THE CYBER TOOLCHAIN

NO. 018

Head

THE WEEK'S TOP PICKS

PUBLISHED OCTOBER 30, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

The week in numbers · Oct 23–29

519
changes we read →

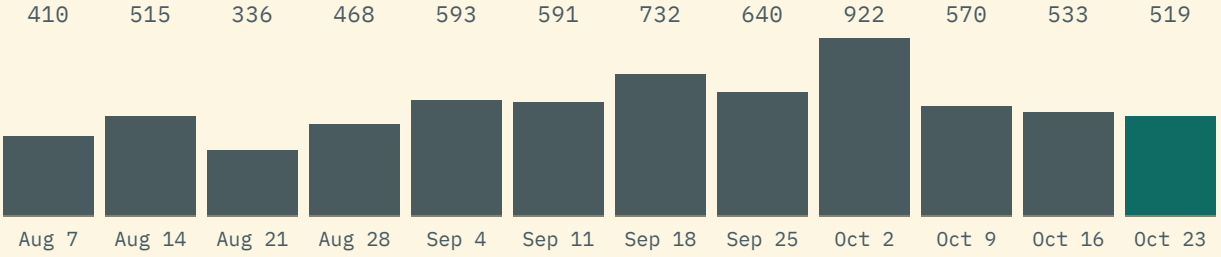
245
tools shipped →

3
product launches →

332
breaking changes →

Changes per week, last 12 weeks · Oct 23–29

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



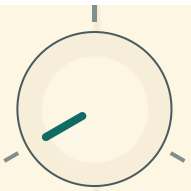
▸ // HOW THESE PICKS ARE MADE SHOW

VIEW

vim ? help

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every change in this week's picks, in full.

all 27 changes

Do you prefer **scroll**? 👍 👎

This week's 5

01	**Aikido Security**	Feature · SAST / SCA / Code Security
02	**CyberArk**	Feature · Privileged Access Management
03	**Aviatrix**	Feature · Firewalls & Network Defense
04	**OpenSSL**	Feature · General Utilities & Libraries
05	**Dreadnode**	Feature · AI-Powered Offensive Security

The picks in full

01 Aikido Security

🚀External Network Pentesting: test your perimeter before someone else does

SAST / SCA / Code Security

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

External Network Pentesting sends AI agents to map your perimeter and exploit it from the outside — discovering subdomains, ports, and services, then chaining findings across hosts into a full attack path, with continuous retesting of newly discovered subdomains. No code access or credentials required.

> Aikido Security launched External Network Pentesting for internet-exposed assets, added a Cloud Security Audit capability for cloud/IaC attack-path analysis, achieved FedRAMP Moderate authorization, and rounded out the release with new API endpoints, a Google SecOps SIEM integration, expanded Zen Firewall language support, and Code Quality/Terraform improvements.

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

01 External Network Pentesting for internet-exposed assets NEW

95

AI agents discover subdomains, open ports, and services from a list of top-level domains, then attempt exploitation using known CVEs, misconfigurations, and exposed secrets — including forgotten subdomains and open admin panels, with evidence-confirmed findings. Agents chain findings across hosts, pivoting using credentials or access found on one system, and the report shows the full attack path from entry point to impact; the Agents tab shows each agent's current thinking, tools in use, and findings in real time. Testing can run in continuous mode so newly discovered subdomains are tested shortly after they go live, with intensity controlled via a maximum requests-per-second per target and optional restriction to inside or outside business hours. A scoping agent checks how much of each subdomain you own, limits testing to owned parts, and routes all agent traffic through a proxy that blocks out-of-scope connections. The output is an auditor-ready report in the same format as Aikido's web app and mobile pentests, covering the attack surface up to the login page.

Start an external network pentest against your perimeter by supplying your top-level domains through the Aikido console.

📍 1. In the Aikido console, navigate to Pentesting › External Network Pentesting. 2. Add one or more top-level domains to scope. 3. Review the discovered subdomains, IP addresses, open ports and services; remove anything to keep out of scope. 4. Set your maximum requests per second per target and optionally

restrict testing to inside or outside business hours. 5. Click `Start Pentest` — the quoted credit cost is shown before anything runs.

Enable continuous mode so newly exposed subdomains are automatically tested shortly after they go live, without manually kicking off a new pentest.

📍 1. In the Aikido console, open your External Network Pentest configuration. 2. Under `Run mode`, select `Continuous`. 3. Confirm the yearly credit price (based on assets found in your first pentest) and save.

— Full console steps, rate-limit and scope controls, mechanism all given

`🚀External Network Pentesting: test you…` `#58 - Cloud Security Audits, External N…`

---

02 **New API endpoints for CWE, invites, assessments and subdomains**

NEW 70

Adds `API - List CWE descriptions` to retrieve CWE metadata, `API - List pending invites` to query outstanding user invitations, `API - List AI code analysis assessments` to fetch AI code analysis results, `API - List pentest assessments` to retrieve pentest findings, `API - Disable user` to deactivate users, and `API - Remove attach surface subdomain` to remove subdomains from the attack surface.

— Endpoint names given but no HTTP method or path

`#58 - Cloud Security Audits, External N…`

---

03 **Google SecOps SIEM integration** NEW 70

New integration to optionally forward Aikido Activity logs to Google SecOps as a SIEM destination, configured via Settings > Integrations > SIEM by enabling `Transmit Activity logs` and selecting `Google SecOps`.

Forward Aikido Activity logs to Google SecOps for centralized SIEM alerting and correlation.

📍 Steps: In Aikido, go to Settings › Integrations › SIEM, enable `Transmit Activity logs`, and select `Google SecOps` as the destination.

— Exact UI navigation path given, but no API/config key

`#58 - Cloud Security Audits, External N…`

THINNER COVERAGE BELOW

04 **Cloud Security Audit for multi-step attack paths** NEW 50

New feature where AI agents reason across cloud and IaC configurations to surface multi-step attack paths and link each finding back to the causative template.

– Mechanism described but no UI path or example given

#58 - Cloud Security Audits, External N…

05 **Terraform container management resources and OpenTofu registry**

NEW 50

Adds Terraform Container management resources and publishes the Aikido provider to the OpenTofu registry.

– Named resource and registry but no usage example

#58 - Cloud Security Audits, External N…

06 **Code Quality scan limits removed and auto-PR creation**

IMPROVED 45

Removes the maximum-rule restriction for Code Quality scans and enables auto-creation of PRs for issues found in full repository scans.

– Describes change but no config key or flag given

#58 - Cloud Security Audits, External N…

07 **Zen Firewall agents for five languages** NEW 40

Adds Zen Firewall agent releases for Dotnet, Python, PHP, Ruby, and Node.

– Lists languages but no usage or behaviour detail

#58 - Cloud Security Audits, External N…

08 **FedRAMP Moderate authorization for Aikido for Government** NEW

40

Aikido for Government achieves FedRAMP Moderate authorization and is listed on the FedRAMP Marketplace, allowing federal agencies to adopt Aikido without a full security review.

– Compliance milestone with no actionable steps #58 - Cloud Security Audits, External N…

09 **Hugging Face added to Aikido Intel** NEW 30

Adds Hugging Face to supported malware ecosystems in Aikido Intel.

– Bare addition with no mechanism or scope given #58 - Cloud Security Audits, External N…

└──▷ **ALSO FROM THESE RELEASES**

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 CyberArk

announced · Oct 28, 2026 Privileged Access Management

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

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

Idira Agent Identity Security reaches GA, issuing SPIFFE SVIDs and short-lived OAuth 2.1 credentials to each AI agent and enforcing tool-level least-privilege on MCP servers, with zero standing privilege for databases and an onboarding flow for shadow agents.

Idira Agent Identity Security is now GA with Prisma AIRS, delivering per-agent SPIFFE/OAuth 2.1 credentials, MCP tool-level ZSP, on-behalf-of audit logging, shadow agent onboarding, and an activity map.

NEW SUPPORTED PLATFORM Provisions SPIFFE Verifiable Identity Documents (SVIDs) and OAuth 2.1 client credentials for each registered agent, issuing short-lived, time-bound access tokens in place of long-lived shared secrets.

Applies fine-grained access policies to specific tools available on an MCP server, enabling least-privilege and zero standing privilege (ZSP) controls per agent.

NEW SUPPORTED PLATFORM Delivers ZSP for databases via the Secure Infrastructure Access MCP server, with ZSP for cloud providers listed as coming soon.

NEW INTEGRATION Supports automatic OAuth app registration for MCP servers that implement Dynamic Client Registration (DCR), capturing audit data for tool interactions.

Brings shadow AI agents into central governance through an unmanaged agent onboarding flow, including agents identified by cloud-provider origin IDs such as AWS ARNs or App IDs.

+4 more

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Aviatrix

announced · Oct 28, 2026

Firewalls & Network Defense

// WHY IT MADE THE LIST

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

Clearance enforces per-workload outbound and lateral containment by workload identity across clouds, Kubernetes, and serverless — down to individual containers, databases, and AI agents — with a shadow-rule monitor mode that proves exactly which connections a policy would block before any enforcement.

> Aviatrix launched Clearance, a per-workload outbound and lateral containment system built on an autonomous policy-drafting agent and a monitor-mode testing workflow, alongside a free assessment and promotional pricing window.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 Outbound and Lateral Containment in Clearance NEW

72

Aviatrix Clearance™, part of the Aviatrix Containment Platform, adds Outbound Containment (limiting each workload to outbound destinations it is explicitly cleared for, named by domain rather than network address) and Lateral Containment (restricting workload-to-workload reachability by identity across clouds and Kubernetes/serverless paths not governed by hub-based inspection). It enforces per-workload policy by identity and protocol across clouds, down to individual containers, serverless functions, databases, and AI agents.



— Names the product and two containment modes but no config surface

launch-20261028-5ed5a331

02 **Containment Assessment and launch pricing** NEW 70

A free, read-only Containment Assessment at `aviatrix.ai/containment-assessment` installs nothing and can reserve AWS Marketplace enrollment. Clearance (Outbound Containment , Lateral Containment , and Intrusion Prevention , no workload cap) is free for all enterprises from November 1, 2026 through May 1, 2027 via AWS Marketplace . A General Availability tier offers the Containment Assessment plus `15` workloads in monitor or enforce mode with no time limit, priced per workload, on Aviatrix Cloud or Aviatrix Enterprise ; Aviatrix Active Defense is available today in free early access for design partners.

— Exact dates, URL and workload count but is pricing not a feature

launch-20261028-5ed5a331

THINNER COVERAGE BELOW

03 **Shadow Rule Set monitor mode** NEW 56

A new Shadow Rule Set monitor mode evaluates a proposed policy alongside the live policy against real traffic and shows exactly which connections would be blocked, letting teams validate a containment policy before enforcing it.

— Clear workflow but no named flag or UI path launch-20261028-5ed5a331

04 **Active Defense autonomous policy agent** NEW 54

Aviatrix Active Defense™ is an autonomous agent layer that models the estate from live traffic, drafts per-application policies without a tagging project or inventory import, proves each policy in monitor mode on real traffic, and keeps policies current as the estate changes.

— Describes mechanism in prose, no concrete interface named launch-20261028-5ed5a331

05 **Single-setting enforcement and backout** NEW 38

Every containment policy can be turned on with a single setting and rolled back with a single action, covering enforcement and backout as one-step operations.

— States the capability but no concrete setting name or UI location

launch-20261028-5ed5a331

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 OpenSSL `openssl-4.1.0`

General Utilities & Libraries

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

DTLS 1.3 per RFC 9147 lands in the most widely used crypto toolkit, alongside RFC 8701 GREASE and AVX-512/VAES post-quantum and AES-CBC performance work.

> OpenSSL 4.1.0 adds DTLS 1.3 (RFC 9147) and GREASE (RFC 8701) support, new IKEv2 KDF and DTLS listener APIs, and performance optimizations for ML-DSA/ML-KEM post-quantum crypto, alongside breaking changes to removed `Configure` options, the `tsget` utility's curl dependency, and dropped Windows build targets.

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

01 Removal of no-ecdsa and no-ecdh Configure options — BREAKING

75

The `no-ecdsa` and `no-ecdh` `Configure` options are removed; use `no-ec` instead to disable elliptic curve support.

— Exact removed flags and replacement flag given `openssl-4.1.0`

02 tsget requires Net::Curl::Easy dependency — BREAKING

75

The `tsget` utility now requires `Net::Curl::Easy` (from the `Net-Curl` CPAN distribution) instead of the abandoned `WWW::Curl::Easy` — install it before upgrading.

— Exact dependency names and clear migration step `openssl-4.1.0`

03 Post-quantum crypto performance optimizations — IMPROVED

70

Adds optimized ML-DSA and ML-KEM NTT operations on `ppc64le`; optimized ML-DSA operations on `x86_64` including AVX-512-optimized SHAKE x4 for ML-DSA and AVX-512/VAES optimizations for AES-CBC decryption.

— Names algorithms and architectures but no benchmark numbers `openssl-4.1.0`

04 DTLS protocol enhancements — NEW

60

OpenSSL 4.1.0 adds DTLS 1.3 support per RFC 9147 (documented in the new ossl-guide-dtlsv13(7) manual page) and adds DTLS listener API support.

– Names RFC and manual page but no usage walkthrough Only named in release summary, no API detail given

openssl-4.1.0

THINNER COVERAGE BELOW

05 **MSVC 2013 build targets added** NEW 45

Adds `VC-WIN32-MSVC2013` and `VC-WIN64A-MSVC2013` build targets for MSVC 2013 C99 bridging support.

– Named targets but no further build instructions openssl-4.1.0

06 **VC-WIN64I and VC-CE build targets dropped** BREAKING 45

The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are dropped from Configurations.

– Named targets removed but no migration guidance openssl-4.1.0

07 **GREASE extension support (RFC 8701)** NEW 35

Adds RFC 8701 GREASE (Generate Random Extensions And Sustain Extensibility) support.

– Named with RFC number but no config or usage detail openssl-4.1.0

08 **Elbrus2000 (e2k) architecture support** NEW 30

Adds initial support for the Elbrus2000 (`e2k`) architecture.

– Bare mention of new architecture, no build detail openssl-4.1.0

09 **IKEv2 KDF support** NEW 20

OpenSSL 4.1.0 adds IKEv2 KDF (key derivation function) support.

– Only named in release summary, no implementation detail openssl-4.1.0

└──▷ **BREAKING ON UPGRADE**

! The `no-ecdsa` and `no-ecdh` Configure options are removed; use `no-ec` to disable elliptic curve support.

! The `tsget` utility now requires `Net::Curl::Easy` (from the Net-Curl CPAN distribution) instead of the abandoned `WWW::Curl::Easy` — install it before upgrading.

! The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are dropped from Configurations.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Dreadnode

New navigation now live for all users

AI-Powered Offensive Security

// WHY IT MADE THE LIST

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

Dreadnode adds white-box AI attack support — GCG-family, embedding, generator, and activation methods run against a local open-weight surrogate — plus three new jailbreak algorithms and two classifier-evasion attacks, letting red teams test model internals rather than just black-box prompts. A deterministic tool_call_attack_landed scorer gives a judge-free success signal for agentic attacks.

> Dreadnode added a tree-search jailbreak attack tool for AI red-teaming, enriched findings exports with audit and compliance data, and introduced session pinning alongside more resilient hosted inference streaming.

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

01 SLIP self-jailbreaking attack tool

NEW DETECTION NEW

75

Adds `SLIP` to the AI red-teaming attack tool list, performing tree-search-based prompt jailbreaks against a target model with configurable depth, branching, and judge models. Launched from the red-teaming UI by selecting a project, choosing `SLIP`, setting the target model, and configuring depth, branching, and judge models before running the campaign.

Run a SLIP tree-search jailbreak campaign against a target model to measure attack success rate with controlled branching.

📍 In the AI red-teaming UI, select your project, choose `SLIP` from the attack tool list, set your target model, configure depth, branching, and judge models, then launch the campaign.

— Names mechanism and config params; UI steps given but no CLI/API

Agentic attack trace parity

02 Session pinning in Chat and Agents

NEW

70

Sessions can now be pinned from the `...` menu in the Chat sidebar or via `Agents › Sessions`. Pinned sessions appear at the top of the list and are visible only to the user who pinned them.

— Exact menu and navigation path given for both surfaces Agentic attack trace parity

03 **Enriched findings parquet export** IMPROVED 60

The findings parquet export gains new columns for audit trail, ASR, compliance tags, adversarial-ML metrics, and project-level rollups, and now reflects the latest operator-reviewed severity and status.

— Lists specific new columns but gives no export command or path

Agentic attack trace parity

THINNER COVERAGE BELOW

04 **Automatic retry for hosted inference streams** IMPROVED 55

Hosted inference SDK streams dropped mid-connection are now retried up to 3 times with exponential backoff (1s, 2s, 4s) to reduce failures on unstable network paths.

— Precise retry mechanism and timing, but no user-facing action

Agentic attack trace parity

NEW DETECTION New structural and long-context jailbreak attacks in the airt SDK

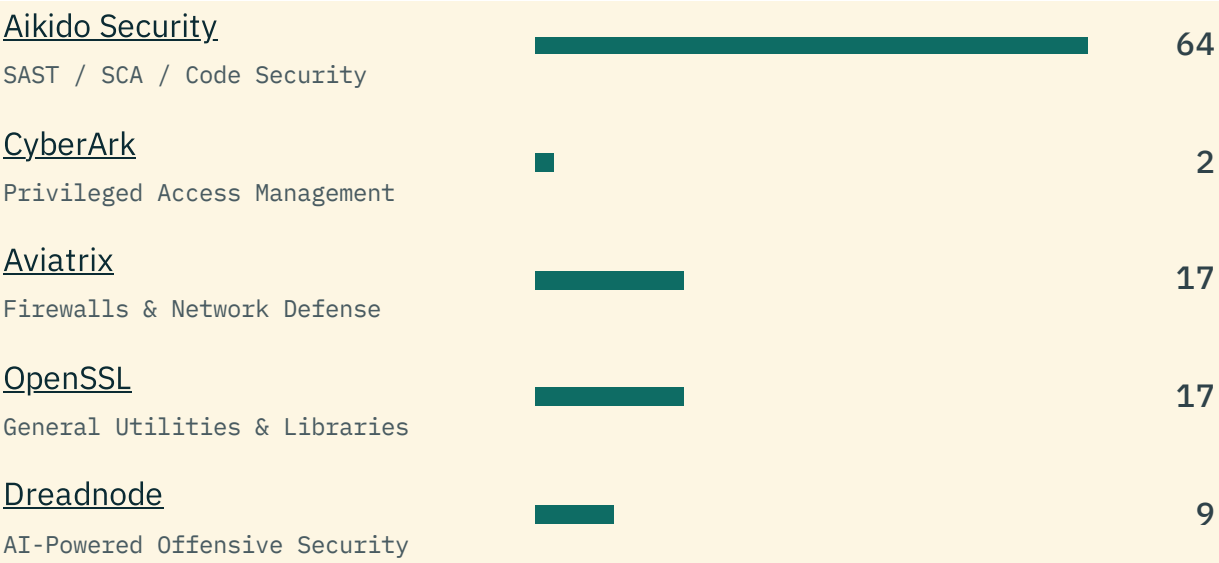
NEW DETECTION New classifier-evasion attacks in the airt SDK

> New navigation now live for all users
>
> Dreadnode adds a `tool_call_attack_landed` scorer for judge-free agentic red teaming and ships a third-party notices file for air-gapped Enterprise installs.
>
> NEW DETECTION Adds `tool_call_attack_landed` scorer for agentic red teaming — measures whether an agent emitted a specific attacker-intended tool call (name and arguments), providing a deterministic, judge-free attack success rate signal.
>
> Enterprise installs now include a third-party notices file downloadable from Account Settings or at `https://<domain>/third-party-notices`, air-gap compatible and version-matched.
>
> Redesigned navigation is now the default for all users on SaaS and Enterprise.
>
> NEW DETECTION New structural and long-context jailbreak attacks in the airt SDK
>
> NEW DETECTION New classifier-evasion attacks in the airt SDK

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

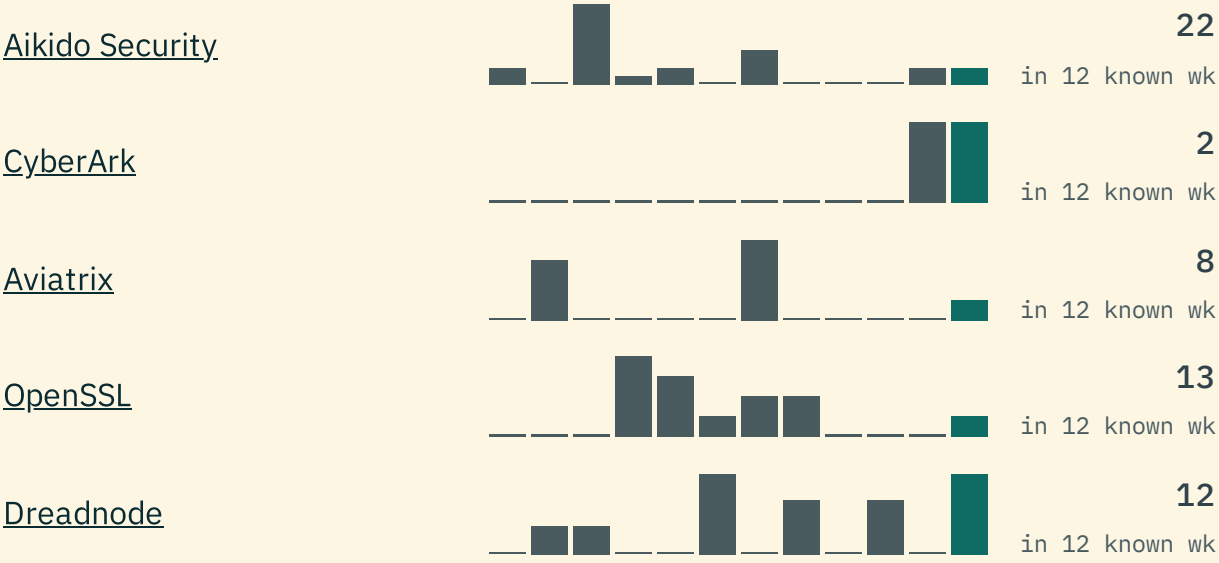
Where each pick stands · Oct 23–29

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



Each pick's release rhythm · Oct 23–29

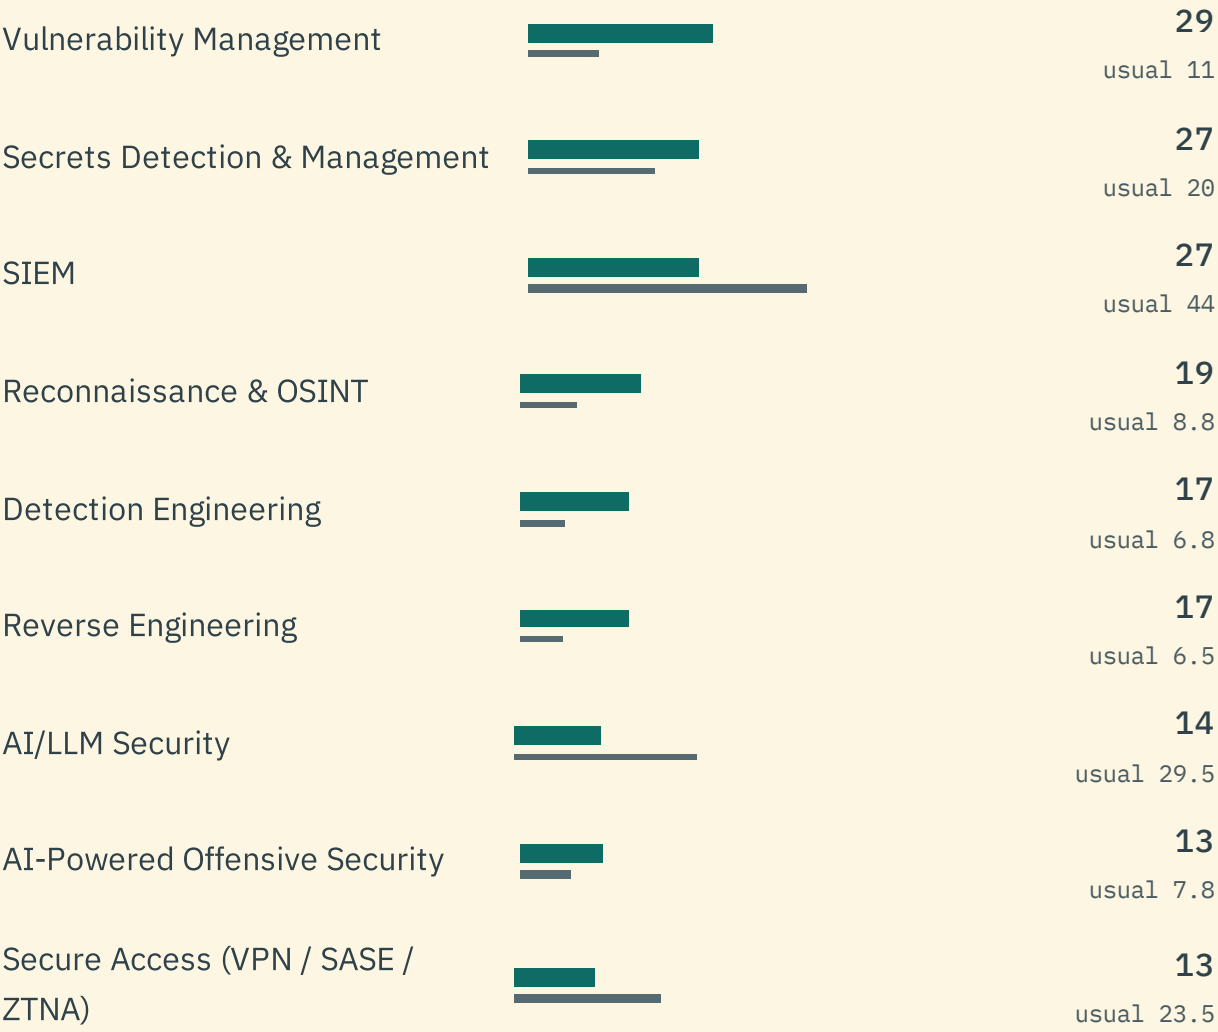
Changes per week for the last 12 weeks. The bright bar is this week; a dash marks a week with no daily.



Which categories moved · Oct 23–29

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 · Oct 23–29

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



// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · October 30, 2026

The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe