THE CYBER TOOLCHAIN

NO. 030

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 30, 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 29 tools

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Codacy ⓘ

Support for Shellcheck configuration file - July, 2026

NOTES ↗

Codacy is an automated code quality and security analysis platform that reviews pull requests, enforces coding standards, and tracks technical debt across repositories.

Codacy now respects ShellCheck configuration files in your repository.

- Supports ShellCheck configuration files, allowing per-repository ShellCheck rule customization without modifying Codacy settings.

WAS THIS USEFUL? 👍 👎

Portkey ⓘ snapshot-20260730 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 adds vault-backed secret references, weekly/endpoint-scoped rate limits, two new guardrail integrations, and broader provider coverage.

└──▷ USE IT

Apply Zscaler Detection Policies to all LLM traffic by wiring Zscaler AI Guard into your Portkey config as a guardrail hook.

json

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

Set a weekly spend cap scoped only to the embeddings endpoint so it doesn't share budget with chat completions.

📍 In the Portkey console, go to Budget & Rate Limit Policies, create a new policy, set the window to `rpw` (weekly), set the endpoint scope to `embeddings`, and enter your desired token or cost limit.

› Introduces weekly rate/budget windows (`rpw`) and endpoint-scoped rate limits, letting teams cap spend per API surface (chat completions, embeddings, admin routes, etc.) independently.

WAS THIS USEFUL? 👍 👎

Veracode ⓘ Integration updates (2026-07-30)

3 RELEASES · 2026-07-29 → 2026-07-30 BLOG ↗ LATEST STABLE

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 integrations gain platform-driven GitHub install/uninstall, scheduled scans, new API pagination, and broader language autopackaging support.

└──▷ USE IT

Enable debug logging for a GitHub or Azure DevOps workflow scan to troubleshoot packaging or scan failures.

```yaml
# In veracode.yml
debug: true
```

Install and schedule Veracode GitHub Integration scans from the platform without touching the GitHub Marketplace.

> 📍 In the Veracode Platform, go to Integrations › GitHub › Install, select your organization, configure credentials, then go to Schedule Scans to set up Static Analysis, SCA, or IaC scan schedules.

Scope a Visual Studio scan to only the active solution for faster, more focused results in large repos.

> 📍 In Visual Studio with the Veracode Scan extension, open Settings › Veracode Scan and set Scan Scope to 'Active Solution', then run a scan.

- Adds a new uninstallation workflow on the Veracode Platform to remove the Veracode app from GitHub organizations.
- Enables installing, configuring, and updating credentials for the Veracode GitHub Integration directly from the Veracode Platform, bypassing manual GitHub Marketplace steps.
- Adds autopackaging support for Apex, VB6, and Perl in the GitHub Workflow Integration.
- Enables scheduling Static Analysis, SCA, and IaC scans from the Veracode Platform for GitHub organizations using the GitHub Workflow Integration.
- Adds scheduled scan support for every organization where the Azure DevOps Workflow

Integration is installed, driven by veracode.yml configuration.

+14 more

└──▷ **BREAKING ON UPGRADE**

! The `Veracode Flaw Importer` task (with spaces in the task name) in the Azure DevOps extension is deprecated and replaced by `VeracodeFlawImporter`; the old task is scheduled for removal in a future release.

▾ **2 more releases in this issue** · 2026-07-29 → 2026-07-30

`Platform updates (2026-07-29)` BLOG ↗ STABLE

Veracode adds Custom PDF REST API, OAuth for XML APIs, six new built-in policies, and expanded Reporting API filters in its 2026 platform updates.

› New Custom PDF REST API enables programmatic generation of custom PDF reports.

› XML APIs now support OAuth Client Credentials authentication, in addition to REST APIs.

› Auto-Update OWASP and CWE Top 25 policy rules now enforce OWASP Top 10:2025 and MITRE CWE Top 25 for 2025 standards.

› Six new built-in security policies added to the Veracode Platform.

› Veracode Analytics upgraded to Looker 26.6 with finer widget sizing, persistent filtering, and tabbed dashboards.

+18 more

└──▷ **BREAKING ON UPGRADE**

! The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns a 404 error; use the `v2/principal` endpoint in the Identity REST API instead.

`Static Analysis updates (2026-07-29)` BLOG ↗ STABLE

Veracode SAST adds Pipeline Scan support for >200 MB artifacts, exploit-level scoring in results.json, and an enhanced Python scanner with PySpark/Flask/Django/FastAPI support.

› Pipeline Scan now supports packaged artifacts larger than 200 MB with no integration changes required.

› Pipeline Scan results.json now includes a numeric exploit level (-2 to 2) for each finding, surfacing exploitability likelihood alongside vulnerability data.

- Static Analysis Engine v2026.07 adds support for Android 17 and enhances .NET 10 support.
- Static Analysis Engine v2026.07 expands JavaScript dependency resolution with pnpm module support.
- Static Analysis Engine v2026.07 expands third-party library detection in Java to include Fluxnova.

+33 more

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical ⓘ v0.162.15 NOTES ↗

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

Infisical v0.162.15 adds SCEP support for Microsoft Intune, NKP Helm chart, PAM connection testing, and improved machine identity template visibility.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.162.15
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.162.15
```

- Adds SCEP (Simple Certificate Enrollment Protocol) support for Microsoft Intune, enabling certificate-based device enrollment flows.
- New Helm chart (`infisical-nkp`) for deploying Infisical in Nutanix Kubernetes Platform (NKP) partner catalog environments.
- Adds connection test capability for PAM (Privileged Access Management) accounts, allowing validation of account connectivity before use.
- Improves machine identity template visibility in the UI, making templates easier to discover and manage.

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine ⓘ

Detection Engine Updates Version 2.168 (2026-07-20) BLOG ↗

GitGuardian Detection Engine scans code repositories and developer environments to automatically detect secrets, credentials, and sensitive data exposures in real time.

GitGuardian Detection Engine 2.168 adds 4 new detectors, 2 new validity checkers, and extends HTML/Markdown scanning to 5 token types.

- New detector for Azure Immersive Reader Key credentials.
- New detector for Alibaba Cloud Model Studio (Qwen / DashScope) Coding Plan API keys.
- New detector for Alibaba Cloud Model Studio (Qwen / DashScope) Token Plan API keys.
- New detector for Prefect Cloud API keys.
- New validity checker for Alibaba Cloud IAM keys.

+4 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Security Onion ⓘ 3.2.0-20260729 NOTES ↗

Security Onion is a free and open platform for threat hunting, enterprise security monitoring, and log management. It includes our own interfaces for alerting, dashboards, hunting, PCAP, detections, and case management.

Security Onion 3.2.0 adds agentic AI (Gemma-backed Investigator/Engineer), Sigma playbooks, ES|QL defaults, data stream lifecycle management, and UEK8 kernel support.

└──▷ GET THIS VERSION

```
$ git clone --branch 3.2.0-20260729 https://github.com/Security-Onion-Solutions/securityonion.git
# already have the repo? check out this version:
$ git checkout 3.2.0-20260729
```

- Adds agentic AI capability with Gemma-backed Investigator and Detection Engineer agents for automated analysis workflows.
- Supports Sigma playbooks (including airgap environments), enabling SigmaHQ rules against Elastic Defend logs.
- Adds data stream lifecycle management support for Elasticsearch index management.
- Upgrades to Elasticsearch 9.3.7.
- Adds UEK8 kernel support across install and grid, with UEK8 set as the boot default.

+4 more

└──▷ BREAKING ON UPGRADE

! The `icsnpp-modbus` parser is removed from defaults.yaml; grids relying on the default Modbus protocol parser will no longer have it enabled after upgrade.

! The Detection Engineer agent configuration key is renamed to `DetectionEngineer`; existing configurations using the old key will need to be updated.

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Fastly ⓘ Next-Gen WAF for Apache 1.11.0 (2026-07-29) BLOG ↗

Fastly is a cloud edge platform that accelerates, secures, and delivers websites, apps, and APIs by caching content closer to users globally.

Next-Gen WAF Apache module now logs failed-open requests for better visibility into bypass events.

› Logs failed-open requests, giving practitioners visibility into traffic that bypassed WAF inspection.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Sysdig Secure ⓘ Secure SaaS - July 29, 2026 BLOG ↗

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

Sysdig CLI Scanner v1.28.0 adds NixOS and Arch Linux package detection plus dnf/microdnf support.

› Adds detection of NixOS (`nix`) packages in the CLI Scanner.

› Adds detection of Arch Linux packages in the CLI Scanner.

› Adds `dnf` and `microdnf` to the package manager detection policy rule.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Iru ⓘ

Auto App: Elgato Stream Deck | 3 RELEASES · seen 2026-07-30 | NOTES ↗

LATEST | STABLE

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

Iru adds Elgato Stream Deck as a new Auto App integration in the catalog.

› Adds Elgato Stream Deck Auto App integration, now available to all customers via the Library Items catalog.

▾ **2 more releases in this issue** · seen 2026-07-30

> Auto App: Screen Studio | NOTES ↗ | STABLE
>
> Iru adds Screen Studio as a new Auto App, now available to all customers in the catalog.
>
> › Adds Screen Studio to the Iru Auto App catalog, making it available for automated deployment to all customers.

> New Windows Auto Apps | NOTES ↗ | STABLE
>
> Iru Endpoint adds 9 new Windows Auto Apps including Cisco Webex, Claude, GitHub CLI, and Node.js.
>
> › Adds Cisco Webex Meetings to the Windows Auto Apps catalog for automated endpoint deployment.
> › Adds Claude to the Windows Auto Apps catalog for automated endpoint deployment.
> › Adds GitHub CLI to the Windows Auto Apps catalog for automated endpoint deployment.
> › Adds Jabra Direct to the Windows Auto Apps catalog for automated endpoint deployment.
> › Adds Microsoft 365 Apps (Core) to the Windows Auto Apps catalog for automated endpoint deployment.
>
> +3 more

WAS THIS USEFUL? 👍 👎

Trend Vision One ⓘ

New resources in Container Security inventory overview (2026-07-29) BLOG ↗

Trend Vision One is a unified cybersecurity platform that integrates extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

Container Security inventory gains resource-level drill-down, 24-hour trend indicators, and direct log links.

└──▷ HOW TO FIND IT

Investigate whether a specific ECS task is generating more findings than yesterday without leaving the inventory view.

> 📍 In the console, go to Cloud Security → Container Security → Inventory, select an ECS service or task, then review the 24-hour trend indicator next to each finding type.

Jump directly from a finding count to the underlying log events for a specific Kubernetes pod.

> 📍 In the console, go to Cloud Security → Container Security → Inventory, select a pod, then select a finding count to open the corresponding log page.

- Drill down to individual Kubernetes nodes/pods or Amazon ECS services/tasks in the Container Security inventory overview, scoping all security findings to that single resource.
- Each security finding type now displays a 24-hour trend indicator comparing the last 24 hours against the previous 24 hours.
- Selecting a finding count in the inventory overview opens the corresponding log page directly for event-level review.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

defectdojo ⓘ 3.1.302 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

DefectDojo 3.1.302 adds login rate limiting to API token auth and expands scanner protocol support.

└──▷ GET THIS VERSION

```
$ git clone --branch 3.1.302 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.1.302
```

- Applies login rate limiter to the API token authentication endpoint, reducing brute-force exposure on token-based logins.
- Adds additional scanner protocols to the whitelist, broadening supported import sources.
- Updates Fortify parsers to extract and populate CWE data from scan results.
- API now returns HTTP 409 instead of 500 on unique-constraint violations, enabling cleaner client-side error handling.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

runZero ⓘ runZero 5.0.260728.1 (2026-07-29) BLOG ↗

runZero is a continuous asset intelligence platform that discovers, inventories, and monitors all connected devices and software across networks in real-time.

runZero 5.0 adds Jira issue-owner sync, Git Smart Protocol scanning, Purdue-level search terms, and multi-sensor Tanium support.

└──▷ HOW TO FIND IT

Filter software inventory to OT assets at Purdue Level 1 (field devices) to scope vulnerability reviews to critical control-layer systems.

> 📍 In the runZero console, go to Inventory › Software and enter `purdue_level:1` in the search bar.

Identify baseline assets not yet meeting a compliance goal so you can prioritize remediation work.

> 📍 In the runZero console, open a Goals widget and click the Pending metric to view the baseline assets that do not match the target query in inventory.

- › New site details page lets organization-level users view site details and subnet definitions.
- › Adds ability to sync issue owners with Jira.
- › Goals widget now includes a Pending metric showing baseline assets that don't match the target query, with click-through to those assets in inventory.
- › Tanium integration now supports specifying multiple sensors.
- › Scanner now supports services using the Git Smart Protocol.

+3 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta ⓘ snapshot-20260730 NOTES ↗

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

- Expands automated sensitive data discovery (regex, dictionary, and column-name regex identifiers) to Immuta's Amazon Redshift viewless integration.
- Snowflake integration error messages are moved out of error banners and into the connections experience.

└──▷ 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 from July 23, 2026.

WAS THIS USEFUL? 👍 👎

DigiCert CertCentral

July 29, 2026 NOTES ↗

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizational infrastructure.

CertCentral automated delegation approval expands to EV TLS, EV Code Signing, Verified Mark, and PKIo products globally and in Europe.

└──▷ HOW TO FIND IT

Delegate approval authority to a CertCentral user so future EV TLS or EV Code Signing orders are issued automatically without manual intervention.

📍 In CertCentral, go to Account › Users. Select the target user. On the user details page, under Order approval delegation, select Request delegation. In the Submit a delegation request window, select the authorized representative or verified contact, then select Submit Request.

- Extends automated delegation approval to EV TLS, EV Code Signing, Code Signing, Verified Mark, and Common Mark certificates in CertCentral global and Europe.
- Adds automated delegation approval support for a broad set of CertCentral Europe products including EU Qualified eSeal (PSD2), EU Qualified Website Authentication Certificate (PSD2), PKIo Personal Organisation, PKIo Qualified eSeal, PKIo Private Person, and multiple PKIo Authentication/Encryption and Private Services Server variants.
- Simplifies the delegation request workflow: a single request now covers all organizations the selected approver represents and all certificate products they are authorized to approve, removing the previous per-organization selection step.
- Enables fully automated certificate issuance via CertCentral UI and Services API once a delegated approver is assigned, eliminating repeated manual approvals per order.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Nudge Security ⓘ

Detect sensitive data shared in Amazon Quick Suite | 2 RELEASES · seen 2026-07-30

NOTES ↗ | LATEST | STABLE

Nudge Security discovers and secures all SaaS apps used across an organization, helping IT teams manage shadow IT and enforce security policies through automated nudges.

Nudge Security AI conversation monitoring now detects sensitive data entered into Amazon Quick Suite chats.

└──▷ HOW TO FIND IT

Enable detection of sensitive data shared in Amazon Quick Suite for your organization.

> 📍 In the Nudge Security console, go to Browser Extension Settings and configure AI conversation monitoring to include Amazon Quick Suite.

› Detects secrets, PII, financial data, and health data shared in Amazon Quick Suite (Amazon's agentic AI platform) via AI conversation monitoring.

▾ 1 more release in this issue · seen 2026-07-30

Discover AI agents built on Amazon Quick | NOTES ↗ | STABLE

Nudge Security now discovers AI agents built on Amazon Quick, surfacing creator, components, and risk insights in your AI agents inventory.

└──▷ HOW TO FIND IT

Enable browser-based agent discovery to start automatically surfacing Amazon Quick AI agents in your inventory.

> 📍 In the browser extension settings, turn on browser-based agent discovery. Then navigate to Identities › AI Agents tab in Nudge Security to view discovered Amazon Quick agents with their creator, components, and risk insights.

› Discovers AI agents your team builds on Amazon Quick (AWS's agentic AI platform) via browser extension, adding them to the AI agents inventory under Identities with creator, components, and risk insights.

WAS THIS USEFUL? 👍 👎

Secure Access (VPN / SASE / ZTNA)

tailscale ⓘ Tailnet creation API (2026-07-29) BLOG ↗

The easiest, most secure way to use WireGuard and 2FA.

New alpha API lets you programmatically create, list, and delete API-only tailnets in your organization.

- Adds a tailnet creation API (alpha) to programmatically create, list, and delete API-only tailnets within an organization.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

Compliance Automation & Trust Management

Mondoo ⓘ Connect Claude and other AI tools to Mondoo with the MCP server

NOTES ↗

Mondoo is a security platform that continuously scans and assesses infrastructure, cloud environments, and software supply chains for vulnerabilities and misconfigurations.

Mondoo launches a hosted MCP server, enabling Claude Desktop, Claude Code, and VS Code to query spaces, findings, and policies in plain language.

- New hosted MCP server lets AI tools (Claude Desktop, Claude Code, VS Code) connect to Mondoo using a read-only API token.
- Enables natural-language exploration of Mondoo spaces, findings, and policies directly from AI-powered clients.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Windsurf ⓘ snapshot-20260730 NOTES ↗

Windsurf is an AI-powered code editor that combines an intelligent coding assistant with agentic capabilities to help developers write, edit, and debug code faster.

Windsurf v3.6.21 adds browser preview for remote ACP agents, Cascade Migration Wizard, Plan mode for Devin Local, and editable command approvals.

└──▷ HOW TO FIND IT

Migrate all Windsurf hooks, skills, and memories to Devin in one guided flow with dry-run previews before committing any changes.

> 📍 Open the Command Palette (⌘P / Ctrl+P) → run "Devin: Open Cascade Migration Wizard" → review each hooks/skills/memories item with per-item opt-in and dry-run preview → confirm to apply.

Suppress 'Generated by Devin' attribution lines from commit messages produced by Devin Local in your repository.

```yaml
# In your Devin Local config file:
attribution: false
```

- › Adds a unified notifications setting that posts native OS notifications when any agent session finishes or needs input.
- › Codemaps now open in their own editor tabs in agent window mode, and codemap @`-mentions` now work in Devin Local.
- › New 'Add to chat' button (⌘L / Ctrl+L) lets users select text in a message transcript and append it to chat input.
- › Adds keyboard shortcuts for permission request actions (always-allow and reject) with shortcut hints shown on buttons.
- › Pull request editor tabs now show state-specific icons (open, draft, merged, or closed); session-create cards show a Devin mode badge (Fast, Ultra, Fusion).

+21 more

└──▷ BREAKING ON UPGRADE

! The status bar is now hidden by default in Agent Command Center mode; set `"workbench.statusBar.visible": true` in user settings to restore it.

- ! Cascade, Devin Local, and every other ACP agent are unavailable while a workspace is open in Restricted Mode, and hooks no longer load or run there.
- ! Models marked 'Devin Local only' (including every GPT-5.6 variant) now appear disabled in Cascade's model picker rather than running degraded.

WAS THIS USEFUL? 👍 👎

Cursor ⓘ Cursor, now on iPad 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 native iPad app with split-screen reviews, Apple Pencil markup, Inbox, and Bitbucket/Azure DevOps support.

.

└──▷ HOW TO FIND IT

Open a PR review alongside an active agent chat on iPad to read diffs and prompt the agent to resolve comments without switching screens.

> 📍 In the iPad app, open a PR from the Inbox, then use Split Screen to place the review alongside your active chat. Tap a comment and prompt the agent to resolve it directly from the review panel.

Annotate a screenshot with Apple Pencil to give the agent precise visual feedback on a UI issue.

> 📍 In a chat, attach a screenshot, then tap a specific point on the image to drop a comment, or draw on it with Apple Pencil before sending to the agent.

- Adds native iPad app (available on all paid plans) with a rebuilt layout featuring pinned sidebar chats for monitoring multiple agents simultaneously.
- Supports split-screen mode on iPad, keeping a PR review open alongside an active chat with full file diff rendering.
- Enables Apple Pencil markup on screenshots: tap to drop a comment at a specific point or draw directly on an image.
- Adds a full PR review surface on iPhone and iPad covering comments, checks, approvals, reviewer management, and agent-driven resolution.
- Introduces an Inbox to track in-progress work, items needing attention, and PRs currently in review.

+3 more

WAS THIS USEFUL? 👍 👎

OpenCode ⓘ v1.18.10 NOTES ↗ CODE ↗

The open source coding agent.

OpenCode v1.18.10 adds automatic Modal model discovery and desktop UX improvements.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.18.10
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.18.10
```

- Automatically discovers available Modal models, removing the need to manually configure them.
- Always shows the new session button in the desktop app for faster session creation.
- Prevents duplicate attachments from being added in the desktop composer.
- Improves toast notification stacking, dismissal, and mobile layout in the desktop app.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.147.0-alpha.2 NOTES ↗ CODE ↗

PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI adds MCP 2026 protocol support, cloud-managed sandbox profiles, persisted thread sections, and model-owned token budget defaults.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.147.0-alpha.2
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.147.0-alpha.2
```

- Adds MCP 2026-07-28 discovery support and completes MCP 2026 client support, enabling use of the latest Model Context Protocol spec.
- Exposes cached MCP tools to inference before the MCP server finishes initializing, reducing startup latency for tool-heavy workflows.
- Routes MCP OAuth requests through configured HTTP clients, enabling proxy-aware MCP authentication.
- Loads cloud-managed profiles for `codex sandbox`, allowing centrally managed sandbox configurations.
- Adds persisted sections for organizing threads in session history.

+11 more

WAS THIS USEFUL? 👍 👎

OpenHands ⓘ v1.7.0 NOTES ↗

OpenHands: AI-Driven Development

OpenHands v1.7.0 adds persistent agent memory, always-visible LLM selector, and secret overwriting from the edit form.

└──▷ GET THIS VERSION

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

- Adds a persistent agent memory toggle so agents can retain context across sessions.
- LLM selector is now always visible in the UI; agent-profile switching moves to the tools menu.
- Enables overwriting an existing secret's value directly from the secret edit form.

WAS THIS USEFUL? 👍 👎

Pi ⓘ v0.83.0 NOTES ↗

AI agent toolkit: unified LLM API, agent loop, TUI, coding agent CLI

Pi v0.83.0 adds credential export commands, headless OpenRouter login, and Claude Opus 5 via GitHub Copilot.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.83.0 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.83.0
```

└──▷ TRY IT

Pipe a bearer token into another tool (e.g., a curl call to an API) without manually copying credentials.

```
$ pi auth print-bearer-token | xargs -I{} curl -H "Authorization:
Bearer {}" https://api.example.com/endpoint
```

Export an API key into an environment variable for a script that calls the provider directly.

```
$ export API_KEY=$(pi auth print-api-key)
```

- Adds `pi auth print-api-key` and `pi auth print-bearer-token` commands to export configured credentials (with automatic OAuth refresh and minimum-validity enforcement) for use by external clients.
- Enables headless OpenRouter login by accepting a manually pasted redirect URL or authorization code when the loopback callback is unavailable (e.g., over SSH).
- Adds Claude Opus 5 support through the GitHub Copilot provider, with adaptive thinking and a 1 million-token context window.
- Adds per-request `fetch` injection for supported text and image provider transports.
- Adds the `"pending"` stop reason for partial streaming messages in the custom provider stream pattern.

└──▷ BREAKING ON UPGRADE

! TypeBox aliases upgraded to 1.3.7, removing deprecated APIs `Type.Base`, `Type.Awaited`, `Type.Promise`, `Type.AsyncIterator`, `Type.Iterator`,

`Type.Options`, and `Value.Mutate` — extensions using any of these APIs must migrate to supported TypeBox APIs.

WAS THIS USEFUL? 👍 👎

Goose ⓘ v1.45.0 NOTES ↗ CODE ↗

an open source, extensible AI agent that goes beyond code suggestions - install, execute, edit, and test with any LLM

Goose v1.45.0 adds tab completion for model switching, air-gapped docs support, disabling built-in skills, and Gemini/Opus5 model support.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.45.0 https://github.com/aaif-goose/goose.git
# already have the repo? check out this version:
$ git checkout v1.45.0
```

└──▷ TRY IT

Point Goose at a local mirror of the docs in an air-gapped environment so help content loads without internet access.

```
$ GOOSE_DOCS_ROOT=/mnt/internal/goose-docs goose session
```

Run the ACP server with scheduled recipe execution enabled (opt-in after the new default disables it).

```
$ goose server --enable-scheduler
```

- Adds tab completion and inline provider switching to the `model` command.
- Supports latest Gemini models and Opus5 (including adaptive thinking).
- Adds `GOOSE_DOCS_ROOT` environment variable for air-gapped / offline documentation access.
- Gates tool-call label enrichment on ACP client capability, improving ACP interoperability.
- Adds stable identity for agent event messages, enabling reliable message tracking across sessions.

└──▷ BREAKING ON UPGRADE

! The ACP scheduler is now disabled by default; deployments relying on scheduled recipe execution must explicitly pass `--enable-scheduler`.

WAS THIS USEFUL? 👍 👎

Cline ⓘ cli-v3.0.47 NOTES ↗ CODE ↗

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

Cline CLI v3.0.47 adds free-model support, persistent settings, and agentic compaction as the new default.

└──▷ **GET THIS VERSION**

```
$ git clone --branch cli-v3.0.47 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout cli-v3.0.47
```

- Adds support for free Cline models end-to-end, displayed with a "(free)" label and a dedicated card showing the reset time when the free usage limit is reached.
- Persists `/settings` general toggles — plan/act mode, tool auto-approve, and compaction mode — across CLI restarts.
- Expands built-in provider coverage, now generated from models.dev, with an updated bundled model catalog.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

holmesgpt ⓘ 0.38.0 NOTES ↗

SRE Agent - CNCF Sandbox Project

HolmesGPT 0.38.0 adds remote tool call approval workflows and cuts prompt token usage by 50–60%.

└──▷ GET THIS VERSION

```
$ git clone --branch 0.38.0
https://github.com/HolmesGPT/holmesgpt.git
# already have the repo? check out this version:
$ git checkout 0.38.0
```

- Implements remote tool call approval workflows, allowing human review and authorization of agent-initiated tool calls before execution.
- Reduces system prompt and tool description sizes by ~50–60%, significantly cutting prompt token consumption per investigation.
- Makes conversation-worker slot exhaustion visible and bounds reconnect sign-in, improving observability of capacity limits.

WAS THIS USEFUL? 👍 👎

CrewAI ⓘ 1.15.9 NOTES ↗

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

CrewAI 1.15.9 surfaces tool failures explicitly, emits FlowFailedEvent on flow errors, and adds progressive disclosure for skills.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.15.9 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.9
```

- Surfaces tool failures as actual failures instead of silently reporting them as success, enabling agents to react correctly to tool errors.
- Emits a new `FlowFailedEvent` when a flow execution fails, allowing event-driven error handling and observability hooks in Flows.
- Adds progressive disclosure for skills, surfacing contextually relevant skill guidance incrementally rather than all at once.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

AI-Infra-Guard ⓘ v4.5.1 NOTES ↗

A full-stack AI Red Teaming platform securing AI ecosystems via OpenClaw Security Scan, Agent Scan, Skills Scan, MCP scan, AI Infra scan and LLM jailbreak evaluation.

AI-Infra-Guard v4.5.1 adds multi-turn jailbreak attacks, 6 new OWASP agent-scan skills, and 4 new MCP security detection rules.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.5.1 https://github.com/Tencent/AI-Infra-Guard.git
# already have the repo? check out this version:
$ git checkout v4.5.1
```

- Adds Many-Shot, PAIR, GOAT, and ActorAttack multi-turn jailbreak attack methods to PromptSecurity.
- Adds 5 new OWASP-aligned Agent-Scan detection skills: agentic-supply-chain, cascading-failure, human-agent-trust, inter-agent-comm, and unexpected-code-execution.
- Adds web-exfiltration-detection skill to Agent-Scan.
- Adds 4 new MCP-Scan security detection rules covering hardcoded secrets and insecure deserialization.

WAS THIS USEFUL? 👍 👎

Knostic ⓘ

Introducing AgentMesh Access Tiers: Building Together, Growing Together (2026-07-27)

BLOG ↗

AgentMesh shifts to freemium with tiered API access, including a free signed-in API trial and paid tiers bundling Kirin endpoint licenses.

- Introduces anonymous (no-login) free tier with access to 100 items per category via the Web UI.
- Adds free signed-in tier unlocking 300 items per category plus trial API access (capped at 20 calls/month, 1/day) for engineers and security researchers.
- Launches Premium tier ($15/user/month) with full catalog access, 300 API calls/month (20/day cap), and a bundled Kirin endpoint agent license.
- Introduces Pro tier ($850/user/month) with unlimited AgentMesh API access and a Kirin IDE license for production pipelines and commercial integrations.
- Adds Enterprise tier with custom pricing, dedicated support, and deployment for 10+ users.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 30, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe