THE CYBER TOOLCHAIN

NO. 020

Head

THE WEEK'S TOP PICKS

PUBLISHED NOVEMBER 14, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

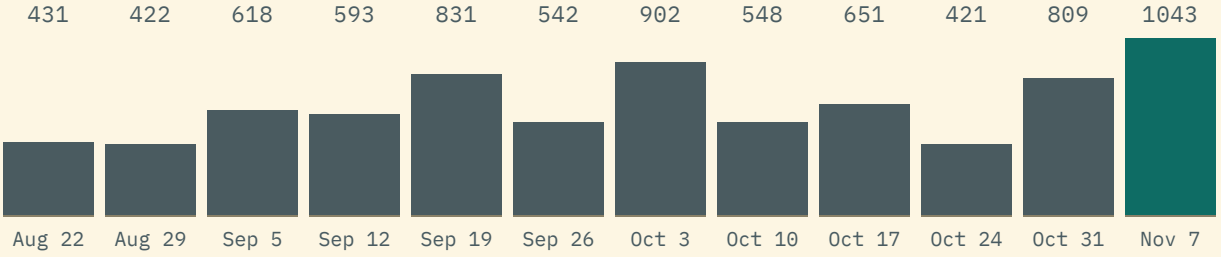
The best tooling updates that shipped this week.

The week in numbers · Nov 7–13

1,043	303	11	274
changes we read →	tools shipped →	product launches →	breaking changes →

Changes per week, last 12 weeks · Nov 7–13

Changes we read in that week's dailies, including docs and API changes. A gap means no daily was published.



Category moves this month

Category moves are held until the measured rollout checks pass.

vim ? help

▸ // HOW THESE PICKS ARE MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every change in this week's picks, in full. all 45 changes

Do you prefer **scroll**? 👍 👎

This week's 5

01	**Semgrep**	Launch · Code Security
02	**Endor Labs**	Feature · Software Supply Chain Security
03	**Zscaler**	Feature · Secure Access (VPN / SASE / ZTNA)
04	**Prowler**	Feature · Cloud Security Posture
05	**teleport**	Launch · Privileged Access Management

The picks in full

01 Semgrep

announced · Nov 12, 2026

Code Security

Semgrep is a static analysis tool that scans code for bugs, security vulnerabilities, and policy violations using customizable rules.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5 NOVELTY 4/5

Semgrep ships Mandoline — a new program analysis engine scoring 100% on the OWASP Benchmark for Java with no AI token usage — alongside Guardian AI-powered detection of IDOR and BOLA at code-generation time, and a Context Layer that surfaces 2x–7x more real vulnerabilities in agentic workflows; together these close the gap between rule-based SAST and the class of business-logic flaws that rule engines structurally cannot find.

> Semgrep shipped a hosted MCP server that lets AI coding agents triage findings, run scans, and manage policies, alongside a major expansion of Semgrep Guardian — org-wide policy enforcement, AI-powered IDOR/BOLA detection, lockfile protection, and installs across Claude Code, Codex, and Copilot — plus new Supply Chain lockfile and registry coverage, broader language parsing, and several breaking changes to CLI platform support and access control.

└──▷ WHAT SHIPPED · 26 FEATURES most completely described first ? what's the number?

01 **Hosted Semgrep MCP server** NEW INTEGRATION NEW 100

Semgrep's hosted MCP server (public beta) lets Claude Code, Codex, Cursor, VS Code, ChatGPT, or any MCP client that supports remote HTTP servers search and filter findings across Semgrep Code, Supply Chain, Secrets, and Agentic Workflows, triage and reopen findings in bulk, read code snippets and finding history, start and retry managed scans, read scan logs, look up advisories and dependencies, dry-run and apply policy changes, manage Guardian rules, pull reports by severity/product/rule, trigger and schedule Agentic Workflows (admins only), create and link Jira/Linear tickets, and check SCM connections — with permissions following your AppSec Platform role. Authenticate via OAuth (browser sign-in scoped to an organization) or an API token: a Web API token allows all read/write operations, a Read-only token allows reads only, and an Agent (CI) token covers read/write on scans (Member tokens aren't supported). A `discover` tool finds the right MCP tool for a task and a `semgrep_account` tool returns the connected organization. The `semgrep mcp` server can also scan files outside the MCP client's workspace root by setting `SEMGREP_MCP_ALLOW_SCAN_OUTSIDE_WORKSPACE=true`. Available to organizations on Teams or Enterprise plans (trials included).

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

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

Connect Cursor (or any MCP client that supports remote HTTP servers) to the Semgrep MCP server so your agent can triage findings and manage policies without leaving the editor.

> 📍 In Cursor, open Settings › MCP and add the Semgrep MCP server URL. Choose OAuth to sign in through your browser and scope access to your organization, or supply a Web API token for full read and write operations. Once connected, your agent can call the `discover` tool to list available capabilities or `semgrep_account` to confirm which organization it is connected to.

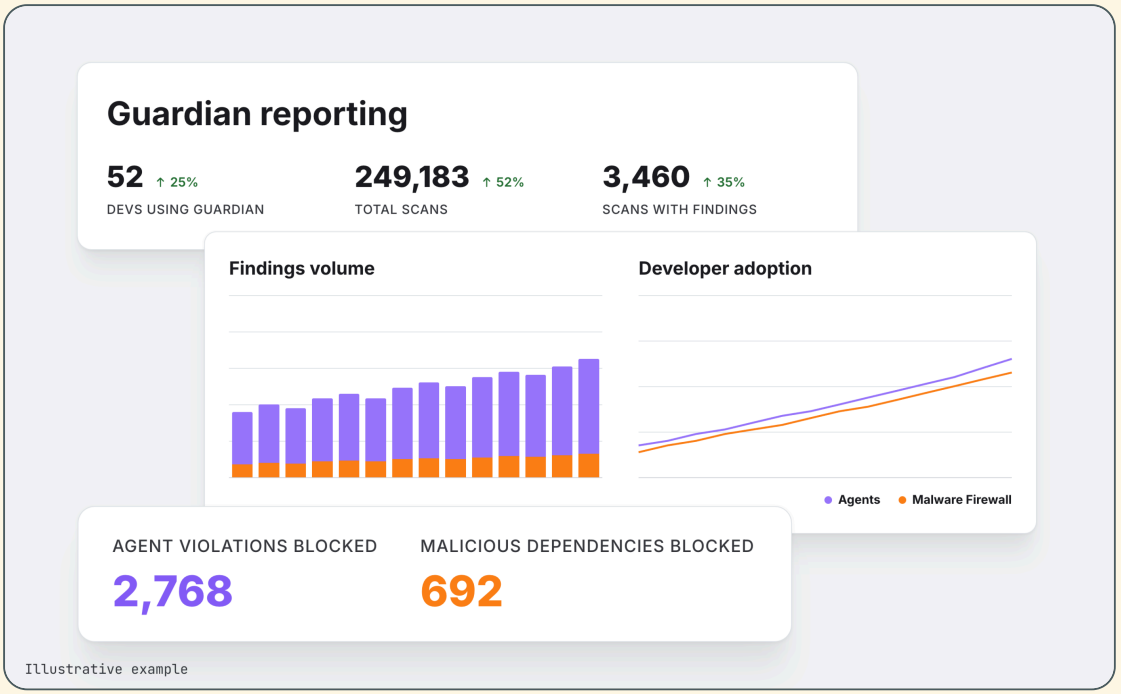
— Exhaustive named tools, tokens, env var and plan gating `launch-20261112-f2d6b2d0` `snapshot-20261112`

---

02 Guardian org-wide policy system and local rules (GA) IMPROVED

95

Security teams can define custom YAML rules (such as `no-legacy-auth`) in a local rules file or repo-wide policy file that Guardian automatically applies to every developer's agent session without manual distribution. Developers can also store local rules in `~/.semgrep/guardian-rules/` or the root of a repository; Guardian runs up to 1,000 organization rules and 50 local developer rules per scan. Guardian Policies and local developer rules are now generally available for all organizations, with per-rule control in the Claude Code remote plugin and the Codex plugin.



Illustrative example

Define a custom org rule that blocks a deprecated authentication library, then commit it to your local rules file so Guardian enforces it in every agent session.

```yaml
rules:
  - id: no-legacy-auth
    languages: [python]
    severity: ERROR
    message: Use auth.v2.verify() — legacy_auth is deprecated.
    pattern: legacy_auth(...)
```

Drop custom Guardian rules into the per-user directory so they run locally in your AI coding agent without affecting CI policy.

```
$ cp my-org-rules.yaml ~/.semgrep/guardian-rules/
```

— Exact rule paths, limits and GA status given `launch-20261112-c889153f` `snapshot-20261107`

03

Guardian coding agent integrations NEW

90

Semgrep Guardian can now be installed directly into coding agents: in Claude Code via `claude plugin install semgrep@claude-plugins-official`, in Codex via `codex plugin marketplace add https://github.com/semgrep/guardian --sparse .claude-plugin --sparse plugin` followed by `codex plugin add semgrep@semgrep-marketplace`, and in Copilot via `copilot plugin marketplace add https://github.com/semgrep/guardian` followed by `copilot plugin install guardian-copilot@semgrep-marketplace`. Guardian also supports Microsoft Copilot alongside Codex and Claude Code, Codex now gets automatic code, dependency, and secret scanning without a Semgrep CLI install, and Semgrep Guardian for GitHub Copilot CLI is in public beta on macOS and Linux.

Bootstrap Guardian in Claude Code so every future session enforces your org's Semgrep policies automatically.

```
$ claude plugin install semgrep@claude-plugins-official
```

Add Guardian to Codex so the agent can detect, fix, and report Semgrep findings without leaving the agent environment.

```
$ codex plugin marketplace add
https://github.com/semgrep/guardian --sparse .claude-plugin --
sparse plugin
codex plugin add semgrep@semgrep-marketplace
```

– Exact install commands given per coding agent `launch-20261112-c889153f` `launch-20261103-a1d24311` `snapshot-20261112`

---

04 **Supply Chain dependency paths and local-build diagnostics**

NEW 90

Supply Chain now shows the dependency path for transitive findings in Go projects when `--allow-local-builds` and dependency paths are turned on. With `--allow-local-builds`, running `semgrep ci --no-suppress-errors` now exits with code `2` when dependency resolution fails for a subproject, listing the affected dependency sources in scan output.

– Exact flags, command and exit code given `snapshot-20261112`

---

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

Set the `SEMGREP_GUARDIAN_OBSERVE_ONLY` environment variable to `1` to let Guardian scan as usual while hiding results from developers and coding agents until you're ready to roll it out. Observe Mode collects metrics in the background and makes no changes to the developer experience during staged rollouts, letting teams validate impact before enabling blocking.

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

```
$ export SEMGREP_GUARDIAN_OBSERVE_ONLY=1
```

– Exact env var and rollout behaviour given `snapshot-20261112` `launch-20261112-c889153f`

---

06 **Detect-2Shell workflow and Agentic Workflow scheduling** NEW 85

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

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

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

— Names workflow, scan surfaces, scheduling requirement `launch-20261103-a1d24311`

07 Supply Chain Bun, requirements.txt, and private registry support

NEW

85

Supply Chain now parses `bun.lock`, giving Bun projects dependency, malware, and vulnerability coverage. Every package in `requirements.txt`, pinned or not, now appears in the dependency inventory and SBOM. Supply Chain can also resolve private Maven or Gradle packages from AWS CodeArtifact for transitive dependency tracking, configured via an IAM role ARN under Settings > Supply Chain > Private Registries.

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

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

— Names lockfile formats and exact settings path `launch-20261103-a1d24311`

08 AI provider settings: bring-your-own-keys expansion NEW

80

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

— Names every provider and configuration field `snapshot-20261112` `launch-20261103-a1d24311`

09 Expanded language parsing across Bash, Clojure, Swift, C#, Terraform, Lua, Objective-C

NEW SUPPORTED PLATFORM NEW

75

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

Terraform provider functions with a namespace (e.g. provider::fn(...)), conditional and loop template directives, Lua 5.4 variable attributes, and Lua 5.5 leading attributes, global declarations, and named variable-argument parameters. Experimental support for Objective-C parses `.m` files and runs search-mode and taint-mode rules with metavariables and ellipsis operators.

— Names every language construct supported `snapshot-20261112` `snapshot-20261107`

10 **Guardian lockfile protection and incident alerts** NEW

70

Guardian blocks agent-initiated changes to lockfiles that would introduce malware or a reachable vulnerability, preventing unsafe dependency changes before they reach the build system. It also pushes advisories to every developer to check whether the version on their machine is still vulnerable when a critical supply chain incident (such as PHP2Shell or Node2Shell) is disclosed.

— Describes blocking mechanism and named incident examples `launch-20261112-c889153f`

11 **Advisories page filtering and matching** NEW

70

A new Advisories with findings filter on the Advisories page shows only advisories affecting your projects, respecting project-level role-based access control and excluding archived projects. An advisory's matching projects now include a 'Potentially matching: declared in a version range' section listing each project, branch, file, and declared version, and Incident Slack notifications report how many such projects exist and which packages they declare.

— Names filter, matching section and Slack detail `snapshot-20261107` `snapshot-20261112`

12 **Intel Mac binaries dropped** BREAKING

70

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

— Exact last supported version and migration path given `snapshot-20261112`

13 **Findings CSV export columns for triage and tickets** NEW

60

Findings CSV exports now include Last triaged by, Ticket key, and Ticket URL columns; ticket columns appear only when a Jira integration is active.

— Exact column names given, no further detail `snapshot-20261107`

14 **Autotriage cost reduction** IMPROVED

60

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

– Exact cost figures and effective date given `snapshot-20261112`

THINNER COVERAGE BELOW

15 **Guardian AI-powered vulnerability detection** NEW DETECTION NEW

55

Guardian's AI-powered detection reviews code at generation time and reasons about complex vulnerabilities including insecure direct object references (IDOR) and broken object-level authorization (BOLA), as well as code written with malicious intent.

– Names vulnerability classes but no config or command `launch-20261112-c889153f`

16 **Self-service Semgrep Secrets purchase** NEW 55

Self-service organizations can now buy Semgrep Secrets standalone or bundled, and manage or cancel subscriptions via the Stripe customer portal by clicking Edit plan.

– Exact purchase flow and UI control named `snapshot-20261107`

17 **Issues API issueUrl field** NEW 50

The issues API now returns an `issueUrl` field containing a link to each issue in AppSec Platform.

– Exact field name given but no endpoint path `snapshot-20261107`

18 **Third-party licenses in CLI --help** NEW 50

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

– Exact flag and link named `snapshot-20261112`

19 **Wiz integration expanded to Azure DevOps and Bitbucket**

IMPROVED 50

The Wiz integration now sends findings from Azure DevOps, Bitbucket Cloud, and Bitbucket Data Center repositories in addition to GitHub and GitLab.

– Names newly supported source platforms `launch-20261103-a1d24311`

20 **Read-only role bulk triage restriction on Secrets findings**

50

BREAKING

Users with the Read-only role and API tokens with the Read-only scope can no longer bulk ignore or reopen Semgrep Secrets findings; bulk triage now requires triage permission.

– Names exact role/scope affected and new requirement `snapshot-20261107`

21 **Managed Scans source restrictions for Free/Teams plans**

50

BREAKING

Free and Teams plan organizations can no longer select GitHub Enterprise, GitLab Self-Managed, or Bitbucket Data Center when setting up Semgrep Managed Scans.

— Names exact plans and restricted source options `snapshot-20261107`

---

22 **Guardian 30-day self-serve trial** NEW 45

New organizations can start a 30-day free trial of Guardian and upgrade in the product without a sales call.

— Trial length and self-serve upgrade path named `launch-20261103-a1d24311`

---

23 **Notes on Agentic Workflows issues** NEW 40

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

— Exact UI control named, no further mechanism `snapshot-20261112`

---

24 **Guardian compromised extension alert** NEW 35

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

— States the alert but no remediation steps `launch-20261103-a1d24311`

---

25 **Guardian auto-update with SLSA verification** IMPROVED 35

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

— Names SLSA level but no further detail `launch-20261103-a1d24311`

---

26 **Guardian PowerShell hook support** NEW 35

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

— Names shells supported, nothing more `launch-20261103-a1d24311`

NEW COMPLIANCE FRAMEWORK Mandoline OWASP Benchmark and NIST Juliet test suite coverage

NEW SUPPORTED PLATFORM EU multi-tenant region

NEW INTEGRATION OpenAI Daybreak Defense Factory import integration

└──▷ BREAKING ON UPGRADE

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

Python on Apple silicon.

Our best guess

Check this claim Claim history

snapshot-20261112

! Users with the Read-only role and API tokens with the Read-only scope can no longer bulk ignore or reopen Semgrep Secrets findings; bulk triage now requires triage permission.

Our best guess

Check this claim Claim history

snapshot-20261107

! Free and Teams plan organizations can no longer select GitHub Enterprise, GitLab Self-Managed, or Bitbucket Data Center when setting up Semgrep Managed Scans.

Our best guess

Check this claim Claim history

snapshot-20261107

announced · Nov 12, 2026

Semgrep adds Mandoline, an agent-oriented static-analysis engine with CLI, MCP, and gRPC interfaces, callgraph/dataflow/taint primitives, and UniFFI bindings for six languages.

Access Mandoline via CLI, MCP, or gRPC interfaces, with callgraph output available in text, JSON, GRAPHVIZ, and protobuf formats.

Semgrep adds Mandoline, a decomposed static-analysis engine built for agentic workflows, exposing callgraph, dataflow, taint paths, API endpoint detection, and dependency reachability as composable primitives.

Mandoline adds native library bindings for six languages via UniFFI, plus WASM support, enabling use in agent containers, edge compute, and the browser or IDE.

Semgrep customers now have access to Mandoline inside Semgrep Agentic Workflows and Semgrep's Context Layer.

NEW SUPPORTED PLATFORM Mandoline language and framework support

NEW COMPLIANCE FRAMEWORK Mandoline OWASP Benchmark and NIST Juliet test suite coverage

announced · Nov 12, 2026

Semgrep adds Mandoline (a millisecond-speed program analysis engine), Context Layer (private beta, persistent codebase context), Agentic Workflows, and Semgrep MCP for agent-native AppSec.

Semgrep adds Mandoline, a next-generation program analysis engine that answers questions about code flows and reachable paths in milliseconds, scoring 100% on the OWASP Benchmark for Java and 98.28% on Python in a single run with no AI token usage.

Semgrep adds Context Layer, now in private beta, giving agentic workflows persistent, dynamic understanding of a codebase and organization across runs, enabling agents to surface 2x–7x more real vulnerabilities compared to agents without it.

Semgrep adds Agentic Workflows, powered by Mandoline and Context Layer, to automate complex detection, triage, and fix tasks following guided paths set by Semgrep's security experts.

NEW INTEGRATION Access the full Semgrep platform from any agent by pointing it at the Semgrep MCP server, available in Codex, Claude Code, Cursor, and CLI.

NEW SUPPORTED PLATFORM Semgrep adds deployment options including bring-your-own-key (BYOK) models, Agentic Workflows in CI, Semgrep Dedicated single-tenant hosting, and an EU multi-tenant region.

NEW DETECTION AI detection for authorization, IDOR, and business logic flaws in Guardian

NEW INTEGRATION OpenAI Daybreak Defense Factory import integration

---

GOOD PICK?  GOOD REASON? 

02 Endor Labs

Introducing AI Sast That Thinks Like A Security Engineer |… | Endor Labs (2026-11-11)

Software Supply Chain Security

Endor Labs is a software supply chain security platform that detects vulnerabilities and malicious dependencies in open source code.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5 NOVELTY 4/5

Endor Labs AI SAST launches a multi-agent pipeline that orchestrates detection, triage, and remediation agents to find IDORs, broken access control, and insecure design across the full OWASP 2025 Top 10 — automatically classifying 95% of findings as false positives before they reach the security team — with C support delivering working exploits and call paths, and BYOK so model requests never leave the customer's own provider account.

> Endor Labs launched AI SAST in private preview, a multi-agent system that triages findings, learns organizational context, and now covers C alongside bring-your-own-key model support.

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

01 Organizational context learning for AI SAST NEW 100

Endor Labs AI SAST now automatically reads `AGENTS.md`, `CLAUDE.md`, skills files, and AI rules files from the repository to extract security-relevant context — such as authentication middleware and user-facing scope — without any configuration. It also ingests IaC files — Dockerfiles, Terraform, CloudFormation, Kubernetes YAML, and Helm charts — to build a deployment model that infers network boundaries and component interactions for every PR and local developer scan. Users can upload architecture and threat model context in plain natural language, which the system verifies against code and classifies into three categories — project context, true positives, and false positives — scopeable to a single project, a set of projects, or the entire tenant. Reviewers can attach thumbs-up or thumbs-down feedback to individual findings during triage, teaching the scanner to adjust severity for entire classes of issues and suppress patterns the team has explicitly ruled on in future scans.

Suppress a class of XSS findings already mitigated by an upstream WAF by marking them as false positives during triage so future scans deprioritize that pattern automatically.

📍 In the AI SAST findings view, open a reflected XSS finding, click the thumbs-down icon, and add a note such as 'Mitigated by CDN-managed WAF ruleset in front

of all public traffic' — repeat for the same class to train severity adjustment across future scans.

Encode organization-wide authentication standards as scan-time enforcement so AI SAST flags any new code that deviates from approved auth libraries.

📍 Upload a plain-language security context document scoped to the entire tenant that states your standard auth libraries and prohibited patterns; AI SAST classifies matching deviations as true positives in every project scan.

— Names exact files, categories, and includes runnable triage steps

Your model, our analysis: BYOK for Endo… | How Endor Labs AI SAST Learns What Matt…

02 **Bring-your-own-key for AI model providers** NEW 80

AI SAST and AI Chat model requests can now be routed through your own OpenAI, Azure OpenAI, Gemini, or Vertex AI account by enabling BYOK, with provider API keys either encrypted in the Endor Labs platform or passed as environment variables at scan time so keys are never stored with Endor Labs. BYOK enforces a provider allowlist and tenant isolation so requests go only to approved providers and your keys and requests remain separate from other tenants.

— Names providers and key-handling mechanism but no exact config key

Your model, our analysis: BYOK for Endo…

03 **C language support in AI SAST** NEW 75

Endor Labs AI SAST now supports C with buildless scans that run in the IDE and on pull requests, covering buffer overflows, use-after-free, integer overflows, and cross-function flaws without requiring a compiler or build to reproduce. C SAST and C SCA now run together in the same platform, covering first-party C code and the open-source it depends on in a single pass. Each C finding includes a call path, a working exploit, and a suggested fix — the same evidence format already used for other supported languages.

— Names scan surfaces and vuln classes but no exact command

Your model, our analysis: BYOK for Endo… | Why Endor Labs AI SAST for C finds what…

04 **AI SAST with multi-agent triage** NEW DETECTION NEW 60

Endor Labs launched AI SAST, in private preview, using multiple agents to detect business logic flaws and IDORs that rule-based tools miss. It automatically classifies every finding as true positive, false positive, or unknown across the full OWASP 2025 Top 10, including broken access control and insecure design, before results reach the security team. An AI Security Code Review agent assesses pull requests for architecture design and business logic flaws affecting security posture, working alongside the SAST detection and triage agents.

— Explains mechanism and scope but no command or UI path given launch-20261111-65428615

NEW INTEGRATION OpenAI Daybreak Defense Factory / Codex Security integration

NEW INTEGRATION OpenAI Daybreak Defense Factory / Codex Security integration

> Introducing AI Sast That Thinks Like A Security Engineer |… | Endor Labs (2026-11-11)
>
> Endor Labs adds AI SAST: a multi-agent system using a Code API, Opengrep, and OPA-backed policies to detect business logic flaws and auto-triage findings.
>
> NEW DETECTION Endor Labs now adds AI SAST — a multi-agent static analysis capability that orchestrates detection, triage, and remediation agents to surface complex vulnerabilities such as IDORs, broken access control, and insecure design flaws.
>
> Triage now runs automatically on every finding, classifying each as true positive, false positive, or unknown before it reaches the security team.
>
> A new Code API builds a semantic model of code on top of which callgraphs, dataflow tracing, and code search operations can be performed, enabling agents to reason across multi-file and multi-function interactions.
>
> AI SAST integrates with Endor Labs' policy engine, powered by Open Policy Agent (OPA), allowing teams to write custom policies to surface only high-confidence findings, suppress false positives, or block risky pull requests.
>
> Custom prompts let organizations teach the system about their own architectural patterns, security requirements, and risk priorities, with user feedback on classifications continuously refining future results.

> Endor Labs + OpenAI: Strengthening the Defense Factory With Accelerated Remediation |
>
> Endor Labs adds an integration with the OpenAI Daybreak Defense Factory, routing reachability analysis, SAST, SCA, and container findings into Codex for automated fix implementation.
>
> NEW INTEGRATION Endor Labs now connects to Codex Security via the OpenAI Daybreak Defense Factory, sending each finding with full-stack reachability, risk, impact, and remediation guidance so Codex can implement and test fixes.
>
> An AI SAST agent (AURI) gives Codex evidence and remediation guidance for vulnerabilities in first-party source code.
>
> An SCA Remediation agent gives Codex reachability evidence and remediation guidance for open source vulnerabilities across direct and transitive dependencies.
>
> Endor Labs now provides Codex with reachability evidence and guidance to remediate vulnerabilities in container base images.
>
> Customer-defined policies let teams prioritize which Endor Labs findings the Daybreak Defense Factory acts on first, based on organization-specific risk criteria.

announced · Nov 12, 2026

Endor Labs adds an integration with the OpenAI Daybreak Defense Factory, routing findings with reachability and remediation guidance to Codex Security for automated fix implementation.

NEW INTEGRATION Endor Labs now connects to Codex Security via the OpenAI Daybreak Defense Factory, sending each finding with full-stack reachability, risk, impact, and remediation guidance so Codex can implement and test fixes.

An AI SAST agent (AURI) supplies Codex with evidence and remediation guidance for vulnerabilities in first-party source code.

An SCA Remediation agent (AURI) supplies Codex with reachability evidence and remediation guidance for open source vulnerabilities across direct and transitive dependencies.

Endor Labs now provides Codex with reachability evidence and guidance to remediate vulnerabilities in container base images.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Zscaler

changelog · Nov 4, 2026

Secure Access (VPN / SASE / ZTNA)

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5 NOVELTY 5/5

Zscaler's Zero Trust Browser now streams thick-client applications — SAP GUI, database utilities, developer IDEs — inside isolated, ephemeral cloud environments rendered directly in a browser tab, with no local VDI client, endpoint agent, or customer-managed infrastructure required; this closes the last major gap in browser-delivered zero-trust access that previously forced organizations to maintain VDI for non-web apps.

> Zscaler's Zero Trust Browser can now stream thick-client applications such as SAP GUI and developer IDEs from isolated, ephemeral cloud environments, eliminating the need for VDI.

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

01 Thick-client app streaming via Application Isolation NEW

65

The Zero Trust Browser now supports thick-client application delivery through Application Isolation, running apps like SAP GUI, database utilities, and developer IDEs inside isolated, ephemeral cloud environments rendered directly in a browser tab, with no local VDI client, endpoint agent, or customer-managed infrastructure required. Users access these thick-client desktop environments from the Zero Trust Browser landing page alongside existing web and private applications. Session persistence is enabled by default, letting users disconnect and reconnect with unsaved work intact, and copy/paste plus file downloads from isolated sessions are governed by policy to keep enterprise data inside the cloud environment.

— Describes mechanism and behaviour but no config/API names or limits 2026-11-04-0029665f

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Prowler 5.49.0

Cloud Security Posture

// WHY IT MADE THE LIST — DEPTH 4/5 — IMPACT 5/5 — NOVELTY 5/5

Prowler adds an emerging threat signature engine — no prior art found — that lets practitioners supply threat signatures per scan to match compute instances across AWS, Azure, and GCP against affected image publishers, name patterns, version ranges, and exposed ports, with Amazon Inspector CVE correlation to automatically clear false positives; three new checks fire on matched instances and store structured threat data in a new `Check_Report.raw` payload.

> Prowler shipped emerging-threat signature matching for AWS, Azure, and Google Cloud compute instances, real-time AWS detection and remediation in Prowler Cloud, and dozens of new checks spanning AWS Bedrock AgentCore, Cloudflare, GitHub, and Vercel.

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

01 Emerging threat signature detection for AWS, Azure, and GCP

NEW — 95

Supply threat signatures per scan with the `threat_signatures` configuration key for AWS, Azure, and Google Cloud to match compute instances against affected image publishers, name patterns, per-branch version ranges, exposed ports, impact, and remediation. Three new checks — `ec2_instance_emerging_threat_image` (AWS), `vm_emerging_threat_image` (Azure), and `compute_instance_emerging_threat_image` (Google Cloud) — flag instances running images matched by a loaded threat signature, with matched-threat data attached to each finding via the new `Check_Report.raw` payload. On AWS, the engine clears a match when Amazon Inspector finds no active CVE finding for the threat, using the new `inspector2:SearchVulnerabilities` permission added to the Prowler additions policy and the CloudFormation scan role. In Prowler Cloud, issues can be filtered by threat with `filter[threat]` or searched by threat name/CVE ID with `filter[search]` on `GET /api/v1/issues` and `GET /api/v1/issues/overview`.

List all Prowler Cloud issues tied to a specific emerging threat to see every affected host in one call.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://api.prowler.com/api/v1/issues?filter[threat]=my-threat-
001'
```

— Names config key, checks, permission, and API endpoints with an example `5.49.0`

02 **Real-time AWS detection and remediation in Prowler Cloud** NEW

90

Prowler Cloud ingests CloudTrail events via EventBridge/SQS or a `LookupEvents` fallback to detect AWS activity in real time, with per-account policy modes `monitor`, `pr`, and `auto` (including dry-run support), a new `/realtime` dashboard, and a Manage Realtime role permission to control who can change the real-time policy. Lighthouse AI and the MCP Server gain tools to manage the event feed, read detections and actions, and revert them.

— Names policy modes, dashboard path and permission, but no numeric limits `5.48.0`

03 **New checks for Bedrock AgentCore, Cloudflare, GitHub, and Vercel**

NEW

85

AWS gains a `bedrockagentcore` service with `bedrockagentcore_full_access_policy_attached`, `bedrockagentcore_payments_process_payment_role_separation`, `bedrockagentcore_runtime_vpc_configured`, plus `cloudtrail_agentcore_data_events_enabled` for AgentCore data event coverage. Cloudflare gains 13 checks: a new `account` service (`account_two_factor_enforced`, `account_member_two_factor_enabled`, `account_super_admins_limited`, `account_api_token_expiration_set`, `account_api_token_ip_restriction_enabled`, `account_api_token_recently_used`), a new `r2` service (`r2_bucket_public_dev_url_disabled`, `r2_bucket_cors_no_wildcard_origin`, `r2_bucket_custom_domain_min_tls_secure`), a new `workers` service (`workers_script_observability_enabled`, `workers_script_workers_dev_disabled`, `workers_script_preview_urls_disabled`, `workers_script_no_secrets_in_plain_text_bindings`), plus `zone_firewall_no_skip_all_traffic_rules`. GitHub gains `repository_secret_scanning_push_protection_enabled` to flag repositories that accept pushes containing secrets. Vercel gains `project_environment_no_secrets_with_public_prefix` (flagging secret-like environment variables exposed via prefixes such as `NEXT_PUBLIC_` or `VITE_`) and `project_linked_to_git_repository` (flagging projects that can only receive unreviewed CLI deployments).

— Enumerates every new check name across four providers `5.48.0`

04 **Env var to suppress Prowler Cloud promotion notice** NEW

70

Set `PROWLER_NO_PROMOTIONS=1` to suppress the Prowler Cloud hand-off note in the OSS CLI dashboard; Prowler Cloud and Private Cloud reports never show it.

Run the local Prowler dashboard without the Prowler Cloud promotion notice, useful in automated or shared terminal environments.

```
$ PROWLER_NO_PROMOTIONS=1 prowler dashboard
```

— Exact env var and runnable command given `5.48.0`

05 **New AWS checks for OAM and Organizations delegated admin** NEW

60

New check `oam_sink_not_publicly_accessible` reports CloudWatch Observability Access Manager sinks whose resource policy allows any AWS account to link to them. New check `organizations_security_services_delegated_admin_not_management_account` verifies that security services with trusted access in the organization use a delegated administrator rather than the management account.

— Names both checks but gives no remediation mechanism `5.49.0`

06 **Compliance framework mappings for Azure GDPR and Cloudflare NCSC**

NEW

60

Prowler adds a GDPR framework for Azure mapping Articles 25, 30, and 32 to existing Azure checks, and a Cloudflare mapping for NCSC Cyber Essentials 3.3.

— Names articles and framework version but no usage steps `5.48.0`

THINNER COVERAGE BELOW

07 **Contextual Fix skill in Lighthouse AI issue panel** NEW

55

A Contextual Fix Lighthouse AI skill on the issue side panel reads open attack-path legs and findings and leads with the single change that closes the issue fastest.

— Describes behaviour but no concrete surface or command `5.48.0`

08 **Early findings display during running first scan** NEW

55

The Findings page now shows findings from a running first scan as they arrive, with a View Early Findings action on executing scans; critical and high findings are stored first.

— Names the UI action and prioritization order `5.48.0`

09 **Issue side panel replaces separate issue page** BREAKING

55

The issue side panel now holds everything the full issue page had — Findings and Resources tabs, leg evidence, compliance frameworks, and Send to Jira — and the separate issue page is removed.

— Names removed page and consolidated tabs clearly 5.48.0

10 **Sample demo data shown before first scan completes** NEW

50

Prowler Cloud shows labelled sample data from a demo AWS account on Overview, Findings, Resources, Compliance, Issues, and Attack Paths until the first scan completes.

— Names affected pages but no mechanism beyond display 5.48.0

11 **Prowler Private Cloud on AWS Marketplace** NEW

50

Prowler Private Cloud is now available on AWS Marketplace as Pay as You Go and 10- and 25-account bundles, installed in the customer's own AWS account.

— Names pricing tiers but no setup steps 5.48.0

12 **Bedrock API key option for Lighthouse AI settings** IMPROVED

45

Prowler Cloud's Lighthouse AI settings now accept a Bedrock API key as an alternative to AWS access keys.

— Simple config change, minimal detail given 5.48.0

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 teleport

announced · Nov 10, 2026

Privileged Access Management

// WHY IT MADE THE LIST

DEPTH 5/5 | IMPACT 4/5 | NOVELTY 5/5

Teleport Beams runs AI agents in micro-VMs with delegated X.509 identities scoped to the current task (auto-expiring in 12 hours), MITRE ATT&CK session scoring, Phases for Git-like filesystem snapshot and fork across runs, agent-initiated JIT access requests, and MCP tool-level RBAC with short-lived certificates — the first tracked tool to give AI agents hardware-isolated, identity-scoped, auditable infrastructure access rather than inheriting a user's full role set.

> Teleport's Beams runtime gained "Phases," a Git-like system for managing container filesystem state with snapshot, fork, rebase, and rebuild workflows, plus a way to run persistent background agent tasks inside a beam.

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

01 Phases: Git-like container filesystem management in Beams

NEW

93

Beams runtime adds Phases, defined in a Kube-style `phase.yaml` file supporting `version`, `metadata.name`, `from`, `files`, `commands`, and `services` fields, giving snapshot, fork, rebase, and periodic rebuild capabilities for a container's filesystem. Create a phase interactively with `tsh phase add <name>`, fork an existing beam's filesystem state with `tsh phase add <name> --fork`, apply a definition with `tsh phase apply -f phase.yaml`, and save or finalize the current `/home` state as a durable snapshot with `tsh phase snapshot` or `tsh phase snapshot --final`. The `services` block can start a pre-start or background service inside a phase via `type: pre-start` or `type: background`, scoped to `context: user` or `context: snapshot`.



Fork a running beam to give a parallel agent its own isolated filesystem and have it commit work independently before saving a snapshot.

```
$ tsh phase add beak-buttons --fork
beamctl start sh -c 'codex exec -C ~/parakeets "Add a beak-
friendly interface. Test it, commit, push, then run tsh phase
snapshot." </dev/null'
```

Define a reproducible dev environment with a local file overlay, an apt+Go install rebuild stage, and a background dev server that starts automatically on each beam launch.

```yaml
version: v1
metadata:
  name: dev-files
files:
  - from: ./dev-files
    to: /opt/dev
---
version: v1
metadata:
  name: environment
from: dev-files
commands:
  - shell: |
      apt-get update
      apt-get install -y curl
  - shell: |
      v=$(curl -fsSL 'https://go.dev/dl/?mode=json' | jq -r '.
[].version | select(startswith("go1.27."))')
      curl -fsSL "https://go.dev/dl/$v.linux-amd64.tar.gz" | tar
-xz -C /usr/local
      ln -sf /usr/local/go/bin/go /usr/local/bin/go
services:
  - name: dev-server
    type: background
    context: snapshot
    shell: |
      export GOCACHE=/tmp/go-build GOPATH=/tmp/go
      go run /home/beams/parakeets/main.go
```

Apply a phase definition and open an interactive beam on top of the resulting environment.

```
$ tsh phase apply -f phase.yaml
tsh phase add dev --from=environment
```

— Names every CLI command and yaml field with runnable examples launch-20261110-ca529eff

02 **Persistent background tasks via beamctl start** NEW

63

Run a persistent agent task inside a beam without staying connected using `beamctl start <command>`, shown forking a beam and kicking off a long-running agent command that commits and snapshots on its own.

Fork a running beam to give a parallel agent its own isolated filesystem and have it commit work independently before saving a snapshot.

```
$ tsh phase add beak-buttons --fork
beamctl start sh -c 'codex exec -C ~/parakeets "Add a beak-friendly interface. Test it, commit, push, then run tsh phase snapshot." </dev/null'
```

— Exact command given but mechanism beyond the example is thin launch-20261110-ca529eff

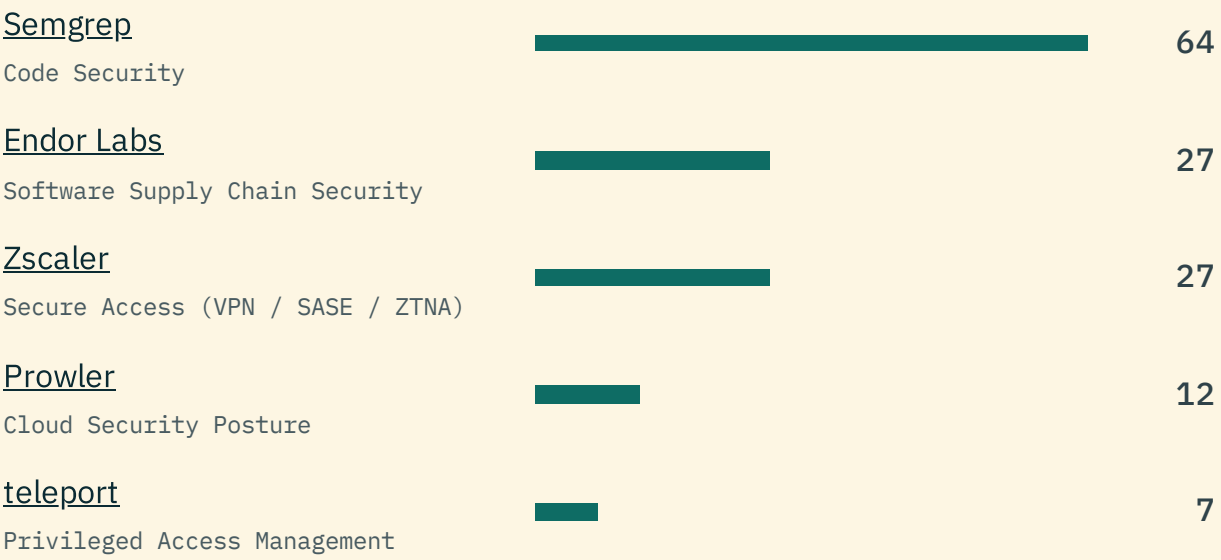
NEW DETECTION MITRE ATT&CK session scoring for AI agent behavior

> announced · Nov 10, 2026
>
> Teleport adds Beams, a secure runtime for AI agents with micro-VM isolation, delegated X.509 identities, phases for state snapshotting, MCP tool-level RBAC, and JIT agentic access requests.
>
> Run AI agents in isolated micro-VMs that auto-expire in 12 hours using `tsh beams add --role=<scoped-role>` to assign a least-privilege delegated X.509 identity scoped to the task instead of inheriting a user's full role set.
>
> Snapshot, fork, and branch a running agent's filesystem state using `tsh phases add`, `tsh phase snapshot`, and `tsh phases add --fork` — giving agents consistent state across runs without rebuilding from scratch.
>
> Create agent-initiated just-in-time access requests for elevated, time-boxed access using `tsh request create --resource=<resource-id> --reason=<reason>`, reviewable by a human or another agent and auto-expiring on task completion.
>
> Connect to MCP servers with tool-level RBAC and short-lived certificates instead of shared API keys using `tsh mcp login <mcp-server-name>`, scoped to only the tools the developer is allowed to use.
>
> Capture every message sent to an inference model — including LLM traces, tool-use, and protocol-level traffic — in a full audit trail for every agentic session.
>
> NEW DETECTION MITRE ATT&CK session scoring for AI agent behavior

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

Where each pick stands · Nov 7–13

Tools with changes in each pick's category over the last 12 weeks, including the pick itself.



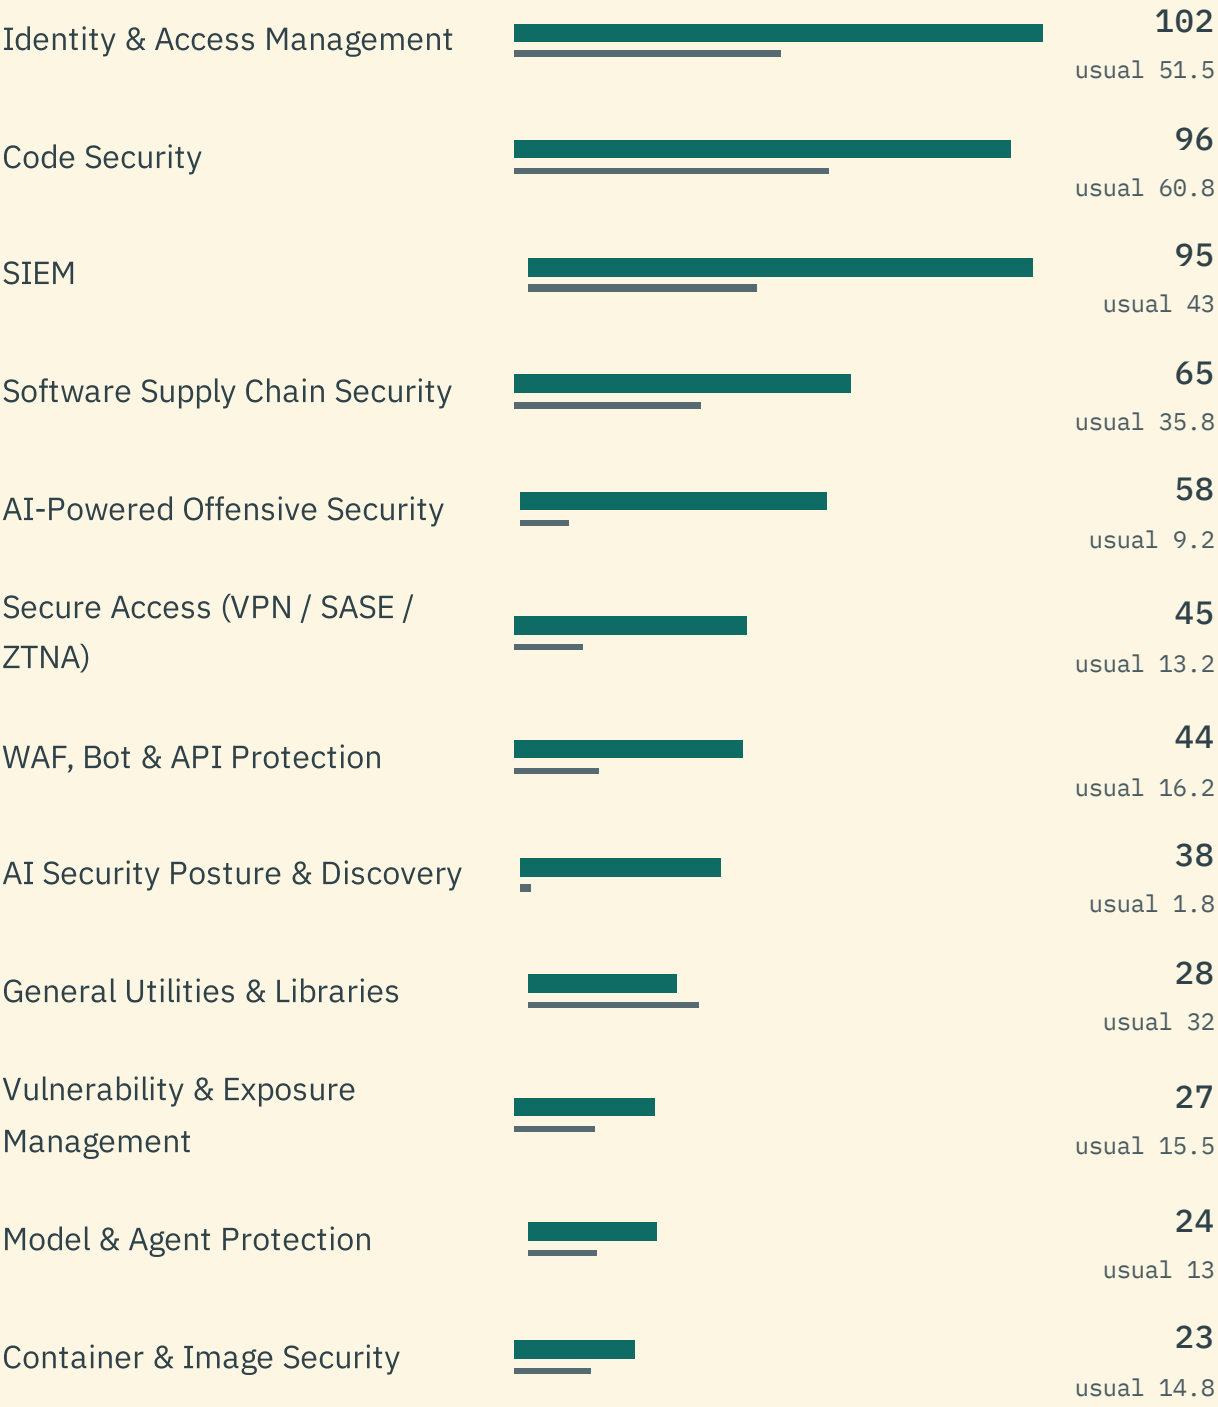
Each pick's release rhythm · Nov 7–13

Changes per week for the last 12 weeks. The bright bar is this week; a dash marks a week with no daily.



Which categories moved · Nov 7–13

The thick bar is this week. The thin bar is the average of the four weeks before, counting only weeks with a daily.



Who shipped the most · Nov 7–13

Changes per tool this week. Volume is not a quality score.



RedAmon		24
code-oss		21
skylos		21

▾ // CONFERENCE CHANGES HIDE

Calls for papers, tickets, dates, talks and recordings this week. Star a conference to hear about the ones you follow.

The talks for Ekoparty 2026 are out.

Upcoming conferences →

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · November 14, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe