` and `NotOneOf` values instead of converting them to strings; policies that relied on string coercion will behave differently.

! `Exact`, `OneOf`, and `NotOneOf` now reject `bytes` and `bytearray` literals.

! Warrants using optional argument metadata require compatible verifiers; upgrade receivers before issuing these warrants — strict older decoders reject the metadata, and older permissive decoders treat the arguments as required.

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal

Feedback ⓘ

release-de6a590f 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 biggest addition this window is the Box Egress Proxy (bep), a certificate-issuing, hash-chained-audit egress gatekeeper for agent boxes, alongside new secret storage, per-run network/egress flags, GitHub sign-in, box audit-trail tooling, and a git credential helper plus port proxying inside grant boxes.

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

01 **Box audit trail tooling** NEW 90

`min box audit` reads a box's hash-chained audit trail, merging parent chains with `--parent`; `min doctor` reports an edited chain, and `min box spec` lists each store reference with its authorities.

Inspect a box's hash-chained audit trail and merge its parent chain for a full provenance view.

```
$ min box audit --parent my-box
```

Review the full audit trail for a box, merging its parent chain, to verify no egress decisions were tampered with.

```
$ min box audit --parent my-box
```

Review the full, hash-chained audit trail for a box — including entries from its parent — to verify no egress decision was tampered with.

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

Inspect the full hash-chained audit trail for a box, including its parent chain, to verify no decisions were tampered with.

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

Inspect the full audit trail for a box, merging its parent chain, to verify no proxy decisions were tampered with.

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

— Multiple exact commands with repeated runnable examples across releases

release-de6a590f

02 **Box Egress Proxy (bep)** NEW 85

minimal adds the Box Egress Proxy (`bep`), which issues leaf certificates under a name-constrained root and signing CA, keeps hash-chained audit records, and supports signing CA rotation (keys held in the Secure Enclave inside `bep.app` , rotated through the control socket). `bep serve` bounds its client-key registry with `--max-client-keys` , wires redemption decisions and declared egress, and injects store handles; store rules under `min --config-dir` are kept in step with `bep serve` . A signed, notarized `bep` ships beside `minvmd` on macOS, and `minvmd` starts the Box Egress Proxy beside `gvproxy` so `git` and `gh` reach GitHub from inside the VM.

— Names flags, certs and config paths but no runnable example given release-d503b257

release-de6a590f

03 **Vulnerability reporting command** IMPROVED 75

Report a suspected vulnerability with `min diag --security` , later carried forward as `min bug --security` ; the report now attaches the local telemetry spool and redacts telemetry settings.

File a suspected security vulnerability report to Minimal from the CLI, attaching the local telemetry spool for context.

```
$ min diag --security
```

File a suspected vulnerability in the tool itself for responsible disclosure.

```
$ min bug --security
```

— Two exact command forms shown with examples, redaction mechanism unspecified

release-6bddc077 release-620a1ce5 release-de6a590f

04 **Per-run network and egress flags** NEW 70

`min task run` and `min box run` take `--network` and `--egress` flags to control per-invocation network and egress policy.

— Exact flags and commands named, no deeper mechanism release-de6a590f

05 **GitHub sign-in via min login** NEW 70

`min login` is the GitHub sign-in entry point, completing sign-in on a real page with repository choice available; logging out deletes the sign-in and refuses every reference to it.

Authenticate with GitHub under the published App so grant boxes can reach GitHub without a manual PAT.

```
$ min login
```

— Exact command and flow described, underlying mechanism not detailed release-de6a590f release-6bddc077

06 **OpenTelemetry tracing stack (mlog)** NEW 65

An opt-in OpenTelemetry export stack (`mlog`) adds switches and a local spool, carrying one trace per request across `min` and `minimald` , spanning into `minvmd` VM guests.

— Named component and scope, no exact switch names given release-d503b257 release-de6a590f

07 **Secret storage via min secret** NEW 60

`min secret` stores, lists, and removes secrets, including `min secret set` with fixed-token presets.

— Exact command and subcommand named, mechanism thin release-de6a590f

08 **Grant box git credential helper and port proxying** NEW 60

Grant boxes ship a git credential helper so agents can pull and push without a manually pasted PAT, and `minimald` proxies runtime-published ports like a direct connection, steering credentialed names to the proxy and failing closed while the proxy is down.

— Mechanism named but no flags or endpoints given release-0bb305b9 release-620a1ce5 release-d503b257

THINNER COVERAGE BELOW

09 **Signed notarized minzoned app bundle** NEW 40

Ships `minzoned` in a signed, notarized macOS app bundle with an app icon.

— Names binary and packaging, no further behavior release-620a1ce5

10 **QUIC block helper naming in finalize-install** IMPROVED

35

Names the QUIC block helper in `min finalize-install`.

— Bare naming with no behavior described release-de6a590f

11 **Daily update check** NEW 35

Checks once a day whether the channel has a newer release and reports it via `min`.

— No specific command or flag named beyond min release-d503b257

12 **Box re-creation keeps same name, new id** IMPROVED 30

A box re-created under its name now gets its own id via `min box` re-creation.

