THE CYBER TOOLCHAIN

NO. 134

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 8, 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.

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 33 tools

OFFENSIVE SECURITY

vim ? help

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ updated · Nov 8, 2026 NOTES ↗

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

Hacktron adds finding IDs, activity timelines, Jira/Linear status sync, Whitebox retest wizards, and Azure DevOps Remediation across this snapshot.

└──▷ HOW TO FIND IT

Keep bot PRs like Dependabot in scope for review while excluding a specific bot author by name.

> 📍 In your repository configuration, add the bot's handle to `exclude.authors` to prevent it from consuming a seat while all other bot pull requests remain eligible for review.

Let Hacktron open and track fix pull requests on an Azure DevOps repository as part of an automated Remediation workflow.

> 📍 In the Hacktron web app, go to the finding detail for an Azure DevOps-hosted repository and start Remediation — Hacktron will open and track the fix pull request on Azure DevOps as it does on GitHub and GitLab.

Filter the findings list to specific findings by pasting two or more finding IDs into the search box, with `Source`, `Reopened`, and `Date range` filters each saved in the URL.

Finding status now syncs with Jira and Linear — admins enable or disable each sync direction separately in the integration settings, Jira updates follow your project's workflow transitions, and triaging a finding in the web app, Slack, a pull request comment, or the API moves the linked Linear issue to the matching workflow state.

Every member can now create, rotate, and revoke their own API keys directly from the unified Personal settings page, without admin involvement.

Bot pull requests (such as Dependabot and Renovate) are now reviewed on GitHub; use `exclude.authors` to exclude specific bots.

Request a retest of a completed Whitebox pentest from the scans list or scan detail page; the retest covers every reported finding except false positives and accepted risk.

+5 more

WAS THIS USEFUL? 👍 👎

xalgorix

Feedback ⓘ v4.6.183 NOTES ↗ LATEST STABLE

Xalgorix autonomous AI pentesting agents — real-time reconnaissance, vulnerability detection, and exploitation orchestration. Go + TypeScript.

xalgorix now defaults `XALGORIX_CONVERGE_FINISH` and `XALGORIX_CONVERGE_RECONCILE` to `true`, automatically closing scans once coverage obligations are met

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v4.6.183 https://github.com/xalgorix/xalgorix.git
# already have the repo? check out this version:
$ git checkout v4.6.183
```

└──▷ **TRY IT**

Restore the old indefinite-scan behaviour in a CI environment where you want the scan to run until manually stopped.

```
$ XALGORIX_CONVERGE_FINISH=false XALGORIX_CONVERGE_RECONCILE=false
xalgorix --web
```

Tune the idle-close window and lead-verification window so long-running scans against large targets wait longer before concluding.

```
$ XALGORIX_CONVERGE_FINISH=true XALGORIX_CONVERGE_FINISH_IDLE=1000
XALGORIX_CONVERGE_FINISH_LEAD_WINDOW=200 xalgorix --web
```

Scans now close automatically after `XALGORIX_CONVERGE_FINISH_IDLE` iterations (default 500) of no new validated findings once the coverage threshold is cleared, with high-confidence open leads getting a `XALGORIX_CONVERGE_FINISH_LEAD_WINDOW-iteration` (default 100) verification window before closure; set `XALGORIX_CONVERGE_FINISH=false` to restore the previous behaviour.

Content-discovery artifacts such as VCS metadata, config and secret files, logs, and OS cruft now owe only information-exposure checks rather than every injection class.

NEW INTEGRATION NVIDIA NeMo Switchyard auto-router strategy

└──▷ **BREAKING ON UPGRADE**

! `XALGORIX_CONVERGE_FINISH` now defaults to `true`; scans that previously ran indefinitely past their coverage threshold will now terminate automatically after idle iterations.

! `XALGORIX_CONVERGE_RECONCILE` now defaults to `true`; exploratory plan tasks without a vulnerability class that previously kept a fully covered scan open indefinitely will now close once all measurable obligations are done.

▾ **3 more releases in this issue** · 2026-11-05 → 2026-11-07

announced · Nov 6, 2026 BLOG ↗

xalgorix adds NVIDIA NeMo Switchyard as a tool-trajectory-aware router strategy, configurable via `switchyard_router_config` YAML, with per-session state, auditable spend logs, and custom tool-class vocabularies.

Configure Switchyard routing via `switchyard_router_config` in `model_list` YAML, specifying `capable` and `efficient` model names, with an optional LLM or decision-model judge for inconclusive turns.

Connect Claude Code to a Switchyard router by setting `ANTHROPIC_BASE_URL`, `ANTHROPIC_AUTH_TOKEN`, and `ANTHROPIC_MODEL` to the router name, with no client-side code changes required.

xalgorix now supports NVIDIA NeMo Switchyard as an `auto_router/switchyard` strategy that routes each agent turn based on the session's rolling tool-call history — escalating to a capable model on failures and stepping back down on recovery — rather than on prompt content alone.

Custom tool vocabularies let internal or MCP tools be declared as `observe`, `mutate`, `plan`, or `new` to contribute to Switchyard's trajectory signals alongside built-in classifications.

Switchyard routing decisions are recorded in spend logs with the selected model and evidence, enabling cost and quality analysis per session.

announced · Nov 6, 2026 BLOG ↗

xalgorix adds `-latest` model aliases so apps automatically track the newest model version without code changes

└──▷ **USE IT**

Call the latest OpenAI model in a family without pinning a version — xalgorix will route to the current best model behind the alias.

```python
from litellm import completion

response = completion(
    model="openai/gpt-luna-latest",
    messages=[{"role": "user", "content": "Say hello."}],
)
```

Use auto-updating `-latest` model aliases — such as `openai/gpt-luna-latest`, `anthropic/claude-sonnet-latest`, or `deepseek/deepseek-pro-latest` — in the Python SDK and AI Gateway so the model behind the name updates without touching application code.

Supported `-latest` aliases span OpenAI (`gpt-luna-latest`, `gpt-sol-latest`, `gpt-terra-latest`, `gpt-astra-latest`, `gpt-mini-latest`), Anthropic (`claude-opus-latest`, `claude-sonnet-latest`, `claude-haiku-latest`, `claude-fable-latest`), xAI (`grok-latest`), Moonshot (`kimi-latest`), Z.AI (`glm-latest`, `glm-flash-latest`), DeepSeek (`deepseek-pro-latest`, `deepseek-flash-latest`), DashScope (`qwen-max-latest`, `qwen-flash-latest`), and MiniMax (`minimax-latest`).

announced · Nov 5, 2026 BLOG ↗

xalgorix adds model deprecation alerts, scheduled email reminders, and per-key usage tracking to help prevent outages during model migrations

Enable model deprecation email reminders by turning on Model Deprecations under Logging & Alerts → Email Alerts → Email Notifications in the UI; xalgorix schedules two reminders at 30 days and 1 day before the provider's deprecation date.

Review per-key traffic during a migration using the Review impact view, which shows a daily requests chart broken down by key alongside affected keys, teams, owners, and tags.

View model deprecation alerts — including provider, model, deprecation date, requests in the last 30 days, and last-seen traffic — under Observability → Alerts → Model deprecations in the UI.

Mark a migration complete with Mark resolved so the alert moves to Resolved; new requests to the old model automatically reopen it.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

ZeroPath

Feedback ⓘ changelog · Nov 8, 2026 CHANGELOG ↗

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

ZeroPath adds directory team sync, secret exposure snapshots, Gerrit CI scanning, restricted service keys, and new CLI commands for contexts, rules, and tags

└──▷ TRY IT

List all scan contexts as JSON to pipe into a downstream script or audit report.

```
$ zeropath contexts list --json
```

Start a Gerrit change scan from CI on self-hosted deployments by calling `scans.scanGerritChange`; switch each server between `Discovery: Polling` and `Discovery: CI-triggered` in Settings.

Manage scan context, natural-language rules, and repository tags with the new CLI commands `zeropath contexts`, `zeropath rules`, and `zeropath tags`, each supporting list, get, create, update, and delete subcommands plus a `--json` flag for scripting.

Retrieve a secret finding's exposure snapshot — including who can read the repository on GitHub or GitLab and which ZeroPath users can open the finding — via the new `issues.exposureSnapshot` API.

Members with API Token Create can now create restricted service keys scoped to selected repositories or all repositories, edit a key's permissions with Edit permissions without rotating it, and only organization admins can grant Full access.

Map identity-provider directory groups to ZeroPath teams under Settings by enabling Enable directory team mappings once a directory is connected.

+5 more

WAS THIS USEFUL? 👍 👎

Mend.io Mend Platform

Feedback ⓘ changelog · Nov 8, 2026

CHANGELOG ↗

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

Mend Platform adds an AI assistant for natural-language risk queries and granular scan-category controls in Scan Profiles

└──▷ HOW TO FIND IT

Reduce false positives from noisy scan categories by excluding File Read and Database Read from your Scan Profile before the next scheduled scan.

> 📍 In the Mend Platform, go to Settings › Scan Profile, expand the low-probability scan categories tree, uncheck File Read and Database Read, then save — the count indicator confirms how many categories remain active. Changes apply on the next scan.

Mend Platform now lets users selectively include or exclude specific low-probability scan categories — such as File Read or Database Read — within the Scan Profile, replacing the previous binary toggle with a configurable selection tree that shows a count of selected categories.

Mend Platform adds the Mend Intelligent Agent, an AI assistant that accepts natural-language questions to explore platform capabilities, navigate to relevant views, and generate executive-level risk summaries across all applications and projects; enabled by default for organizations with a signed AI consent addendum.

WAS THIS USEFUL? 👍 👎

Clover Security

Feedback ⓘ announced · Nov 4, 2026 BLOG ↗

Clover Security automates security design reviews and threat modelling from specs, flags design-to-implementation drift, and feeds security context to coding agents.

Clover Security adds a Notion Custom Agent template that runs a security review inline on a spec page in about a minute



└──▷ **HOW TO FIND IT**

Run a security review on a spec before any coding agent turns it into code, without leaving Notion.

📍 1. Add the `Clover Spec Security Review` template to your Notion workspace and connect it to your Clover account. 2. Open your spec page in Notion. 3. Type @Clover Spec Security Review on the page. 4. Clover reads the design and writes a `Clover Security Notes` block inline on the page.

Adds `Clover Spec Security Review`, a Notion Custom Agent template that lets you trigger a security review by typing `@Clover Spec Security Review` on any spec page, writing a Clover Security Notes block inline.

Security notes written into the Notion spec are read by coding agents alongside the design, so security requirements shape generated code from the first line.

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗

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

Semgrep adds self-serve Guardian trials, Codex and PowerShell support, Detect-2Shell workflow, Bun lockfile parsing, and AWS CodeArtifact private registry resolution

.

└──▷ **HOW TO FIND IT**

Trigger the Detect-2Shell workflow to find all Node.js and PHP runtime versions across your repositories and identify vulnerable or unpinned ones.

> 📍 In the Semgrep platform, go to Agentic Workflows, select the `Detect-2Shell` workflow, choose your target projects, and click Run.

Allow Supply Chain to resolve private Maven or Gradle packages from AWS CodeArtifact when scanning Java projects with transitive dependency tracking.

> 📍 In the Semgrep platform, go to Settings › Supply Chain › Private Registries, add an AWS CodeArtifact registry entry, and configure the IAM role ARN for authentication.

Semgrep Guardian now supports Codex with automatic code, dependency, and secret scanning — no Semgrep CLI install needed.

Guardian hooks now run on PowerShell in addition to Bash, covering developers on Windows.

Guardian now alerts when a compromised version of the Claude Code extension for Cursor containing an SSH backdoor is detected.

Guardian updates itself automatically on new releases, with each download verified to SLSA Level 3.

In Semgrep Agentic Workflows, the new `Detect-2Shell` workflow finds Node.js and PHP versions across Docker images, `.nvmrc` files, infrastructure as code, and CI — marking each as vulnerable, unpinned, fixed, or unknown and using AI to resolve versions set in variables or base images — and Agentic Workflows can now be scheduled to run daily, weekly, or monthly, requiring Unified Policies, the default for new organizations.

+5 more

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

dep-scan

Feedback ⓘ v6.3.1 NOTES ↗

OWASP dep-scan is a next-generation security and risk audit tool based on known vulnerabilities, advisories, and license limitations for project dependencies. Both local repositories and container images are supported as the input, and the tool is ideal for integration.

dep-scan adds `--fail-on-error` to abort scans on failures and `DEPSCAN_SERVER_TRUSTED_HOST_NAMES` to allowlist server hostnames

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/owasp-dep-scan/dep-scan:v6.3.1
```

pipx

```
$ pipx install owasp-depscan==6.3.1
# already installed? upgrade to this version:
$ pipx install --force owasp-depscan==6.3.1
```

git

```
$ git clone --branch v6.3.1 https://github.com/owasp-dep-scan/dep-scan.git
# already have the repo? check out this version:
$ git checkout v6.3.1
```

published under the project's GitHub account · PyPI attestation · checked 8 Nov Wrong command?

└──▷ TRY IT

Fail CI immediately if dep-scan encounters any error during scanning, rather than continuing and producing a partial report.

```
$ depscan --src $PWD --reports-dir $PWD/reports --fail-on-error
```

Run dep-scan in server mode on a host reachable by a non-loopback name, trusting a specific internal hostname without requiring an API key.

```
$ DEPSCAN_SERVER_TRUSTED_HOST_NAMES=scanner.internal depscan --server
```

Abort a scan immediately on failures with the new `--fail-on-error` flag.

Allowlist server hostnames beyond localhost by setting the new `DEPSCAN_SERVER_TRUSTED_HOST_NAMES` environment variable; without an API key, the `Host` header must match localhost, an IP address, or a name in this list.

└─▷ **BREAKING ON UPGRADE**

- ! `--server-allowed-paths` no longer blocks scans that generate a BOM.
- ! On Windows, a custom batch file set through `CDXGEN_CMD` or a similar variable no longer receives arguments containing `&`, `|`, `<`, `>`, `^`, `"`, `%` or `!`.
- ! Without `DEPSCAN_SERVER_API_KEY`, the `Host` header must be `localhost`, an IP address, or a name listed in `DEPSCAN_SERVER_TRUSTED_HOST_NAMES`; a server reached by a non-loopback name or address accepts `path` and `url` only in a JSON body.
- ! A `path` parameter must be an existing file or directory and cannot look like an option, URL or package URL; on Windows, UNC and device paths are refused.
- ! Project types must start with a letter or digit.

WAS THIS USEFUL? 👍 👎

◆ Application Security Posture

Arnica

Feedback ⓘ announced · Nov 5, 2026 BLOG ↗

Arnica is an agentic application security platform that secures AI-generated code with SAST, SCA, IaC, secrets detection, agent governance, and automated remediation.

Arnica adds Agentic Pentesting to test live applications and confirm whether findings are exploitable using code, cloud, runtime, and business context

Arnica now includes Agentic Pentesting, a new capability that tests live applications to confirm or debunk whether SAST, AI SAST, SCA, secrets, and IaC findings are actually exploitable at runtime.

A private connector lets agents test internal apps behind a firewall by making an outbound connection to Arnica — no VPN required, with only allow-listed apps reachable.

Agents work from a context graph built from four sources — source code (SAST, AI SAST, SCA, secrets, IaC), cloud (AWS, Azure, GCP, Kubernetes, containers), runtime (live traffic, DAST), and business context (architecture diagrams, design docs, threat models) — to surface risks invisible to any single scanner.

Playbooks let teams customize the agent's scope and priorities per application, using built-in playbooks or writing their own.

A live browser view lets practitioners watch agent activity in real time and take over at any point during a test run.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Mobile App Security

MobSF Feedback ⓘ CVE-2026-68922 CVE ↗ STABLE

MobSF is an automated mobile application pen-testing, malware analysis and security assessment framework for Android, iOS and Windows apps, with static and dynamic analysis.

`CVE-2026-68922` : Arbitrary File Read via Path Traversal in ZIP Uploads (CVSS 5.5, affects `<= v4.4.6`)

▾ **16 more releases in this issue** · 2024-03-22 → 2026-07-05

> `CVE-2026-68923` `CVE ↗` `STABLE`
>
> `CVE-2026-68923` : Regression: CSRF checks not enforced after Django migration (CVSS 6.5, affects `<= 4.5.0`)

> `CVE-2026-68924` `CVE ↗` `STABLE`
>
> `CVE-2026-68924` : Zip Bomb Denial of Service via Per-File Size Limit Bypass in ZIP/APK Extraction (CVSS 4.9, affects `<= 4.5.0`)

> `CVE-2026-33545` `CVE ↗` `STABLE`
>
> `CVE-2026-33545` : SQL Injection in SQLite Database Viewer utils (CVSS 5.3, affects `<= 4.4.5`)

> `CVE-2026-24490` `CVE ↗` `STABLE`
>
> `CVE-2026-24490` : Stored XSS via Manifest Analysis - Dialer Code Host Field (CVSS 8.1, affects `4.4.4`)

> `CVE-2025-58161` `CVE ↗` `STABLE`
>
> `CVE-2025-58161` : Path Traversal in GET `/download`/<filename> using absolute filenames in MobSF data directory (CVSS 4.3, affects `4.4.0`)

> `CVE-2025-58162` `CVE ↗` `STABLE`
>
> `CVE-2025-58162` : Arbitrary File Write (AR-Slip) via Absolute Path in .a Extraction (CVSS 6.5, affects `4.4.0`)

CVE-2025-46730 CVE ↗ STABLE

`CVE-2025-46730` : Web Server Resource Exhaustion via ZIP of Death Attack (CVSS 6.8, affects `<=4.3.2`)

CVE-2025-46335 CVE ↗ STABLE

`CVE-2025-46335` : Stored Cross Site Scripting (XSS) via malicious SVG Icon Upload (CVSS 5.4, affects `<=4.3.2`)

CVE-2025-31116 CVE ↗ STABLE

`CVE-2025-31116` : SSRF on assetlinks_check with DNS Rebinding (CVSS 4.4, affects `4.3.1`)

CVE-2025-24804 CVE ↗ STABLE

`CVE-2025-24804` : Partial Denial of Service due to strict regex check in iOS report view URL (CVSS 4.3, affects `4.3.0`)

CVE-2025-24805 CVE ↗ STABLE

`CVE-2025-24805` : Local Privilege escalation due to leaked REST API key in web UI (CVSS 5.5, affects `4.3.0`)

CVE-2024-53999 CVE ↗ STABLE

`CVE-2024-53999` : Stored Cross-Site Scripting Vulnerability in Recent Scans Diff or Compare (CVSS 8.1, affects `<=4.2.8`)

CVE-2024-43399 CVE ↗ STABLE

▶ Watch video

`CVE-2024-43399` : Zip Slip Vulnerability in .a Static Library Files (CVSS 8.0, affects `<=4.0.6`)

CVE-2024-41955 CVE ↗ STABLE

`CVE-2024-41955` : Open Redirect in Login Redirect (CVSS 5.2, affects `<=4.0.4`)

CVE-2024-31215 CVE ↗ STABLE

`CVE-2024-31215` : SSRF in firebase database check (CVSS 6.3, affects `<=3.9.7`)

CVE-2024-54000 CVE ↗ STABLE

.

`CVE-2024-54000` : Potential Fix Bypass - [SSRF Vulnerability on assetlinks_check(act_name, well_knowns)] (CVSS 7.5, affects `<=3.9.6`)

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-11-05 CODE ↗

CAPEv2 is a malware Configuration And Payload Extraction

CAPEv2 adds static shellcode analysis for extracted and submitted shellcode payloads

└──▷ **USE IT**

Turn on shellcode analysis for both CAPE-extracted payloads and directly submitted shellcode samples.

conf/processing.conf

```
[static_shellcode_check]
enabled = yes
cape_payloads = yes
submitted_shellcode = yes
```

Enable static shellcode analysis on CAPE-extracted payloads or directly submitted shellcode by setting `enabled`, `cape_payloads`, and `submitted_shellcode` under `[static_shellcode_check]` in the processing config.

CAPEv2 now runs shellcode analysis on extracted and submitted shellcode, surfacing framework family detection via the new `shellcode_utils` integration module.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

hrtng

Feedback ⓘ announced · Nov 4, 2026 BLOG ↗

Hrtng is an IDA Pro plugin with a rich set of features: decryption, deobfuscation, patching, lib code recognition and various pseudocode transformations

hrtng adds the Malware Analysis Add-on for IDA Pro 9.5, bringing automated unpacking, full-system emulation, and sandbox trace overlay into the disassembly workflow

└──▷ **USE IT**

Unpack a malware sample headlessly in an idalib script so that downstream string extraction or code indexing operates on the inner payload rather than the packer stub.

python

```
start_emulator_with_backend("speakeasy")
# IDA will emulate, detect OEP, rebuild IAT, and update the IDB
automatically
# Then use ida_dbg APIs for further stepping or breakpoint control
```

Invoke the full-system Windows emulator headlessly from an idalib script or AI agent with start_emulator_with_backend('speakeasy'), then control it via the existing `ida_dbg` APIs.

Adds automated malware unpacking that integrates Bochs, Sogen, and Speakeasy emulator backends alongside remote physical hardware and virtual machine snapshot backends, rebuilding the IAT and recovering the original entry point (OEP) in-place.

Adds sandbox trace overlay that pulls VMRay analysis archives directly from VMRay's API and displays API calls, parameters, return values, and code coverage inline in the disassembly and pseudocode views.

Adds a 'Sandbox > Select events in this function' context menu action to highlight all sandbox trace rows whose addresses fall within the currently viewed function.

Adds an Apply all API names/types in function context menu action to resolve indirect calls by applying names, parameter types, and return types from the sandbox trace into both the disassembly and pseudocode views.

+2 more

WAS THIS USEFUL? 👍 👎

fickling

Feedback ⓘ v0.1.13 NOTES ↗

Fickling is a Python pickling decompiler and static analyzer

fickling adds detection of private stdlib imports, subinterpreter gadgets, and underscore attributes, and now hooks `pickle._load`, `pickle._loads`, and `pickle._Unpickler`

└──▷ **GET THIS VERSION** ▾

pipx

```
$ pipx install fickling==0.1.13
# already installed? upgrade to this version:
$ pipx install --force fickling==0.1.13
```

git

```
$ git clone --branch v0.1.13
https://github.com/trailofbits/fickling.git
# already have the repo? check out this version:
$ git checkout v0.1.13
```

PyPI attestation · checked 8 Nov Wrong command?

Detect private stdlib imports, including `subinterpreter` modules that can execute arbitrary code, in scanned pickle files.

Extend private import detection to underscore attributes and qualified names in pickle files.

Hook `pickle._load`, `pickle._loads`, and `pickle._Unpickler` so that fickling's safety checks also cover Python-internal pickle entrypoints.

Analyze PyTorch checkpoints with PyTorch's own ZIP reader to reduce risks from parsing differentials between fickling and PyTorch.

Reject inconsistent `FRAME` boundaries in pickle streams that could hide executable opcodes from analysis, and recompute frame lengths when serializing edited pickles.

+2 more

└──▷ **BREAKING ON UPGRADE**

! Python 3.10 is no longer supported; the minimum is now Python 3.11.

! The `max_acceptable_severity` parameter has been removed from `FicklingContextManager`; any code that passes it will break.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

ecapture

Feedback ⓘ v2.10.0 NOTES ↗

Ecapture is a capture SSL/TLS plaintext with eBPF—no MITM proxy or custom CA installation. Supports Linux and Android on x86_64\arm64 and loong64.

ecapture adds QUIC v1/v2 verification, GnuTLS 3.8.10–3.8.13 support, and live-decryption options for HTTP/1.1, HTTP/2, and HTTP/3

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.10.0 https://github.com/gojue/ecapture.git
# already have the repo? check out this version:
$ git checkout v2.10.0
```

└──▷ TRY IT

Capture and reconstruct full HTTP/2 and HTTP/1.1 messages from live TLS traffic on port 443, writing structured events to a file instead of stdout.

```
$ sudo ecapture tls -m text --text-source=pcapng -i eth0 --eventaddr
/tmp/http.events 'tcp port 443'
```

Inspect decrypted MQTT 3.1.1 traffic inside TLS by selecting the MQTT application parser alongside live pcapng decryption.

```
$ sudo ecapture tls -m text --text-source=pcapng --text-protocol mqtt
-i eth0 'tcp port 8883'
```

Capture a Go TLS client that upgrades SMTP to TLS via STARTTLS, saving the raw pcapng for later analysis while streaming events to stdout.

```
$ sudo ecapture gotls -m text --text-source=pcapng --text-tls-upgrade
smtp --elfpath /usr/bin/go-mailer -i eth0 --pcapfile /tmp/raw.pcapng
'tcp port 587'
```

Reconstruct complete HTTP/1.1 and HTTP/2 messages from live captured traffic using `-m text --text-source=pcapng` with the `tls`, `gotls`, or `gnutls` subcommands; output goes to `--eventaddr` (stdout by default).

Filter reconstructed output to raw TLS/QUIC stream bytes with `--text-raw`, or to incremental HTTP/1.x, HTTP/2, HTTP/3, SSE, and HTTP/2 gRPC events with `--text-http-events`.

Limit per-message gzip expansion with `--text-max-decoded-body N` (requires `--text-http-events`; zero keeps content encoding).

Declare a STARTTLS/SSLRequest/CONNECT negotiation with `--text-tls-upgrade smtp|imap|pop3|ldap|postgres|mysql|connect|https-connect`; `https-connect` includes one inner TLS layer.

Parse DTLS 1.2 ECDHE-RSA AES-128/256-GCM records over UDP with `--text-udp-protocol dtls` (default UDP protocol is QUIC).

+2 more

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ announced · Nov 4, 2026 BLOG ↗

Sysdig Secure is a cloud-native security platform that detects threats, monitors runtime activity, and scans containers and Kubernetes environments for vulnerabilities and compliance issues.

Sysdig Secure adds AI Defense, a runtime security and governance layer for AI agents, MCP servers, and local models across workstations, Linux, and Kubernetes

.
.
.
.

Sysdig AI Defense now discovers every AI agent, MCP server, skill, and plugin running across employee workstations, production Linux, Kubernetes, and managed agent platforms — including unapproved and shadow instances — and ties each to the account and person behind it.

Block, Confirm, and Kill process policy responses now intercept an agent's tool call before it runs or end a process the moment it starts, with the agent receiving the policy name and an exception link.

Adds a new Explorer view that expands each identity into the departments, devices, harnesses, MCP servers, models, skills, and threats connected to it, and a new Threats view that shows behavior chains, the sessions behind a detected incident, and the identity the agent acted under.

Session replay now steps through every tool call in a flagged session in order, with each detection naming the rule, policy, identity, harness, and model that fired.

Sysdig AI Defense now rates each person's model tier against the demand of their work by team, flagging light work on expensive models and heavy work on cheap ones.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Buildah

Feedback ⓘ announced · Oct 21, 2026 BLOG ↗

Buildah is a tool that facilitates building OCI images.

Buildah v1.46.0 adds `--umask`, `--include`, `--required-path`, Git-commit-SHA checksums, and `sharing=private` cache mounts

└──▷ TRY IT

Run a container step with a specific umask to tighten file-creation permissions inside the build.

```
$ buildah run --umask 0027 <container> -- install -m 640 config.yaml
/etc/app/config.yaml
```

Copy only selected files from a large directory tree and assert that a critical path was actually present.

```
$ buildah copy --include='src/**' --required-path='src/main.go'
<container> . /app/
```

Give a build step its own private cache directory so concurrent builds do not share or corrupt each other's cache.

```
$ buildah build --layers -f - . <<'EOF'
FROM golang:1.23
RUN --mount=type=cache,target=/go/pkg/mod,sharing=private go mod
download
EOF
```

Add a `--umask` flag to `buildah run` to control the umask for the executed command.

Adds `--include` to `buildah add` and `buildah copy` as the complement to the existing `--exclude` flag, and `--required-path` (which must be used with `--include`) to assert that specific paths were actually copied.

Allow the `--checksum` option in `ADD` Containerfile instructions and in `buildah add`, `buildah build`, and `buildah copy` to accept Git commit SHAs (or a prefix) for Git sources.

Cache mounts now accept `sharing=private`, giving a command its own exclusive cache directory.

Suppress Containerfile step output when `--log-level error` or `--log-level fatal` is passed to `buildah build` or `buildah push`.

+4 more

└─▷ **BREAKING ON UPGRADE**

! An `ADD` Containerfile command from Git sources no longer includes the `.git` directory by default; `--keep-git-dir` or `--keep-git-dir=true` is now required to include it.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ **SIEM**

Databricks Feedback ⓘ

New jobs compute defaults to Auto access mode (2026-11-09) BLOG ↗

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, and machine learning built on Apache Spark.

Databricks now defaults new jobs compute to Auto access mode, switching to Standard unless ML runtime, GPU, or Runtime < 14.3 is selected

└─▷ **BREAKING ON UPGRADE**

! New jobs compute created in the UI defaults to Auto access mode instead of Dedicated; workspaces with the compliance security profile enabled are affected in December 2026. Existing jobs are not changed, and a cluster policy enforcing a different access mode still takes precedence.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Microsoft Security

Feedback ⓘ changelog · Nov 5, 2026

BLOG ↗

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

Microsoft Security adds Dynamic Permissions for agent task-scoped access, Microsoft Execution Containers for local AI agent containment, and On-Premises Baseline Security Mod

.

Dynamic Permissions now gives AI agents authority scoped to their current task, with Dynamic Agent Risk assessing whether the agent's authority aligns with changing behavior and Dynamic Conditional Access narrowing permissions, restricting resource access, or requiring human approval for higher-impact actions.

Microsoft Execution Containers (MXC), now generally available, applies OS-enforced least-privilege containment for AI agents on Windows across files, networks, processes, and user interactions, with additional isolation options for higher-risk scenarios and upcoming integration with Microsoft Intune, Microsoft Entra, and Microsoft Agent 365.

Microsoft Intune Endpoint Privilege Management now supports system settings elevation, allowing standard users to change IPv4, IPv6, and DNS server values on network adapters without broad administrator rights, with administrator-controlled requirements for business justification or Windows authentication before changes are applied.

Secure Now now pairs information about recent threats with focused initiatives across security domains in Microsoft Security Exposure Management, connecting observed attack patterns to foundational security controls teams can prioritize.

Emergency Mitigation Service 2.0 for Exchange Server and Skype for Business Server, planned for December 2026, will enable the service by default and expand the security mitigations it can apply, with extension to SharePoint Server planned for January 2027.

+1 more

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ changelog · Nov 5, 2026

BLOG ↗

Microsoft Defender is Microsoft's security suite that protects devices, identities, and cloud resources against malware, phishing, and other threats.

Microsoft Defender adds a consolidated Defender Experts MDR overview page and local AI agent/MCP discovery on Windows and macOS; enhanced SQL agent now available in Azure Government.

Defender Experts MDR now offers a consolidated overview page with updated navigation and visibility into active and completed hunts, findings, and items requiring follow-up.

NEW DETECTION Microsoft Defender now discovers more than 50 supported local AI agents and their configured model context protocols (MCPs) on Windows and macOS, surfacing agent, device, account, and MCP context together for security teams.

NEW SUPPORTED PLATFORM The enhanced agent for Defender for SQL Servers on machines is now generally available for resource- and subscription-level enablement in Azure Government cloud, using existing SQL infrastructure without requiring the Azure Monitor Agent.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Endpoint Management (UEM/RMM)

NinjaOne

Feedback ⓘ announced · Nov 5, 2026 BLOG ↗

NinjaOne is an improve IT efficiency with an autonomous patch management solution. Automate updates and secure endpoints across your organization with NinjaOne.

NinjaOne 16.0 adds a technician mobile app, real-time device chat for Windows, and early-access self-service OS updates for end users.

NinjaOne Assist now includes a technician-facing interface alongside the existing end-user experience in the same app, letting technicians handle remote support and device management from their phones.

NinjaOne Chat lets technicians initiate real-time chat with the person logged in to a managed Windows device, replacing ticket back-and-forth with in-context messaging.

NEW INTEGRATION Early Access self-service OS updates let end users install policy-approved OS updates from the NinjaOne End User Portal without opening a ticket or waiting for a maintenance window.

NEW SUPPORTED PLATFORM ChromeOS and ChromeOS Flex device support

▾ 2 more releases in this issue · 2026-11-04 → 2026-11-05

announced · Nov 4, 2026 BLOG ↗

NinjaOne Vulnerability Management 16.0 adds consolidated patch scanning, cutting scan time by ~90%, and expands third-party patch coverage to 146 new products

.

Third-party patch catalog adds 146 new products (55 for Windows, 91 for macOS) added from April through September 2026, with an average of 12.3 patches published per day over the last 30 days.

announced · Nov 4, 2026 BLOG ↗

NinjaOne 16.0 adds unified browser policy configuration covering startup pages, bookmarks, HTTPS enforcement, credential storage, notifications, and browser AI agent controls

└──▷ HOW TO FIND IT

Roll out a compliance-focused browser baseline that blocks password saving, disables payment info storage, and enforces HTTPS across a pilot group before expanding org-wide.

📍 In the NinjaOne console, go to Browser Management › Policies, create a new policy, enable HTTPS enforcement, set password storage to disabled, set payment info

storage to disabled, and assign the policy to your pilot device group.

Standardize startup pages and bookmarks for new hires so their browser opens to internal tools with the right links already pinned — no per-machine scripting required.

📍 In the NinjaOne console, go to Browser Management › Policies, configure startup pages and homepage URLs, add your bookmark list with pinned entries, and push the policy to your onboarding device group.

Set an organization-wide policy for built-in browser AI agents before they proliferate, to prevent unmanaged access to page content across managed devices.

📍 In the NinjaOne console, go to Browser Management › Policies, locate the browser AI agent controls section, configure the allowed or restricted state, and apply the policy to all managed devices.

NinjaOne Browser Management now controls startup pages and homepages, letting admins define what opens when the browser launches and where the home button points, all within a single browser policy.

NinjaOne Browser Management now controls bookmarks, allowing admins to manage bookmark toolbar behavior, pin bookmarks, and push custom bookmark lists to users across the organization.

NinjaOne Browser Management now enforces HTTPS connections across managed devices via browser policy, replacing reliance on user awareness.

NinjaOne Browser Management now sets policy on password saving, payment info storage, and cookie handling across all managed devices from a single policy.

NinjaOne Browser Management now controls browser notification permissions, letting admins decide whether sites can prompt users with notifications.

+1 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Privileged Access Management

Palo Alto Networks CyberArk

Feedback ⓘ

announced · Nov 8, 2026 BLOG ↗

CyberArk is a protect your organization with Idira®. Modernize identity governance, automate access reviews, and secure every identity across your hybrid cloud environment.

CyberArk's Idira platform adds X.509 SVIDs, automatic OAuth registration for MCP servers, Google Agent Identity attestation, AWS/Kubernetes node attestation, and Azure management group onboarding for Secrets Hub.

NEW SUPPORTED PLATFORM Secure Workload Access now issues X.509-SVIDs alongside JWT-SVIDs, with trust bundles published in PEM, DER, or SPIFFE bundle format and configurable certificate lifetime per trust domain — enabling mutual TLS and AWS IAM Roles Anywhere access without a separately managed CA.

Two new SWA node attestation methods — `aws_iid` (AWS Instance Identity Document on EC2) and `k8s_jwt` (projected service account token on Kubernetes) — join the existing `x509pop` and `k8s_psat` methods, removing the need for pre-provisioned bootstrap credentials.

NEW INTEGRATION A new `gemini_enterprise` attestor on the SWA Server verifies a Google Agent Identity token and issues a JWT-SVID or X.509-SVID for Gemini Enterprise agents, requiring no SWA Agent on the host and no pre-placed key material.

NEW SUPPORTED PLATFORM Secure AI Agents now automatically registers OAuth apps for MCP servers that support dynamic client registration (DCR), managing credential rotation so agents connect through the Identity Broker without holding MCP server credentials.

Unmanaged agents discovered in the AI agents inventory can now be remediated in a guided flow — registered as a new managed agent (pre-filled from discovery metadata) or linked to an existing managed registration — clearing the risk finding within minutes.

+3 more

NEW SUPPORTED PLATFORM Secrets Rotation Service region expansion to Israel (Tel Aviv)

└──▷ BREAKING ON UPGRADE

! Legacy tenant role-based authentication for Secrets Hub's AWS organization onboarding is deprecated and will be removed in February 2027; redeploying the template through Connect cloud environments migrates to AWS External ID automatically.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

iDenfy

Feedback ⓘ changelog · Nov 8, 2026 CHANGELOG ↗

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

iDenfy: iDenfy's Android SDK now checks CPF (Brazilian taxpayer identification) during Bank Card Verification.

iDenfy's iOS SDK now checks CPF (Brazilian taxpayer identification) during Bank Card Verification.

iDenfy's iOS SDK now supports iPhone Duo during Bank Card Verification.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

OpenSSH

Feedback ⓘ changelog · Nov 8, 2026 CHANGELOG ↗

Portable OpenSSH

OpenSSH 10.6p2 adds the post-quantum hybrid signature scheme ssh-mldsa44-ed25519 as an allowed CA signature algorithm and adds PubkeyAuthOptions max-pk-ok:nnn to cap per-key

OpenSSH now supports `ssh-mldsa44-ed25519` as an allowed CA signature algorithm, enabling post-quantum hybrid signatures for certificate authorities.

WAS THIS USEFUL? 👍 👎

Zscaler

Feedback ⓘ changelog · Nov 4, 2026 BLOG ↗

Zscaler is a cloud-based security platform that provides secure internet access, zero-trust network access, and threat protection for users and applications.

Zscaler's Zero Trust Browser now streams thick-client applications like SAP GUI and developer IDEs through isolated, ephemeral cloud environments, removing the need for VDI i

Zscaler's Zero Trust Browser now supports thick-client application delivery through Application Isolation, running apps like SAP GUI, database utilities, and developer IDEs inside isolated, ephemeral cloud environments rendered directly in a browser tab.

Users access thick-client desktop environments from the Zero Trust Browser landing page alongside existing web and private applications, with no local VDI client, endpoint agent, or customer-managed infrastructure required.

Session persistence is enabled by default, allowing users to disconnect and reconnect to a running session with unsaved work intact.

Copy/paste and file downloads from isolated thick-client sessions are governed by policy, keeping enterprise data inside the cloud environment.

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Veeam Backup & Replication

Feedback ⓘ

changelog · Nov 6, 2026 BLOG ↗

Veeam Backup & Replication is backup, recovery, and replication software for virtual, physical, cloud, and SaaS workloads.

Veeam Backup & Replication v13.1.2 adds automatic best-practice repository settings on Data Domain add, plus Instant VM Recovery direct from deduplicated storage.

UPDATED SUPPORTED PLATFORM NEW COMPLIANCE FRAMEWORK Veeam Backup & Replication now automatically applies recommended best-practice repository settings when you add a Dell Data Domain system as a backup repository, reducing manual configuration.

Veeam Backup & Replication can now drive DD Boost Managed File Replication between Data Domain systems, replicating backup files appliance-to-appliance in deduplicated form for offsite copies while tracking those copies for restore.

UPDATED INTEGRATION Dell Data Domain (DD3410 and all higher current-line models) now holds the Veeam Ready Repository Primary Target certification, enabling Instant VM Recovery to run directly from Data Domain without a landing zone, as of Veeam Backup & Replication v13.1.2.

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

warden

Feedback ⓘ announced · Nov 6, 2026 BLOG ↗

Warden is the secure gateway connecting AI agents to enterprise systems.

Warden adds an agent identity registry with sponsorship, governed-only roles, an OIDC auth method for operators, and a deny mode for unclaimed agents

└──▷ TRY IT

Register a new agent, attach a team sponsor, and place it under governance so it can hold roles through static groups.

```
$ warden identity entity create billing-agent -json '{"description":
"Reconciles supplier invoices against purchase orders", "profile":
{"purpose": "billing", "runtime": "langgraph"}}'
warden identity entity sponsor propose billing-agent -json '{"group":
"billing-team"}'
```

Create a static group granting specific roles to governed agents, then claim the agent's native SPIFFE identity so virtual groups no longer apply to it.

```
$ warden identity group create billing-operators -json '{"type":
"static", "description": "Billing agents governed by Warden",
"members": ["billing-agent"], "roles": ["invoice-reader", "invoice-
matcher"]}'
warden identity alias add -json '{"mount": "auth/spiffe/", "name":
"spiffe://actor/billing-team/agent-1", "to": "billing-agent"}'
```

After registering and sponsoring all known agents and verifying the audit log is quiet, switch the namespace to deny mode so any unclaimed agent is refused before role evaluation.

```
$ warden identity entity list -ungoverned
# Review audit log until quiet, then:
warden identity config write -unknown-aliases=deny
```

Operators can now sign in through their company identity provider using the new `oidc` auth method with `warden login`, holding roles through virtual groups over IdP claims.

Agent discovery now runs under namespace built-in policies and is audited.

Agents can call Microsoft Graph as the signed-in person they act for through Entra on-behalf-of, and Azure apps can authenticate with a certificate instead of a secret.

Amazon Bedrock now works keyless, with no model API key stored anywhere.

Token and credential caches are now sized in real bytes with a configurable budget.

└──▷ **BREAKING ON UPGRADE**

- ! Auth mounts no longer bind roles — any configuration relying on this will break on upgrade.
- ! Agent mounts now require an audience — existing mounts without one will fail on upgrade.
- ! There are twelve breaking changes in total; the upgrade guide must be read before upgrading from v0.21.0.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Vanta

Feedback ⓘ changelog · Nov 8, 2026 CHANGELOG ↗

Vanta is a compliance automation platform that continuously monitors a company's cloud, code, and identity systems to help achieve and maintain security certifications like SOC 2, ISO 27001, and HIPAA.

Vanta adds MCP-based controls management, a native DocuSign integration, rebuilt Jira/Linear/Vercel connectors, and 27 new risk library scenarios

Delete custom controls through the Vanta Agent or MCP, with associated mappings removed under the same authorization rules as the web app.

Through authorized `MCP` `-connected` agents, update common control fields, retrieve common control subdomain information, retrieve common control details and current status, retrieve deactivated control details without changing status (also via the `Vanta Agent`), and deactivate common controls while preserving existing permissions and business rules (Enterprise, Professional packages).

Vanta now exposes shared, permission-aware APIs for creating, viewing, updating, deactivating, and reactivating controls across Agent, MCP, and GraphQL surfaces (Professional package).

Mark audits complete and upload audit reports directly through the Vanta Agent and MCP.

Connect DocuSign directly to Vanta and ingest eligible contracts into Customer Commitments without manual uploads.

+5 more

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v4.2.2 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 MCP findings tools, asset lifecycle tracking, EBIOS RM study duplication, and new compliance frameworks including Swift CSCF v2026/v2027 and SAMA BCM

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.2.2 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v4.2.2
```

CISO Assistant adds MCP tools for findings assessment and writing findings, expanding the MCP surface with read and write operations on findings.

CISO Assistant adds an instance time zone setting that controls timestamps in reports and exports.

CISO Assistant adds an asset lifecycle feature to track assets through defined lifecycle states.

CISO Assistant adds a new cumulative-based outcome model to support Essential Eight maturity logic.

CISO Assistant now lets workflows read requirement assignments and automatically promote their actors to respondents.

+5 more

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ announced · Nov 4, 2026 BLOG ↗

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

Runlayer adds Artifacts (stable versioned links) and Apps (agent-built internal tools with owner-approved actions, a built-in database, scoped MCP access, and an audit log).

Runlayer now publishes agent-generated outputs as Artifacts — stable, versioned links requiring company sign-in that update in place when you ask the agent to revise them, with sharing controlled by people, groups, teams, or roles.

Runlayer now supports Apps — Artifacts with an owner-approved action list, a built-in database, and live access to business systems through existing MCP connectors, with no API keys stored in the app and each call running under the viewer's own permissions.

An Admin view lets app owners review every action before it runs, showing which MCP connectors it calls and whose access it uses; a changed action waits for re-approval while the last approved version continues to serve.

Every call an App makes is written to the Runlayer audit log alongside existing agent and MCP activity, recording which app, which action, which connector, and on whose behalf.

NEW INTEGRATION MCP connector backend for Apps

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

stride-gpt

Feedback ⓘ v0.21.0 NOTES ↗

An AI-powered threat modeling tool that leverages OpenAI's GPT models to generate threat models for a given application based on the STRIDE methodology.

stride-gpt adds per-subsystem `--max-llm-calls` and `--max-tool-calls` caps, service-scoped agent planning, and a DFD editor in the web UI

└──▷ GET THIS VERSION ▾

pipx

```
$ pipx install stride-gpt==0.21.0
# already installed? upgrade to this version:
$ pipx install --force stride-gpt==0.21.0
```

git

```
$ git clone --branch v0.21.0 https://github.com/mrwadams/stride-gpt.git
# already have the repo? check out this version:
$ git checkout v0.21.0
```

PyPI attestation · checked 8 Nov Wrong command?

Caps LLM calls per subsystem with `--max-llm-calls` and tool calls per subsystem with `--max-tool-calls`, and plans subsystems by service while scoping each agent worker to its own files so analysis does not bleed across service boundaries.

Show a generated or parsed data flow diagram in the web UI editor after threat modelling.

Print the checkpoint path and resume command on Ctrl-C so interrupted runs can be restarted without digging through logs.

Improves LLM and tool call reliability by sending `json_object` instead of `json_schema` response format for DeepSeek models so structured calls succeed, and depending on `litellm[mcp]` so fresh installs can run tool calls without a separate MCP installation step.

Improves threat modelling by treating external services as trust-boundary components and covering hosted inference serving in the GenAI reference card, and reporting a threat chain that spans subsystems as an emergent cross-cutting threat rather than a simple aggregation.

+1 more

WAS THIS USEFUL?  

Minimal Software Research minimal

Feedback ⓘ

release-6a01d5ae NOTES ↗ LATEST STABLE

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 adds a Box Egress Proxy, secret management, OpenTelemetry tracing, and a vulnerability reporting flag across CLI and daemon components

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-6a01d5ae
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-6a01d5ae
```

└──▷ TRY IT

Manage a secret token inside a box — store it once, then list to confirm it is registered.

```
$ min secret set MY_API_TOKEN
min secret list
```

Inspect a box's proxy-enforced audit trail and merge its parent chain for a full history.

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

Submit a suspected vulnerability report from the local telemetry spool.

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

Adds `min box audit` to read a box's audit trail, `min box audit --parent` to merge parent chains, `min doctor` to report an edited audit chain, and `min box spec` to list each store reference with its authorities.

Report a suspected vulnerability to the project with `--security` on the `diag` command.

Sign in to GitHub and finish the OAuth flow on a real page with `min login`, with repository choice offered during sign-in.

Adds the `bep` crate implementing a Box Egress Proxy that issues leaf certificates under a name-constrained root and signing CA, holds proxy keys in the host Secure Enclave, maintains a hash-chained audit log per proxy decision, rotates the signing CA through a control socket, starts beside `gvproxy` so `git` and `gh` reach GitHub from inside VM boxes, and wires a client-key bound into `bep serve` via `--max-client-keys` to cap the registry size.

Name the QUIC block helper in `min finalize-install` for post-install network configuration.

+4 more

▾ **8 more releases in this issue** · 2026-10-15 → 2026-11-07

release-1e236070 NOTES ↗ STABLE

minimal: Store, list, and remove secrets inside a box with `min secret` (subcommands: `set`, `list`, `remove`).

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch release-1e236070
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-1e236070
```

└──▷ **TRY IT**

Store a GitHub token as a named secret so a box can retrieve it at runtime without it ever appearing in the toml.

```
$ min secret set GH_TOKEN
```

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

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

Limit how many client keys the Box Egress Proxy will register before refusing new ones, to bound the registry surface.

```
$ bep serve --max-client-keys 50
```

Carry an edited-audit-chain warning into `min doctor`.

Carry telemetry spool and redact telemetry settings automatically when running `min bug`.

release-3c2b2a73 NOTES ↗ STABLE

minimal: Control the client-key registry limit in `bep serve` with `--max-client-keys`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-3c2b2a73
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-3c2b2a73
```

└──▷ TRY IT

Store a secret in a box so an agent can access it without a manually pasted credential.

```
$ min secret set MY_API_KEY
```

Inspect a box's full audit trail, merging in its parent chain, to verify no decisions were tampered with.

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

Review every store reference a box declares along with the authorities bound to each, to audit egress policy before activating a session.

```
$ min box spec
```

Bundle the telemetry spool and redacted telemetry settings in `min bug` reports.

Sign in to GitHub through the published App and manage credentials with `min login`; logout deletes the sign-in and refuses every reference to it.

Ship a git credential helper inside grant boxes so agents can reach GitHub without a manually pasted PAT.

Block UDP port 443 for boxes whose QUIC stance is blocked, enforced both at the host and from a blocked cohort's own leaves.

release-46d9e530 NOTES ↗ STABLE

minimal: Introduces the new `bep` crate shipping the GitHub module and sealed-envelope proxy, with `bep serve` as the runtime entry point for egress control in the Box

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-46d9e530
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-46d9e530
```

└──▷ TRY IT

Store a GitHub token as a named secret in a box so agents can use it without it being pasted into the session environment.

```
$ min secret set GH_TOKEN
```

Review the tamper-evident audit trail for a box, merging its parent's history, to verify every proxy egress decision made during an agent session.

```
$ min box audit --parent <box-name>
```

Issue leaf certificates under a name-constrained root and signing CA; rotate the signing CA under an unchanged root, including rotation through the control socket.

Report non-Rust dependency pins and bump the `locked_commit` for packages in `scripts`.

release-fad7c721 NOTES ↗ STABLE

minimal: Refuse Minimal's own infrastructure domains in the box CA and store rules.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-fad7c721
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-fad7c721
```

└──▷ TRY IT

Store a secret token for a box so the agent can retrieve it without a manual paste.

```
$ min secret set MY_API_KEY
```

Inspect the tamper-evident audit trail for a specific box after a session ends.

```
$ min box audit <box-name>
```

List all store references and their declared authorities for a box to review egress policy.

```
$ min box spec <box-name>
```

Record a revocation when a live box is destroyed, and check minted box IDs against retained revocations.

release-5378e56c NOTES ↗ STABLE

minimal adds opt-in OpenTelemetry tracing across the CLI, daemon, and VM host, plus a new `bep` proxy subsystem with certificate issuance, audit chaining, and egress enforcem

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-5378e56c
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-5378e56c
```

└──▷ TRY IT

Inspect the full proxy decision history for a specific box to verify egress enforcement.

```
$ min box audit
```

Adds the `bep` crate for proxied access control, including a GitHub module and sealed-envelope support, storing and publishing proxy key identities from the host keychain, hash-chained audit records for every proxy decision, leaf certificate issuance under a name-constrained root and signing CA, wiring redemption decisions and each box's declared egress into `bep serve`, refusing undeclared egress, host mismatches, off-port, off-set authorities, and expired values, and

accepting only box connections while terminating TLS for module hosts over HTTP/1.1.

release-d67d78ee NOTES ↗ STABLE

minimal adds opt-in OpenTelemetry tracing across the CLI, daemon, and VM host, plus a signed macOS app bundle for minzoned

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-d67d78ee
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-d67d78ee
```

└──▷ HOW TO FIND IT

Opt out of OpenTelemetry telemetry export by setting the `DO_NOT_TRACK` environment variable before running any min command.

📍 `DO_NOT_TRACK=1` min session activate `--attach` .

minimal now exports OpenTelemetry spans from `mlog` with opt-in switches and a local spool, grounded in `DO_NOT_TRACK` opt-out support.

minimal now traces each request end-to-end with one trace per request across both `min` and `minimald` .

release-6c6c39ae NOTES ↗ STABLE

minimal adds opt-in OpenTelemetry export, `min diag upload` , `min finalize-install` , and staged-version installs via `--version`

└──▷ GET THIS VERSION ▾

```
$ git clone --branch release-6c6c39ae
https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout release-6c6c39ae
```

└──▷ TRY IT

Install a specific staged release of minimal rather than the current stable channel.

```
$ curl --proto "=https" --tlsv1.2 -fsSL
'https://go.minimal.dev/stable' | sh -s -- --version <version>
```

Run the new network finalisation step after install, reviewing the script before it applies.

```
$ min finalize-install --show --script
```

Install a specific staged version of minimal with `--version` on the installer.

Upload diagnostic bundles directly to the portal for analysis with `min diag upload`.

Replace `min net setup` with `min finalize-install` for host DNS and network setup; `--show --script` replaces `--print`.

Warn within a minute when the switch datapath is lost.

End a task run's box daemon-side.

+1 more

└──▷ **BREAKING ON UPGRADE**

- ! The `[harness]` table in `minimal.toml` no longer loads a stack and now fails to load; rename it to `[stack]`.
- ! `--token` is removed from `min bug`, `min diag collect`, and `min diag upload`; passing it is now a parse error. The `Status:` line is no longer printed after an upload.
- ! `min net setup` is removed with no alias; run `min finalize-install` instead (`--show --script` replaces `--print`).
- ! The own-ip deny-all egress default is now enforced for sessions and `minvmd`.

v0.7.0 NOTES ↗ STABLE

minimal adds per-session network isolation with egress rules, named addresses, port publishing, a root-setup command, multi-VM support, and CLI bug reporting

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.7.0 https://github.com/gominimal/minimal.git
# already have the repo? check out this version:
$ git checkout v0.7.0
```

└──▷ **TRY IT**

Activate a session that can only reach github.com and nothing else — useful for an agent that must pull code but must not exfiltrate data.

```
$ min session activate --network own_ip --allow-dns-hosts github.com
--attach
```

Forward port 8080 from a running session to localhost:8080 so you can open a dev server in your browser without installing anything in the session.

```
$ min net forward mysession 8080:8080
```

Upload a diagnostic bundle immediately after a hang so the team can investigate without manual log collection.

```
$ min bug --upload
```

Adds session network control: `min session activate --network <none|host_ip|own_ip>` sets the network mode (`host_ip` shares the host network, `own_ip` gives the session its own IP, `none` removes all network access), `--allow-dns-hosts <host>` , `--allow-subnets <cidr>` , `--deny-subnets <cidr>` , `--allow-protocols <tcp|udp|icmp>` , and `--deny-all-egress` flags restrict what an `own_ip` session can reach, and `min session policy <session>` (also shown in the detail pane of `min dash`) inspects the effective network policy for a running session.

Adds session port management: `--ingress <external>:<internal>` publishes ports when a session starts, `--dynamic-ingress <allow|ask|deny>` and `--dynamic-range <lo-hi>` control runtime port requests, `min net expose <port>` requests a port from inside a session, and `min net forward <session> <local>:<port>` forwards a session port to localhost from outside.

Every session gets a `<name>.min.internal` hostname reachable via a local HTTP proxy, and via full DNS resolution after `min finalize-install` .

Collect and upload a diagnostic bundle in one step with `min bug --upload` , which prints a shareable report link and a delete command; upload a previously collected bundle with `min diag upload <path>` , and remove it later with `min diag delete <id> <token>` (bundles auto-expire after 7 days).

Collect all root-requiring setup into one reviewable `sudo` call with `min finalize-install` ; preview what is missing with `--show` , print the exact script with `--show --script` , and undo the installation with `--undo` .

`+1 more`

└──▷ **BREAKING ON UPGRADE**

- ! `[harness]` in `minimal.toml` no longer loads — rename the table to `[stack]`.
- ! `min ssh-forward` is removed — use `min net forward <session> <local>:<port>` instead.
- ! The daemon's HTTPS reverse proxy and client certificates are removed, and `min login` no longer issues a certificate — reach sessions via `<session>.min.internal` or `min net forward`.
- ! Session names must now be valid hostnames (letters, digits, hyphens only; case-insensitive; `host`, `local`, `localhost`, and `minimald` are reserved; names shaped like a session id are refused) — rename non-conforming sessions with `min session rename`.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

PyPA python-pipx

Feedback ⓘ 1.18.16 NOTES ↗

The python-pipx CLI installs and runs Python applications in isolated environments while exposing their commands on the user's PATH.

pipx's `ensurepath` command now reports accurately when PATH directories are already configured

└──▷ GET THIS VERSION ▾

pipx

```
$ pipx install pipx==1.18.16
# already installed? upgrade to this version:
$ pipx install --force pipx==1.18.16
```

git

```
$ git clone --branch 1.18.16 https://github.com/pypa/pipx.git
# already have the repo? check out this version:
$ git checkout 1.18.16
```

PyPI attestation · checked 8 Nov Wrong command?

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Breachrx

Feedback ⓘ announced · Nov 8, 2026 BLOG ↗

BreachRx is an agentic AI CIRM platform that elevates cyber incident response to a governed enterprise-wide process.

BreachRx adds Slack and Microsoft Teams integrations to the Rex Platform, creating a private incident channel and direct-message action alerts per incident

└──▷ TRY IT

Raise an incident from an ongoing Slack conversation without leaving the channel, so responders can begin formal tracking immediately.

```
$ /breachrx create
```

Connect Slack or Teams to the Rex Platform so private incident channels and action-owner DMs are created automatically for every new incident.

📍 In the Rex Platform, go to Integrations, select Slack or Microsoft Teams, enable automatic channel creation, choose the roles to invite, and set notification rules for action assigned, updated, due soon, and overdue events.

BreachRx now creates a private Slack channel named `breachrx-{identifier}` for each new incident or exercise automatically, or on demand from the incident's menu.

Open a new incident from any Slack or Microsoft Teams channel using the `/breachrx create` slash command, which presents an intake form with incident name, type, severity, and description fields.

BreachRx adds a Microsoft Teams integration driven by the BreachRx Bot, replacing the older webhook-based connection, that creates private per-incident channels and auto-adds invited members to the team before the channel.

Direct messages now notify action owners in Slack when an action is assigned, updated, due soon, or overdue, with configurable notification rules.

Restricted incidents in Slack auto-create channels limited to the Admin role, the default channel member, and the incident creator; in Teams, restricted incidents invite only BreachRx organization administrators.

└──▷ BREAKING ON UPGRADE

! The Microsoft Teams integration replaces the older webhook-based connection with the BreachRx Bot; existing webhook-based Teams setups will need to be reconfigured.

WAS THIS USEFUL? 👍 👎

Arcade

Feedback ⓘ announced · Nov 5, 2026 BLOG ↗

Arcade is the actions runtime for enterprise AI agents â one control point to enforce, execute, and govern every action, every time.

Arcade adds Enterprise MCP Hosting so self-hosted customers can deploy custom MCP servers with the `arcade deploy` command

└──▷ **TRY IT**

Ship a custom MCP server to your self-hosted Arcade environment from a CI pipeline without filing tickets for a host or cluster access.

```
$ arcade deploy
```

Deploy custom MCP servers to self-hosted Arcade environments with the `arcade deploy` command from a local checkout or CI pipeline, without needing cluster access.

Every deployed MCP server automatically runs behind Arcade's authorization, policy enforcement, and audit logging from its first call.

Operators can set deployment limits per organization and per project; when a team hits its limit, the error message states the limit plainly.

Monitor deployed servers through the dashboard's Servers page, which shows versioned releases, status, logs, and restart and removal controls.

Arcade adds Enterprise MCP Hosting support for Python servers built with the Arcade MCP framework, with FastMCP support planned next.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 8, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue