THE CYBER TOOLCHAIN

NO. 021

Head

THE WEEK'S TOP PICKS

PUBLISHED NOVEMBER 20, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

EDITIONS tail | grep | head | ls | diff | uniq

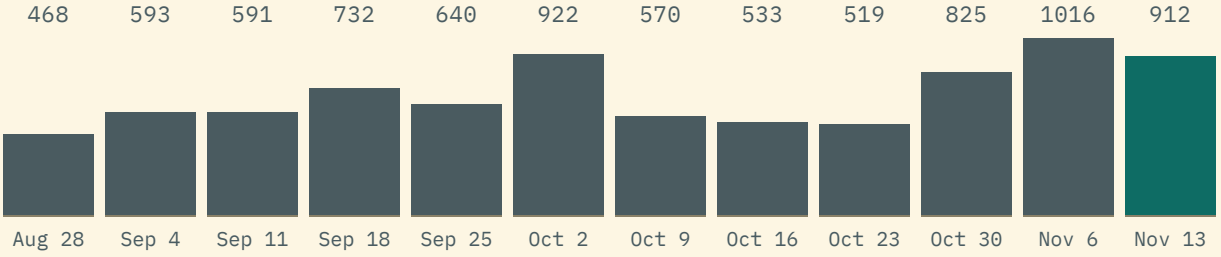
The best tooling updates that shipped this week.

The week in numbers · Nov 13–19

912	273	15	399
changes we read →	tools shipped →	product launches →	breaking changes →

Changes per week, last 12 weeks · Nov 13–19

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



Category moves this month

Plays could not be checked. This does not mean there is nothing to do. Retry

vim ? help

▸ // HOW THESE PICKS ARE MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every change in this week's picks, in full. all 9 changes

Do you prefer **scroll**? 👍 👎

This week's 5

01	**Humanix**	Launch · Firewalls & Network Defense
02	**Zenity**	Feature · AI Security Posture & Discovery
03	**Wallarm**	Feature · WAF, Bot & API Protection
04	**Auth0**	Feature · Identity & Access Management
05	**StepSecurity**	Feature · Software Supply Chain Security

The picks in full

01 Humanix

announced · Nov 17, 2026

Firewalls & Network Defense

Humanix detects and stops social engineering attacks as they unfold across email, chat, voice, and help desk workflows.

// WHY IT MADE THE LIST

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

Humanix adds Complete Call Context, spawning a micro-agent per sensitive service-desk request that correlates live call behaviour with IdP and ITSM records in real time — flagging mismatched managers, non-existent ticket numbers, or unverified identity claims before the call ends, closing a gap no other tracked tool addresses.

Humanix adds Complete Call Context: per-request micro-agents correlate live service-desk calls with IdP and ITSM records to detect vishing in real time.

NEW DETECTION NEW INTEGRATION Adds Complete Call Context, which spawns a micro-agent per sensitive service-desk request (such as an MFA reset or password change) to correlate call behaviour with identity provider and IT service management records, then raises a severity-rated alert with a plain-language explanation when any discrepancy or missing evidence is found. Not checked yet

Micro-agents now check caller claims against directory, ticketing, and identity provider records in real time — flagging conflicts such as a mismatched manager, a non-existent ticket number, or an unverified identity claim — rather than waiting until after the call. Not checked yet

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Zenity

Securing the Next Control Plane for AI Agents: Introducing Zenity for Sandboxes

AI Security Posture & Discovery

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

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

Zenity now treats agent sandboxes as a first-class security surface: continuous discovery of every sandbox and template, posture checks that flag real API keys or wildcard egress, active red-team testing inside each sandbox profile, and runtime enforcement that evaluates actions against the full session context — no other tracked tool covers this attack surface.

Zenity adds a full sandbox security layer: discovery, posture checks with active red-team testing, runtime protection with taint and intent tracking, four containment levels, a trajectory view, and a new Sandboxes policy page.

NEW INTEGRATION A new Sandboxes page lets security teams write per-agent file, network, and runtime policies — including specific blocked paths, egress allow-lists, and environment variable listings by name — that the sandbox enforces directly. Not checked yet

New sandbox discovery continuously inventories every sandbox and sandbox template, linking each to its agent, credentials, workspace, egress policy, owner, image, and mid-session-installed packages, skills, and MCP servers, including shadow sandboxes and sandboxes without a managed platform. Not checked yet

Posture checks now flag misconfigured sandboxes — such as real API keys in the environment or image, wildcard egress rules, internet-reachable sandboxes, or root-running containers — as posture violations with risk descriptions and remediation guidance. Not checked yet

Posture is now also tested actively: a red-team agent runs inside each sandbox profile attempting to escape, reach denied hosts, read credentials, or tamper with hooks, and every path found becomes a hardening rule for that profile. Not checked yet

NEW DETECTION NEW INTEGRATION UPDATED INTEGRATION Runtime protection now evaluates each action against the full session context — combining harness intent, taint from untrusted content, and low-level process, file, and network telemetry — blocking actions that conflict with the agent's declared task. Not checked yet

+6 more

NEW INTEGRATION Sandbox platform integrations

NEW SUPPORTED PLATFORM Vendor-operated cloud coding agents

NEW INTEGRATION Claude Enterprise in Microsoft 365 — inline prevention, telemetry, and findings

NEW INTEGRATION Microsoft Copilot Cowork — telemetry and findings via Cowork transcripts

UPDATED INTEGRATION Microsoft Sentinel integration upgraded from raw-log export to high-fidelity alerts

▾ **1 more release this week** · 2026-11-18 → 2026-11-19

> Zenity Now Secures Agents Across Your Microsoft Ecosystem
>
> Zenity adds actor-identity resolution via Microsoft Entra ID for GitHub Enterprise, where GitHub settings hide actor names, enabling full attribution of coding-agent actions.
>
> NEW INTEGRATION UPDATED INTEGRATION Where enterprise settings hide actor names, Zenity uses Microsoft Entra ID to resolve them, so the security team can see whose agent made a change. Not checked yet
>
> NEW INTEGRATION GitHub Copilot coding agent — telemetry, findings, and inline prevention
>
> NEW INTEGRATION Microsoft Copilot Cowork — telemetry and findings via Cowork transcripts
>
> UPDATED INTEGRATION Microsoft Sentinel integration upgraded from raw-log export to high-fidelity alerts

---

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Wallarm

Your security team just got bigger (2026-11-19)

WAF, Bot & API Protection

// WHY IT MADE THE LIST

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

Wallarm ships seven autonomous Agentic Operators — Radar, Sift, Trim, Tune, Lens, Bard, and Wally — that run the full WAAP operations loop (attack-spike investigation, hourly false-positive clearing, rule-set pruning, policy tuning, inventory review, posture briefing, and on-demand configuration) without prior art among tracked tools.

> Wallarm introduced seven AI Agentic Operators that autonomously triage alerts, prune rules, and tune policies, manageable via an MCP server and an audit-logged autonomous mode.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 Seven Agentic Operators for triage and policy tuning

Not checked yet NEW 70

Wallarm added seven named Agentic Operators: Radar, which detects attack spikes, investigates them, and suggests a response; Sift, which clears false positives on an hourly cycle so alerts remain actionable; Trim, which tidies the active ruleset to prevent rule sprawl; Tune, which adjusts protection policies to match observed application behaviour; Lens, which reviews the API inventory and flags risky endpoints; Bard, which converts the current security posture into a human-readable briefing; and Wally, which answers questions and handles configuration requests on demand.

— Names all seven operators and their functions but no setup steps here

Your security team just got bigger (202…

02 Agentic Operators managed via MCP server

Not checked yet NEW 60

All Agentic Operators can be managed through Wallarm's MCP server, letting them plug into external agentic workflows without returning to the console; a customer success team enables Agentic Operators, after which an agentic workflow connects to the MCP server to invoke operators such as Sift, Radar, and Wally programmatically.

Plug Wallarm Agentic Operators into an external agentic workflow via the MCP server so your own agents can trigger Sift, Radar, or Wally without opening the Wallarm console.

📍 Contact your customer success team to enable Agentic Operators, then connect your agentic workflow to Wallarm's MCP server to invoke operators such as

Sift, Radar, and Wally programmatically.

— Has concrete enablement and connection steps via MCP server

Your security team just got bigger (202…

THINNER COVERAGE BELOW

03 **Autonomous mode with audit logging** Not checked yet NEW

40

Agentic Operators support an autonomous mode where routine approved actions run without manual confirmation, with every finding logged to the audit trail.

— Describes behaviour but no config to enable or scope it

Your security team just got bigger (202…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Auth0

Auth0 Security Agent is now in Early Access (2026-11-19)

Identity & Access Management

Auth0 is an identity and access management platform that provides authentication, authorization, and user management services for applications via APIs and SDKs.

// WHY IT MADE THE LIST

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

Auth0's Security Agent detects credential stuffing attacks that stay under rate limits by computing per-tenant traffic baselines and alerting when failed logins deviate from normal for a sustained period — a class of attack that threshold-based controls miss entirely — then triages the attack and recommends specific Attack Protection and Tenant ACL configuration changes an admin can apply in one step.

> Auth0 shipped three distinct capabilities this window: general availability of OpenID Connect for third-party apps, a beta in-app sign-in flow called Embedded Authorize, and a reworked Email Templates editor with live preview and agent-prompt handoff.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Embedded Authorize for in-app sign-in (Beta)** Not checked yet

NEW

95

A new `POST /e/authorize` endpoint lets apps build in-app sign-in without redirecting to Universal Login; it returns the allowed next steps for the app to render until Auth0 issues an authorization code. During Beta it supports email and phone one-time-code sign-in and second-factor MFA (email OTP, phone OTP, TOTP, push, recovery code), and works out-of-the-box with Auth0 Universal Components for web, iOS, and Android.

Kick off an in-app sign-in flow — no browser redirect required — and receive the first allowed step your UI should render.

```
$ curl -X POST https://<your-tenant>.auth0.com/e/authorize \
  -H 'Content-Type: application/json' \
  -d '{"client_id": "<client_id>", "redirect_uri": "
<redirect_uri>", "response_type": "code", "scope": "openid
profile email"}'
```

— Runnable endpoint with method, path, and example request

Embedded Authorize (Beta): in-app sign-…

02 **OpenID Connect for third-party apps (GA)** Not checked yet NEW

80

Third-party apps can now use OpenID Connect, with per-app control over which standard scopes (`profile` , `email` , `address` , `phone`) they may request via OpenID Connect Settings in the dashboard. Dynamically registered apps (via Client ID Metadata Documents) can't receive OIDC settings directly, so OIDC is instead configured through Application Profiles, which all apps registered under that profile inherit; third-party apps now get profile claims in the ID token while the `/userinfo` endpoint returns only `sub` for them.

Turn on OIDC for a third-party app and restrict it to email and profile claims only.

📍 In the Auth0 dashboard, go to Applications › your third-party app › Settings › OpenID Connect Settings, enable Allow ID tokens, and check only profile and email under allowed scopes.

Enable OIDC for a dynamically registered app that cannot receive OIDC settings at registration time.

📍 In the Auth0 dashboard, go to Applications › Application Profiles › your profile › OpenID Connect Settings and enable OIDC there, so all apps registered via that

Client ID Metadata Document inherit the setting.

— Names exact scopes, endpoint behavior, and UI path

OpenID Connect for Third-Party Apps is …

03 **Email Templates editor with live preview and agent prompt**

Not checked yet NEW 72

A dedicated editor at Emails > Templates in the dashboard lists all templates with status and opens each on its own detail page, where editor and live preview tabs show the From line, Subject line, and rendered email. A new Variables panel lets users search and copy template variables, and a Copy agent prompt action copies the template's variables, current body, and deploy steps for handoff to a coding agent.

When iterating on a branded email template, switch to the live preview tab to see exactly how the rendered email, From line, and Subject look before deploying.

> 📍 In the Auth0 dashboard, go to Emails › Templates, select the template you want to edit, then click the Preview tab to see the rendered output alongside From and Subject fields.

When delegating email template authoring to a coding agent, use the Copy agent prompt button to give the agent full context — variables, body, and deploy steps — in one action.

> 📍 In the Auth0 dashboard, go to Emails › Templates, select a template, then click Copy agent prompt to copy a prompt containing the template's variables, current body, and deploy steps for use with a coding agent.

— Clear UI navigation path but no API/config surface

Email Templates Editor UX Improvements …

NEW DETECTION Credential stuffing attack detection

└──▷ **ALSO FROM THESE RELEASES**

Auth0 Security Agent is now in Early Access (2026-11-19)

Auth0 Security Agent EARLY ACCESS adds guided triage and remediation: it recommends Attack Protection and Tenant ACL changes and lets admins apply them after authenticating.

The Security Agent triages detected attacks and recommends Attack Protection and Tenant ACL configuration changes, letting an Auth0 tenant administrator review and apply each one after authenticating. Not checked yet

Re-run triage after applying recommendations to check the current status of an ongoing attack. Not checked yet

NEW DETECTION Credential stuffing attack detection

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 StepSecurity

announced · Nov 18, 2026

Software Supply Chain Security

// WHY IT MADE THE LIST

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

StepSecurity ships two first-of-kind supply chain controls this week: Secure Registry routes npm/pnpm/Bun installs through a cooldown proxy that blocks known-compromised packages and holds new versions for a configurable window, integrated into Harden-Runner with a single flag; and Harden-Runner now enforces network, DNS, process, and HTTPS monitoring inside Kata Containers and Amazon EKS Nitro Sandbox microVMs, closing a gap where CI hardening stopped at the VM boundary.

> StepSecurity's biggest release this window is a free community tier for its Secure Registry, which intercepts npm/pnpm/Bun installs in Harden-Runner to block compromised packages and enforce a cooldown; it also shipped a new fork-PR approval audit control and remediation tracking for Threat Center.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Secure Registry community tier in Harden-Runner** Not checked yet

NEW SUPPORTED PLATFORM NEW 95

Enable protection for any GitHub Actions job by adding `secure-registry: true` to a `step-security/harden-runner@v2` step, which routes npm, pnpm, and Bun 1.4.4+ installs through `community.registry.stepsecurity.io/javascript/` and blocks known-compromised packages via StepSecurity SOC flagging (always on). Installs are held back for a configurable cooldown of 0–30 days (default 2 days) using the `secure-registry` block's `npm.cooldown.days` key, with exemptions for specific packages/versions such as `typescript@*` or `[email protected]`; a custom block list can also refuse specific packages or versions. Disable the registry for a job (e.g. one running `npm publish`) with `npm: mode: off`, and inspect held/blocked versions in the job summary (with a ready-to-paste allow snippet) or review every request on the new Packages tab in Harden-Runner Insights, which shows request type, status, and evaluated controls per download.

Add Secure Registry protection to a CI job so all npm installs are checked for the two-day cooldown and known-compromised packages before they reach your build.

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: true
```

Extend the cooldown to three days and exempt a pinned TypeScript version so the registry holds other new releases but lets a known-good version through.

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: |
      npm:
        cooldown:
          days: 3
          exemptions:
            - "typescript@*"
            - "@typescript/*"
```

Turn off Secure Registry in a publish job so `npm publish` can reach the public registry directly without being intercepted.

```yaml
- name: Harden the runner
  uses: step-security/harden-runner@v2
  with:
    egress-policy: audit
    secure-registry: |
      npm:
        mode: off
```

Enable Secure Registry on a public repo to block compromised npm packages and apply a 7-day cooldown before new versions reach your CI build.

```yaml
- name: Harden Runner
  uses: step-security/harden-runner@v2.23.0
  with:
    secure-registry: true
```

— Names exact config keys, defaults, cooldown range, and runnable YAML

Harden-Runner GitHub Action v2.23.0: Se… | launch-20261118-cefeb282

02 **Remediation tracking in Threat Center** Not checked yet NEW

75

Threat Center findings can now be marked remediated individually (via row menu, with an optional comment) or all at once in the current view via Remediate all ; the incident API endpoint now returns each finding's remediation status, and new endpoints let you mark findings remediated or reopen them for one organization or tenant-wide. Matched components and Matched IOCs tables show status, who remediated each finding, when, and their comment, with All , Active , and Remediated filter tabs plus source-based search; incidents move to a Remediated tab once fully cleared and move back automatically if a finding reappears, flagged with a 'Reappeared:' title prefix and a notification through existing threat intel channels. A summary banner shows outstanding vs. fully remediated findings, every action is recorded in audit logs, and remediating requires Read & write access to the Threat Center permission.

Track remediation progress through the Threat Center UI, using the status tabs to separate active from resolved incidents.

> 📍 In Threat Center, go to the Affecting You tab and open an incident. Use the All / Active / Remediated tabs in the findings table to filter by status, then select Remediate all to close every active finding in the current view, or choose Remediate from an individual finding's row menu to add a comment before marking it done.

— Rich behaviour described but no exact API endpoint paths given

Threat Center: Remediation Tracking (20…

03 **Fork PR approval control audits org policy hourly** Not checked yet

NEW DETECTION NEW

65

A new 'Fork pull request workflows should require approval' control, rated Critical and listed under Code Injection in the Overview dashboard, checks every GitHub organization and public repository's fork PR workflow approval policy every hour and links directly to the fix in GitHub's Settings > Actions > General. Findings are also surfaced through the existing control tools in the StepSecurity MCP server, and the check requires the StepSecurity GitHub App to have Administration: read permission on the organizations and repositories it scans.

Review the new fork PR approval control findings for your organization via the StepSecurity Overview dashboard after granting the required permission.

> 📍 In StepSecurity, go to Overview dashboard › All Controls › Code Injection, then select 'Fork pull request workflows should require approval' to see which organizations and public repositories are on the least restrictive policy, and click `Open Settings` on any finding to jump directly to Settings › Actions › General in GitHub.

— Names control, category, permission, and dashboard path but no API

`launch-20261118-fff0dbd5`

`NEW SUPPORTED PLATFORM` Conda package inventory and configuration audits

`NEW SUPPORTED PLATFORM` Java package inventory and configuration audits

`NEW SUPPORTED PLATFORM` Swift package inventory and configuration audits

`NEW SUPPORTED PLATFORM` Pi coding-agent package and extension inventory

`NEW SUPPORTED PLATFORM` Support for pooled and multi-session Windows virtual desktops

> `announced · Nov 18, 2026`
>
> Harden-Runner agent v1.9.7 adds runtime monitoring and egress policy enforcement inside microVM-isolated ARC runner pods, covering network, DNS, process, and HTTPS events.
>
> `NEW SUPPORTED PLATFORM` Set `runtimeClassName: isolated-sandbox` on runner pods to enable Harden-Runner monitoring and egress enforcement inside Amazon EKS Nitro Sandbox microVMs. `Not checked yet`
>
> HTTPS visibility inside microVMs now covers requests from Go programs and HTTP/2 traffic — including calls from `gh` and other Go tools that carry their own TLS and bypass system OpenSSL — plus Docker-in-Docker image pulls.
> `Not checked yet`

> `Dev Machine Guard v1.19.1 (2026-11-17)`
>
> Dev Machine Guard adds stable device identity and per-session scanning for pooled virtual desktops, device tagging and retirement on uninstall, virtualization detection, and prompt-free agent configuration via `configure` flags.
>
> `NEW SUPPORTED PLATFORM` Set `device_id` to `{hostname}` in agent config to keep a stable device record and history when a session host is rebuilt from an image, suiting Azure Virtual Desktop, Citrix, and Horizon host pools. `Not checked yet`
>
> Exclude an imaging machine from scanning by running `configure --image-builder` or setting MSI property `IMAGEBUILDER=1` during deployment.
> `Not checked yet`

Label a device for its host pool using `device_tag`, and remove it from the dashboard and license count on teardown with `--retire` or by setting `retire_on_uninstall`. Not checked yet

Adds per-session scanning on Windows Enterprise multi-session and Remote Desktop Session Host servers, reporting each signed-in user's inventory and findings independently under the one device record for that machine. Not checked yet

Every scan now reports whether the machine is virtual or physical and which hypervisor or cloud it runs on. Not checked yet

+2 more

NEW SUPPORTED PLATFORM Pi coding-agent package and extension inventory

New Control: Fork Pull Request Workflows Should Require Approval (2026-11-13)

StepSecurity adds a Critical-severity dashboard control that flags weak fork PR approval policies at both org and repo level, with hourly refresh.

Each finding distinguishes whether the weak policy comes from the organization's default or a repository-level override, so repositories that deviate from the org setting are caught separately. Not checked yet

NEW DETECTION Findings update on the next hourly scan after the setting changes; to remediate, set the policy to Require approval for all external contributors under Settings > Actions > General in GitHub. Not checked yet

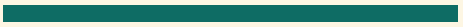
GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

Where each pick stands · Nov 13–19

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

Humanix Firewalls & Network Defense		19
Zenity AI Security Posture & Discovery		5
Wallarm Category unknown		–
Auth0 Identity & Access Management		34

StepSecurity
Software Supply Chain Security

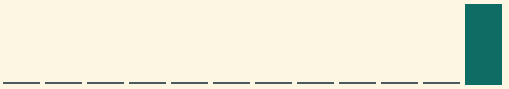


28

Each pick's release rhythm · Nov 13–19

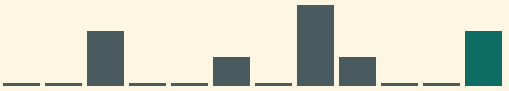
Changes per week for the last 12 weeks. The bright bar is this week; a dash marks a week with no daily.

Humanix



1
in 12 known wk

Zenity



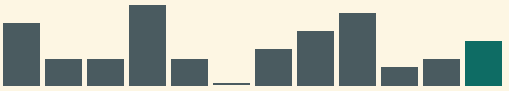
9
in 12 known wk

Wallarm



4
in 12 known wk

Auth0



53
in 12 known wk

StepSecurity



53
in 12 known wk

Which categories moved · Nov 13–19

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

Software Supply Chain Security



110
usual 46.2

Code Security



109
usual 69.8

SIEM



53
usual 54

Secrets Detection & Management



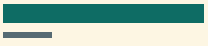
49
usual 19.5

Identity & Access Management



43
usual 66.8

Detection Engineering

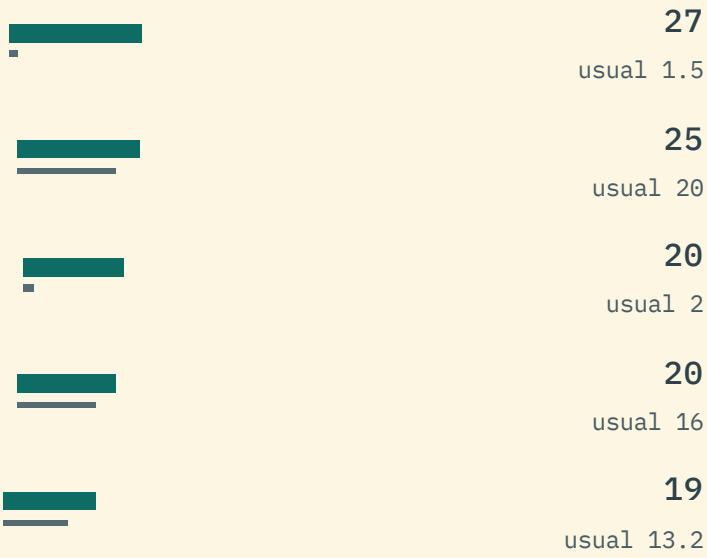


41
usual 9.8

Secure Access (VPN / SASE / ZTNA)



32
usual 22.2



Who shipped the most · Nov 13–19

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



Upcoming conferences →

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · November 20, 2026

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

▷ Subscribe