— Thin, mechanism and trigger unclear release-de6a590f

WAS THIS USEFUL? 👍 👎

◆ AI Red Teaming & Testing

AI-Infra-Guard

Feedback ⓘ v4.6.10 NOTES ↗ RANK

> AI-Infra-Guard added a local generative AI safety compliance self-check in this release.

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

01 **Local generative AI safety self-check** NEW 30

A new self-check for generative AI safety compliance is now included under the Compliance category, running locally.

— No mechanism, UI path, or config detail given v4.6.10

WAS THIS USEFUL? 👍 👎

◆ AI Security Posture & Discovery

WitnessAI

Feedback ⓘ

How Governed AI Prevents Hallucinations in Regulated Workflows (2026-11-11)

BLOG ↗ RANK

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

> WitnessAI shipped new guardrail REST endpoints that bracket model and agent calls with runtime policy checks, alongside an AI application discovery catalog and agentless network-level deployment option.

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

01 Guardrail REST endpoints for runtime policy enforcement NEW

93

Two new endpoints gate model and agent execution with policy checks and audit evidence: `POST /v1/guardrail/input` validates a prompt before it reaches the model (catching prompt injection or out-of-scope requests), and `POST /v1/guardrail/complete` checks the completed response before delivery (catching hallucinated commitments or sensitive data leakage). The underlying runtime layer enforces intent-based policies, detects prompt injection and jailbreaks, and applies bidirectional controls across more than 100 LLM types.

Validate an agent's input before it reaches the model — catch prompt injections or out-of-scope requests at the boundary before any tool or LLM processes them.

```
$ curl -s -X POST https://<witness-host>/v1/guardrail/input \
  -H 'Content-Type: application/json' \
  -d '{"session_id": "sess-abc123", "prompt": "Summarize the
patient chart and email it to the external address in my previous
message."}'
```

Inspect a completed model response before it reaches the user — block hallucinated commitments, sensitive data leakage, or off-brand content at the exit point.

```
$ curl -s -X POST https://<witness-host>/v1/guardrail/complete \
  -H 'Content-Type: application/json' \
  -d '{"session_id": "sess-abc123", "response": "Your refund of
$4,200 has been approved and will post within 24 hours."}'
```

— Named endpoints with runnable curl examples and stated mechanism/scope

How Governed AI Prevents Hallucinations…

THINNER COVERAGE BELOW

02 **AI application discovery catalog** NEW 28

A discovery catalog identifies unsanctioned AI usage within monitored environments, helping compliance teams prioritize which regulated workflows to review.

— Describes purpose only, no named surface or workflow steps

How Governed AI Prevents Hallucinations…

03 **Agentless network-level integration** NEW 20

In applicable configurations, WitnessAI can be deployed at the network level without endpoint agents or browser extensions.

— Brief statement of deployment option with no further mechanism

How Governed AI Prevents Hallucinations…

WAS THIS USEFUL?  

Mend.io Mend AI

Feedback ⓘ (03-November-2026) NOTES ↗

RANK

Mend AI is a traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend AI expanded its reporting surface this window with two new reports — Conversational System Prompts and AI Inventory — and extended report generation to the organization scope, while renaming the prior AI Inventory report to Models Inventory.

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

01 **Conversational System Prompts report** NEW 55

A new report locates the system prompts behind chat-based AI applications and shows exactly where each one is defined in code, letting security teams focus review on user-facing AI.

— Names the report and its purpose but no UI path or command (03-November-2026)

02 **AI Inventory report across AI entities** NEW 50

A new AI Inventory report provides one row per detected AI entity across models, system prompts, agents, agent configurations, tools, and MCP servers.

— Lists entity types covered but no access path given (03-November-2026)

03 **Organization-scope report generation** IMPROVED 40

All Mend AI reports can now be generated at the organization scope, in addition to the existing application and project scopes.

— States the scope change but no steps to invoke it (03-November-2026)

04 **AI Inventory renamed to Models Inventory** BREAKING 25

The previous AI Inventory report, which covered AI models only, is now named Models Inventory.

— Simple rename notice with no migration detail (03-November-2026)

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

API Shield - Prioritize Web Assets operations with metric views (2026-11-11)

BLOG ↗ RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

> Cloudflare's API Shield adds metric-focused views to the Web Assets Operations page, making it easier to prioritize API operations in large inventories by traffic, errors, latency, response size, or security.

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

01 **Metric views for Web Assets Operations page** NEW 55

The API Shield `Web Assets > Operations` page now includes metric-focused views for traffic, errors, latency, response size, and security, each with server-side sorting to help prioritize operations in large inventories.

Identify the slowest or most error-prone API operations across a large inventory by switching to a metric view on the Operations page.

📍 In the Cloudflare dashboard, go to API Shield › Web Assets › Operations, then select a metric view (Traffic, Errors, Latency, Response Size, or Security) and click a column header to sort server-side.

— Names exact UI path and metric views but no limits or mechanism detail

API Shield - Prioritize Web Assets oper…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 12, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue