THE CYBER TOOLCHAIN

NO. 019

Head

THE WEEK'S TOP PICKS

PUBLISHED NOVEMBER 6, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

The week in numbers · Oct 30–Nov 5

825
changes we read →

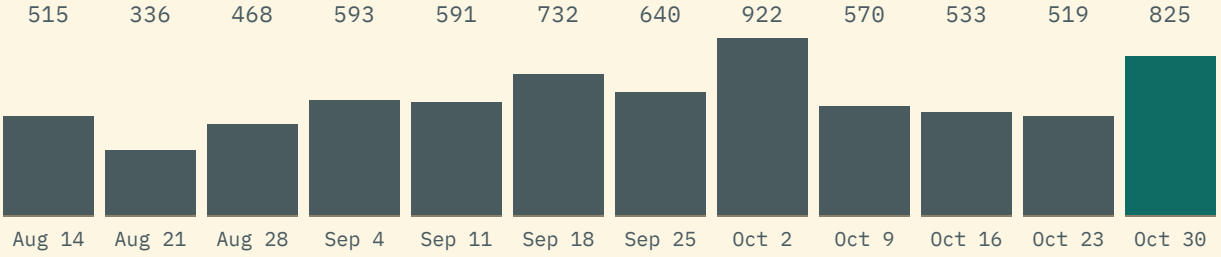
283
tools shipped →

2
product launches →

343
breaking changes →

Changes per week, last 12 weeks · Oct 30–Nov 5

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



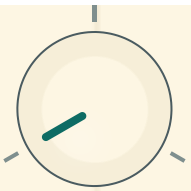
▸ // HOW THESE PICKS ARE MADE SHOW

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 30 changes

Do you prefer **scroll**? 👍 👎

This week's 5

01 **SentinelOne** — Feature · Endpoint Security (EDR/XDR)

02 **GitLab** — Feature · Code Security

03 **Joe Security (Joe Sandbox)** — Feature · Malware Analysis

04 **Sysdig Secure** — Launch · Cloud Security Posture

05 **OWASP ZAP** — Feature · Web Application Security

The picks in full

01 SentinelOne

announced · Oct 30, 2026

Endpoint Security (EDR/XDR)

SentinelOne's Singularity platform protects endpoints, cloud, identity and data with AI-driven detection, Purple AI and an agentic SOC layer.

// WHY IT MADE THE LIST

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

Singularity Agentic Detection and Response builds a live inventory of AI agents, tools, connectors and MCP servers with their permissions, and can contain an individual agent without halting the wider workflow — a control for an attack surface most EDRs do not yet touch. It also links AI-driven actions to subsequent processes to flag when an agent deviates from its task.

> SentinelOne launched Singularity Agentic Detection and Response (ADR) to discover, monitor, and contain AI agents across the enterprise, and added Anthropic's Claude Mythos 5.1 model to its Wayfinder Frontier AI Services for exploitability-validated vulnerability assessment.

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

01 Singularity Agentic Detection and Response (ADR)

NEW INTEGRATION NEW

80

ADR builds an inventory of AI agents, tools, connectors, and Model Context Protocol (MCP) servers along with their permissions and associated users, supporting agents built on Salesforce Agentforce, Microsoft Copilot Studio, OpenAI ChatGPT workspace agents, and Anthropic Claude managed agents. It monitors agent activity through event traces linking AI-driven actions to subsequent system processes to detect deviation from assigned tasks, lets security teams set controls on which agents can operate and which resources they can access, and can contain an individual AI agent without stopping the wider workflow. ADR detects and blocks threats including prompt injection and credential exfiltration before agent actions execute, and integrates agent activity into the Singularity Platform's existing investigation and response workflows.

— Names mechanisms and integrations but no exact command or UI path

launch-20261030-2d04f863

THINNER COVERAGE BELOW

02 Claude Mythos 5.1 in Wayfinder Frontier AI Services IMPROVED

55

Wayfinder Frontier AI Services now includes Anthropic Claude Mythos 5.1, a restricted-access model, alongside existing OpenAI and Anthropic models to improve vulnerability

discovery and exploitability confirmation. Findings are filtered through exploitability testing so customers receive validated results and prioritised remediation guidance rather than raw vulnerability lists.

— Names the model and benefit but no concrete usage steps launch-20261030-e5b68b0c

NEW INTEGRATION Tenable OPEN for AI partner integrations via MCP connectors and direct APIs

> announced · Oct 28, 2026
>
> SentinelOne adds Singularity ASOC, a new autonomous SOC suite with agentic investigation, hyperautomation, AI SIEM, IOC management, and a credits-based consumption model.
>
> NEW INTEGRATION Singularity ASOC bundles Agentic Investigation, Hyperautomation, AI Data Pipelines, AI SIEM, and threat intelligence and IOC management into one suite on the Singularity Platform.
>
> Singularity Credits now govern AI-powered capability consumption across the platform, with administrators setting guardrails to keep automated actions within organisation-defined limits.

> announced · Oct 28, 2026
>
> NEW INTEGRATION Tenable OPEN for AI partner integrations via MCP connectors and direct APIs

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 GitLab

GitLab Orbit: Find vulnerability blast radius without the manual hunt (2026-10-29)

Code Security

// WHY IT MADE THE LIST

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

Orbit indexes projects, pipelines, security findings, and ownership into one relationship graph you can query in plain language to map a CVE's blast radius across a whole group, and external agents such as Claude Code and Codex reach it over MCP. It replaces manual correlation with a live, queryable context graph of the software lifecycle.

> GitLab shipped Orbit, a live relationship graph connecting code, merge requests, pipelines, vulnerabilities, work items and ownership across projects, reaching general availability for Premium and Ultimate customers with plain-language blast-radius queries, MCP/REST agent access, exact code search alongside the graph, and usage-based query-pack pricing.

└──▷ **WHAT SHIPPED · 7 FEATURES** most completely described first ? what's the number?

01 Agent access to Orbit via MCP and REST NEW

95

MCP-compatible agents — Claude Code, Codex, Cursor, OpenCode, Gemini CLI, Antigravity — connect to the Orbit context graph by running `glab auth login` then `glab orbit setup`; custom agents and internal tooling can connect through REST. Agents on Duo Agent Platform query Orbit natively with nothing to configure, and GitLab's built-in Security Analyst Agent now uses Orbit directly to traverse the graph and answer blast-radius questions from indexed project data.

Wire up an MCP-compatible agent such as Claude Code or Cursor to Orbit so it can query live codebase context instead of crawling files.

```
$ glab auth login
glab orbit setup
```

— Names exact CLI commands and every supported agent.

GitLab Orbit is GA: Higher agent accura… GitLab Orbit: Find vulnerability blast …

02 Orbit relationship graph reaches general availability NEW

78

GitLab adds Orbit, a live, queryable graph connecting code, merge requests, pipelines, vulnerabilities, work items and ownership across all projects in a top-level group — indexing projects, merge requests, pipelines, work items, security findings and source code into one graph. It indexes every push, merge and deploy in real time, handling more than 1.5 million projects and 8 TB of graph data in under 2 hours at GitLab scale, and is now generally available for Premium and Ultimate customers on `GitLab.com`, with Self-Managed and GitLab Dedicated support coming in the following weeks and a free 30-day trial available through account teams.

.

— Mechanism, scale numbers and GA rollout are detailed.

GitLab Orbit: Find vulnerability blast … | GitLab Orbit is GA: Higher agent accura…

03 Exact code search with regex mode alongside Orbit NEW

75

Exact code search at the top-level group now supports regular expression mode for finding container base images, e.g. `lang:Dockerfile ^FROM.*php` or `file:composer\.json$`, used alongside Orbit graph queries.

Find all Dockerfiles that pull a PHP base image across the entire top-level group using exact regex search alongside Orbit.

```
$ lang:Dockerfile ^FROM.*php
```

— Exact runnable query syntax provided. GitLab Orbit: Find vulnerability blast …

04 Plain-language blast-radius queries in Duo Agentic Chat NEW

70

Query blast radius in plain language via GitLab Duo Agentic Chat — for example asking it to map the blast radius of PHP in a group, listing every project with PHP code grouped by subgroup with owners, or to collect every container scanning finding for specific CVEs across the group from the vulnerability report (this requires Ultimate and the Security Manager role).

Map every project in a group that uses PHP, grouped by subgroup with owners named, during a vulnerability triage.

> 📍 In GitLab Duo Agentic Chat, type: Use Orbit to map the blast radius of PHP in this group. List every project that contains PHP code, grouped by subgroup, with owners.

Collect every container scanning finding for a specific CVE across the group via Agentic Chat (requires Ultimate and Security Manager role).

📍 In GitLab Duo Agentic Chat, type: Collect every container scanning finding for CVE-2026-109618 and CVE-2026-109619 across the group from the vulnerability report.

— Gives exact chat prompts but no API or config surface.

GitLab Orbit: Find vulnerability blast …

05 **Orbit query-pack pricing** NEW 60

Orbit is sold as query packs — a shared pool per organization — priced at $30 in GitLab Credits per user per month for 1,000 queries, with overages billed at 33 queries per GitLab Credit (about $0.03 per query).

— Concrete pricing figures but no purchase mechanism described.

GitLab Orbit is GA: Higher agent accura…

THINNER COVERAGE BELOW

06 **Role-based access and auditability for Orbit queries** NEW 35

Orbit access follows existing GitLab roles, so people and agents see only what they are permitted to see, and the graph queries behind each answer can be audited.

— States the control exists but no config or endpoint named.

GitLab Orbit: Find vulnerability blast …

07 **Data Explorer for Orbit graph queries** NEW 35

Engineers can query the Orbit graph through the Data Explorer.

— Names a UI surface with no further detail. GitLab Orbit is GA: Higher agent accura…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Joe Security (Joe Sandbox)

The Sandbox That Fights Back: Introducing Adaptive Detonation (2026-11-03)

Malware Analysis

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

Adaptive Detonation adds an AI agent that runs a sample, identifies what blocked execution — an evasion check, missing prerequisite, or failure — then supplies missing launch arguments, adjusts the guest clock or environment, and re-runs the unchanged sample to surface behavior the first detonation missed. It automates defeating sandbox evasion that previously stalled analysis entirely.

> Joe Sandbox introduces Adaptive Detonation, an AI agent built into Joe Sandbox AI that diagnoses why a sample stalled and actively intervenes — supplying missing arguments, adjusting the guest clock, or preparing files and dependencies — before re-running the sample to expose behavior the initial pass missed.

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

01 Adaptive Detonation AI agent for sandbox evasion NEW

78

Joe Sandbox adds Adaptive Detonation, an AI agent that runs a sample, identifies what blocked execution (evasion check, missing prerequisite, or failure), applies a targeted intervention, and re-analyzes the sample under the new conditions. Interventions include supplying missing launch arguments — e.g. recovering and injecting the required `S1=1` MSI installer property — without modifying the sample package; adjusting the guest machine's date or advancing its clock (demonstrated by a two-day advance to bypass a file-age loader delay) before re-running the unchanged sample in Joe Sandbox Cloud; and running preparation code to create required files, directories, or registry values, or install a missing dependency such as Python in a location the sample already searches. Analysis reports now include the agent's findings: what blocked execution, the runtime and code evidence, the recommended action, alternative explanations, and remaining uncertainty. It is available as part of Joe Sandbox AI and has been running in Joe Sandbox Cloud Basic.

.
.
.
.

— Mechanism and a named MSI property detailed, but no runnable command given.

The Sandbox That Fights Back: Introduci…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Sysdig Secure

announced · Nov 6, 2026 · Cloud Security Posture

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

Sysdig AI Defense is a runtime layer that blocks unvetted MCP servers, skills, and plugins before they run and holds risky credential reads for review, with every agent session recorded as an immutable, audit-grade action graph tied to an accountable human. Built from runtime telemetry across Claude Code, Codex, Copilot, and Cursor, it lets a team stop and prove agent actions from endpoint to cloud rather than just inventory them.

> Sysdig Secure's biggest addition this window is AI Defense, a new runtime security and governance layer for AI agents covering inventory, policy enforcement, audit trails, and cost attribution; alongside that, Vulnerability Findings gained grouped notifications and the CLI Scanner reached v1.32.0 with new OS support and detection improvements.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 CLI Scanner excludes standalone apps on older on-prem collectors

85

BREAKING

When scanning images against on-premises collectors 7.5.0 to 7.9.x that contain standalone binary applications, the CLI Scanner now excludes applications from the SBOM and scan results; upgrading the on-premises backend to 7.10.0 or later restores application inclusion.

— Names exact collector versions and the upgrade fix · Secure SaaS - November 4, 2026

02 **CLI Scanner v1.32.0 detection and OS support updates** IMPROVED

75

Sysdig CLI Scanner v1.32.0 adds detection of Claude Code installed via its native installer during image scanning; adds support for scanning images on SUSE Linux Enterprise Server 16, `openSUSE Leap 16`, and `Ubuntu 26.10` (including end-of-life information for SUSE Linux Enterprise Server 16); marks vulnerabilities with a known exploit as exploitable even when the exploit has no publication date or reference link; and improves PHP detection in official Docker Hub `php` images by also reading the `PHP_VERSION` environment variable.

— Names OS versions and env var but no CLI flags given Secure SaaS - November 4, 2026

03 **Runtime policy enforcement on AI agent actions** NEW

65

AI Defense enforces runtime policy against agent actions, blocking unvetted MCP servers, skills, and plugins before they run, and holding or blocking risky credential reads and behavior drift for human review, with exception paths for developers.

— Describes enforcement mechanism but no named config or flag launch-20261106-55e985f3

04 **Grouped notifications for Vulnerability Findings** NEW

65

Vulnerability Findings automations can now group findings by resource or image, sending one Slack or email notification — or creating one Jira ticket — per resource or image instead of per finding (Preview).

— Names channels and grouping mechanism, marked Preview Secure SaaS - October 29, 2026

05 **Live inventory of AI agents and their owners** NEW 60

AI Defense provides a live inventory of AI agents, harnesses, models, MCP servers, skills, and plugins built from runtime telemetry, with every agent traceable to an accountable human across platforms including Claude Code, Codex, GitHub Copilot, and Cursor.

— Names agent types and platforms but no concrete UI/API path launch-20261106-55e985f3

06 **Fedora subpackage vulnerability matching change** BREAKING

60

After upgrading, scans of Fedora images may show additional vulnerabilities because Fedora subpackages are now matched against their source package instead of the previous matching logic.

— Explains before/after behavior but no remediation step `Secure SaaS - November 4, 2026`

THINNER COVERAGE BELOW

07 **AI model call attribution and spend tracking** NEW 55

AI Defense attributes every model call to an agent, team, and individual — including calls that bypass an AI gateway — to surface real-time AI spend and identify where premium models are doing routine work.

— Clear capability but no dashboard or command referenced `launch-20261106-55e985f3`

08 **Audit-grade action graph for agent sessions** NEW 50

AI Defense records every agent session as an immutable, audit-grade action graph showing each step in order, which policies were applied, when, and who signed off.

— Explains the artifact but not how to access it `launch-20261106-55e985f3`

09 **OWASP-mapped policy templates for AI Defense** NEW 50

AI Defense ships with policy templates that include rules mapped to OWASP Top 10 for LLM Applications, with policies starting in alert mode and moving to enforcement when teams are ready.

— Names the standard but not specific rule names `launch-20261106-55e985f3`

10 **AI Defense general availability for macOS and Windows** NEW

20

Sysdig AI Defense is generally available for macOS and Windows.

— Bare availability statement with no further detail `launch-20261106-55e985f3`

NEW COMPLIANCE FRAMEWORK CIS Kubernetes Benchmark v2.1.0 control 1.2.1 node audit

NEW COMPLIANCE FRAMEWORK CIS Kubernetes Benchmark v2.1.0 control 4.1.11 node audit

└──▷ BREAKING ON UPGRADE

! After upgrading, scans of Fedora images may show additional vulnerabilities because Fedora subpackages are now matched against their source package instead of the previous matching logic.

! When scanning images against on-premises collectors 7.5.0 to 7.9.x that contain standalone binary applications, the CLI Scanner now excludes applications from the SBOM and scan results. Upgrade the on-premises backend to 7.10.0 or later to restore application inclusion.

> Secure SaaS - November 3, 2026
>
> **NEW COMPLIANCE FRAMEWORK** CIS Kubernetes Benchmark v2.1.0 control 1.2.1 node audit
>
> **NEW COMPLIANCE FRAMEWORK** CIS Kubernetes Benchmark v2.1.0 control 4.1.11 node audit

---

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 OWASP ZAP

announced · Nov 3, 2026

Web Application Security

// WHY IT MADE THE LIST

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

AI Assisted Authentication uses an LLM to drive a real browser through multi-step logins that Browser Based Authentication heuristics cannot handle, then converts the successful session into a reusable Zest replay script — closing a long-standing gap where complex auth flows left apps unscannable. The same release adds OAuth2 client_credentials and password-grant authentication.

> OWASP ZAP's headline addition is AI Assisted Authentication, which uses an LLM to drive a real browser through complex multi-step login flows and auto-generate a reusable Zest script; the same window also added OAuth2 authentication, an in-browser Zest script editor, LLM provider configuration, and a passive rule for server-side technology CVEs.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 AI Assisted Authentication for complex login flows NEW

96

A new `ai` authentication method in the Automation Framework `authentication` block (parameters `loginPageUrl`, `browserId`, `loginPageWait`, optional `hint`) uses an LLM to drive a real browser through multi-step login flows that Browser Based Authentication heuristics cannot handle, then automatically converts the successful session into a reusable Client Script Authentication Zest script. It can also be triggered programmatically via the `aiAssistedAuthentication` action in the authentication API, steered with free-text `hint` guidance (e.g. 'Set the domain dropdown to Project1 before clicking Next.'), monitored step-by-step in the new AI Assisted Auth tab inside the LLM Chat tab, and tried interactively via Tools > Authentication Tester (or Ctrl+T) by selecting AI Assisted as the Auth Method.

Run a scan against an app with a complex multi-step login that Browser Based Authentication cannot handle, letting the LLM work out the flow once and generate a reusable script.

```yaml
env:
  contexts:
    - name: "ai-auth-context"
      urls:
        - "https://www.example.com"
      authentication:
        method: "ai"
        parameters:
          loginPageUrl: "https://www.example.com/login"
          browserId: "firefox-headless"
          loginPageWait: 5
          hint: "Set the domain dropdown to Project1 before
clicking Next."
      verification:
        method: "autodetect"
      sessionManagement:
        method: "autodetect"
      users:
        - name: "test-user"
          credentials:
            username: "[email protected]"
            password: "${AUTH_PASSWORD}"
jobs:
  - type: llm-config
    parameters:
      defaultProvider: "my-provider"
      defaultModel: "my-model"
      providers:
        - name: "my-provider"
          type: "OPENAI_COMPATIBLE"
          apiKey: "${LLM_API_KEY}"
          endpoint: "https://llm.example.com/v1"
          models:
            - "my-model"
  - type: "requestor"
    parameters:
      user: "test-user"
      requests:
        - url: "https://www.example.com/"
          method: "GET"
          responseCode: 200
```

Open the Authentication Tester interactively to test a complex login page with AI assistance and optional hints before committing to automation.

📍 1. Ensure the LLM Support add-on is configured with a provider. 2. Open Tools › Authentication Tester (or press Ctrl+T / Cmd+T on macOS). 3. Set Auth Method to `AI Assisted`. 4. Enter the Login URL, Username, and Password. 5. In the Hints field, enter any guidance, e.g. 'Set the domain dropdown to Project1 before clicking Next.' 6. Press Test and watch progress in the results panel and the AI Assisted Auth tab inside the LLM Chat tab.

— Names method, params, API action, UI path and runnable example.

`launch-20261103-b7ebd607`

---

02 **Insights job exitAutoLevel threshold** IMPROVED 90

ZAP's Insights add-on (v0.7.0) adds an `exitAutoLevel` option (accepting `medium` or `high`) to the insights Automation Framework job, preferred over the now-secondary `exitAutoOnHigh` field.

Restrict a CI scan to exit non-zero only when high-severity findings exceed the threshold, using the new exitAutoLevel setting in the Insights Automation Framework job.

```yaml
- type: insights
  parameters:
    exitAutoLevel: high
```

— Exact config key, accepted values, and a runnable YAML example.

`launch-20261102-b931fc36`

---

03 **Client Spider as AJAX Spider stand-in with modern-app targeting**

IMPROVED 85

ZAP's Client Side Integration add-on (v0.32.0) adds a `runOnlyIfModern` option to the `spiderClient` Automation Framework job to run the Client Spider only when a `Modern Web Application` alert has been raised, and now stands in for the AJAX Spider's `spiderAjax` automation job and `ajaxSpider` API using the Client Spider when the AJAX Spider add-on is not installed.

Run the Client Spider only on pages ZAP has already flagged as modern web apps, avoiding unnecessary crawl overhead on traditional server-rendered sites.

```yaml
- type: spiderClient
  parameters:
    runOnlyIfModern: true
```

– Names both jobs, the API and the exact option, with a runnable example.

launch-20261102-b931fc36

04 **OAuth2 authentication method** NEW 74

ZAP's Authentication Helper add-on (v0.43.0) adds an OAuth2 authentication method supporting the `client_credentials` and `password` grant types, configurable via the desktop UI, the Automation Framework (`oauth2` method, added in Automation Framework v0.61.0), and the API.

– Names method, grant types and three config surfaces, no example.

launch-20261102-b931fc36

05 **Diagnostics job failure retention controls** NEW 74

ZAP's Authentication Helper add-on (v0.43.0) adds `type` and `count` parameters to the diagnostics automation job, retaining only the last failed attempt with `auth_on_failure` or a rolling window of failures with `auth_failure_rolling` .

– Names exact parameters and option values, lacks worked example.

launch-20261102-b931fc36

06 **LLM provider configuration via llm-config job** NEW 65

ZAP's LLM Support add-on (v0.1.0) adds an `llm-config` Automation Framework job to configure an LLM provider, specifying `defaultProvider` , `defaultModel` , and a `providers` list with `type` , `apiKey` , `endpoint` , and `models` fields; it underpins AI Assisted Authentication.

– Names config job and fields but lacks its own standalone example.

launch-20261103-b7ebd607 launch-20261102-b931fc36

07 **In-browser Zest Script Viewer** NEW 65

ZAP's Browser Extension (v0.3.0, included in Client Side Integration v0.34.0) adds a Script Viewer that lets you see, edit, replay (including step-by-step), load, and save a Zest script without leaving the browser.

– Describes capability and versions but no exact UI path. launch-20261102-b931fc36

08 **Passive rule for server-side technology CVEs** NEW 63

ZAP's Passive scanner rules (v77) add a Server Side Technology Version with Known Vulnerabilities rule that checks technology versions disclosed in response headers against known vulnerabilities, initially covering PHP versions via php-src security advisories.

– Named rule and initial scope, but it runs automatically with no config shown.

launch-20261102-b931fc36

THINNER COVERAGE BELOW

09 **Network add-on cross-add-on notifications** NEW 46

ZAP's Network add-on now notifies other add-ons when an outbound proxy is enabled (v0.30.0) and when requests fail to send, such as DNS errors, connection refused, or timeouts (v0.32.0).

— Internal developer-facing event with no end-user config shown. launch-20261102-b931fc36

10 **Automation Framework plan run duration limit** NEW 20

ZAP's Automation Framework (v0.61.0) now allows defining how long a plan can run.

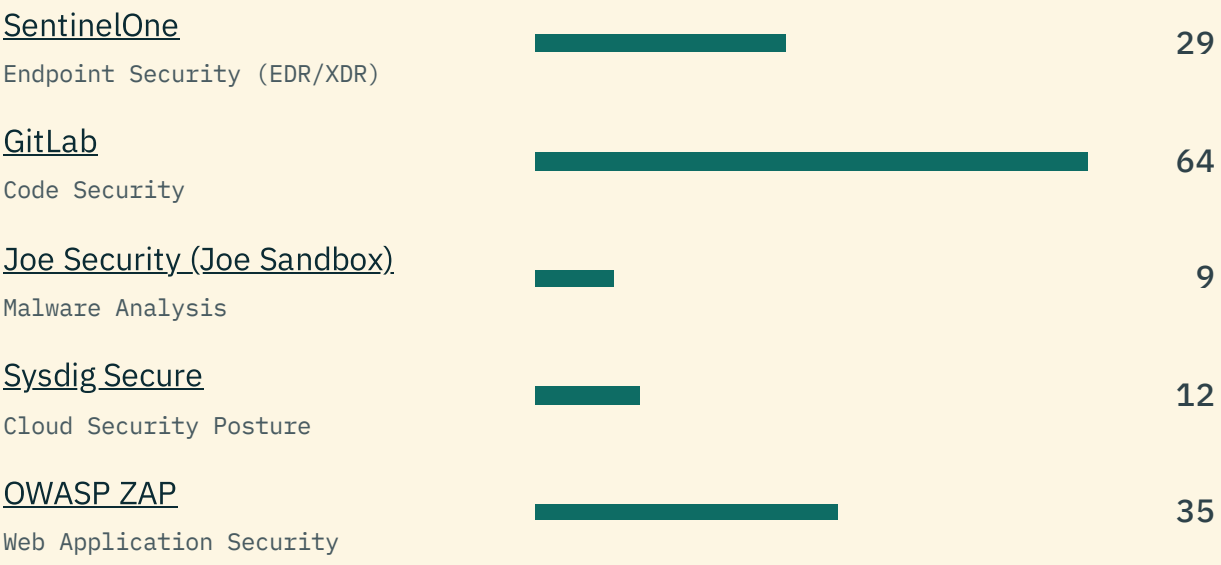
— No named flag or config key given, purely descriptive. launch-20261102-b931fc36

└──▷ **ALSO FROM THESE RELEASES**

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

Where each pick stands · Oct 30–Nov 5

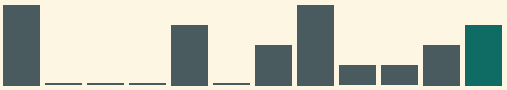
Tools with changes in each pick's category over the last 12 weeks, including the pick itself.



Each pick's release rhythm · Oct 30–Nov 5

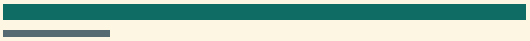
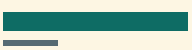
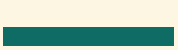
Changes per week for the last 12 weeks. The bright bar is this week; a dash marks a week with no daily.



GitLab		41 in 12 known wk
Joe Security (Joe Sandbox)		3 in 12 known wk
Sysdig Secure		20 in 12 known wk
OWASP ZAP		7 in 12 known wk

Which categories moved · Oct 30–Nov 5

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

General Utilities & Libraries		80 usual 16.2
Identity & Access Management		68 usual 62
SIEM		61 usual 39.2
Software Supply Chain Security		47 usual 23.8
Threat Intelligence		46 usual 25.2
Code Security		43 usual 80.8
Container & Image Security		36 usual 12
Web Application Security		34 usual 15.5
Reverse Engineering		33 usual 9.2
Endpoint Management (UEM/RMM)		27 usual 8
Secrets Detection & Management		26 usual 18.5

Network Analysis & Monitoring		24 usual 9

Who shipped the most · Oct 30–Nov 5

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

authentik	25
Databricks	24
OpenCTI	23
joplin	22
Wazuh	19
kubernetes-helm	17
openbao	14
GitHub	13

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · November 6, 2026

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

▷ Subscribe