THE CYBER TOOLCHAIN NO. 008

Head

THE WEEK'S TOP PICKS

PUBLISHED AUGUST 21, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

// HOW THESE PICKS ARE MADE

Every feature release from the 841 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and pick the few updates worth your attention, on two questions: **how deep and complete is the single best capability in the release**, and **how much it changes what you can actually do**.

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

// THE PICKS

01	**Bugcrowd**	Exploitation & C2
02	**Plerion**	CNAPP
03	**Endor Labs**	SAST / SCA / Code Security
04	**GitLab**	SAST / SCA / Code Security
05	**Black Duck**	SAST / SCA / Code Security

01 Bugcrowd

1 RELEASE · 2026-08-17

Exploitation & C2

Bugcrowd Customer Changelog is a feed that documents updates, new features, and improvements made to Bugcrowd's crowdsourced security platform for customers.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Savant Pathseeker launches autonomous penetration testing that runs authenticated tests behind login, including external IdPs, and pushes validated findings straight to the Security Inbox — surfacing exploitable paths a team would otherwise wait for a manual engagement to find.

> Bugcrowd launched Savant Pathseeker in Early Access, a new autonomous penetration testing product that handles authenticated tests and feeds validated findings into existing triage workflows.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first

? what's the number?

THINNER COVERAGE BELOW

01 Savant Pathseeker autonomous penetration testing (Early Access)

NEW

53

Bugcrowd introduced Savant Pathseeker, a new autonomous penetration testing product now available in Early Access. It supports authenticated tests behind login, including tests against external identity providers, allows adding assets mid-setup during an autonomous test run, and pushes validated findings to the Security Inbox for triage and workflow integration.

— Describes product capabilities but no setup steps, API, or config given.

Savant Pathseeker Early Access (1.0) (20…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Plerion

3 RELEASES · 2026-08-17 → 2026-08-21

CNAPP

Plerion is a cloud security platform that detects and responds to infrastructure vulnerabilities and misconfigurations across multi-cloud environments.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Builds an external access grant inventory across 97 AWS resource types and four sharing mechanisms — resource policies, role trust policies, RAM shares, and cross-account attributes — surfacing every external principal that can reach your resources, most of which teams had no consolidated view of. In-product access reviews let you assign owners, record keep/remove decisions, and raise a finding for every untrusted grant.

> Plerion shipped an external access grant inventory and review workflow for AWS, opened detection exemptions to the API, and rebuilt all 40 compliance frameworks against the current detection catalog while retiring 23 superseded editions.

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

01 External access grant inventory and reviews NEW

90

Plerion now builds an external access grant inventory for every tenant, analyzing 97 AWS resource types across four sharing mechanisms: resource policies, role trust policies, Resource Access Manager (RAM) shares, and cross-account permissions set directly as an attribute. Grants can be exported to a spreadsheet for offline review, or reviewed end-to-end in-product: assign an accountable team, record a `Keep`, `Remove`, `Review later`, or `Trust until review` decision, add notes, and set a future review date, with a full audit history per grant and a `Past due` preset to surface overdue reviews. A finding is raised for every asset granting access to an external principal not confirmed as trusted, with trusted principals configured under `Settings > Profiles > Trusted principals`.

Suppress the first-scan wave of untrusted-access findings by pre-loading known-good external principals before reviewing findings one by one.

> 📍 In Plerion, go to Settings › Profiles › Trusted principals and add the AWS account IDs, federated identities, or AWS services that your organization intentionally grants access to.

— Names decisions, config path, and mechanisms but not an API surface

`Review external access to AWS with resou…`

02 Compliance frameworks rebuilt with new mappings BREAKING

85

Plerion rebuilt 40 compliance frameworks from scratch against the current detection catalog, producing 15,250 control-to-detection mappings across AWS, Azure, GCP, and now Kubernetes (29 frameworks). Eighteen frameworks were promoted to their publisher's latest edition, including PCI DSS 4.0.1, NIST SP 800-53 Rev 5 Release 5.2.0, CSA CAIQ v4.1, UK NCSC CAF 4.0, and FedRAMP Rev 5 CR26 baselines. CIS Amazon Web Services Foundations Benchmark v7.0.0 was added with all 70 recommendations mapped, including six new AWS Organizations controls (centralized root access, authorization guardrails, organizational unit structure, delegated administrators) and a check for reaching AWS services over VPC endpoints, each tagged CIS Level 1 or Level 2. Twenty-three superseded framework editions were retired — including PCI DSS 4.0 and 3.2.1, NIST CSF v1.1, ISO/IEC 27001:2013, CIS AWS Foundations Benchmark v4.0.1, and FFIEC CAT (with no successor) — removing them from the framework list while preserving existing mappings; as a result, compliance posture percentages will shift in both directions across all 40 rebuilt frameworks.

— Thorough on scope and named frameworks, but no direct user action beyond awareness

`Plerion’s compliance frameworks, rebuilt…`

---

03 **API for detection exemption management** NEW

78

New API endpoints let teams read, write, and clear detection exemptions on any detection without using the platform UI. Exemption matching supports asset name, tags, region, principal, or API Gateway route as criteria, and existing findings on the dashboard are re-evaluated immediately when an exemption is written rather than waiting for the next scan. Profiles should be listed via the API first, since the default profile only covers integrations with no profile of their own.

— Describes match criteria and behavior but no exact endpoint paths given

`Detection exemptions, now over the API:`

└──▷ BREAKING ON UPGRADE

! Compliance posture percentages will shift in both directions across all 40 rebuilt frameworks because the control-to-detection mappings were reconstructed rather than patched.

! 23 retired framework editions (including PCI DSS 4.0, PCI DSS 3.2.1, NIST CSF v1.1, ISO/IEC 27001:2013, CIS AWS Foundations Benchmark v4.0.1, and FFIEC CAT) are removed from the framework list; FFIEC CAT has no successor framework.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Endor Labs

2 RELEASES · 2026-08-18 → 2026-08-19

SAST / SCA / Code 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 4/5 IMPACT 5/5

Endor Labs AI SAST now scans C buildless in the IDE and on pull requests, pairing a call-graph/reachability engine with LLM reasoning to trace cross-function buffer overflows, use-after-free, and integer overflows that pattern-based rules miss — no compiler hook or reproducible build required.

> Endor Labs extended its AI SAST engine with full C language support and open-sourced complete example patches showing how to reproduce, verify, and install its dependency fixes.

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

01 **Open-sourced Endor Patches examples repository** NEW 95

Endor Labs published two complete Endor Patch examples in the `endor-patches` repository with reproducible builds, attestations, and auditable diffs. Patched JARs in `patches/<library>/artifacts/` can be installed to a local Maven repository via `mvn install:install-file` and declared as a dependency using the `-endor-YYYY-MM-DD` version string, with transitive override support via `<dependencyManagement>` in Maven or `resolutionStrategy` in Gradle; the attestation JSON can be inspected to verify security diffs against upstream commit SHAs, and the artifact can be rebuilt from source and checksum-verified.

Independently verify a patch's security diffs by inspecting the attestation file and matching upstream commit SHAs before installing.

```
$ cat patches/com.fasterxml.woodstox/woodstox-core/security_attestation.json | jq '.[] | select(.fix_type == "FIX_TYPE_VULN")'
```

Reproduce the patched artifact from scratch and verify its checksum matches the published artifact.

```
$ # Extract the source tarball, then build with Docker and Bazel
per the included README.md
tar -xf patches/com.fasterxml.woodstox/woodstox-
core/source.tar.gz
cd woodstox-core-patched && docker run --rm -v "$PWD":/build
<hermetic-image> bazel build //...
sha256sum output/woodstox-core-6.3.1-endor-2024-09-03.jar
diff <(sha256sum output/woodstox-core-6.3.1-endor-2024-09-03.jar
| awk '{print $1}') SHA256SUM.txt
```

Install the patched JAR into your local Maven repository and override a transitive dependency so nothing else in the build needs to change.

```
$ mvn install:install-file \
  -Dfile=patches/com.fasterxml.woodstox/woodstox-
core/artifacts/woodstox-core-6.3.1-endor-2024-09-03.jar \
  -DgroupId=com.fasterxml.woodstox \
  -DartifactId=woodstox-core \
  -Dversion=6.3.1-endor-2024-09-03 \
  -Dpackaging=jar
```

— Names the repo path, exact commands, and version-string format

When you can't upgrade: open source exam…

---

02 **C language support in AI SAST** NEW 70

Full C language support was added to Endor Labs AI SAST, with buildless scans that run in the IDE and on pull requests with no compiler hook or reproducible build required. It pairs a deterministic call-graph and reachability engine with LLM reasoning to trace data flow across functions and files, catching buffer overflows, use-after-free, integer overflows, and other cross-function memory-safety flaws that pattern-based rules miss. C SAST runs alongside C SCA in the same platform, covering first-party C code and its open-source dependencies together, and is also surfaced through AURI by Endor Labs, giving an independent security harness outside the coding agent with policy enforced across any agent, model, or CI stage.

— Explains mechanism and scope but gives no exact command or config key

Why Endor Labs AI SAST for C finds what …

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 GitLab

4 RELEASES · 2026-08-18 → 2026-08-20

SAST / SCA / Code Security

GitLab is a web-based DevOps platform providing Git repository hosting, CI/CD pipelines, and collaboration tools for software development teams.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

GitLab 19.3 adds bulk-select agentic SAST False Positive Detection and Vulnerability Resolution against an existing backlog, so triage and fix-MR creation run across many findings at once — including SARIF-imported third-party scanner results, with progress and cancellation controls.

> GitLab's 19.3 window centers on agentic security and automation: bulk agentic triage and remediation for SAST backlogs, a Flow Creator agent that generates Custom Flow YAML from plain language, an AI Gateway now hosted inside GitLab Dedicated single-tenant environments, and a new Git Clone Override Policy mechanism for enforcing optimized shallow/partial clones repo-wide.

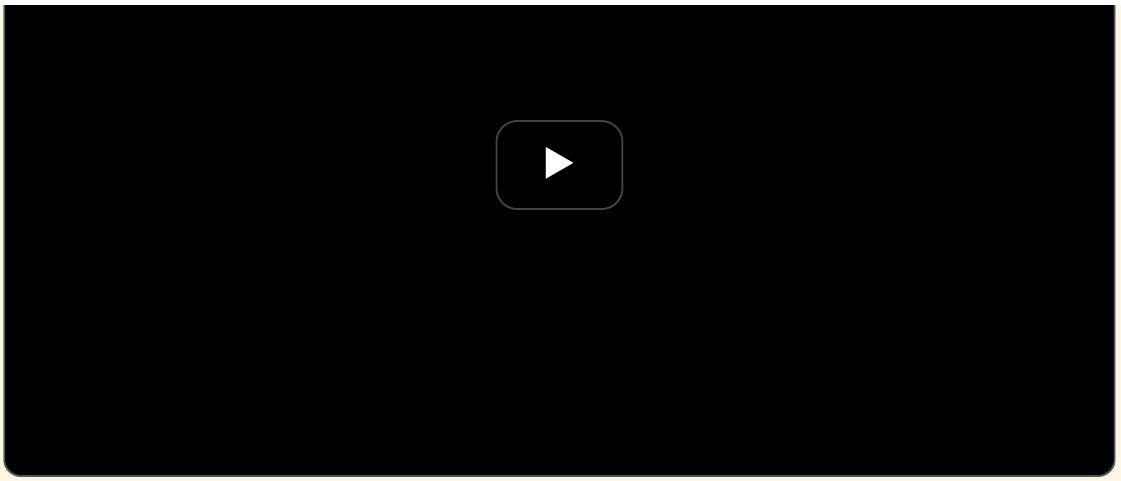
└──▷ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Git Clone Override Policy for automated shallow clones NEW

100

Adds `.afullhistorycloneoverridepolicy.toml` (TOML, `schema_version = 2`) at the repo root to declaratively configure clone optimization, including `[git_config]`, `[fetch]`, `[sparse_checkout]`, `exclude_extensions`, `conemode`, `includes`, `excludes`, `case_variants`, and `[policyinfo]` fields (`description`, `OptimalForAIAgents`). A self-contained Go binary, `git-clone-override-policy`, intercepts bare `git clone <URL>` calls and runs a fixed 13-step shallow-and-partial-clone sequence before applying the repo's policy. `[fetch]` supports flags like `--depth=1` and `--filter=blob:none`; `[sparse_checkout]` supports `conemode = 'auto'` (resolving to `cone` or `no-cone` at runtime), `includes` / `excludes`, and `exclude_extensions` covering 30+ binary file types; `[git_config]` scopes large-transfer tuning (`core.compression`, `http.postBuffer`, `http.lowSpeedLimit`, `http.lowSpeedTime`, `pack.windowMemory`, `pack.packSizeLimit`, `pack.threads`) to the cloned repo only, without touching global git config. The binary runs on Linux, macOS, and Windows for `amd64` and `arm64` (CMD, PowerShell, Git Bash, WSL) with only a `git` binary on `PATH` required, and is bypass-safe: clones already carrying `--depth`, `--filter`, or any other flag, and repos without a policy file, pass through unchanged.







Commit a policy file to a monorepo so every bare clone — by developers, CI jobs, or AI agents — automatically gets a shallow partial clone that skips binaries, without anyone remembering the flags.

```toml
# .afullhistorycloneoverridepolicy.toml at repo root
schema_version = 2

[policyinfo]
description = "Optimized clone for CI and AI agents"
OptimalForAIAgents = true

[git_config]
"core.compression" = 0
"http.postBuffer" = "1024M"
"pack.threads" = 4

[fetch]
flags = ["--depth=1", "--filter=blob:none"]

[sparse_checkout]
enabled = true
conemode = "auto"
includes = ["/*"]
exclude_extensions = [
  "png", "jpg", "gif", "svg", "pdf",
  "zip", "tar", "gz", "mp4", "exe",
  "dll", "woff", "woff2", "ttf", "psd"
]
case_variants = "both"
```

Run the override binary in place of `git clone` so the policy is enforced for a one-off zero-footprint clone without any global git config changes.

```
$ ./git-clone-override-policy https://example.com/group/my-repo.git
```

— Full config schema and runnable binary command both given

Avoid the massive end-to-end tax of defa…

02 **Flow Creator agent for natural-language automation** NEW

88

The Flow Creator agent, accessible via GitLab Duo Agentic Chat, converts plain-language descriptions into complete, runnable Flow Registry YAML definitions with no YAML schema knowledge required. It reads current Flow Registry documentation before every response, applies rules encoding known failure patterns, and runs a pre-output checklist catching issues like missing `project_id`, broken human-in-the-loop gates, wrong

`sends_response_to` targets, and missing stopping instructions. Generated definitions can be pasted into the configuration editor when creating a flow from a project or the AI Catalog with private, restricted, or public visibility; enabling a flow still requires Maintainer role or higher, Group Owners and instance administrators can disable Custom Flows entirely, and every flow runs under a scoped service account with composite identity bounded to the top-level group.

Use Flow Creator in Agentic Chat to instantly generate a triage automation without writing any YAML schema.

> 📍 Steps: 1. Open GitLab Duo Agentic Chat in your GitLab 19.3 instance. 2. Invoke Flow Creator and enter: 'When a work item is created with the label security-triage, check it against our known false-positive patterns, find the affected dependency and its owner, and post a severity assessment as a comment.' 3. Respond to any clarifying questions (e.g., target project, whether a human-approval step is needed). 4. Copy the generated flow definition YAML. 5. Go to your project › AI Catalog › Create Flow, paste the YAML into the configuration editor, set visibility (private / restricted / public), and add the desired triggers.

— Detailed mechanism and named fields, but no exact runnable command

`Build custom flows in minutes with the F…`

---

03 **Bulk SAST triage and remediation for vulnerability backlogs** NEW

85

Adds a bulk-select UI in the vulnerability report to run SAST False Positive Detection and Agentic SAST Vulnerability Resolution against multiple findings at once, replacing one-at-a-time manual triage. False Positive Detection analyzes surrounding code and returns a confidence score and explanation; SARIF-imported findings from third-party scanners into the unified vulnerability report are also eligible for these bulk agentic flows. Pipelines can run both automatically on Critical and High severity new findings so triage and fix MR creation happen before a developer opens the finding, and bulk jobs show a progress indicator with a cancellation control plus concurrency limits to avoid competing with pipeline build capacity.

Triage and remediate your entire SAST backlog at once instead of opening each finding individually.

> 📍 In the GitLab vulnerability report, select all SAST findings using the bulk-select checkbox, then choose 'Run SAST False Positive Detection' followed by 'Run Agentic SAST Vulnerability Resolution' from the bulk-action menu. A progress indicator tracks the job; click Cancel to stop it at any point.

— UI flow and mechanism named, but no API/CLI surface given

`When your backlog outgrows your team, Gi…`

---

04 AI Gateway hosted inside GitLab Dedicated tenant NEW

60

GitLab Dedicated now hosts the AI Gateway for the Duo Agent Platform inside each customer's single-tenant AWS environment, keeping AI inference within the customer's AWS region and data residency boundary. It supports Amazon Bedrock as a model backend and adds bring-your-own-model configuration so customers can connect preferred model providers under their own data classification rules, unlocking all Duo Agent Platform capabilities — including automated custom code reviews, fixing failed pipelines, and security analysis — against customer-selected models on Dedicated deployments.

— Describes architecture and backends but no config steps given

Run agentic software delivery inside the…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Black Duck

1 RELEASE · 2026-08-19

SAST / SCA / Code Security

Black Duck is a software composition analysis tool that identifies and manages open source components, vulnerabilities, and licensing risks in applications.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 3/5

Black Duck Polaris adds bidirectional triage, status, and fix-by-date sync across Jira, ServiceNow, GitHub/GitLab Issues, and Azure Boards, plus CI-native AI-assisted SAST fix pull requests and central rule governance enforced across pipelines, CLI, and integrations.

> Black Duck's Polaris update centers on two-way bug tracker synchronization across five systems and AI-assisted SAST fix pull requests, alongside secure tunnels to self-hosted SCMs, central rule governance, and expanded Rapid Scan Static and fAST SCA engine coverage.

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

01 Two-way bug tracker synchronization across five systems NEW

83

Polaris adds bidirectional status and triage sync across GitLab Issues, GitHub Issues, Jira, ServiceNow, and Azure Boards — triage states, ticket statuses, and fix-by dates now propagate automatically in both directions, with fix-by date changes in Polaris updating due dates in Jira, ServiceNow, and Azure Boards and vice versa. GitLab Issues, GitHub Issues, and ServiceNow become new native bug-tracking targets supporting individual and bulk export, single-ticket bundling for ServiceNow, and policy-driven issue creation. Configurable Jira field mapping lets teams map Polaris issue data to standard or custom Jira fields and customize ticket descriptions, and configurable per-project bug-tracker integration mappings can be saved as reusable templates scoped to all branches or the default branch only. Linked tickets close automatically when findings are dismissed or absent in a subsequent scan, and a reappearing issue creates a new ticket to preserve audit trails.

— Names all five systems and mapping mechanics but no API endpoint given

Polaris release update: Two-way bug trac…

02 AI-assisted SAST fix pull requests NEW

83

Polaris now generates AI-assisted fix PRs for SAST findings, triggerable CI-natively via the Bridge CLI from GitHub Actions, GitLab CI, Bitbucket Pipelines, Jenkins, or Azure DevOps pipeline configuration with no server-side scheduling required. Fix PRs can also be created on-demand or in bulk from the UI, or automatically via issue policy actions,

with per-branch or per-run caps on open PRs and filtering by CWE, issue type, language, and file pattern. A new 'View issue policies' action on a SAST finding surfaces which policies, compliance requirements, and governance rules are attached to it.

— Names CLI, pipelines and filters but no exact command syntax

Polaris release update: Two-way bug trac…

03 **Central rule configuration governance** NEW 77

Security teams can now define rule profiles at organization, application, project, or branch level with hierarchical inheritance, enforced automatically across CI pipelines, CLI scans, and integrations. The feature adds RBAC, audit trails, and version awareness, with visibility into which rule profile was used for each scan.

Apply a mandatory central rule profile at the organization level so every scan — CI, CLI, or integration — enforces the same ruleset without per-repo YAML.

```
📍 # In Polaris UI: Settings › Rule Profiles › New Profile # Set scope: Organization # Enable 'Mandatory' and configure inheritance for application/project/branch overrides # All pipelines pick up the profile automatically on next scan — no per-pipeline changes needed.
```

— Includes a concrete UI navigation example for a mandatory org profile

Polaris release update: Two-way bug trac…

04 **Secure tunnels to self-hosted SCM and Jira servers** NEW 70

Polaris adds outbound-only secure tunnel connectivity (no inbound firewall changes required) to self-hosted Azure DevOps Server, GitLab Self-Managed, Bitbucket Server and Data Center, and on-premises Jira, built on Bridge CLI and Teleport infrastructure. Multitunnel support for on-premises Jira allows multiple concurrent tunnels, with tunnel setup restricted to org admins via RBAC and full audit logging of tunnel activity.

— Names mechanism and products but gives no setup command

Polaris release update: Two-way bug trac…

05 **Rapid Scan Static engine checks and language support** NEW 68

The Rapid Scan Static engine extends seven existing Java JWT checks (algorithm confusion, nonexpiring tokens, ignored expiration/start times, missing claims validation, untrusted decoding) to Scala, adds six new Scala web security checks covering permissive and credentialed CORS origins, excessive preflight age, missing TLS, and disabled security headers across http4s and Play, adds a new crypto bits private key pattern to detect embedded cryptographic private keys in source and configuration files, and adds support for Go 1.26 and Dart 3.11.

— Names checks and languages but no flag to enable them

Polaris release update: Two-way bug trac…

06 **fAST SCA package manager and ecosystem coverage** NEW

68

The fAST SCA engine adds new package manager detectors for UV (Python), Rush monorepo tooling, a tree-based Conda detector, and Ant/Ivy dependency extraction. It also extends ecosystem coverage to Poetry, pnpm, SBT, SetupTools, Bazel 9, `.slnx` solution files for NuGet, buildless `go.mod` analysis for Go, and Cargo workspace support with include and exclude controls.

— Names many ecosystems and formats but shows no command

Polaris release update: Two-way bug trac…

07 **Webhook-triggered scanning for self-hosted SCMs** NEW

60

Polaris adds webhook-triggered SAST and SCA scanning on push, branch creation, and pull/merge request activity for self-hosted SCMs, with PR/MR decoration including severity breakdowns and Code Insights reports for Bitbucket.

— Names trigger events and Bitbucket Code Insights but no config steps

Polaris release update: Two-way bug trac…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · August 21, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe