THE CYBER TOOLCHAIN

NO. 028

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 28, 2026 · EVERY WEEKDAY

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

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 38 tools

OFFENSIVE SECURITY

◆ Post-Exploitation

SpecterOps BloodHound ⓘ v9.5.0 NOTES ↗

Six Degrees of Domain Admin

BloodHound v9.5.0 adds OpenGraph data quality stats, webhooks UI, markdown-rich entity panels, and new management operation endpoints.

└──▷ GET THIS VERSION

```
$ git clone --branch v9.5.0
https://github.com/SpecterOps/BloodHound.git
# already have the repo? check out this version:
$ git checkout v9.5.0
```

- Adds `/data-quality-stats` endpoint and aggregations endpoint for querying OpenGraph node and relationship counts by environment.
- Adds `environment_kind_id` and `environment_kind_display_name` fields to the available-domains API response.
- New Webhooks table and webhook details modal on the Delivery Page for managing webhook configurations.
- Renders markdown-based KindInfo for OpenGraph extensions directly in the entity panel.
- Adds `KindInfo` to the `api/v2/nodes/{node_id}` response and to extension upsert operations.

+8 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Veracode ⓘ changelog · Jul 28, 2026 CHANGELOG ↗

Veracode is a cloud-based application security platform that helps developers identify and remediate vulnerabilities in their code through static, dynamic, and software composition analysis.

Veracode adds CLI `--exclude` flag, DAST custom host-to-IP mappings, EU Package Firewall, enhanced Python scanner, and Android 17 / .NET 10 static analysis support.

└──▷ TRY IT

Exclude a legacy JAR or noisy internal package from a Java scan to focus results on application code you own.

```
$ veracode package --exclude com.example.legacy:legacy-lib.jar
```

Configure a custom host-to-IP mapping for a DAST scan targeting an internal staging environment not resolvable via public DNS.

> 📍 In the Veracode Platform, open your DAST scan configuration, navigate to the Custom Host settings, and add a host-to-IP address mapping for your target host.

- Adds `--exclude` flag to Veracode CLI v2.51.1 to prevent specified Java packages or JAR files from being scanned.
- Enables custom host-to-IP address mappings for DAST scans when DNS lookup is unavailable or not desired.
- Enhances Python static scanner with taint tracking, cross-file analysis, and pattern recognition; adds framework support for PySpark, Flask, Django, and FastAPI.
- Adds support for Android 17 in the Static Analysis Engine v2026.07.
- Enhances support for .NET 10 in the Static Analysis Engine v2026.07.

+4 more

WAS THIS USEFUL? 👍 👎

Snyk ⓘ changelog · Jul 28, 2026 DOCS ↗ BLOG ↗

Snyk finds and fixes security vulnerabilities in code, open-source dependencies, and containers throughout the development pipeline.

Snyk DAST gains AI-powered authenticated scan setup and native Bruno collection support for API targets.

└──▷ HOW TO FIND IT

Set up an authenticated DAST scan without manually recording login flows — use the AI-powered configuration to handle login sequences and 2FA in minutes.

> 📍 In the Snyk console, go to your DAST target settings, select 'Automatic Target Authentication Configuration', enter your credentials in the dedicated credential fields, and optionally provide natural-language instructions for complex login scenarios.

Onboard a Bruno-based API collection as a Snyk API Target without converting files to Postman or OpenAPI.

> 📍 In the Snyk console, navigate to API & Web › API Targets › Create Target, select 'Bruno Collection', and upload your collection as a ZIP file or folder.

- Adds Automatic Target Authentication Configuration for DAST: AI-driven setup automatically generates login sequences, handles 2FA, and maps CSS selectors, with natural-language instructions for complex scenarios.
- Supports native Bruno collection imports for API Targets — upload via ZIP or folder without converting to Postman or OpenAPI formats first.
- Syncs Bruno collection environments automatically and warns when variables are missing required secret values during import.

WAS THIS USEFUL? 👍 👎

Portkey ⓘ snapshot-20260728 NOTES ↗

Portkey is an AI gateway and observability platform that helps developers manage, monitor, and optimize LLM API calls with routing, caching, and reliability features.

Portkey March 2026: vault-backed credentials, weekly rate limits, cross-cloud GCS logging, and four new guardrail integrations.

└──▷ USE IT

Apply Zscaler AI Guard detection policies to all LLM inputs and outputs by wiring in your existing `policyId`.

```json
{
  "guardrails": [
    {
      "type": "zscaler",
      "hooks": ["beforeRequestHook", "afterRequestHook"],
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    }
  ]
}
```

- Adds Secret References to pull credentials from AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault at runtime via `secret_mappings`, keeping secrets out of Portkey config.
- Supports weekly budget and usage windows (`rpw`) for rate limits, enabling caps aligned to weekly planning and review cycles.
- Adds endpoint-scoped rate limits so chat completions, embeddings, and admin routes can carry independent limits instead of a single global rule.
- Enables GCP Workload Identity Federation log delivery from AWS-hosted gateways to GCS via `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL`, eliminating long-lived cross-cloud keys.
- Allows organization admins to include archived workspaces in analytics graphs and summaries; saved filters now accept workspace slugs alongside IDs.

+11 more

WAS THIS USEFUL? 👍 👎

Wiz Code ⓘ changelog · Jul 28, 2026 BLOG ↗

Wiz Code is a developer security tool that identifies and remediates code vulnerabilities and misconfigurations in applications.

Wiz introduces Atlas, an autonomous AI SAST system that validates every finding with a working exploit and scans pipelines for Secure by Design compliance.

- Introduces Atlas, an autonomous AI vulnerability research system that uses specialized agents to map attack surfaces, hunt in parallel, and adversarially validate findings — each confirmed with a dynamically generated, working exploit.
- Adds AI SAST at scale: Atlas performs continuous, cost-efficient deep code scanning across entire enterprise codebases using a code property graph (CPG) for call- and data-flow–grounded analysis.
- Adds shift-left tracing that links vulnerabilities discovered in production back to their originating code repository or container image, enabling fixes at the source.
- Extends pipeline scanning to AI application configurations and model deployments, in addition to container images, software packages, and IaC, to enforce Secure by Design guardrails.

WAS THIS USEFUL? 👍 👎

StepSecurity ⓘ

`Secure Registry: Maven Support in Beta, Custom Block List, and Tarball URL Rewriting (2026-07-24)`

BLOG ↗

StepSecurity is a platform that secures CI/CD pipelines by detecting and preventing supply chain attacks, hardening GitHub Actions workflows, and monitoring runtime behavior.

StepSecurity Secure Registry adds Maven (Beta) support, per-ecosystem Custom Block List, and npm Tarball URL Rewriting.

└──▷ HOW TO FIND IT

Enforce an internal ban on all 4.x versions of lodash and an entire scoped namespace across your npm policy.

> 📍 In the Secure Registry console, go to Policy › npm › Custom Block List and add entries: `lodash@4*` and `@rejected-scope/*`, then save.

Route all Maven artifact downloads through Secure Registry for evaluation instead of Maven Central.

> 📍 In the Secure Registry console, go to Setup Guide › Registry selector › Maven, select your integration path (e.g., Direct), then point your Maven client at `https://registry.stepsecurity.io/java` as the mirror URL.

Ensure npm package tarballs download through Secure Registry rather than the public npm registry.

> 📍 In the Secure Registry console, go to npm Registry Settings and enable the Tarball URL Rewriting toggle.

- › Adds Maven (Java) ecosystem support (Beta) to Secure Registry — route all Maven metadata and artifact downloads through `https://registry.stepsecurity.io/java` with Cooldown Period, Compromised Packages, and Custom Block List controls at launch.
- › Supports four Maven integration paths: JFrog Artifactory, Google Artifact Registry, Sonatype Nexus Repository, and Direct, with setup steps switchable via a new Package manager selector on the Setup Guide tab.
- › Introduces Custom Block List, a new per-ecosystem control that blocks specific packages or versions using glob patterns (e.g., `lodash@4*`, `@scope/*`), evaluated independently of Cooldown, Compromised Packages, and Typosquatting Protection.

› Adds Tarball URL Rewriting setting for npm, rewriting tarball URLs in metadata responses so package downloads route through Secure Registry instead of the public npm registry.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

DataDome

Go module v2.5.0 (2026-07-28) · 3 RELEASES · 2026-07-28

BLOG ↗ · LATEST · STABLE

DataDome is a bot protection and fraud prevention platform that detects and blocks malicious traffic, scraping, and cyberattacks in real time.

DataDome Go module v2.5.0 adds a GraphQLEndpoint parameter to configure GraphQL path protection.

└──▷ USE IT

Override the default GraphQL path when your API exposes GraphQL at a non-standard route.

```go
GraphQLEndpoint: "/api/graphql"
```

› Adds `GraphQLEndpoint` parameter to specify custom GraphQL endpoint paths (defaults to `/graphql`).

▾ 2 more releases in this issue · 2026-07-28

Google Cloud Load Balancer module 1.7.0 (2026-07-28)

BLOG ↗ · STABLE

DataDome's Google Cloud Load Balancer module adds configurable GraphQL endpoint path support.

└──▷ USE IT

Override the default GraphQL endpoint path when your API exposes GraphQL at a non-standard route.

```yaml
GraphQLEndpoint: /api/graphql
```

› Adds `GraphQLEndpoint` parameter to specify custom paths for the GraphQL endpoint (defaults to `/graphql`).

Python ASGI module 1.2.0 (2026-07-28)

BLOG ↗ · STABLE

DataDome Python ASGI module 1.2.0 adds configurable MCP endpoint support and expands cookie length to 512 chars.

› Adds `mcp_endpoint` attribute to `DataDomeMiddleware` to specify a custom MCP endpoint path (defaults to `/mcp`).

> Increases DataDome cookie length limit from 128 to 512 characters to support future capabilities.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS ⓘ

Track your GitHub check runs from the dashboard (2026-07-28) BLOG ↗

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

GitGuardian adds a dedicated GitHub Check Runs dashboard page with filtering, KPIs, re-run controls, and incident linking.

└──▷ HOW TO FIND IT

Audit scan coverage and unblock a stuck pull request by finding and re-running a failed check run without leaving GitGuardian.

> 📍 In the dashboard, go to Perimeter › GitHub check runs. Filter by repository and set Status to 'Failed', locate the run for the relevant commit SHA, then click Re-run.

Jump straight to remediation when a check run catches a secret, without switching to GitHub.

> 📍 In the dashboard, go to Perimeter › GitHub check runs. Find the check run with detected secrets, then click the linked incident to open the incident details and begin remediation.

- New GitHub Check Runs page in the Perimeter section surfaces every check run with status, repository, commit SHA, pull request, and timing in one place.
- Filter and search check runs by repository, status, or date — including workspaces with hundreds of thousands of runs.
- KPI strip at the top of the Check Runs page summarizes activity to surface failures and trends at a glance.
- Re-run or skip a check run directly from the dashboard without switching to GitHub.
- Navigate from a check run directly to linked incidents when secrets are detected, enabling faster remediation.

+2 more

WAS THIS USEFUL? 👍 👎

infisical ⓘ v0.162.14 NOTES ↗

Infisical is the open-source platform for secrets, certificates, and privileged access management.

Infisical v0.162.14 adds Google Workspace service template, Windows domain account rotation, enriched secret reminder emails, and instance version update checks.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.162.14
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.162.14
```

- Adds Google Workspace as a proxied service template (available as a popular template in the service proxy catalog).
- Tracks and displays a last-used timestamp for proxied services.
- Adds an instance version update check with in-UI display so self-hosted admins can see when a newer release is available.
- New machine identity creation sheet includes policy templates to accelerate identity provisioning.
- Enriches secret reminder emails with secret details and a direct 'View secret' link.

+1 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

tcpreplay ⓘ v4.6.0 NOTES ↗ CODE ↗

Edit and replay captured network traffic at arbitrary speeds — a suite of pcap tools for *NIX and Windows.

tcpreplay 4.6.0 adds `AF_XDP` and io_uring injection backends, `--loss` simulation, `--raw` socket path, and a static libtcpreplay library.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.6.0 https://github.com/appneta/tcpreplay.git
# already have the repo? check out this version:
$ git checkout v4.6.0
```

└──▷ TRY IT

Saturate a 1GigE link using the `AF_XDP` kernel-bypass backend; fall back automatically if the adapter lacks native XDP support.

```
$ tcpreplay --xdp -i eth0 capture.pcap
```

Benchmark `AF_XDP` on a specific queue without silent fallback masking results — fail loudly if XDP can't be used.

```
$ tcpreplay --xdp --xdp-no-fallback --xdp-queue 2 -i eth0
capture.pcap
```

Simulate a lossy link by randomly dropping 10% of packets during replay to stress-test application resilience.

```
$ tcpreplay --loss 10 -i eth0 capture.pcap
```

- Adds `--xdp` `AF_XDP` injection backend, now reaching ~941 Mbps (305k pps) line rate on 1GigE hardware — a 23× improvement over the previous implementation.
- Adds `--xdp-queue` to bind the `AF_XDP` socket to a specific adapter queue instead of always using queue 0.
- Adds `--xdp-no-fallback` to fail hard rather than silently fall back to the default injector — useful for benchmarking where a silent injector change would invalidate

results.

- Adds `--xdp-batch-size` for tuning `AF_XDP` batch depth on very high-speed links (100GigE+) with small packets.
- Adds `--io-uring` injection backend that batches sends 64 deep via Linux io_uring, reducing per-packet syscall count 64× for ~11% less CPU per packet.

+5 more

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Rapid7 InsightCloudSec

changelog · Jul 28, 2026

DOCS ↗

BLOG ↗

Rapid7 InsightCloudSec is a cloud-native security platform that continuously monitors, detects, and automatically remediates security risks and compliance violations across multi-cloud environments.

InsightCloudSec 26.7.28 adds AI Root Cause Analysis for runtime findings, HITRUST CSF v11.7.1 compliance packs, and AWS Tag Policy resource harvesting.

└──▷ HOW TO FIND IT

Investigate why a container runtime detection fired by reviewing AI-generated root causes and when the underlying posture risks first appeared.

📍 In InsightCloudSec, open a container runtime detection finding › AI Analysis panel › Root Cause Analysis tab to view the AI summary, Prevention Window, Primary Causes with confidence scores, and Contributing Factors.

- Adds a Root Cause Analysis tab to the AI Analysis panel for container runtime detection findings, surfacing an AI-generated summary, a Prevention Window showing when posture risks were first detected relative to the runtime incident, Primary Causes with confidence scores, and Contributing Factors.
- New HITRUST CSF v11.7.1 e1 Essential and i1 Intermediate compliance packs available.
- Adds harvesting support for AWS Organizations Tag Policy resources across AWS, AWS GovCloud, and AWS China environments via the new TagPolicyHarvester, enabling visibility into tag enforcement policies.

WAS THIS USEFUL? 👍 👎

Qualys TotalCloud ⓘ changelog · Jul 27, 2026 BLOG ↗

Qualys TotalCloud is a cloud security platform providing vulnerability management, compliance monitoring, and threat detection across cloud infrastructure.

Qualys TotalCloud adds automated vulnerability scanning for AWS Lambda, extending serverless security beyond CSPM to workload-level risk.

└──▷ HOW TO FIND IT

Deploy the Qualys Lambda scanning integration to your AWS environment to enable automated, event-driven vulnerability scanning across accounts and regions.

> 📍 Deploy the provided CloudFormation template in your AWS account (same-account or cross-account) to configure event-driven Lambda function scanning. Once deployed, Qualys automatically discovers and scans Lambda functions on creation or update — no manual triggering needed.

- Adds automated vulnerability scanning for AWS Lambda functions, detecting vulnerable open-source packages and outdated libraries inside function code — beyond what CSPM posture checks can see.
- Enables event-driven scanning that triggers automatically when Lambda functions are created or updated, with no manual intervention required.
- Supports cross-account and multi-region AWS Lambda environments, including same-account and cross-account configurations deployed via CloudFormation templates.
- Combines Lambda vulnerability findings with cloud security context (permissions, exposure) to power hyper-prioritization via automated TruRisk Scoring and unified attack path analysis merging CSPM and real-time detections.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Aikido Security ⓘ

#52 - DSPM, Aikido Libraries and Repository Code Quality scans NOTES ↗

Aikido Security is a developer-first security platform that identifies and fixes vulnerabilities in code, dependencies, and infrastructure.

Aikido adds DSPM, backport library patches, code quality scans, and new agent/IDE releases in one drop.

└──▷ HOW TO FIND IT

Point DSPM at a repository to discover where sensitive data is exposed in code — no database credentials or access grants needed.

> 📍 In the Aikido console, navigate to your repository, then open the Data Exposure Audit (DSPM) scan to review sensitive data flow findings and remediate at the source.

Merge a backported CVE fix for an existing dependency version without upgrading to a new major version.

> 📍 In the Aikido console, go to Aikido Libraries, select an open-source CVE affecting your current dependency version, and open the generated PR to apply the backported patch.

- New Data Exposure Audit (DSPM) scans repos to track how sensitive data flows and where it gets exposed — no database connections or data access grants required.
- New Aikido Libraries capability delivers backported CVE fixes for the exact dependency version already in use, without requiring major-version upgrades or refactors.
- New Code Quality Repository Scans surface existing code quality issues across repos to help identify and prioritize technical debt over time.
- Noise reduction for container images in Cloud scans.
- Bitbucket Server support added for PR Checks.

+7 more

WAS THIS USEFUL? 👍 👎

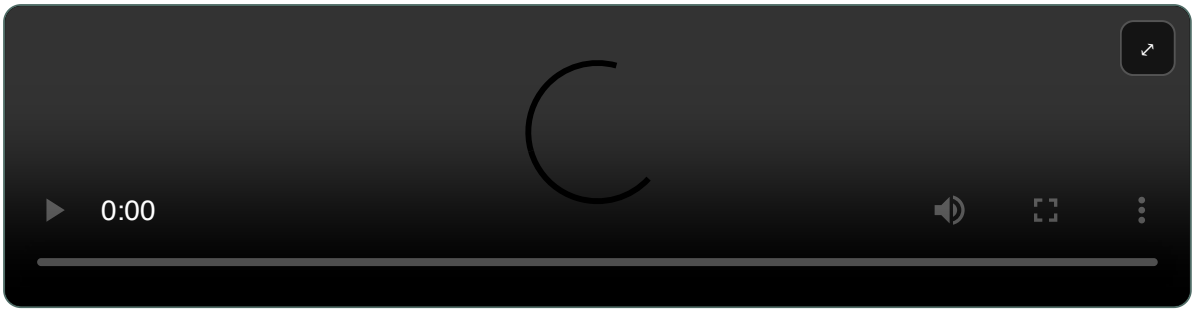
DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Iru ⓘ New Apple Beta Enrollment Assignment Rule NOTES ↗

Manage devices and everything they touch. See how security, IT, and compliance get easier when unified around devices with Iru.

Iru adds Assignment Rule to target Library Items at devices based on Apple beta program enrollment status.



› New Assignment Rule on Assignment Maps lets you assign Library Items to devices based on Apple beta program enrollment — either any beta enrollment or specific beta programs.

WAS THIS USEFUL? 👍 👎

XFA ⓘ A New Look for the XFA Desktop App NOTES ↗

XFA is a developer security tool that scans and monitors open-source dependencies for vulnerabilities, helping teams manage software supply chain risks efficiently.

XFA desktop app gains a full security dashboard so end users can see every check and self-remediate without IT.

› New full desktop app UI (macOS, Windows, Linux): users can open the menu-bar icon to see a complete device security overview.

› Per-check pass/fail detail view shows every security check XFA runs with step-by-step remediation guidance for end users.

WAS THIS USEFUL? 👍 👎

Huntress ⓘ Microsoft Usage Location Inference Live in Managed ITDR

2 RELEASES · seen 2026-07-28 NOTES ↗ LATEST STABLE

Huntress is a managed threat detection and response platform that helps MSPs and enterprises detect and respond to security threats.

Huntress Managed ITDR now infers Microsoft 365 home country from tenant data, eliminating false-positive Unexpected Country escalations.

- Automatically infers a Microsoft 365 identity's home country from the tenant when no Entra usage location is set, requiring zero partner configuration.
- Inferred country value serves as the Expected Country for Unexpected Country escalation evaluation, reducing false-positive alerts for users logging in from their home country.

▾ **1 more release in this issue** · seen 2026-07-28

> Onboarding Escalation Throttle Now Live in Managed ITDR NOTES ↗ STABLE
>
> Managed ITDR now throttles and prioritizes escalations during a new tenant's first week to reduce onboarding noise.
>
> - Adds onboarding escalation throttle to Managed ITDR, pacing and prioritizing alert flow during a new tenant's first week while behavioral baselines build.

WAS THIS USEFUL? 👍 👎

Microsoft Defender

changelog · Jul 27, 2026

BLOG ↗

Microsoft Defender is an antivirus and security tool protecting Windows devices from malware, viruses, and cyber threats.

Microsoft Defender's Project Perception launches public preview with coordinated red/blue/green AI agents for closed-loop autonomous defense.

- Introduces Project Perception, a continuously learning defense system that coordinates three classes of specialized AI agents — red team (attack path discovery), blue team (investigation and risk reasoning), and green team (corrective action and hardening) — operating in a closed-loop across identities, endpoints, applications, data, clouds, and AI systems.
- Launches software vulnerability management as the first Project Perception scenario, integrating MAI-Cyber-1-Flash into MDASH (software vulnerability multi-model agent team) for automated vulnerability discovery and prioritization.
- MDASH with MAI-Cyber-1-Flash achieves 96% on the CyberGym benchmark (+12 points above the previous Mythos configuration) while delivering ~50% cost reduction versus the current production MDASH configuration.

WAS THIS USEFUL? 👍 👎

Microsoft Security

changelog · Jul 27, 2026

BLOG ↗

Microsoft Security provides comprehensive cybersecurity solutions, threat protection, and security monitoring tools for organizations and enterprises.

Microsoft launches Project Perception agentic security system and the External Red Team Alliance (EXTRA) for global AI safety research.

- Introduces Project Perception, a new agentic security system with a multi-layer Cyber Stack — signals/sensors, security context, multi-model architecture, agent harness, and actuators — designed to continuously perceive, reason, and act at machine speed.
- Adds a continuously updated security context layer that transforms signals, threat intelligence, and asset/identity/relationship data into a shared, near-real-time representation for AI agents, reducing token cost and improving reasoning accuracy.
- Adopts a multi-model architecture that dynamically selects the best-fit model (frontier or specialized) per security task, optimizing for quality, latency, reliability, and cost without locking customers to a single model.
- Introduces actuators that translate AI reasoning directly into real-time protections and actions across integrated security workflows, closing the gap between insight and enforcement.
- Launches EXTRA (External Red Team Alliance), a formalized global program funding 18 university labs across six continents with unrestricted gifts to advance independent AI safety and red-teaming research.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Lookyloo ⓘ v1.40.1 NOTES ↗

Lookyloo is a web interface that allows users to capture a website page and then display a tree of domains that call each other.

Lookyloo v1.40.1 adds WebSocket stream visibility, private captures with token access, and background tree-building scripts.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.40.1 https://github.com/Lookyloo/lookyloo.git
# already have the repo? check out this version:
$ git checkout v1.40.1
```

└──▷ HOW TO FIND IT

Restrict a sensitive capture so only admins or a shared temporary token can view it.

> 📍 In the capture view, open the visibility form and set the capture to Private; share the generated temporary token with authorised recipients.

- Adds private captures mode — captures visible only to admins or holders of a temporary token.
- Adds a UI form to change the visibility of an existing capture (e.g. public → private).
- Adds background scripts to build trees and index archived captures without blocking the main process.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder ⓘ

AI Pentesting - expert-level web app pentesting without the wait (2026-07-28)

BLOG ↗

Intruder is a cloud-based vulnerability scanner that continuously monitors and identifies security weaknesses in internet-exposed systems, APIs, and web applications.

Intruder launches AI Pentesting: on-demand web app pentests with PDF reports, whitebox code repo analysis, and logic-flaw detection.

.

└──▷ HOW TO FIND IT

Start an AI-driven pentest against your web app directly from the Intruder portal.

> 📍 In the portal, go to AI Pentesting in the left nav, click 'Start new pentest', then choose the 'AI Pentest' option.

- New AI Pentesting capability delivers full web app penetration tests on demand, producing a PDF report with executive summary, severity ratings, reproduction steps, and remediation guidance.
- Supports whitebox testing via code repository integration, enabling deeper analysis beyond surface-level scanning.
- Detects access control, business logic, and auth bypass vulnerabilities that require multi-step chaining — issues missed by static or purely dynamic scanning.
- Allows practitioners to define custom entry points, credentials, and application context to scope tests to their real environment.

WAS THIS USEFUL? 👍 👎

Censys

July 27, 2026 (2026-07-27) BLOG ↗

Censys is an internet intelligence platform that continuously scans and indexes internet-connected devices, helping security teams discover and monitor exposed assets and vulnerabilities.

Censys adds OAuth MCP server setup, 10 new software fingerprints, and 23 new protocol/application scanners.

└──▷ HOW TO FIND IT

Connect the Adversary Investigation MCP server to your account using the new OAuth workflow instead of manual credential setup.

> 📍 In the Platform web console, go to Connected Applications, then follow the OAuth workflow to authorize the Adversary Investigation MCP server.

- New Connected Applications page in the Platform web console lets users review and manage OAuth connections.
- Adds 10 new software fingerprints covering Apple HTTPServer, Control Web Panel, Dynatrace, LinuxServer SWAG, Meteor, Microsoft ASP.NET MVC, Microsoft WCF, Oracle Mojarra, Weaver e-cology, and ZKoss ZK Framework.
- Adds 23 new protocol, application, and endpoint scanners including `CORTEX_XDR_P2P`, `DENON_HEOS`, `DOTNET_NEGOTIATE_STREAM`, `ELASTICSEARCH_TRANSPORT`, GIOP, INFLUXDB, `ISC_DHCP_OMAPI`, JUPYTER, `KCODES_NETUSB`, KEYCLOAK, LANGFLOW, MYSQLX, `NGINX_PROXY_MANAGER`, `PERFORCE_P4D`, `PROXMOX_VE`, `QDRANT_GRPC`, `R1SOFT_BUAGENT`, `SYNCTHING_BEP`, `SYNOLOGY_DSM`, TARANTOOL, `TELEXPER_TLXP`, `TWAMP_CONTROL`, and VAULT.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Ory ⓘ

Ory Network, Ory Hydra, Ory Kratos v26.3.3 released NOTES ↗

Ory is an open-source identity and access management platform providing authentication, authorization, and user management infrastructure for developers building secure applications.

Ory Network v26.3.3 adds configurable legal footer links (ToS, Privacy Policy, Contact) to Account Experience UI.

- Adds configurable Terms of Service, Privacy Policy, and Contact links to the Account Experience legal footer.

WAS THIS USEFUL? 👍 👎

Passbolt ⓘ

Sunshine | Browser Extension & API NOTES ↗

Passbolt is an open-source password manager designed for team collaboration, enabling secure password sharing and management across organizations.

Passbolt 5.14 adds smarter credential suggestions, stronger MFA policy enforcement, and email-free SSO setup and recovery.

- Improves visibility over shared access for resources.
- Delivers more relevant in-context credential suggestions.
- Strengthens MFA policy enforcement.
- Expands email-free SSO setup and recovery workflows.
- Makes subscription changes easier to assess.

WAS THIS USEFUL? 👍 👎

ConductorOne

July 24, 2026 (2026-07-27)

BLOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance across enterprise applications and infrastructure.

ConductorOne adds AI agent policy steps, group/role access reviews, per-user MCP auth, and scoped personal API keys.

└──▷ HOW TO FIND IT

Delegate a policy step to the C1 AI agent so it automatically evaluates access requests and logs each decision for audit — no human reviewer needed for routine approvals.

> 📍 In the console, go to Policies › [your policy] › Edit, add or select a workflow step, choose Assign to agent, enter any custom instructions, and save. Audit log entries for that step will appear under the C1 System agent identity with the AI-agent action label.

Scope a personal API key to read-only to limit blast radius when integrating with a third-party reporting tool.

> 📍 In the console, go to Settings › Personal API Keys › Create API Key, open the Limit Scopes selector, choose the read-only role, and generate the key.

Run a quarterly access review that covers group and role grants, not just individual user grants, to satisfy broader compliance requirements.

> 📍 In the console, go to Access Reviews › Create Campaign, open the Scope tab, find the Grant types selector, choose 'Users and resources', set a custom recurrence interval (e.g., every three months), configure the campaign duration, and launch.

- Assign a workflow step in a request or review policy to the C1 AI agent, which evaluates each request using built-in guidance plus custom instructions; every evaluation is attributed to the C1 System agent identity, labeled as an AI-agent action in the audit log, and recorded as a reviewable conversation.
- Request catalog search now includes an 'Entitlements matching' section showing entitlements that indirectly grant access to a searched resource (e.g., a GitHub team whose repo permissions include the target repo), with a 'View access graph' link showing the full grant chain.
- Access review campaigns can now review grants held by groups or roles in addition to users, using a new 'Grant types' selector (Users only, Resources only, or Users and resources) on the Scope tab.
- Campaign templates now support custom recurrence intervals (e.g., every two days, every three months) and configurable campaign durations, beyond the previous fixed

presets.

- IAM policy documents on AWS resources now render as a readable, collapsible JSON tree via a 'View policy' button on resource-level access reviews and in task details.

+10 more

WAS THIS USEFUL? 👍 👎

SailPoint ⓘ

Introducing the connectivity agent for simplified web services configuration in SailPoint (2026-07-27)

BLOG ↗

SailPoint is an identity security platform that helps organizations manage and govern user access to applications, data, and systems across the enterprise.

SailPoint introduces an AI-guided connectivity agent wizard to configure Web Services sources without deep API expertise.

- Adds a connectivity agent: an AI-powered, wizard-based setup experience for configuring Web Services connector sources directly within the SailPoint application onboarding workflow.
- Supports AI-recommended configuration for authentication endpoints, account aggregation with data preview, pagination rules for large datasets, attribute mapping for accounts and entitlements, and identity-to-account correlation.
- Enables detection of compatible REST and SCIM public APIs for any target application, compressing multi-day Web Services configuration into hours.
- Includes provisioning endpoint recommendations covering Get Object, Create Account, Account Update, and Add/Remove Entitlement.
- Each AI recommendation is validated in-UI with a live data preview before the administrator accepts — no changes are applied without explicit approval.

+1 more

WAS THIS USEFUL? 👍 👎

Clerk ⓘ

Composable UserProfile and OrganizationProfile components (2026-07-27) BLOG ↗

Clerk is a user authentication and identity management platform that provides developers with pre-built UI components, APIs, and SDKs for sign-up, sign-in, and user management.

Clerk introduces composable, experimental building-block APIs for UserProfile and OrganizationProfile pages.

- New experimental composable API lets developers build custom UserProfile and OrganizationProfile pages from Clerk's individual UI building blocks instead of using the pre-built monolithic components.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Progress ShareFile ⓘ snapshot-20260728 NOTES ↗

Progress ShareFile is a cloud-based file sharing and storage platform enabling businesses to securely send, receive, and manage documents and client communications.

ShareFile adds Tasks workspace, Request List export, Outlook signature requests, session management, AI document validation, and a Security Insights dashboard.

- New Export a Request List capability lets users export Request Lists to Excel (.xlsx) or CSV (.csv) from the three-dot menu, reflecting currently active columns.
- New ShareFile Tasks provides a centralized workspace for viewing and managing work items across projects in a single location.
- New Create Signature Request in ShareFile for Outlook Online add-in lets users send documents for e-signature or self-sign directly from Outlook without leaving the inbox.
- ShareFile for Windows now provides access to Projects, Client Hub & Solutions files and folders for all Premium plan users (v26.03.12.0+).
- New Bulk client entity creation in Client Hub allows admins and employees to create multiple client entities and primary contacts at once.

+16 more

└──▷ BREAKING ON UPGRADE

! Upgrading to ShareFile for Mac 26.04.20 requires all users to re-authenticate.

! Upgrading to ShareFile for Windows 26.3.12.0 requires all users to re-authenticate.

! ShareFile for Windows 25.6.5.0 is now signed with Progress Software Corporation's digital certificate; environments restricting software by publisher may block this update until 'Progress Software Corporation' is added to the trusted publishers list.

WAS THIS USEFUL? 👍 👎

DNSFilter ⓘ

07/28 July 28 – Web Application and API Update (2026-07-28) BLOG ↗

DNSFilter is a cloud-based DNS security platform that blocks malicious websites, filters unwanted content, and protects networks from cyber threats in real time.

DNSFilter adds Brave SafeSearch enforcement, AppAware Sites filter, clickable MSP widgets, and Query Log 'Uncategorized' filtering.

└──▷ HOW TO FIND IT

Enforce SafeSearch on Brave (and all other supported engines at once) in a Filtering Policy.

> 📍 In the console, go to Filtering Policies › select a policy › Privacy tab › SafeSearch section › click 'Enforce All' to enable every supported engine including Brave.

Find unclassified DNS queries that may need custom categorization or blocking.

> 📍 In the console, go to Query Log › add a filter for 'Uncategorized' to display only queries DNSFilter has not yet classified.

- Adds Sites filter to AppAware Reporting, giving Core plan and mixed Core/Pro environments query data scoped by site.
- Supports SafeSearch enforcement on Brave browser in Filtering Policies, joining Google, Bing, DuckDuckGo, Ecosia, and Yandex.
- Adds 'Enforce All' shortcut in the redesigned SafeSearch policy editor to enable every SafeSearch engine at once; enforcement toggles are no longer nested under 'Block all search engines'.
- Enables filtering DNS Query Log by 'Uncategorized' to surface unclassified queries.
- Makes Collections and Users widgets on the MSP Overview page clickable, linking directly to their respective management pages.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Immuta ⓘ snapshot-20260728 NOTES ↗

Immuta is a data security platform that automates access control, privacy policy enforcement, and compliance monitoring across cloud data environments.

Immuta adds separate access/masking request forms, Redshift viewless auto-discovery, managed groups for policies, and OAuth 2.0 for Alation.

└──▷ HOW TO FIND IT

Write a subscription policy that allows every registered identity (agents and users) to subscribe to non-sensitive data without manually maintaining group membership.

> 📍 In the Immuta policy editor, create a subscription policy and set the target group to the built-in managed group 'All Identities' to automatically include all agent and non-agent user accounts.

- Adds separate access and masking exception request forms for data products, allowing stewards to configure independent questions and data use agreements per request type.
- Expands automated sensitive data discovery to Immuta's Amazon Redshift viewless integration, automatically tagging matching columns using regex, dictionary, and column-name regex identifiers.
- Introduces three built-in Immuta managed groups — All Agents, All Users, and All Identities — available for authoring subscription and data policies without manual group management.
- Adds OAuth 2.0 authentication support for Immuta's Alation integration alongside existing API key authentication, eliminating recurring token refresh requirements.
- Extends Snowflake manual policy resync to cover both data policies and subscription policies when triggered from a data source health check.

+2 more

└──▷ BREAKING ON UPGRADE

- ! Immuta policy role names used to grant users access to Snowflake data sources will be renamed as part of the Snowflake integration architecture update rolling out beginning July 23, 2026; policy roles will be recreated under new names.
- ! Integration error banners for Snowflake integration validation failures will no longer be displayed; errors will only be available through the connections experience.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Cursor ⓘ Cursor Start NOTES ↗

Cursor is an AI-powered code editor that helps developers write, edit, and debug code faster using integrated large language model assistance.

Cursor launches India-local 'Start' plan at ₹649/month with UPI billing, cloud agents, and iOS remote control.

└──▷ HOW TO FIND IT

Sign up as a new user in India and select the Start plan during onboarding.

> 📍 Visit cursor.com/signup and select the Start plan during onboarding to enroll at ₹649/month.

- New Cursor Start plan (₹649/month, INR billing with UPI or card) targets developers in India with local pricing.
- Includes always-on cloud agents that build, test, and ship code in the background while the developer keeps working.
- Adds Cursor for iOS with remote control to launch and steer agents from a mobile device.
- Bundles access to Grok 4.5 and Composer models alongside plugins, MCP servers, hooks, and skills.

WAS THIS USEFUL? 👍 👎

OpenCode ⓘ v1.18.8 NOTES ↗ CODE ↗

The open source coding agent.

OpenCode v1.18.8 upgrades the MCP client SDK to v2 and adds a project selector keyboard shortcut.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.18.8
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.18.8
```

- Upgrades the MCP client SDK to v2, improving compatibility with newer MCP servers and OAuth flows.
- Adds a `mod+shift+o` shortcut to open the project selector from the new-session page in the desktop app.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.146.0-alpha.14 NOTES ↗ CODE ↗

PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI alpha.14 adds cloud-managed profiles for `codex sandbox`, network policy callbacks for remote exec, and configurable developer instructions for v2 subagents.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.146.0-alpha.14
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.14
```

- Loads cloud-managed profiles for `codex sandbox`, enabling permission profiles and managed requirements to be sourced from cloud config bundles.
- Adds configurable developer instructions for v2 subagents, allowing per-subagent instruction customization in multi-agent workflows.
- Routes curated plugins by authentication mode, so plugin recommendations are matched to the active auth context.
- Prepares MCP and plugin recommendations concurrently, reducing latency when loading tool suggestions.
- Terminates Windows non-TTY processes on interrupt, routing Ctrl-C through the termination callback for pipe-backed processes.

WAS THIS USEFUL? 👍 👎

Cline ⓘ desktop-v0.0.6 2 RELEASES · 2026-07-27 → 2026-07-28 NOTES ↗

CODE ↗ LATEST STABLE

Autonomous coding agent as an SDK, IDE extension, or CLI assistant.

Cline desktop v0.0.6 adds message queuing UI, persistent sidebar update indicator, and connector session persistence.

└──▷ GET THIS VERSION

```
$ git clone --branch desktop-v0.0.6
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout desktop-v0.0.6
```

- Adds a queued-messages panel above the composer showing a count with expand controls to edit, send immediately, or delete individual queued turns.
- Adds a persistent sidebar update indicator that displays the new version and a one-click restart button after an update is downloaded, surviving toast dismissal.
- Displays the full workspace name and current git branch in the header and sidebar, with stable branch tracking that survives transient git lookups.

▾ 1 more release in this issue · 2026-07-27 → 2026-07-28

desktop-v0.0.5 NOTES ↗ CODE ↗ STABLE

Cline desktop v0.0.5 adds a Cline API key onboarding path, cancellable browser sign-in, and consistent MCP server cards with install guidance.

└──▷ GET THIS VERSION

```
$ git clone --branch desktop-v0.0.5
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout desktop-v0.0.5
```

- Adds a Cline API key path to onboarding so users can authenticate without a browser.
- Enables cancellation of a pending browser sign-in instead of waiting indefinitely.
- Unifies MCP server cards across marketplace views with a single uninstall action and setup guidance on installed servers.

> › Delivers a major performance overhaul: animated background locked at 60 fps, composer keystroke stutter reduced from 245 slow keystrokes to 3, streaming responses coalesce per-token re-renders, and app boot fetches the provider catalog once instead of three times.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI ⓘ 1.15.8 NOTES ↗

Framework for orchestrating role-playing, autonomous AI agents. By fostering collaborative intelligence, CrewAI empowers agents to work together seamlessly, tackling complex tasks.

CrewAI 1.15.8 adds WaitTool to let agents pause execution while waiting on long-running async jobs.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.15.8 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.8
```

› Adds `WaitTool` to pause agent execution while waiting on long-running background jobs.

WAS THIS USEFUL? 👍 👎

holmesgpt ⓘ 0.37.0 NOTES ↗

SRE Agent - CNCF Sandbox Project

HolmesGPT 0.37.0 adds JSON log scraping, Langfuse/OpenTelemetry tracing, and pre-acquired Azure AD token support.

└──▷ GET THIS VERSION

```
$ git clone --branch 0.37.0
https://github.com/HolmesGPT/holmesgpt.git
# already have the repo? check out this version:
$ git checkout 0.37.0
```

- Adds JSON logging format support for log scrapers, enabling structured log ingestion pipelines.
- Adds Langfuse integration via OpenTelemetry tracing for LLM observability and trace export.
- Supports pre-acquired `AZURE_AD_TOKEN` for Entra ID authentication, allowing token reuse without re-authentication.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

vMLX ⓘ v1.6.18 NOTES ↗

vMLX - JANGTQ Uber Compressed MLX Models - L2 Disk Cache (survives restart) + L1 Paged (super fast ttft) + Hybrid SSM Scheduler + Cont Batching + etc!

vMLX 1.6.18 adds a process-wide SSD cache budget with LRU eviction and detailed cache reporting.

└──▷ GET THIS VERSION

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

- Adds a process-wide SSD cache budget with coordinated LRU eviction and off-request-path publication for improved prefix reuse management.
- Introduces detailed cache reporting covering attempted prefixes, applied reuse, uncached suffixes, reconstruction, fallback prefill, and aggregate disk usage.
- Supports MiniMax tool arguments emitted as request-schema XML children.
- Retains media across Responses tool continuations and enables local attachment playback in the Electron app.

└──▷ BREAKING ON UPGRADE

! This release requires JANG 2.5.34.

WAS THIS USEFUL? 👍 👎

◆ AI Model & Data Infrastructure

emisar ⓘ v0.35.0 2 RELEASES · 2026-07-26 → 2026-07-27 NOTES ↗ LATEST

STABLE

An MCP that lets AI tools securely connect to your infrastructure, write IaaS code, debug issues, and assist during incidents - without risking production stability. Built for security teams to approve and infrastructure teams to experience like magic.

Emisar v0.35.0: hostname-based runner identity, retention safety for connected runners, and faster behavior matrix CI.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.35.0
https://github.com/AndrewDryga/emisar.git
# already have the repo? check out this version:
$ git checkout v0.35.0
```

- Runners without an explicit `runner.id` now use their hostname as identity, so reboots reconnect the same runner and replaced ephemeral hosts auto-enroll without a persisted identity file.
- Pack retention no longer sweeps versions still advertised by a connected runner, preventing `pack_untrusted` dispatch failures after idle-window cleanups.
- Hosted admin runner bootstraps from a version-pinned release download with retries, eliminating GitHub API anonymous rate-limit failures on control-plane VM start.
- Behavior cases for high/critical pack actions now run in disposable Compose projects as a non-root identity against real services, with network-path readiness probes enforced before cases start.
- The snmp pack's OSPF adjacency fixture supports AppArmor-confined hosts via network namespace construction.

+2 more

▾ 1 more release in this issue · 2026-07-26 → 2026-07-27

> v0.34.0 NOTES ↗ STABLE
>
> emisar v0.34.0 adds streaming run output via MCP cursor paging and authoring-time pack argument validation.
>
> └──▷ GET THIS VERSION

```
$ git clone --branch v0.34.0
https://github.com/AndrewDryga/emisar.git
# already have the repo? check out this version:
$ git checkout v0.34.0
```

- New `wait_for_run` MCP call accepts an output cursor and returns the next cursor, enabling agents to stream run output incrementally as it arrives rather than polling.
- Curl-backed API actions now fail on 4xx/5xx HTTP responses instead of reporting transport success, making pack error handling reliable.
- Pre-run dispatch rejections (contract changes, refusals, rate limits) are now logged with bounded, allowlisted fields, making refused MCP calls visible in operations.
- Pack-trust conflict errors now name the specific runners that disagree about an action instead of failing generically.
- Operator input (approval notes, grant scopes, policy overrides) now survives console re-renders triggered by co-approver broadcasts or refused submits.

+2 more

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA SkillSpector ⓘ v2.5.0 NOTES ↗

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, and security risks.

SkillSpector v2.5.0 adds inspection-ledger reporting with execution-completeness fields in JSON/SARIF and v2 baseline fingerprints.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.5.0
https://github.com/NVIDIA/SkillSpector.git
# already have the repo? check out this version:
$ git checkout v2.5.0
```

└──▷ TRY IT

Check whether a scan completed reliably in CI — block the pipeline if `execution_successful` is false rather than silently accepting a zero-finding report.

```
$ skillspector scan ./my-skill/ --format json --output report.json &&
jq '.execution_successful' report.json
```

Regenerate a v2 baseline after upgrading from v1 so fingerprint suppressions are accepted again.

```
$ skillspector baseline ./my-skill/ -o .skillspector-baseline.yaml
```

Surface ledger exceptions in CI to diagnose why an incomplete scan was blocked.

```
$ skillspector scan ./my-skill/ --format json --output report.json;
jq '.analysis_completeness.ledger_exceptions' report.json
```

- Adds canonical inspection-ledger accounting across static and LLM analysis stages, including per-component coverage and explicit out-of-scope records.
- New `execution_successful` top-level field and `analysis_completeness.ledger_exceptions` diagnostics in JSON output let automation block incomplete or failed scans.

- JSON and SARIF reports now include execution-completeness information and analyzer status so a zero-finding report can be distinguished from a partial or failed scan.
- Recursive scans now propagate child-scan failures to the combined report and return a failure when any child scan fails.
- CLI now exits with code 2 for fatal execution or accounting failures, even when a JSON report was produced.

+1 more

└──▷ BREAKING ON UPGRADE

- ! Baseline files with version 1 fingerprints are no longer accepted; run `skillspector baseline <path>`, review the generated version 2 entries, and commit the replacement baseline (rules-only version 1 baselines remain supported with a warning).
- ! JSON integrations must treat invalid or missing output, a nonzero process exit, or `execution_successful: false` as a blocking validation error and surface `analysis_completeness.ledger_exceptions` for diagnosis.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 28, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe