THE CYBER TOOLCHAIN

NO. 011

Head

THE WEEK'S TOP PICKS

PUBLISHED SEPTEMBER 11, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

// HOW THESE PICKS ARE MADE

Every feature release from the 916 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and choose our top picks, judged on two questions: **how deep and complete is the single best capability in the release**, and **how much it changes what you can actually do**.

A tool is judged on everything it shipped that week, so a project that releases daily gets credit for the sum — and still only takes one slot. We would rather run a short list than a padded one.

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 Opus 5

What stands out in this week's picks, grouped by what it lets you do. Every tool named links to its pick below.

Three unrelated things stand out today: Aurascape can police live AI coding sessions before secrets or destructive commands leave the workstation; GitHub can block merges on unresolved secret alerts; Tenzir can run stock Sigma rules on OCSF telemetry without hand-mapping fields.

PROTECT

Put guardrails around AI coding agents while they are working

Aurascape inspects live Claude Code, GitHub Copilot, and Cursor sessions for secrets, destructive commands, prompt injection, and supply-chain risk, so teams can enforce policy before generated actions become commits or shell damage. Cloudflare extended AI agent sandboxing to Cursor Cloud Agents, reducing the amount of trust placed in agent-run code paths.

Aurascape · Cloudflare

PROTECT

Stop risky code from merging instead of finding it after the fact

GitHub added REST APIs for AI Scan on pull requests and merge blocking on unresolved secret scanning alerts, turning findings into enforceable gates rather than advisory comments. The same release adds enterprise controls for Copilot agents and JetBrains sandboxes, giving platform teams a central way to constrain where automated development help can operate.

GitHub

DETECT

Run stock detection content against normalized telemetry with less field plumbing

Tenzir now matches SigmaHQ rules to OCSF -normalized events automatically, removing the brittle manual field-mapping step that often stops shared detection content from running in a local pipeline. Elastic Security Labs added Linux fileless-execution telemetry and simplified DLL search-order hijacking detection into a single-field rule, making two common detection paths cheaper to test and maintain.

Tenzir · Elastic Security Labs

PROTECT

Give AI and human access shorter lifetimes

CrowdStrike's Agentic Identity Provider gives AI agents cryptographic identities and brokers short-lived access, so automation does not need to inherit standing human credentials. Its expanded just-in-time and zero-standing-privilege access across SaaS apps, endpoints, private apps, and AWS also narrows how long an abused identity remains useful.

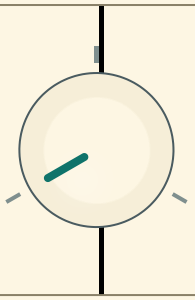
CrowdStrike Falcon Identity Protection

Does **GPT-5.5** read better? 👍 👎

VIEW

ISSUE VIEW

full issue



FULL ISSUE | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **full issue**? 👍 👎

// THE PICKS

01	**Aurascape**	AI/LLM Security
02	**powershell-empire**	Exploitation & C2
03	**GitHub**	SAST / SCA / Code Security
04	**Recorded Future**	Threat Intelligence
05	**RAPTOR**	AI-Powered Offensive Security

DENSITY | BRIEFED | **EVERYTHING** | Every change in this week's picks, in full.

all 76 changes

01 Aurascape

1 RELEASE · 2026-09-04

AI/LLM Security

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Coding Assistant Guardrails inspect and enforce policy on live AI coding sessions across Claude Code, Copilot, and Cursor — blocking secrets pasted into prompts, destructive shell commands and git pushes, and prompt injection hidden in READMEs and MCP outputs before the client acts on them. It decodes SSE/WebSocket/gRPC/MCP streams inline, catching a class of exposure that only exists in flight and never touches an artifact scanner.

> Aurascape launched Coding Assistant Guardrails, a new solution that inspects and enforces policy on live AI coding-assistant sessions across Claude Code, GitHub Copilot, and Cursor, adding wire-level detection for secrets, destructive commands, prompt injection, and supply-chain risks alongside audit logging, coaching, and natural-language querying.

└──▷ **WHAT SHIPPED · 11 FEATURES** most completely described first ? what's the number?

01 Inline decoding of coding-assistant streaming protocols NEW

65

Decodes coding-assistant streaming protocols — Server-Sent Events, WebSockets, Protobuf, gRPC, and MCP — inline without buffering, preserving native streaming speed while enabling content inspection.

— Explains mechanism and names five protocols, no user action

Secure the Coding Session, Not Just the …

02 Destructive command and git push interception NEW 60

Intercepts destructive shell commands and unauthorized git pushes at the wire level — in the tool call the model streams back — before the client executes them, with configurable block or confirmation-checkpoint policy.

— Describes wire-level mechanism and configurable policy modes

Secure the Coding Session, Not Just the …

THINNER COVERAGE BELOW

03 Coding Assistant Guardrails for live sessions NEW 55

A new solution that inspects and enforces policy on live AI coding-assistant sessions — covering Claude Code, GitHub Copilot, and Cursor — acting before any artifact exists to scan.

— Names three covered assistants but no config surface

Secure the Coding Session, Not Just the …

04 **In-flight secrets and credential detection** NEW 55

Detects and blocks or redacts secrets and credentials pasted into prompts before the model receives them, with the event captured as auditable evidence.

— Clear behavior described, no specific config or command

Secure the Coding Session, Not Just the …

05 **Continuous discovery of IDE plugins and MCP servers** NEW 55

Provides continuous discovery of IDE plugins, extensions, and MCP servers as they appear, with new or uncataloged apps receiving signatures on a 48-hour SLA — removing the need for hand-maintained allowlists.

— Names concrete 48-hour SLA and allowlist removal

Secure the Coding Session, Not Just the …

06 **Customizable in-context coaching templates** NEW 45

Offers customizable coaching templates that surface in-browser, in-IDE, or in-terminal explanations and safer paths when a policy action fires.

— Names three surfaces where coaching appears, no config detail

Secure the Coding Session, Not Just the …

07 **Auditable conversation logs with RBAC** NEW 40

Delivers auditable conversation logs — capturing prompts, responses, connections, and tool calls behind each policy flag — under role-based access control for security investigations.

— Names RBAC and log contents, no access path given

Secure the Coding Session, Not Just the …

08 **Prompt injection detection in docs and MCP outputs** NEW 35

Detects prompt injection hidden in READMEs, documentation, and MCP outputs, neutralizing hijacked instructions before the agent acts on them.

— States what it does, no mechanism or config detail

Secure the Coding Session, Not Just the …

09 **Typosquatted package and spoofed link detection** NEW 35

Flags typosquatted packages and spoofed links recommended by a model, surfacing supply-chain attempts as visible, reportable events.

— Describes outcome, no mechanism or threshold given

Secure the Coding Session, Not Just the …

10 **Natural-language queries via Auri** NEW 35

Enables natural-language queries over AI usage, risk, and policy data through Auri.

— Names Auri but no query examples or scope given

Secure the Coding Session, Not Just the …

11 **MCP tool call parameter visibility** NEW 30

Exposes MCP tool calls down to their parameters for visibility and policy enforcement.

— Brief statement, no mechanism or example given Secure the Coding Session, Not Just the …

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Recorded Future

1 RELEASE · 2026-09-09

Threat Intelligence

Get started with the Recorded Future API. Covers authentication, base URLs, API key setup, and how to integrate threat intelligence into your security workflows.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Launches Digital Risk Protection unifying malicious-site, impersonation, code-repository, dark-web and identity-exposure monitoring, surfacing compromised employee credentials from active infostealer logs within hours and lookalike phishing domains within hours of registration, with an AI triage agent and one-click takedown.

> Recorded Future launched Digital Risk Protection, a unified workspace bringing together five external threat-monitoring surfaces plus an AI Triage Agent, configurable risk thresholds, and one-click takedown coordination.

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

01 **Malicious Site Monitoring** NEW 60

Flags newly registered lookalike domains and phishing infrastructure within hours of registration, using logo detection and image OCR to catch impersonation beyond keyword matching.

— Names mechanism and timing but no config or UI path `launch-20260909-ba3a1237`

02 **Identity Exposure Monitoring** NEW 60

Surfaces compromised employee credentials from active infostealer logs within hours of exfiltration, with VIP monitoring available for executive accounts.

— Concrete timing and VIP scope, no config path `launch-20260909-ba3a1237`

03 **AI Triage Agent for high-interest detections** NEW 60

For Malicious Site Monitoring and Dark Web Brand Monitoring, an AI Triage Agent automatically evaluates high-interest detections and delivers an explicit verdict with reasoning, surfacing only critical, high-priority Alerts.

— Explains mechanism and scope across two surfaces `launch-20260909-ba3a1237`

THINNER COVERAGE BELOW

04 **Dark Web Brand Monitoring** NEW 55

Surfaces brand mentions, planned attacks, and compromised data across underground forums, marketplaces, and ransomware extortion sites.

– Names sources monitored but no mechanism launch-20260909-ba3a1237

05 **Impersonation Monitoring** NEW 45

Detects fake executive profiles and company impersonation accounts across social and professional networking platforms.

– States what it does, little mechanism detail launch-20260909-ba3a1237

06 **Code Repository Monitoring** NEW 45

Continuously scans public code repositories for exposed credentials, PII, and source code.

– Brief description of scope, no depth on scanning method launch-20260909-ba3a1237

07 **Customer Credential Monitoring add-on** NEW 45

A paid add-on that extends credential exposure visibility to an organization's customer base.

– Names add-on and scope, no further detail launch-20260909-ba3a1237

08 **Configurable risk thresholds per use case** NEW 45

Teams can set configurable risk thresholds per use case to control which detections escalate to Alerts without losing underlying detection coverage.

– Describes a settable control but no exact UI or key launch-20260909-ba3a1237

09 **One-click takedown coordination** NEW 40

One-click takedown coordination is available across all five Digital Risk Protection use cases from day one.

– Clear single action but no further mechanism given launch-20260909-ba3a1237

└──▷ **ALSO FROM THESE RELEASES**

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 powershell-empire v7.0.2

Exploitation & C2

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Empire v7 lands 30+ BOF modules for in-memory recon, LDAP/AD enumeration, service abuse, LSASS dumping via static syscalls, and LAPS extraction, giving operators post-exploitation tradecraft that runs without dropping to disk. It also adds AES-256-GCM across the stack, chunked uploads past the old 1MB cap, and a one-command self-update.

Empire v7 adds 30+ BOF modules, AES-256-GCM, chunked uploads, persistent chdir, and a new `ps-empire update` subcommand

└──▷ **GET THIS VERSION**

```
$ git clone --branch v7.0.2 https://github.com/BC-SECURITY/Empire.git
# already have the repo? check out this version:
$ git checkout v7.0.2
```

└──▷ **TRY IT**

Upgrade an Empire server to the latest release, automatically confirming the database backup prompt and applying any pending schema migrations.

```
$ ./ps-empire update -y
```

Serve an already-built Starkiller from a local directory for an air-gapped install, bypassing the GitHub clone step at startup.

config.user.yaml

```
starkiller:
  directory: /opt/starkiller/dist
```

- Adds `./ps-empire update` subcommand that moves the checkout to the latest release tag, refreshes `config.yaml`, fast-forwards Starkiller, the plugin registry, and the Empire Compiler binary, and applies pending Alembic migrations; use `-y` to auto-confirm the backup prompt.
- Adds `starkiller.directory` config option to serve an already-built Starkiller instead of cloning from GitHub, enabling air-gapped and distro-packaged installs.

› Adds `POST /api/v2/agents/{id}/tasks/chdir` endpoint and `TASK_CHDIR` (opcode 44) for persistently changing an agent's working directory across PowerShell, Python, and IronPython agents.

› Adds chunked file uploads to agents — files split into 512KB chunks drip-fed across checkins — removing the previous 1MB upload limit with backward-compatible protocol detection.

› Adds AES-256-GCM across the stack: an `AES256GCM` AEAD class in `encryption.py` (server), `agent/stagers/common/aesgcm.py` for Python/IronPython stagers, and a .NET 4.x `AesGcmHelper` for PowerShell stagers.

+45 more

└──▷ **BREAKING ON UPGRADE**

! Password hashing switched from bcrypt to PBKDF2-HMAC-SHA256 (600K iterations); existing bcrypt hashes are incompatible — reset passwords or recreate the database. Non-default bcrypt users must be manually reset via `PUT /api/v2/users/{id}/password`.

! Staging key normalization replaced MD5 with SHA-256; agents staged under the old algorithm must be re-staged.

! CSPRNG generation replaced `random` with `secrets` (server) and random.SystemRandom() (agent) across listeners, stagers, agents, and utilities; all deployed agents must be re-staged.

! `./ps-empire update` overwrites the base `config.yaml` on every run — move local customizations to `config.user.yaml` before upgrading.

! Debian 11 (bullseye) is no longer a supported install platform; it has been removed from the install script, CI matrix, and docs.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 GitHub

20 RELEASES · 2026-09-03 → 2026-09-10

SAST / SCA / Code Security

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

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5

Code Quality agentic autofix lets you bulk-assign up to 25 findings to Copilot to fix on a branch, self-validate, and open a reviewed PR. Paired with a new ruleset rule that blocks PR merges while secret scanning alerts stay open and enterprise-managed per-operation permissions for Copilot agent shell, file, and network actions.

> GitHub's biggest push this window is security governance: new REST APIs to manage AI Scan for pull requests and to block merges on unresolved secret scanning alerts, enterprise-managed permission controls for Copilot agents and JetBrains sandboxes, and GHES 3.22 reaching GA with ruleset bypass-by-user and required-reviewers rules. Alongside that, CodeQL 2.27.0 and 2.26.4 shipped a long run of query and language coverage updates, Copilot gained several new models (Claude Fable 5.1, Gemini 3.8 Flash, GPT-6 Astra, and the HydraFusion orchestration preview), and Actions/npm got supply-chain hardening via `cache-mode` and multiple npm trusted publishing configurations.

└──▷ **WHAT SHIPPED · 40 FEATURES** most completely described first ? what's the number?

01 **cache-mode control for Actions cache access** NEW 93

GitHub Actions adds `cache-mode` to scope least-privilege cache access at the workflow or job level, with values including `read`, `write`, and `write-only`; job-level settings override workflow-level ones, and a called reusable workflow cannot receive more cache access than its caller granted. A warning annotation is emitted when `cache-mode` is explicitly set to `write` or `write-only` for low-trust events such as `pull_request_target`, where `read` is the default.

Lock a PR workflow to read-only cache access to prevent a contributor's code from poisoning the cache used by trusted builds.

yaml

```yaml
on:
  pull_request:

permissions:
  contents: read

cache-mode: read

jobs:
  build:
    runs-on: ubuntu-latest
    steps:
      - uses: actions/checkout@v4
      - uses: actions/cache@v4
        with:
          path: ~/.npm
          key: ${{ runner.os }}-node-${{ hashFiles('**/package-lock.json') }}
      - run: npm ci
```

Allow a deployment job to save a freshly built artifact to the cache while preventing it from restoring — useful when you want to seed the cache without consuming a potentially tainted entry.

yaml

```yaml
on:
  push:
    branches: [main]

jobs:
  seed-cache:
    runs-on: ubuntu-latest
    cache-mode: write-only
    steps:
      - uses: actions/checkout@v4
      - uses: actions/cache@v4
        with:
          path: dist/
          key: ${{ runner.os }}-dist-${{ github.sha }}
      - run: npm run build
```

— Named config key, defaults, and two runnable workflow examples.

Control GitHub Actions cache access with…

02 **AI Scan management API for pull requests** NEW 90

New REST endpoints let admins read and update AI Scan for pull requests enablement programmatically: `GET` / `PATCH /orgs/{org}/code-scanning/ai-scan` at the organization level and `GET` / `PATCH /repos/{owner}/{repo}/code-scanning/ai-scan` at the repository level, in public preview. This enables at-scale rollout of AI-powered PR security detections without configuring each repository through the UI.

Enable AI Scan for pull requests on a specific repository without touching the GitHub UI.

```
$ curl -X PATCH \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  -d '{"enabled": true}' \
  https://api.github.com/repos/{owner}/{repo}/code-scanning/ai-scan
```

Check whether AI Scan for pull requests is enabled at the organization level before rolling out to individual repos.

```
$ curl -X GET \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  https://api.github.com/orgs/{org}/code-scanning/ai-scan
```

— Exact endpoints, methods, and runnable curl examples given.

AI Scan for pull request APIs in public …

---

03 **Block pull requests with unresolved secret scanning alerts**

NEW 87

Rulesets can require secret scanning alerts to be resolved before a PR can merge, configurable via the 'Require secret scanning alerts are resolved' rule under Repository > Rulesets, the REST API's `require_secret_scanning_alert_resolution` rule type (with a `secret_types` parameter), or the GraphQL `REQUIRE_SECRET_SCANNING_ALERT_RESOLUTION` rule type — available at repository, organization, and enterprise level, and configurable to cover custom or generic secret patterns beyond default provider patterns.

Enable the secret scanning alert resolution rule via the UI to protect your default branch without writing any API calls.

> 📍 1. Go to your repository (or organization/enterprise) Settings. 2. Navigate to Repository › Rulesets. 3. Create or edit a ruleset targeting your protected branches.

> 4. Enable 'Require secret scanning alerts are resolved'. 5. Optionally expand 'Secret types' to include custom or generic patterns. 6. Save the ruleset.

— Names REST field, GraphQL type, and exact UI path.

`Block pull requests with exposed secrets…`

---

04 **Runner deprecation API endpoint** NEW 84

New `GET /actions/runners/deprecations/{version}` endpoint (available at repository, organization, and enterprise level) returns `runner_version`, `runtime_deprecates_at`, and `registration_deprecates_at` so teams can plan runner upgrades before support ends.

Check when a specific runner version loses runtime and registration support, so you can schedule upgrades before deprecation.

```
$ curl -H "Authorization: Bearer <token>" \
  -H "Accept: application/vnd.github+json" \
  https://api.github.com/actions/runners/deprecations/2.315.0
```

— Exact endpoint path, method, and response fields with runnable curl.

`GitHub Actions: Early September 2026 upd…`

---

05 **New CodeQL queries and taint-tracking coverage across languages**

NEW 83

CodeQL 2.27.0 and 2.26.4 expand query and sink coverage across languages: the new `rust/command-line-injection` query detects uncontrolled command lines in Rust; `PQexec`, `PQexecParams`, `PQprepare`, `PQsendQuery`, `PQsendQueryParams`, and `PQsendPrepare` are added as SQL injection sinks for `cpp/sql-injection` and related C/C++ queries; taint tracking is added for OData action parameter binding in C# (`ODataActionParameters`, `Delta`); SQL injection sink models are added for Spring R2DBC `DatabaseClient` and the R2DBC SPI in Java/Kotlin; Python taint flow now propagates through `list.extend` and `list.insert` (matching existing `list.append` coverage); Java/Kotlin taint propagation is added through String.valueOf(Object) when the argument is a `CharSequence`; JavaScript/TypeScript gains support for regular expressions using the `d` flag and the React Native Worklets `'worklet'` directive; and `EnvironmentCheck` can now be specified via a models-as-data model for custom environment sanitizer definitions.

— Every new query, sink, and function named across both releases.

`CodeQL 2.27.0 adds support for Linux ARM…` `CodeQL 2.26.4 improves GitHub actions se…`

---

06 **Automatic Dependabot access to GitHub-hosted registries** NEW

80

Dependabot's `GITHUB_TOKEN` now automatically requests the `packages: read` permission and sends that token when pulling from `*.pkg.github.com` and `ghcr.io`, so no PAT-based registry entries in `dependabot.yml` are needed. Dependabot reuses existing 'Manage Actions access' grants on package settings pages to authenticate, covering every GitHub Packages ecosystem Dependabot supports.

Grant Dependabot read access to a private GitHub Package so it can resolve updates without a PAT — replaces any PAT-based registry block you previously added to dependabot.yml.

> 📍 1. Go to your organization's or personal account's Packages tab and open the target package. 2. Under 'Manage Actions access', click 'Add repository'. 3. Select the repository that runs Dependabot and set access to 'Read'. 4. Remove the corresponding PAT-based entry from dependabot.yml (no other changes needed).

— Names permission, token, registries, and config file with migration steps.

`Automatic Dependabot access to GitHub-ho…`

---

07 CodeQL per-platform CLI distribution and Linux ARM64 support

80

BREAKING

CodeQL 2.27.0 adds native Linux ARM64 support for the CLI and bundle via the `linux-arm64` per-platform release assets. The generic multi-platform `codeql.zip` CLI distribution will be removed in a future release, so users must switch to the per-platform `.zip` for their platform; the CLI now emits a warning when run from an all-platforms distribution, which can be suppressed by setting `CODEQL_ALLOW_ALL_PLATFORMS_DIST=true`.

— Names asset, env var, and required migration action.

`CodeQL 2.27.0 adds support for Linux ARM…`

---

08 Reusable workflow job context identity properties NEW

78

Four new `job` context properties — `job.workflow_ref`, `job.workflow_sha`, `job.workflow_repository`, and `job.workflow_file_path` — let a reusable workflow determine its own source identity at runtime; not available on GitHub Enterprise Server.

Surface the exact source file and commit of a reusable workflow at runtime for audit logging or conditional logic.

```yaml
jobs:
  called-job:
    runs-on: ubuntu-latest
    steps:
      - name: Print reusable workflow identity
        run: |
          echo "Workflow ref: ${{ job.workflow_ref }}"
          echo "Workflow SHA: ${{ job.workflow_sha }}"
          echo "Repository:   ${{ job.workflow_repository }}"
          echo "File path:    ${{ job.workflow_file_path }}"
```

— All four property names given with a runnable usage example.

GitHub Actions: Early September 2026 upd…

---

09 **New Copilot models: Claude Fable 5.1, Gemini 3.8 Flash, GPT-6 Astra**

NEW 78

Claude Fable 5.1 is now available for Copilot Pro+, Max, Business, and Enterprise users. Gemini 3.8 Flash rolls out to Copilot Pro, Pro+, Max, Business, and Enterprise users across Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot cloud agent, the Copilot app, JetBrains IDEs, Xcode, and Eclipse, with admin control via the model policy's 'default model enablement' toggle. GPT-6 Astra, OpenAI's agentic long-horizon coding model, is now generally available and selectable across VS Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, the Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse, with Enterprise/Business admins able to enable or restrict it via the same model policy control.

▶ Watch video

Enable or restrict Gemini 3.8 Flash for your organization before the gradual rollout reaches all users.

> 📍 In GitHub Copilot settings, go to Settings › Copilot › Model policy › Default model enablement and explicitly enable or disable Gemini 3.8 Flash for your Enterprise or Business plan.

— Names three models, plans, surfaces, and the admin toggle.

GitHub Copilot weekly releases — August … GPT-6 Astra is generally available in Gi… Gemini 3.8 Flash is now available in Git…

---

10 **vulnerability-alerts permission for GITHUB_TOKEN** NEW

75

A new `vulnerability-alerts` permission for `GITHUB_TOKEN` supports `read` and `none` values, enabling least-privilege read-only access to Dependabot alerts without broader scopes.

Grant a workflow read-only access to Dependabot alerts using least-privilege scoping instead of broad repository permissions.

```yaml
permissions:
  vulnerability-alerts: read

jobs:
  audit-alerts:
    runs-on: ubuntu-latest
    steps:
      - name: List Dependabot alerts
        run: gh api repos/${{ github.repository
}}/dependabot/alerts
```

— Named permission key with concrete values and a runnable YAML example.

GitHub Actions: Early September 2026 upd…

---

11 **Agentic autofix for Code Quality findings** NEW 75

An 'Assign to Copilot' bulk action lets you select up to 25 standard Code Quality findings and have Copilot fix them agentically on a branch, validate its own changes, and open a pull request for review. Available for repositories with GitHub Code Quality enabled on GitHub Team and GitHub Enterprise Cloud (including data residency), governed by the existing enterprise Code Quality policy with no separate policy to manage.

Burn down a backlog of code quality findings by assigning a batch to Copilot, which opens a PR with validated fixes for each one.

> 📍 In your repository, go to Code Quality › Findings. Select up to 25 standard findings using the checkboxes, then choose 'Assign to Copilot' from the bulk-action menu. Copilot will fix them on a new branch and open a pull request for your review.

— Exact limit, UI path, and plan availability given.

Remediate Code Quality findings with age…

---

12 **Ruleset bypass by user and required reviewers rule** NEW 73

Repository rulesets now support bypass by individual users, letting administrators add specific accounts (e.g. a service account) to a bypass list without creating a dedicated

role or team. A new required reviewers rule supports pattern matching on branches, files, and folders, plus a configurable minimum number of reviews per team.

Require the security team to review all changes to the default branch by adding a required reviewers rule to a repository ruleset — enforcing sign-off from teams beyond CODEOWNERS.

> 📍 In your repository, go to Settings › Rules › Rulesets › New ruleset, set the target branch to 'main', add a 'Required reviewers' rule, select your security team, and set the minimum number of required reviews.

Grant a service account ruleset bypass rights without creating a dedicated role or team — useful for automated merge workflows in CI/CD.

> 📍 In your repository, go to Settings › Rules › Rulesets, open the target ruleset, navigate to Bypass list, click 'Add bypass', search for the service account by username, and save.

— Two named ruleset capabilities with exact UI navigation steps.

`GitHub Enterprise Server 3.22 is now gen…`

---

13 **Multiple trusted publishing configurations for npm** NEW

72

npm packages can now have multiple trusted publishing (OIDC) configurations, each with independent repository, workflow, and environment criteria — a publish or stage is authorized if the incoming OIDC token matches any one configuration. Direct publishing is now opt-in per configuration, with staging-only as the default and every configuration able to stage without extra setup. The staged publishing queue's approval button is disabled while malware scanning is in progress and re-enables automatically once it completes, refreshing scan status every minute. The versions tab on npmjs.com now shows a per-version staged history, including whether each version was approved, rejected, or remains staged.

— Explains authorization logic, defaults, and UI behavior in detail.

`Multiple trusted publishing configuratio…`

---

14 **Refreshed pull requests listing page** NEW

72

A redesigned repository PR listing page enters public preview with content assist to help find and apply filters, advanced search supporting `AND` / `OR` keywords and nested searches, a collapsible sidebar with quick filters like 'Authored by me' and 'Involves me', a compact presentation mode, and richer per-PR context including status check counts, stack indicators, and unread-update indicators. Opt in via the 'Preview' badge on the pull requests tab, and switch back to the classic view the same way.

Opt in to the refreshed PR listing page and switch back to the classic view if needed.

📍 1. Navigate to the pull requests tab of any repository. 2. Click the 'Preview' badge at the top of the page to enter the refreshed experience. 3. To return to the classic view, click the 'Preview' badge again and choose 'Switch back'.

— Names every UI element and gives the opt-in path.

`Refreshed repository pull requests page …`

---

15 **Enterprise-managed sandbox policies and diagnostics in Copilot for JetBrains**

NEW 70

Administrators can centrally control Copilot's sandbox enablement, filesystem and network access, proxy settings, developer-tool access, and macOS Keychain access via managed restrictions that override user settings, visible under `GitHub Copilot > Sandbox` when the Editor Preview feature flag is on or a managed setting is configured (public preview). Enterprise policy diagnostics let admins verify that Copilot policies are correctly detected and enforced on a device.

— Names the settings location and covered access surfaces.

`Enterprise-managed sandbox in Copilot fo…`

---

16 **/ide command in Copilot CLI for JetBrains** NEW 70

The `/ide` command in GitHub Copilot CLI (public preview) connects a terminal session to JetBrains IDE context, including selections, diagnostics, and file references. Shell commands run in that session can consume IDE terminal environment variables and the project's configured Python interpreter, and can activate the project's local Python virtual environment.

— Names the command and exact behaviors it unlocks.

`Enterprise-managed sandbox in Copilot fo…`

17 GitHub Actions ControlCheck query refinements BREAKING

70

CodeQL changes how it evaluates author-association and actor-field checks in GitHub Actions event payloads (e.g. `github.event.pull_request.user.login`), treating them as protection only when the relevant field is actually populated for that event — this may surface additional alerts for workflows relying on previously ineffective checks, and environments are no longer treated as a sufficient sanitizer. The `actions/unpinned-tag` query also now detects mutable references to reusable workflows, extending pinning enforcement beyond actions to workflow calls.

— Explains before/after behavior and alert impact clearly.

CodeQL 2.27.0 adds support for Linux ARM… | CodeQL 2.26.4 improves GitHub actions se…

18 HydraFusion multi-model orchestration in Copilot NEW

70

HydraFusion, a research preview model selectable like any other model in GitHub Copilot, dynamically orchestrates multiple models across providers using Single, Cascade, or Critique execution patterns to balance quality, cost, and latency per task. Feedback on the preview can be submitted directly from the command line with `/feedback` in Copilot CLI.

Submit feedback on HydraFusion task quality directly from the terminal without leaving your workflow.

```
$ /feedback
```

— Names execution patterns and a runnable feedback command. launch-20260904-1f78a8ec

19 Enterprise-managed permissions for Copilot agent operations

NEW

69

Enterprise admins can centrally block, require human approval for, or auto-allow specific shell commands, file reads/edits, and network domains that Copilot agents perform, enforced across the GitHub Copilot app, GitHub Copilot CLI, and Visual Studio Code Agent Host sessions. Policies can be scoped per enterprise team, and managed restrictions override and cannot be weakened by user settings, workspace settings, auto-approval rules, or previously saved approvals.

— Mechanism and scope described, but no exact settings path given.

Enterprise managed permissions for GitHu…

20 ASP.NET Core CodeQL query improvements IMPROVED

65

CodeQL improves ASP.NET Core MVC controller and action discovery to more closely match runtime behavior, covering application parts, endpoint mappings, inherited actions, and exclusions (2.27.0); the `cs/web/missing-token-validation` query now recognizes enabled ASP.NET Core `RequireAntiforgeryToken` attributes when antiforgery middleware is used (2.26.4).

— Names the query and specific attribute it now recognizes.

CodeQL 2.27.0 adds support for Linux ARM… | CodeQL 2.26.4 improves GitHub actions se…

THINNER COVERAGE BELOW

21 **Micronaut framework modeling in CodeQL** NEW 59

CodeQL 2.27.0 adds modeling for the Micronaut framework in Java/Kotlin, covering HTTP controllers, WebSocket endpoints, configuration injection, data access, security annotations, and HTTP client sinks.

— Names framework and covered surfaces but no usage command.

CodeQL 2.27.0 adds support for Linux ARM…

22 **Private registry authentication for code scanning default setup**

NEW 52

GitHub code scanning default setup can now authenticate with container registries or the GitHub API using an organization's private registry configurations, supporting custom content pulled from private Git sources and Docker registries.

— Describes capability but no exact config surface named.

CodeQL 2.27.0 adds support for Linux ARM…

23 **Air-gapped Copilot CLI support for GitHub Enterprise Server**

NEW 52

GHES 3.22 (technical preview) supports configuring Copilot CLI to work in disconnected or air-gapped environments — admins set up a model provider once in GHES and end users authenticate with their GHES credentials.

— Explains setup flow but no exact config keys given.

GitHub Enterprise Server 3.22 is now gen…

24 **Unified customer portal at help.github.com** NEW 52

A new portal at `help.github.com` consolidates support, entitlements, docs, learning, community, and account resources into a single entry point, with Copilot-powered search across all portal content available without sign-in, and product status and important announcements surfaced directly in the portal.

— Names the URL and consolidated content areas. New customer portal help.github.com (202…

25 **VS Code Copilot chat and agent session enhancements** NEW

50

VS Code gains experimental multi-root workspaces support, bringing Copilot and Claude agent sessions to every folder in a workspace; experimental chat backgrounds to personalize the Agents window with built-in patterns or custom images; and chat sessions that organize related chats into a hierarchy and surface which ones need attention.

— Names three distinct experimental additions in one area.

GitHub Copilot weekly releases — August …

---

26 **Agent Merge in VS Code** NEW

49

Agent Merge enters public preview in VS Code 1.136, resolving review feedback, failed checks, and merge conflicts to get pull requests merge-ready.

— Names the VS Code version and what it resolves.

GitHub Copilot weekly releases — August …

---

27 **Rust data flow alert location fix in CodeQL** BREAKING

48

CodeQL 2.26.4 changes Rust data flow query alert locations to reflect actual source and sink nodes; as a result some alerts will appear at new locations — showing as new alerts while the previously, differently-located alerts close.

— Explains before/after but no affected query names given.

CodeQL 2.26.4 improves GitHub actions se…

---

28 **Sort secret scanning bypass and dismissal requests** IMPROVED

48

Secret scanning push protection bypass requests and alert dismissal requests can now be sorted by date (ascending or descending) via the filter bar at repository, organization, and enterprise levels.

— Names the exact UI control and levels it applies at.

GitHub Enterprise Server 3.22 is now gen…

---

29 **Privacy-safe star history API endpoint** NEW

45

A new REST API endpoint returns historical repository star counts with timestamps without exposing individual stargazer identities.

— Describes the endpoint's data but not its exact path.

New API endpoint provides privacy-safe s…

---

30 **Xcode 27 runner image on macOS 27** IMPROVED

42

The GitHub-hosted `xcode-27` and `xcode-27-xlarge` runner images (arm64, public preview) now run on macOS 27, up from macOS 26, enabling Apple app validation against the latest OS.

– Names the runner labels and OS versions but no workflow example.

Xcode 27 runner image now runs on macOS …

---

31 **Release and contributor context in issue/PR lists** NEW

41

The issue sidebar now shows a 'Latest release' or 'Pre-release' badge when a linked pull request has been included in a release, and the pull request list view for public repositories shows contributor role labels ('First-time contributor', 'Contributor', 'Member') directly in the list.

– Names both UI surfaces and label text shown. GitHub Enterprise Server 3.22 is now gen…

---

32 **Global project context in Copilot chat** NEW 40

Chat in Copilot for JetBrains gains global project context, letting you include global files and folders as persistent context across a project rather than re-adding them per session.

– Describes mechanism with a screenshot of the dialog.

Enterprise-managed sandbox in Copilot fo…

---

33 **Cross-file cursor jumps in next edit suggestions** NEW

38

Copilot's next edit suggestions can now move your cursor directly to the relevant location in another file when a suggested change continues there, in Copilot for JetBrains.

– Describes the behavior but not how it's triggered.

Enterprise-managed sandbox in Copilot fo…

---

34 **Enterprise teams generally available** IMPROVED 36

Enterprise teams reach general availability, letting enterprise owners manage users and access across all organizations and repositories from a single centralized team structure.

– States scope and GA status, no navigation path given.

GitHub Enterprise Server 3.22 is now gen…

---

35 **Content exclusions honored in Copilot app and CLI** IMPROVED

31

The GitHub Copilot app and CLI now honor content exclusions, keeping sensitive code out of context across agentic workflows.

– Bare statement with no config surface named. GitHub Copilot weekly releases – August …

36 **Session model selection for Copilot subagents** NEW 30

Copilot for JetBrains adds support for selecting the session model used by built-in subagents in the Copilot agent harness.

– Bare capability statement, no path or example given.

Enterprise-managed sandbox in Copilot fo…

37 **Go 1.27 language support in CodeQL** NEW 25

CodeQL 2.26.4 adds support for analyzing Go 1.27 code.

– Bare statement of added language version support.

CodeQL 2.26.4 improves GitHub actions se…

38 **OpenTelemetry settings GA in Copilot Chat** IMPROVED 25

OpenTelemetry settings in 'GitHub Copilot – Chat' are promoted to general availability for all users in Copilot for JetBrains.

– Bare GA promotion notice with no configuration detail.

Enterprise-managed sandbox in Copilot fo…

39 **Copilot agent harness GA in JetBrains** IMPROVED 25

The GitHub Copilot harness reaches general availability in Copilot for JetBrains, enabling faster feature delivery and better code quality.

– Bare GA announcement without mechanism detail. GitHub Copilot weekly releases – August …

40 **Enterprise-managed sandbox policies and diagnostics in Copilot for JetBrains**

NEW 0

placeholder

– duplicate-removed Enterprise-managed sandbox in Copilot fo…

└──▷ **BREAKING ON UPGRADE**

! The generic multi-platform `codeql.zip` CLI distribution will be removed in a future release; users must switch to the per-platform `.zip` for their platform. The CLI now emits a warning when run from an all-platforms distribution.

! Rust data flow query alert locations now reflect actual source and sink nodes — some alerts will change location, appearing as new alerts while previous alerts close.

! GitHub Actions `ControlCheck` `-based` queries may produce more alerts because actor fields from event payloads are no longer treated as protection for events that do not populate them, and environments are no longer a sufficient sanitizer.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 RAPTOR

1 RELEASE · 2026-09-07

AI-Powered Offensive Security

RAPTOR chains static analysis, binary analysis, LLM vulnerability validation, exploit generation, and patch writing into one autonomous security research workflow.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

RAPTOR now generates an agentic checker per confirmed finding and adds a full SCA supply-chain pipeline, plus OCI-registry, Dockerfile, and tar scan substrates so it hunts and validates vulnerabilities across source and container images autonomously. Multi-LLM dispatch and Frida reachability wiring make the confirmed findings a team would otherwise miss the headline, not a faster UI.

> RAPTOR v3.1.0 adds container/tarball scan substrates (OCI , Dockerfile, tar), a full software composition analysis pipeline, multi-LLM provider support, and Frida -based reachability integration, alongside new compiler-analyzer and sandbox-audit controls.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 Compiler-analyzer channel for static analysis NEW 86

Adds `--compiler-scan` / `--no-compiler-scan` flags to enable or disable a compiler-analyzer channel using gcc `-fanalyzer` or clang `--analyze`, and `--compiler-scan-max-tus <n>` to cap the number of translation units analyzed (default 2000).

Run the compiler-analyzer channel on a C/C++ codebase alongside Semgrep, capping the number of translation units to keep the scan tractable on a large repo.

```
$ /scan /path/to/c-project --compiler-scan --compiler-scan-max-tus 500 --build-command 'make -j4'
```

— Flags, tool backends, and a default value all named. v3.1.0

02 Sandbox audit mode verbosity and budget controls NEW 78

Adds `--audit-verbose` for strace-style verbose output during sandbox audit runs, and `--audit-budget <n>` to cap the maximum number of audit events collected per run.

Audit a target's sandbox behavior before enforcing — log every syscall that would have been blocked without actually blocking anything, with a cap to avoid log flooding.

```
$ /scan /path/to/target --audit --audit-verbose --audit-budget
500
```

— Two exact flags with a runnable example command. v3.1.0

03 **Expanded-Semgrep for macro-heavy C/C++** NEW 70

Adds `--expanded-semgrep` to re-run Semgrep rules over preprocessor-expanded views of macro-heavy C/C++ translation units, catching bugs that only appear after macro expansion.

Scan a macro-heavy C/C++ codebase and catch bugs that only appear in expanded preprocessor output, which plain Semgrep misses.

```
$ /scan /path/to/c-project --expanded-semgrep --languages c,cpp
```

— Named flag with a runnable example and clear mechanism. v3.1.0

THINNER COVERAGE BELOW

04 **Reachability analysis: framework registration and Frida integration**

NEW 52

Adds JS and Go function-as-argument framework registration for reachability analysis, and a Frida evidence pipeline with full consumer wiring integrated into the reachability pipeline.

— Names languages and Frida but no invocation detail. v3.1.0

05 **Container and tarball scan substrates** NEW 50

Adds OCI registry, Dockerfile, and tar as scan substrates, enabling RAPTOR to analyze container images and tarballs directly.

— Names three substrates but no command shown. v3.1.0

06 **Web correlation tier for vulnerability analysis** NEW

50

Adds a web correlation tier with expectations, check legs, and a poisoned-render oracle for web vulnerability analysis.

— Names distinct components but no usage instructions. v3.1.0

07 **Multi-LLM provider support** NEW 48

Adds full multi-LLM provider support for the analysis dispatch layer, with Mantle and InvokeModel coexistence.

— Names providers but no config surface given. v3.1.0

08 **CVE environment modeling package** NEW 45

Adds `cve-env` as an integrated package (`packages/cve_env`) for CVE environment modeling (Phase 1).

— Names the package path but phase-1 scope unexplained. v3.1.0

09 **Mechanical sink discovery in understand module** NEW 45

Adds mechanical sink discovery and framework API detection (MAP-5f) to the `understand` module.

— Named module and internal code but no usage path. v3.1.0

10 **Per-function annotations in validate module** NEW 43

Adds per-function annotations emitted across validation stages A, B, and D via the `validate` module.

— Names module and stages but no mechanism detail. v3.1.0

11 **Corpus runner for rule calibration** NEW 38

Adds a corpus runner with rule evaluation and label linting for calibrating detection rules.

— Names components but no usage path. v3.1.0

12 **Software composition analysis pipeline** NEW 35

Adds a full SCA (software composition analysis) mechanical pipeline with a supply-chain detector and first-party hardening.

— Named capability but no mechanism or usage detail. v3.1.0

13 **Agentic checker synthesis** NEW 35

Adds agentic checker synthesis: the system now generates a checker per confirmed finding.

— Describes behavior but no named surface or usage. v3.1.0

14 **Persistent verdict-bearing coverage store** NEW 35

Adds a persistent verdict-bearing coverage store wired into the live pipeline.

— Thin description, no mechanism or access path given. v3.1.0

15 **Proxy-hosts auto-calibration for CodeQL** NEW 27

Adds proxy-hosts auto-calibration for CodeQL analysis.

— One-line addition with no further mechanism. v3.1.0

16 **Source intelligence substrate** NEW 25

Adds a source intelligence substrate with a minimal axis-1 rule for source-level intel gathering.

— Bare description, minimal scope stated as 'minimal'. v3.1.0

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · September 11, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe