THE CYBER TOOLCHAIN

NO. 031

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 31, 2026 · EVERY DAY

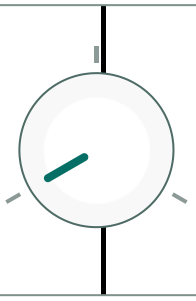
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 53 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ changelog · Jul 31, 2026 CHANGELOG ↗

Bugcrowd is a crowdsourced security platform that connects organizations with ethical hackers to identify and fix vulnerabilities.

Bugcrowd adds managing and assigning researcher credentials for customers.

Enables customers to manage and assign researcher credentials within the Bugcrowd platform.

WAS THIS USEFUL? 👍 👎

shellfire

Feedback ⓘ v0.15 NOTES ↗

An exploitation shell focusing on exploiting command injection vulnerabilities, eg., LFI, RFI, SSTI, etc.

shellfire v0.15 adds a built-in intercepting proxy to capture live browser sessions and seed exploit configs automatically.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.15 https://github.com/unix-ninja/shellfire.git
# already have the repo? check out this version:
$ git checkout v0.15
```

└──▷ **TRY IT**

Capture a real authenticated browser request to a vulnerable target and instantly seed your exploit config — no manual URL/cookie transcription needed.

```
$ (config)> proxy start 8080
(config)> proxy list
(config)> proxy use 0
(config)> url https://example.com/view?file={}
(config)> shell
```

Inspect the full contents of a captured request before committing it to your exploit config.

```
$ (config)> proxy show 1
```

Adds a built-in HTTP/HTTPS intercepting proxy (`proxy start`) to observe real browser sessions without manually transcribing URLs, cookies, or POST bodies.

New `proxy list` command enumerates all requests captured during a proxy session for review.

New `proxy show <index>` command displays the full details of a captured request.

New `proxy use <index>` command seeds the active exploit config (URL, method, headers, cookies, POST body) directly from a captured request.

New `proxy ca` command reprints the CA certificate path needed to intercept HTTPS traffic.

+1 more

WAS THIS USEFUL?  

APPLICATION SECURITY

◆ Web Application Security

Acunetix 360

Feedback ⓘ changelog · Jul 31, 2026 CHANGELOG ↗

Acunetix 360 is a cloud-based web application security testing platform that identifies vulnerabilities through automated scanning and penetration testing.

Acunetix 360 adds agent-page alerts for failed vulnerability database updates.

WAS THIS USEFUL?  

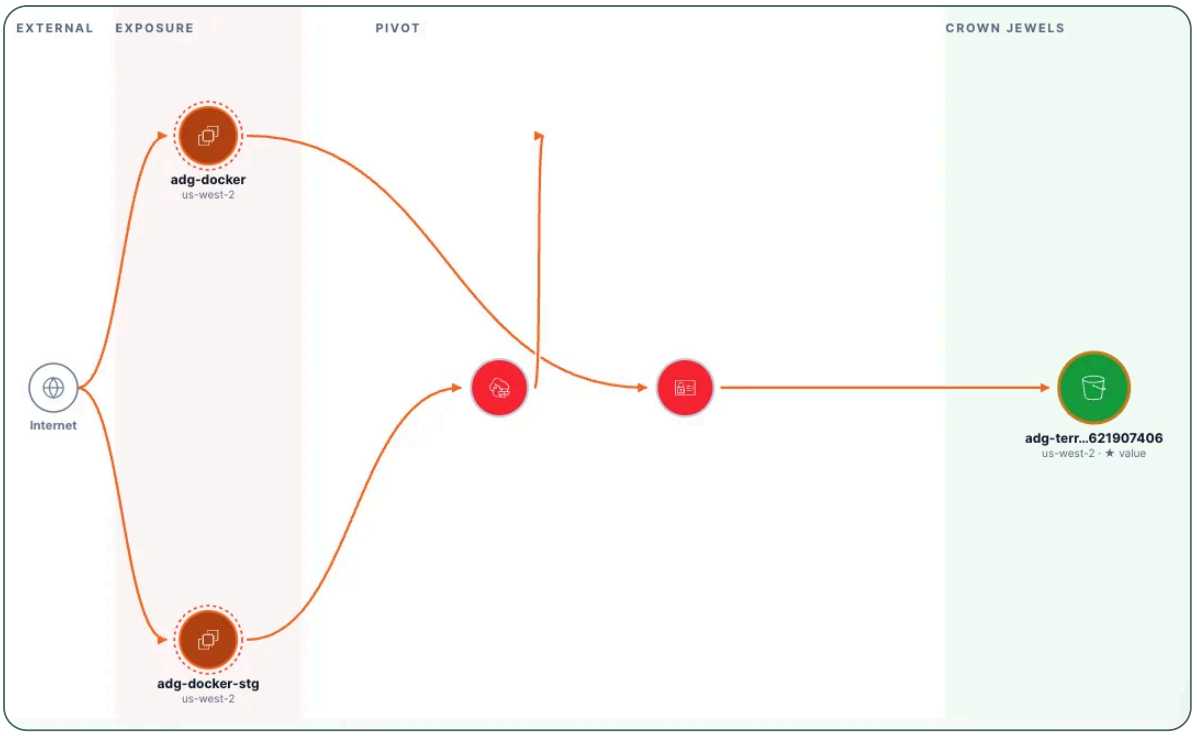
Wallarm

Feedback

See how an attacker would reach your critical AWS assets (2026-07-30) BLOG ↗

Wallarm is a cloud-native API security platform that detects and prevents attacks on APIs and web applications in real-time.

Wallarm adds Attack Paths to Infrastructure Discovery, mapping how attackers could reach critical AWS assets via IAM and network relationships.



└──▷ **HOW TO FIND IT**

Navigate to the new Attack Paths view to see ranked attacker routes through your connected AWS accounts.

📍 In the Wallarm console, go to Infrastructure Discovery → Attack Paths to view the ranked path map for your connected AWS accounts.

Define your own crown jewels so path severity scoring reflects your business priorities rather than default assumptions.

📍 In the Wallarm console, go to Infrastructure Discovery → Attack Paths, select an asset on the map, and use Pin / Unpin to mark it as a crown jewel. Path scores recalculate immediately.

New Attack Paths feature in Infrastructure Discovery visualizes end-to-end attacker routes through AWS estates, built from configuration and IAM relationships rather than CVE findings.

Paths are ranked by severity and displayed as a map showing entry points, pivot nodes, choke points, and crown jewels with plain-English explanations.

Crown jewels are user-defined: pin and unpin assets from the console to reprioritize path scoring around your own high-value targets.

Identity-aware path analysis follows IAM chains including role chaining, pass-role, policy attachment, self-escalation, and effective-admin through AWS managed policies.

Choke point identification surfaces nodes through which the most attack paths run, highlighting where a single remediation has the greatest blast-radius reduction.

+2 more

WAS THIS USEFUL? 👍 👎

Invicti

Feedback ⓘ Release 20260723 (2026-07-23) BLOG ↗

Invicti is an automated application security testing platform that identifies and reports vulnerabilities in web applications and APIs.

LSR reports now capture full HTTP request/response traffic with automatic token obfuscation for each recorded login step.

Adds full HTTP request and response details for each step in Login Sequence Recorder (LSR) reports, with automatic obfuscation of sensitive tokens in headers and bodies.

Adds a direct link to the Invicti Azure Pipelines plugin in the Azure Marketplace within the CI/CD integration setup, enabling plugin-based integration alongside the existing script-based method.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Self-Hosted

Feedback ⓘ

changelog · Jul 31, 2026 CHANGELOG ↗

GitGuardian Self-Hosted is an on-premise security tool that detects and prevents secrets leakage in Git repositories.

GitGuardian Self-Hosted 2026.7 adds Public Monitoring, native Gerrit scanning, Smart Notifiers, MCP server, and 9 new secret detectors.

└──▷ USE IT

Enable the Public Monitoring worker so your self-hosted instance processes public perimeter secret-leak incidents.

```yaml
celeryWorkers:
  public-monitoring:
    enabled: true
```

Deploy the GitGuardian MCP server in your cluster to expose the beta MCP endpoint.

```yaml
mcpServer:
  enabled: true
```

Disable migration-log sidecar collection on a Kubernetes cluster older than 1.29 to prevent upgrade failures.

```yaml
logCollector:
  collectMigrationLogs: false
```

Adds Public Monitoring on self-hosted instances (license-gated): monitor your public perimeter for secret leaks and route public incidents through webhooks, Slack, Microsoft Teams, ServiceNow, and email alerts filterable by public incident severity.

Adds native Gerrit VCS source with full historical scanning, real-time detection via the webhook plugin, and read-replica support for cloning.

Introduces Smart Notifiers: define per-channel rules to filter alerts, tickets, or emails by severity, validity, secret type, or GitGuardian tag across internal and public monitoring.

GitHub check runs now scan oversized pull requests by cloning the repository branch instead of calling the GitHub API commit-by-commit, with configurable size thresholds and the ability to skip merge commits to protect shared API rate limits.

Supports routing external traffic through the Kubernetes Gateway API instead of the legacy Ingress resource, with Kubernetes HPA or KEDA autoscaling for webapps.

+7 more

└─▷ **BREAKING ON UPGRADE**

! `logCollector.collectMigrationLogs` is enabled by default and requires Kubernetes 1.29 or later; clusters running older Kubernetes versions will break unless `logCollector.collectMigrationLogs: false` is set before upgrading.

! Developer Endpoint Protection is not functional on self-hosted instances in this release; installations relying on Endpoint Protection will lose that capability on upgrade.

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

Feedback ⓘ changelog · Jul 31, 2026

CHANGELOG ↗

GitGuardian SaaS is a secrets detection platform that scans code repositories and developer workflows to identify and remediate exposed credentials and sensitive data.

GitGuardian adds a dedicated GitHub Check Runs dashboard page, bulk perimeter actions, and Bitbucket source health management.

└──▷ HOW TO FIND IT

Bulk-unmonitor a set of deprecated sources and remove them from all team assignments in one operation.

> 📍 On the perimeter page, select the sources you want to remove from monitoring, then choose Unmonitor from the bulk actions menu. Confirm the popup that appears when more than 30 sources are selected.

New dedicated GitHub Check Runs page in the dashboard lets security teams browse all check run statuses, filter by repository/status/date, view KPI metrics, re-run or skip checks, and jump directly to linked incidents — without leaving GitGuardian.

Fine-grained GitHub personal access tokens can now be used to monitor private members of a GitHub organization (previously only classic tokens were supported).

Customizable column display on the perimeter table lets users surface the source attributes most relevant to their workflow.

Bulk actions on perimeter sources now support Monitor, Unmonitor, and Assign Team operations on multiple sources at once, with a confirmation popup when more than 30 sources are selected.

Historical scans now have a dedicated tab on the perimeter page for tracking past and in-progress scan statuses in one place.

> └──▷ BREAKING ON UPGRADE
>
> ! The `Display as Tree` view has been removed from integration source pages; only the `Display as List` view remains.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Snyk Code

Feedback ⓘ changelog · Jul 31, 2026 BLOG ↗

Snyk Code is a developer security tool that identifies and fixes vulnerabilities in source code during development.

Snyk Code adds COBOL SAST, Python LangChain & Lambda support, Spring Data taint tracking, and new framework coverage.

Adds generally available SAST scanning for fixed-format COBOL, bringing mainframe applications into the same CLI, IDE, SCM, and PR-check workflows used for other languages.

Adds Python LangChain support, flagging unsafe data flows through LangChain applications.

Adds Python serverless support for applications running on AWS Lambda.

Adds weak-hash detection and detection of debug features left enabled.

Adds Java Spring Data taint-source tracking, enabling injection, SSRF, path-traversal, and XSS findings for data flowing from Spring Data repositories.

+1 more

└──▷ **BREAKING ON UPGRADE**

! The COBOL preview flag is removed at general availability; COBOL scanning now runs automatically across all repositories, and new findings may appear in repositories not previously scanned.

! Fix examples no longer appear in Fix analysis in the Snyk Web UI and no longer appear in SARIF output.

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

CLI 1.20.0 no longer requires Python (2026-07-30) BLOG ↗

Cloudsmith is a cloud-native, universal package management platform that enables teams to securely host, manage, and distribute software artifacts across multiple formats.

Cloudsmith CLI 1.20.0 ships as a standalone binary, dropping the Python runtime dependency on Linux, macOS, and Windows.

Ships the Cloudsmith CLI as a self-contained binary for Linux, macOS, and Windows — no local Python installation required.

Adds a new one-line installer for the standalone binary.

Refreshed Homebrew tap and Docker image for the standalone CLI.

Updated GitHub Actions, Azure DevOps, and CircleCI integrations to use the standalone binary.

WAS THIS USEFUL? 👍 👎

OWASP dependency-track

Feedback ⓘ 5.0.4 NOTES ↗

Dependency-Track is an intelligent Component Analysis platform that allows organizations to identify and reduce risk in the software supply chain.

Dependency-Track 5.0.4 adds configurable query timeouts and enforced execution timeouts for background workflow activities.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 5.0.4
https://github.com/DependencyTrack/dependency-track.git
# already have the repo? check out this version:
$ git checkout 5.0.4
```

Makes DEX query timeout configurable for tuning background task responsiveness.

Enforces execution timeouts for workflow activities with a defined default, preventing runaway background tasks.

Enables automatic heartbeating for long-running activities to prevent premature timeout detection.

Improves workflow task scheduling performance for deep backlogs using concurrency keys.

Executes multiple deletion batches per maintenance cycle, increasing throughput for cleanup operations.

WAS THIS USEFUL? 👍 👎

Wiz Code

Feedback ⓘ changelog · Jul 30, 2026 BLOG ↗

Wiz Code is a developer security tool that identifies and remediates code vulnerabilities and misconfigurations in applications.

Wiz Code adds AI SAST (Private Preview) and a three-tier agentic scanning system with Atlas deep-analysis and Google CodeMender integration.

└──▷ HOW TO FIND IT

Enroll your organization in broad, continuous AI code scanning to catch business logic flaws on every PR without custom configuration.

> 📍 In the Wiz console, go to Wiz Code settings and join the Wiz AI SAST Private Preview to enable continuous AI scanning within your existing Wiz Code workflows.

New Wiz AI SAST (Private Preview): continuous AI-powered first-party code analysis that models application entry points, trust boundaries, and data flows to surface complex business logic flaws missed by pattern-matching scanners.

New Wiz Rule-Based SAST (GA): deterministic scanning that automatically blocks known vulnerability patterns, insecure code, and hardcoded secrets before code is merged.

Atlas vulnerability research system is being productized as a first-party deep-analysis engine for targeted agentic pentesting of high-value applications and critical code.

New orchestration of Google CodeMender as a third-party agentic engine for complementary code analysis and remediation.

New open, multi-engine orchestration layer routes analysis across first- and third-party agentic scanning engines, with all findings correlated on the Wiz Security Graph and flowing through existing AppSec validation and remediation workflows.

+1 more

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

CHIPSEC

Feedback ⓘ 2.0.7 NOTES ↗ CODE ↗

Platform Security Assessment Framework

CHIPSEC 2.0.7 adds HOB listing/dumping, Q35/QEMU platform support, and PCH detection in banner output.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.0.7 https://github.com/chipsec/chipsec.git
# already have the repo? check out this version:
$ git checkout 2.0.7
```

Adds `uefi` subcommands to list and dump Hand-Off Blocks (HOBs) from UEFI firmware.

Adds Q35 platform support for QEMU-based testing environments.

Displays detected PCH (Platform Controller Hub) information in the CHIPSEC banner at runtime.

WAS THIS USEFUL? 👍 👎

pwndbg

Feedback ⓘ 2026.07.29 NOTES ↗ CODE ↗

Exploit Development and Reverse Engineering with GDB & LLDB Made Easy

pwndbg 2026.07.29 adds exithandlers inspection, four kernel task commands, glibc 2.43 heap support, and heap allocation call-site tracking.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 2026.07.29 https://github.com/pwndbg/pwndbg.git
# already have the repo? check out this version:
$ git checkout 2026.07.29
```

.
.
.
.

└──▷ **TRY IT**

Inspect registered glibc exit handlers during exploitation to identify hijack opportunities and see the pointer_guard cookie.

```
$ exithandlers
```

Track heap allocations in a running process and annotate every malloc/free event with the nearest symbol that triggered it.

```
$ track-heap enable --where
```

Step execution forward until a specific syscall is reached, useful for quickly reaching an interesting I/O or exec call.

```
$ stepsyscall write
```

Inspect the seccomp filter of a kernel task by PID when debugging a kernel with qemu-system.

```
$ kseccomp 1234
```

Adds `exithandlers` command (alias `exitfuncs`) to list registered glibc exit handlers by walking `__exit_funcs` and `tls_dtor_list`, demangling `PTR_MANGLE`d function

pointers, and printing the `pointer_guard` cookie.

Adds four new Linux kernel task introspection commands without requiring typeinfo: `kstack [pid]`, `kseccomp [pid]`, `ksighand [pid]`, and `knamespace [pid]`.

Adds `percpu` GDB function to resolve per-CPU variables inside `print` / `break` expressions during kernel debugging.

Adds `v2p` and `p2v` GDB function helpers to convert between virtual and physical addresses inside expressions.

Adds support for breaking at the start of a JIT-compiled eBPF program.

+22 more

└──▷ **BREAKING ON UPGRADE**

- ! The `mallocng-dump`, `mallocng-vis`, `jemalloc-heap`, `jemalloc-extent-info`, `knft-list-tables` (and similar dash-separated) commands are replaced by subcommands: use `mallocng dump`, `mallocng vis`, `jemalloc heap`, `jemalloc extent-info`, `knft list-tables` etc. instead.
- ! The `cyclic` command now generates a de Bruijn sequence consistent with `pwn cyclic` from pwntools; previously generated sequences will differ.
- ! Developer API: `pwndbg.aglib.disasm.arch` is renamed to `pwndbg.aglib.disasm.assistant`.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Aikido Security

Feedback ⓘ

#52 - DSPM, Aikido Libraries and Repository Code Quality scans NOTES ↗

Aikido Security is a developer-first security platform that identifies and fixes vulnerabilities in code, dependencies, and infrastructure.

Aikido adds DSPM, backport CVE patches via Aikido Libraries, code quality repo scans, and new language firewall agents.

└──▷ HOW TO FIND IT

Point Aikido at a repository to audit sensitive data flows and exposures at the source — no database credentials needed.

📍 In the Aikido console, go to your repository, open the Data Exposure Audit (DSPM) scan, and review where sensitive data is exposed in code.

Gate pull requests on code quality severity so high-severity issues block merges automatically.

📍 In the Aikido console, go to Code Quality settings for your repository and enable PR/MR Gating by Severity, then select the minimum severity threshold to block.

Assign multiple cloud accounts to a team in bulk rather than one at a time.

📍 In the Aikido console, go to Settings › Clouds, select the cloud accounts you want to assign, and use Bulk assignment to map them to a team.

New Data Exposure Audit (DSPM) feature tracks how sensitive data flows and where it gets exposed, reading entirely from source code — no database connections or data access grants required.

New Aikido Libraries capability delivers PRs that backport open-source CVE fixes to the exact version already running, avoiding major-version bumps or breaking changes.

New Code Quality Repository Scans surface existing code quality issues across repos to help find and prioritize technical debt over time.

Noise reduction for container images in cloud scanning reduces alert fatigue.

Bitbucket Server support added for PR Checks.

+8 more

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Buildah

Feedback ⓘ v1.45.0 NOTES ↗ CODE ↗

A tool that facilitates building OCI images.

Buildah v1.45.0 adds `--no-follow-symlinks` for COPY/ADD, `--timestamp` for manifest push, symbolic chmod support, and new SFTP/SSHFS CLI commands.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.45.0 https://github.com/podman-container-tools/buildah.git
# already have the repo? check out this version:
$ git checkout v1.45.0
```

└──▷ TRY IT

Preserve symlinks as-is during a COPY step instead of following them to their targets — useful when the symlink itself is the artifact.

```
$ buildah copy --no-follow-symlinks mycontainer ./src/mylink /app/mylink
```

Push a manifest list with all timestamps clamped to a fixed epoch value for a fully reproducible, content-addressable image.

```
$ buildah manifest push --timestamp 1700000000 mylist docker://registry.example.com/myimage:latest
```

Adds `--no-follow-symlinks` flag to `buildah add` and `buildah copy` to copy symlinks themselves rather than their targets.

Adds `--timestamp` flag to `buildah manifest push` to set timestamps on new manifest content to a fixed epoch-relative value for reproducible builds.

Adds support for symbolic chmod options (e.g. `u+x`, `go-w`) in `buildah add` and `buildah copy` `--chmod`.

Adds Copier.Symlink() API for creating symlinks with inline targets.

Adds `CreateDestPath` option to copier `PutOptions` to automatically create destination paths.

+4 more

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Wiz Cloud

Feedback ⓘ changelog · Jul 30, 2026 BLOG ↗

Wiz Cloud is a cloud security platform that provides full-stack visibility, risk assessment, and threat detection across multi-cloud environments.

Wiz launches ASM and Red Agent in GA: continuous attack-surface monitoring plus AI-driven autonomous penetration testing.

New Wiz ASM (GA): continuously identifies exposed assets, weak credentials, and misconfigurations, linking risks to application and code ownership so teams know where risk is introduced and where it becomes exposed.

New Wiz Red Agent (GA): autonomously tests live environments and applications using AI to determine which exposures lead to real, exploitable attack paths.

WAS THIS USEFUL? 👍 👎

Microsoft Defender for Cloud

Feedback ⓘ

changelog · Jul 30, 2026 BLOG ↗

Microsoft Defender for Cloud monitors and protects cloud infrastructure, applications, and data across Azure, AWS, and on-premises environments against threats.

Defender for Cloud CSPM now covers serverless containers on Azure Container Apps, ACI, and AWS ECS on Fargate.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Elastic Security Labs Feedback ⓘ

Elastic goes all-in on Hacker Summer Camp at Black Hat and DEF CON in Las Vegas (2026-07

3 RELEASES · 2026-07-31 BLOG ↗ LATEST STABLE

Elastic Security Labs publishes threat research, vulnerability analysis, and security insights to help organizations understand emerging cyber threats.

Elastic Security adds autonomous attack investigation, auto-drafted detection rules, Windows on ARM endpoint protection, and plain-language workflow authoring.

Attack Discovery now performs autonomous analyst-style investigation: threat-hunting raw events, checking entity risk, corroborating across data sources, and classifying events as validated attacks before they reach the queue.

Attack Discovery auto-drafts detection rules when it finds gaps not covered by existing rules, routing the draft to an analyst for approval.

New alert analysis workflow differentiates likely false positives from true positives upstream, reducing noise before Attack Discovery investigates.

Elastic Defend gains always-on vulnerable driver coverage: threat research continuously monitors VirusTotal, loldrivers.io, and Microsoft's blocklist and automatically generates and instantly deploys YARA rules as new drivers are disclosed.

Elastic Defend now provides full endpoint protection for Windows on ARM devices (e.g., Surface laptops), bringing them into the same protection tier as the rest of the fleet.

+6 more

▾ 2 more releases in this issue · 2026-07-31

Alert Zero: AI-driven alert triage and attack investigation for the agentic SOC (2026-

BLOG ↗ STABLE

Elastic Security 9.5 ships Alert Zero: AI-driven alert triage, attack chain investigation, and SOC workflow integration.

.
.
.
.

└──▷ HOW TO FIND IT

Schedule Attack Discovery to run automatically inside an existing Elastic Workflow so investigations start without waiting for an analyst.

📍 In the Elastic Security console, go to Workflows, open or create a playbook, add a step using the Attack Discovery skill (retrieves alerts, generates discoveries, validates results), then configure a schedule trigger so the pipeline runs continuously throughout the shift.

Enable AI-driven first-pass triage on specific detection rules and review classifications before optionally turning on auto-close.

📍 In Elastic Security 9.5, go to the Security alert analysis workflow, select the detection rules to enroll, choose the AI agent and model, set the confidence threshold, and leave auto-close disabled initially — review the notes and tags added to alerts to validate accuracy before enabling auto-close for high-confidence false positives.

Review an AI-drafted ES|QL detection rule surfaced by Attack Discovery's gap analysis, then explicitly approve it to close a coverage gap found during investigation.

📍 In the Attacks page under Detections, open an Attack Discovery result, navigate to the detection gap analysis section, review the drafted ES|QL rule, and click Approve to create the detection — no rule is added until you explicitly confirm.

Adds the Security alert analysis workflow for first-pass AI triage of alerts generated by selected rules, classifying them as true positive, false positive, or inconclusive with optional auto-close above a configurable confidence threshold.

Adds a new Attacks page under Detections that surfaces Attack Discovery results alongside Alerts, making correlated attack chains a first-class operational object.

Extends Attack Discovery with a skill-backed investigation pipeline (powered by Elastic Agent Builder and Elastic Workflows) that hunts for related activity, inspects entity context, and searches raw logs before generating an attack narrative.

Adds detection gap analysis inside Attack Discovery: when the investigation uncovers activity missed by existing detections, it drafts an ES|QL rule for analyst review and explicit approval before creation.

Enables Attack Discovery to run on a schedule or be triggered from within an Elastic Workflow, making deep investigation always-on without requiring manual analyst initiation.

+3 more

What's new in Elastic Defend: 800+ vulnerable driver rules, automated troubleshooting,

BLOG ↗ STABLE

Elastic Defend adds 800+ auto-generated vulnerable driver rules, ARM Windows support, and an AI troubleshooting skill in Agent Builder.

└──▷ **HOW TO FIND IT**

Verify the open, automatically generated vulnerable driver rules before applying them to your environment.

📍 Browse `https://github.com/elastic/protections-artifacts` to review auto-generated vulnerable driver YARA rules published by Elastic Security Labs Threat Command.

Continuously monitors VirusTotal, LOLDrivers catalog, and Microsoft's Vulnerable Driver Block List to automatically generate and instantly deploy YARA-based endpoint protection rules for newly disclosed vulnerable drivers — decoupled from any release cycle and published openly in the elastic/protections-artifacts repository.

Expands vulnerable driver detection coverage from 65 rules (2023) to 800+ known vulnerable drivers, growing automatically as new drivers are disclosed.

Adds Automatic Troubleshooting as a conversational skill in Elastic Agent Builder, enabling natural-language diagnosis of endpoint policy failures, third-party antivirus conflicts, and performance issues with specific remediation guidance.

Extends Elastic Defend support to Windows on ARM, covering Snapdragon laptops, Copilot+ PCs, and ARM workstations with the same detection rules and telemetry as x64 endpoints, enrollable under existing policies.

WAS THIS USEFUL? 👍 👎

Google SecOps

Feedback ⓘ changelog · Jul 31, 2026 CHANGELOG ↗

Google SecOps is a cloud-based security operations platform that helps organizations detect, investigate, and respond to cybersecurity threats using AI-powered analytics and threat intelligence.

Google SecOps adds prebuilt parser preview visibility and append/replace control for repeated fields in code snippet extensions.

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Create Monitors with Mobot (Search) (2026-07-31)

2 RELEASES · 2026-07-30 → 2026-07-31 BLOG ↗ LATEST STABLE

Sumo Logic is a cloud-native platform that collects, analyzes, and visualizes machine data and logs for security, observability, and operational intelligence.

Sumo Logic's Mobot can now create and update logs monitors from plain-language prompts, skipping the monitor form entirely.

└──▷ HOW TO FIND IT

When you want to spin up a logs monitor quickly without navigating the monitor form.

> 📍 In Mobot, type a plain-language prompt such as "Create a monitor that alerts when error rate exceeds 5% in the last 15 minutes for the prod environment", review the suggested configuration, then confirm to create it.

Enables creation and updating of logs monitors in Mobot via plain-language prompts, with a review-and-confirm workflow before saving.

▾ 1 more release in this issue · 2026-07-30 → 2026-07-31

Mobot Now Generally Available with a Unified Conversational Interface (Search) (2026-0

BLOG ↗ STABLE

Mobot, Sumo Logic's AI conversational assistant, is now GA with unified log analysis, content creation, and Cloud SIEM investigation in one interface.

└──▷ HOW TO FIND IT

Investigate a spike in failed logins across your logs without writing any query syntax — Mobot plans and runs the analysis and surfaces notable findings inline.

> 📍 In the Sumo Logic console, open Mobot and type: "Show me failed login attempts in the last 24 hours and identify any unusual source IPs."

Create a monitor through conversation so you can set up alerting without navigating the UI manually — Mobot confirms before anything goes live.

> 📍 In the Sumo Logic console, open Mobot and type: "Create a monitor that alerts me when error rate exceeds 5% over a 15-minute window."

Continue a Cloud SIEM investigation in Mobot to enrich an insight with log context without switching tools.

📍 In Cloud SIEM, open an insight and select `Continue in Mobot` to carry the full insight context into a Mobot conversation for deeper log analysis.

Mobot is now generally available as a unified conversational interface — ask plain-language questions across log investigation, platform how-tos, and Cloud SIEM without selecting an agent or writing query syntax.

Log Analysis Agent performs multi-step log analysis, returns inline results with notable findings, and recommends next steps automatically.

Enables creation and management of monitors and dashboards through conversation, with a confirmation step before any changes go live.

Conversational playbook management for Automation Service is rolling out within Mobot.

Cloud SIEM integration lets SOC analysts continue a SOC Analyst Agent investigation inside Mobot with full insight context.

`+1 more`

WAS THIS USEFUL? 👍 👎

Wazuh

Feedback ⓘ `v4.10.4` NOTES ↗

Wazuh - The Open Source Security Platform. Unified XDR and SIEM protection for endpoints and cloud workloads.

Wazuh v4.10.4 adds detection of suppressive Audit rules that could blind FIM whodata on Linux.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v4.10.4 https://github.com/wazuh/wazuh.git
# already have the repo? check out this version:
$ git checkout v4.10.4
```

Adds detection of the `-a never,task` Audit rule in FIM whodata for Linux, alerting when this rule would suppress file integrity monitoring events.

WAS THIS USEFUL? 👍 👎

CrowdStrike Falcon Next-Gen SIEM

Feedback ⓘ

undated-d8abab8d BLOG ↗

CrowdStrike Falcon Next-Gen SIEM collects, analyzes, and correlates security data to detect and respond to threats across enterprise environments.

Falcon Next-Gen SIEM gains IOA correlation, Claude Code session visibility, and AIDR findings integration for agentic AI security.

Adds Falcon Platform Indicators of Attack (IOAs) into Next-Gen SIEM to surface new threat detections correlated with existing telemetry.

Enables correlation of Falcon AIDR findings with EDR, identity, and network telemetry in Next-Gen SIEM for unified AI activity risk visibility.

Supports tying AI security events — including prompt interception and dangerous tool calls — back to specific Claude Code sessions and users within Next-Gen SIEM.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ Trusted Device Suppression Live in Managed ITDR

NOTES ↗

Huntress is a managed threat detection and response platform that helps MSPs and enterprises detect and respond to security threats.

Huntress Managed ITDR adds Trusted Device Suppression to cut unexpected-login escalations by 40–50%.

New Trusted Device Suppression in Managed ITDR automatically suppresses location and VPN escalations for logins from corporate-managed devices, configurable per account or per organization.

Supports granular trust-level selection — compliant, domain-joined, or registered — giving partners precise control over which device states qualify for suppression.

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ changelog · Jul 30, 2026

BLOG ↗

Microsoft Defender is an antivirus and security tool protecting Windows devices from malware, viruses, and cyber threats.

Microsoft Defender adds agentic SOC workflows, prompt injection email protection, and unified cloud-agent posture in July 2026.

Introduces Project Perception, a multi-agent autonomous defense system combining red, blue, and green team agents for continuous end-to-end security workflows.

New prompt injection protection (preview) identifies and isolates emails containing malicious AI instructions before delivery.

Unified posture and runtime protection for cloud agents across Microsoft Foundry, Microsoft Copilot Studio, and third-party agents in Microsoft Agent 365.

Microsoft Defender Threat Intelligence convergence and an enhanced Threat Intelligence Agent bring additional out-of-the-box intelligence and automation into the unified SecOps workflow.

New interconnected experiences between Defender and Microsoft Entra enable SOC teams to disable compromised identities directly via an RBAC mode enforcing least privilege.

+2 more

WAS THIS USEFUL? 👍 👎

Microsoft Security

Feedback ⓘ changelog · Jul 30, 2026

BLOG ↗

Microsoft Security provides comprehensive cybersecurity solutions, threat protection, and security monitoring tools for organizations and enterprises.

Microsoft Intune Suite advanced endpoint management now included in Microsoft 365 E5 (and select E3) at no added cost.

Includes Microsoft Intune Suite capabilities in Microsoft 365 E5 (and select capabilities in Microsoft 365 E3) at no additional cost as of July 1, 2026.

Enables reduction of standing admin rights, modernized certificate management, accelerated issue resolution, and streamlined app delivery through advanced endpoint management.

Integrates Microsoft Security Copilot into Intune workflows for AI-assisted endpoint management.

WAS THIS USEFUL? 👍 👎

CrowdStrike

Feedback ⓘ undated-3cd143b1 BLOG ↗

CrowdStrike is a cybersecurity platform that provides cloud-native endpoint protection, threat intelligence, and real-time detection to defend organizations against cyberattacks and breaches.

Falcon AIDR extends AI visibility and control to Microsoft Copilot Studio, Claude Code, and the Falcon browser extension.

└──▷ HOW TO FIND IT

Manage Falcon AIDR browser extension coverage across your fleet without a separate workflow by aligning policy to host groups you already use.

> 📍 In the Falcon console, navigate to Falcon AIDR policy settings, assign AIDR policies to existing host groups, and enable the Falcon browser extension to cover browser-based AI activity.

Integrates Falcon AIDR with Microsoft Copilot Studio as an external threat detection provider, checking tool names and input parameters against policy and returning allow/block decisions before an agent executes a tool call.

Connects Falcon AIDR to Claude Code's hook event system to inspect and block prompts and tool activity in real time, configured via a JSON block in the Claude Code settings file with no agent install required.

Adds Falcon AIDR capability to the Falcon browser extension, enabling monitoring and control of browser-based AI activity managed through the Falcon console with policy assignment aligned to existing host groups.

Feeds Copilot Studio, Claude Code, and browser extension AI activity data into Falcon Next-Gen SIEM to support agentic SOC workflows.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ MCP support for Apps (2026-07-31) BLOG ↗

Tines is a no-code automation platform that helps security and operations teams build workflows to automate repetitive tasks and incident response processes.

Tines MCP server now supports Apps, enabling clients to modify, preview, and publish Apps programmatically.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Node v6.9.0-rc.1: Native Apache Iceberg output (2026-07-31) BLOG ↗

PRE-RELEASE

Tenzir is an open-source security data pipeline platform that helps teams collect, route, shape, and analyze large-scale security telemetry efficiently.

Tenzir v6.9.0-rc.1 adds native Apache Iceberg table output via REST catalog with auto-evolving schemas.

New `to_iceberg` operator writes events directly into Apache Iceberg tables through a REST catalog, automatically creating and evolving table schemas as heterogeneous streams arrive.

Optionally validate batches read from persisted stores for data integrity checks.

Faster, lower-memory catalog loading at node startup for large catalogs.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

ConductorOne

Feedback ⓘ changelog · Jul 31, 2026 CHANGELOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance across enterprise applications and infrastructure.

ConductorOne adds AI agent policy steps, group/role access reviews, per-user MCP auth, and scope-limited API keys in its July 2026 release.

└──▷ HOW TO FIND IT

Restrict a new personal API key to read-only access so it can be shared with a reporting script without granting full account permissions.

> 📍 In the console, go to your personal API key creation screen and use the Limit Scopes selector to select a read-only role before saving the key.

Require each user to authenticate with their own credential when connecting to an external MCP server instead of sharing an admin credential.

> 📍 In the console, open the MCP server's Configure server settings panel, enable per-user authentication, and select the credential type (OAuth2, bearer token, custom header, or basic auth).

Assigns AI agent (C1 AI) to workflow steps in request or review policies, with custom instructions, audit-log attribution to a C1 System identity, and reviewable evaluation conversations.

Extends request catalog search to surface an `Entitlements matching` section — showing entitlements that indirectly grant access to the searched resource — plus a `View access graph` link showing the full grant chain.

Adds a `Grant types` selector to access review campaigns so reviews can now target grants held by groups or roles, not just users (options: Users only, Resources only, Users and resources).

Enables custom recurrence intervals (e.g., every two days, every three months) and configurable campaign durations on campaign templates.

Renders IAM policy documents as a collapsible JSON tree via a `View policy` button on resource-level access reviews and task details, replacing raw unformatted JSON.

+10 more

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ changelog · Jul 31, 2026

CHANGELOG ↗

Nudge Security discovers and secures all SaaS apps used across an organization, helping IT teams manage shadow IT and enforce security policies through automated nudges.

Nudge Security adds Amazon Quick Suite AI monitoring, agent discovery, and enforced SSO login restrictions.

└──▷ HOW TO FIND IT

Enforce SSO-only access to Nudge Security so all logins inherit your IdP's MFA and conditional access policies.

> 📍 In the console, go to Settings › Organization and enable login restrictions.

Turn on AI conversation monitoring to flag sensitive data entered into Amazon Quick Suite chats.

> 📍 In the console, go to Browser Extension Settings and configure AI conversation monitoring detection.

Enable browser-based agent discovery to populate your AI agents inventory with Amazon Quick agents your team builds.

> 📍 In the console, go to Browser Extension Settings and enable browser-based agent discovery, then view discovered agents under Identities › AI Agents.

Detects sensitive data (secrets, PII, financial, health) shared in Amazon Quick Suite AI conversations via browser extension monitoring.

Discovers AI agents built on Amazon Quick and surfaces them in the AI agents inventory under Identities, with creator, components, and risk insights.

Enables organizations to enforce SSO-only login, disabling Google and Microsoft sign-in methods so all authentication routes through the configured IdP (e.g., Okta) with its MFA and conditional access controls.

WAS THIS USEFUL? 👍 👎

Material Security Feedback ⓘ

Map Sensitive File Sharing in Drive and Simplify Customization in Material

NOTES ↗

Material Security protects organizations from email-based threats like phishing and account takeovers through AI-powered security solutions.

Material Security adds a Drive sensitive file sharing dashboard, AI-powered file search, and a unified custom detections modal.

└──▷ HOW TO FIND IT

Use the Drive sensitive file sharing dashboard to quickly check whether source code or confidential M&A documents are being shared outside the company.

> 📍 In the Material console, navigate to the Drive sensitive file sharing dashboard, select your content category (e.g., Source Code or Confidential), and review files shared publicly or externally.

During an employee offboarding, use AI Prompt File Search to locate specific sensitive files the departing employee may still have access to.

> 📍 In the Material console, open AI Prompt File Search and enter a plain-English description such as "spreadsheet with social media logins shared with [employee name]" to retrieve matching files for remediation.

Rapidly deploy a new custom detection for a targeted threat without navigating multiple screens.

> 📍 In the Material console, go to Custom Detections, click New Detection, and use the single modal to define detection criteria and response actions across email and files, then save and deploy.

New Drive sensitive file sharing dashboard maps how sensitive files are shared publicly, externally, company-wide, and with trusted external parties, with drill-down by content category.

New AI Prompt File Search lets users describe files in plain English to locate sensitive documents in Drive without building structured queries.

New Quick File Search provides structured fields to filter Drive files across public, external, and trusted external shares.

Redesigned custom detections UX consolidates creation and editing of detections and responses for email and files into a single modal.

WAS THIS USEFUL?  

Auth0

Feedback ⓘ Global Search in Cmd+K is now in Beta (2026-07-31)

5 RELEASES · 2026-07-30 → 2026-07-31 BLOG ↗ LATEST STABLE

Auth0 is a cloud-based identity platform that provides authentication and authorization services, enabling developers to securely add login and access control to applications.

Auth0 adds Global Search to the Command Palette (Cmd+K), enabling real-time tenant-wide entity lookup in Beta.

└──▷ HOW TO FIND IT

Quickly locate a user by email or ID during an incident without leaving your current page.

> 📍 Press Cmd+K, type the user's email or client ID, and select the matching result to navigate directly to that entity or copy its identifier.

New Global Search in Command Palette (Cmd+K) lets you search Applications, APIs, Organizations, and Users by name, ID, email, or identifier in real time.

One-click identifier copying from search results lets you instantly grab entity IDs without navigating to detail pages.

View more links from search results navigate directly to filtered entity lists for deeper exploration.

▾ 4 more releases in this issue · 2026-07-30 → 2026-07-31

> Agents as Principal now in Early Access! (2026-07-30) BLOG ↗ STABLE
>
> Auth0 launches Agents as Principal in Early Access, giving AI agents their own identity, credentials, and audit trail.
>
> Introduces Agents as Principal: every AI agent gets a unique identity distinct from human users and shared clients, with its own credentials and audit trail.
>
> Enables end-to-end delegation tracking: when an agent acts on a user's behalf, the full delegation chain (who delegated, to which agent, through which app) is encoded in the token and auditable at every hop.
>
> Supports associating existing external agent identifiers (from your own registry, an MCP server, or elsewhere) with an Auth0 client, avoiding re-provisioning of already-built agents.

> Curated Blocklists in Tenant Access Control Lists is Now GA! (2026-07-30) BLOG ↗ STABLE
>
> Auth0 Curated Blocklists for Tenant ACLs are now GA, enabling automated IP threat blocking via Auth0-managed threat intelligence.

└──▷ **TRY IT**

Block all TOR exit node traffic at the authentication scope to prevent anonymous abuse via TOR.

```
$ curl -X POST https://<tenant>.auth0.com/api/v2/network-acls \
  -H 'Authorization: Bearer <mgmt_api_token>' \
  -H 'Content-Type: application/json' \
  -d '{
    "description": "Block TOR exit nodes",
    "active": true,
    "priority": 2,
    "rule": {
      "action": { "block": true },
      "match": { "auth0_managed": ["auth0.tor"] },
      "scope": "authentication"
    }
  }'
```

Adds Curated Blocklists to Tenant Access Control Lists (ACLs), letting Attack Protection customers delegate IP risk management to Auth0 Threat Intel instead of manually maintaining IP/ASN lists.

Supports five Auth0-managed threat categories: `auth0.low_reputation`, `auth0.tor`, `auth0.proxy`, `auth0.vpn`, and `auth0.icloud_relay_proxy` for precise traffic-origin blocking.

Introduces the `auth0_managed` array within the Network ACLs Management API `match` object to reference dynamically updated curated lists in ACL rule definitions.

Enterprise Connect - Beta (2026-07-30) BLOG ↗ STABLE

Auth0 Enterprise Connect (Beta) adds modular B2B identity federation with SSO, SCIM provisioning, and self-service onboarding.

Federate existing SAML or OIDC authorization servers (including Okta) to Auth0 as a modular B2B identity layer.

Layer enterprise SSO and self-service onboarding/setup flows on top of federated connections for B2B customers.

Provision users outbound via SCIM wired to federated connections.

Inject user profile claims into ID tokens using Actions code during federated authentication.

Subscribe to connection lifecycle events (`connection.created`, `connection.updated`, `connection.deleted`) to keep local domain maps in sync as customers onboard.

Token Vault Privileged Worker is now Early Access! (2026-07-30) BLOG ↗

STABLE

Auth0 Token Vault Privileged Worker (Early Access) lets background agents fetch user tokens without an active session.

Adds Token Vault Privileged Worker, enabling autonomous/background agents to retrieve a user's third-party tokens (Gmail, Drive, Slack, etc.) from Token Vault without requiring an active user session.

Supports Private Key JWT and mTLS authentication for trusted worker identities authorizing token exchanges.

Enables pinning each worker credential to specific connections and scopes to limit blast radius of a compromised credential.

WAS THIS USEFUL? 👍 👎

BeyondTrust Feedback ⓘ

iOS Representative Console 3.1.11 release notes (2026-07-30) BLOG ↗

BeyondTrust is a privileged access management platform that secures, controls, and monitors privileged accounts, credentials, and remote access across enterprise environments.

BeyondTrust iOS Representative Console 3.1.11 adds audio support during Jump Client sessions.

WAS THIS USEFUL? 👍 👎

Microsoft Entra

Feedback ⓘ changelog · Jul 30, 2026 BLOG ↗

Microsoft Entra is a cloud-based identity and access management platform that secures user authentication and authorization across applications and resources.

Microsoft Entra adds tenant governance, passkey-default auth, Defender SOC integration, and Purview network-layer data protection.

New Tenant Governance capability enables centralized discovery, management, and governance of tenants with cross-tenant delegated administration.

Passkeys are now the default authentication experience in Microsoft Entra ID, reducing reliance on SMS/voice and improving phishing-resistant security ahead of Microsoft telecom delivery retirement in 2027.

New interconnected Entra and Microsoft Defender experiences unify identity and SOC teams with shared user experience, RBAC, and agentic workflows.

Microsoft Purview integrates with Microsoft Entra Internet Access to extend data security controls to the network layer, enabling real-time protection of sensitive data shared with unmanaged cloud and AI apps.

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Weekly release 2026.07.3 (2026-07-29) BLOG ↗

Okta Identity Engine is a platform that manages user authentication and identity verification through flexible, secure API-driven workflows.

Okta Identity Engine adds editable issuer URLs for AI agent resource connections and resilient AI agent import.

Supports editing the issuer URL for AI agent resource connections.

AI agent import now skips failed entries instead of halting, improving resilience during bulk operations.

└──▷ **BREAKING ON UPGRADE**

! The Application Cross App Access Connections API has been removed.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

New Capability: US Region Deployment for SAAM (2026-07-27) BLOG ↗

SailPoint is an identity security platform that helps organizations manage and govern user access to applications, data, and systems across the enterprise.

SailPoint SAAM now available in AWS US East (us-east-1), enabling US customers in-region data residency and full app discovery-to-governance workflows.

└──▷ **HOW TO FIND IT**

Set up the Browser Extension connector in ISC to surface US-region discovered apps with risk and exposure data on the Discovered Applications page.

> 📍 In ISC, go to Admin › Connections › Discovery Connectors → Select Create Connector → choose Browser Extension category → Continue → Configure → Start Connector Setup → Run Discover to populate Admin › Connections › Sources › Discovered Applications.

Access the full Application Visibility portal via the App Switcher to connect an IdP and deploy the browser extension for app inventory and risk analytics.

> 📍 In ISC, select the SailPoint Solutions Center icon (App Switcher) → Application Visibility tile → Connect Identity Provider (Microsoft Entra ID, Okta, or Other IdP) → Deploy browser extension via MDM (Intune / Jamf / Workspace ONE).

Adds AWS US East (us-east-1) region hosting for SAAM, giving US customers in-region data residency and lower latency.

Enables the App Switcher Application Visibility portal in the US region for full SAAM configuration, discovery, risk & usage analytics, and privacy controls.

Supports Browser Extension (BE) connector in ISC for US-region customers, syncing SAAM discovery and risk data into ISC's Discovered Applications page with risk score, sensitivity, exposure, and credential/MFA signal columns.

Provides a direct discovery-to-governance path within ISC, enabling access reviews, certifications, and app onboarding without leaving the platform.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

DigiCert CertCentral

Feedback ⓘ changelog · Jul 31, 2026

CHANGELOG ↗

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizational infrastructure.

CertCentral adds DCV polling status to APIs, ACME URL delegation/transfer, API key transfer, and broader automated delegation approvals.

└──▷ HOW TO FIND IT

Transfer an API key tied to a departing user to a Service User before deleting the account, keeping integrations uninterrupted.

> 📍 In CertCentral, go to the API Keys page, find the API key and select its link. In the Update API Key window, use the owner menu to select the target Service User, then click Update API Key.

Create an ACME Directory URL for a new Service User so ACME credentials are not tied to an individual employee account.

> 📍 In CertCentral, go to the ACME Directory URLs page, select Add ACME Directory URL, enter a Name for the URL, choose a product from the product menu, add the new Service User's details, and finish configuring the URL.

Enables Administrators to transfer API key ownership to another CertCentral user or Service User, keeping transferred keys valid until revoked or deactivated.

Simplifies delegation request workflow by removing the organization selection step — a single delegation request now applies to all organizations the approver represents across all authorized certificate products.

Extends automated delegation approval support in CertCentral Europe to a broader set of PKIo and EU Qualified certificate products.

WAS THIS USEFUL? 👍 👎

Varonis

Feedback ⓘ Block unauthorized tenant access

7 RELEASES · SEEN 2026-07-31 NOTES ↗ LATEST STABLE

Varonis is a data security platform that detects and stops insider threats, data breaches, and ransomware attacks across cloud and on-premises environments.

Varonis SaaS tenants can now restrict UI logins and API requests to approved IP addresses only.

▾ **6 more releases in this issue** · seen 2026-07-31

> Discover open shares across CIFS environments NOTES ↗ STABLE
>
> Varonis adds Automatic Share Detection to find open shares across monitored and unmonitored CIFS environments.

> Enhanced alert routing and MDDR escalations NOTES ↗ STABLE
>
> Varonis adds segment information to Syslog messages, event publishing workflows, and MDDR escalation emails for richer alert routing.
>
> Includes segment information in Syslog messages to improve alert routing and correlation in downstream SIEM and SOAR tools.
>
> Adds segment context to event publishing workflows for richer automation across security operations pipelines.
>
> Surfaces segment data in MDDR escalation emails to provide analysts with clearer origin context during incident response.

> Export Events and Audit data at scale NOTES ↗ STABLE
>
> Varonis adds Export API to extract events and audit logs at scale for reporting, compliance, and SIEM ingestion.
>
> New Export API enables bulk extraction of events and audit logs as downloadable exports for compliance reporting and long-term analysis.

> Gain more control over label lifecycles NOTES ↗ STABLE
>
> Varonis admins can now configure automatic removal or retention of labels when triggering conditions no longer apply.
>
> New label lifecycle controls let administrators choose to automatically remove Varonis-applied labels when their triggering conditions are no longer met, or preserve them for compliance continuity.

Preserve higher priority Microsoft Purview labels automatically

NOTES ↗

STABLE

Varonis now automatically preserves higher-priority Microsoft Purview labels to prevent accidental downgrades.

Adds automatic protection against Microsoft Purview label downgrades: Varonis now compares an existing label to its recommended label and only applies a change when it results in stronger protection, honoring manually applied or third-party labels.

Secure Multi-Tenant VAST Environments with Greater Visibility

NOTES ↗

STABLE

Varonis adds tenant-level visibility and monitoring for multi-tenant VAST Data environments.

Supports tenant-level visibility and monitoring for multi-tenant VAST Data environments, including tenant-specific activity, access patterns, and events.

Enables consistent security controls across shared VAST infrastructure while maintaining strong tenant separation.

WAS THIS USEFUL? 👍 👎

Microsoft Purview

Feedback ⓘ changelog · Jul 30, 2026

BLOG ↗

Microsoft Purview is a unified data governance platform that discovers, manages, and protects data across cloud and on-premises environments.

Microsoft Purview gains network-layer DLP for SaaS/AI apps, Copilot email controls, and a unified Insider Risk alert triage experience.

└──▷ HOW TO FIND IT

Block uploads of sensitive data to shadow AI apps at the network layer without deploying third-party proxies.

> 📍 In the Microsoft Purview portal, go to Data Loss Prevention › Policies › Create policy, select the `Microsoft Entra Internet Access` scope, define a condition for sensitive info types, and set the action to Block for unmanaged cloud and AI app destinations.

Prevent Copilot from using external emails as grounding data to keep AI responses anchored to vetted internal content.

> 📍 In the Microsoft Purview portal, go to Data Loss Prevention › Policies › Create policy, choose `Microsoft 365 Copilot` as the location, add a condition for `Sender is external`, and set the action to exclude matched emails from Copilot grounding.

Integrates with Microsoft Entra Internet Access to extend data loss prevention to the network layer, detecting and blocking sensitive data shared with unmanaged cloud and AI apps in real time.

New DLP policy for Microsoft 365 Copilot (preview) lets admins exclude external-sender emails from being referenced, summarized, or used as grounding data in Copilot responses.

New centralized Insider Risk Management alert experience consolidates agentic alert triage, enriched user details, case notes, and combined classic/agent filtering into a single view.

Data Security Triage Agent's advanced AI reasoning layer — performing multi-step analysis across user, device, and data activity signals — is now generally available.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Zscaler Internet Access

Feedback ⓘ

changelog · Jul 30, 2026 BLOG ↗

Zscaler Internet Access is a cloud-based security platform that inspects and controls all internet traffic to prevent threats and enforce policies.

ZIA adds DLP log filters, time-bound file type rules, and per-channel DLP engine scoping in the July 2026 GovCloud release.

└──▷ HOW TO FIND IT

Scope a DLP engine to only the Endpoint DLP channel to reduce noise from irrelevant matches in other channels.

> 📍 In the ZIA console, go to Policies › Data Protection › Common Resources › DLP Dictionaries & Engines › DLP Engines, then create, edit, or clone a DLP engine and select the specific channels (e.g., Endpoint DLP, DSPM) under the channel selection option.

Set a temporary exception for a file type control rule that automatically expires after an operational window.

> 📍 In the ZIA console, navigate to your File Type Control policy rule, open the rule settings, and set an expiration date under the rule configuration to auto-expire the exception at the end of the allowed window.

Adds DLP Policy Action and DLP Confirmation Message filter columns to Web Insights Logs and Endpoint DLP Insights Logs for improved audit and incident investigation.

Enables expiration dates on File Type Control policy rules for time-limited enforcement windows and temporary exceptions.

Allows custom and predefined DLP engines to be scoped to specific channels (e.g., Endpoint DLP, DSPM), reducing false-positive noise from broad engine application.

Risk360 is now included in most ZIA and platform SKUs for organizations with more than 500 users.

WAS THIS USEFUL? 👍 👎

Zscaler Private Access

Feedback ⓘ

changelog · Jul 30, 2026 BLOG ↗

Zscaler Private Access provides zero-trust network access to internal applications without exposing them to the public internet.

ZPA adds multi-Private Cloud support for LSS-based log streaming in Business Continuity scenarios.

Extends LSS-based Log Streaming in Business Continuity to stream logs across multiple Private Clouds, maintaining observability and compliance logging during failover or multi-cloud operations.

Releases Manager software version 26.54.4 with updated App Connector and Private Service Edge RPM packages for RHEL 8.x and 9.x, and a new Private Cloud Controller RPM package for RHEL 9.x, downloadable from the Zscaler repository.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Mondoo

Feedback ⓘ Analytics Brings Self-Serve Reporting to the Console

NOTES ↗

Mondoo is a security platform that continuously scans and assesses infrastructure, cloud environments, and software supply chains for vulnerabilities and misconfigurations.

Mondoo Analytics is GA: self-serve security reporting with 7 built-in reports, pivots, custom reports, and PNG/PowerPoint export.

└──▷ HOW TO FIND IT

Share a full security posture report with leadership by exporting it as a PowerPoint deck directly from the Analytics module.

📍 In the Mondoo console, go to Analytics, open a built-in or custom report, then select Export › Download as PowerPoint.

New Analytics module is generally available in the Mondoo console, offering self-serve exploration of live security data.

Includes seven built-in reports covering key security metrics out of the box.

Supports configurable pivots to slice and dice security data across different dimensions.

Enables creation of custom reports tailored to specific organizational needs.

Export individual report pages as PNG images or full reports as PowerPoint decks for stakeholder sharing.

WAS THIS USEFUL? 👍 👎

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.37.69 (2026-07-31) BLOG ↗

RegScale is a compliance operations platform that automates and continuously monitors regulatory compliance, security controls, and risk management for organizations.

RegScale CLI 6.37.69 adds asset-free Qualys/Prisma sync and Tanium ports/protocols integration.

Adds option to sync Qualys and Prisma vulnerabilities and issues without creating or updating assets.

Adds Tanium open ports and protocols synchronization to assets and the system Ports, Protocols, and Services table.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

OpenAI Codex CLI

Feedback ⓘ rust-v0.147.0-alpha.3 NOTES ↗

CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI alpha.3 adds parallel MCP tool execution, cloud-managed MCP server loading, forked chat naming, and subagent direct-input reporting.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.147.0-alpha.3
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.147.0-alpha.3
```

Allows unrelated tools to run concurrently while MCP servers are still starting up, eliminating gate-blocking delays.

Loads cloud-managed MCP servers in MCP CLI commands (`codex mcp list`, `codex mcp get`).

Enables naming forked chats directly from the TUI.

Exposes MCP read-only hints in tool call items.

Denies network access when an allow amendment fails, tightening sandbox network policy.

+2 more

└──▷ BREAKING ON UPGRADE

! The `--full-auto` flag is removed from `codex exec`; legacy handling has been dropped.

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ cli-v3.0.48 2 RELEASES · 2026-07-31 NOTES ↗

CODE ↗ LATEST STABLE

Autonomous coding agent as an SDK, IDE extension, or CLI assistant.

Cline CLI v3.0.48 brings in-TUI history management, resilient connector threads, and per-server MCP timeouts.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch cli-v3.0.48 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout cli-v3.0.48
```

Adds `cline history` inside the existing TUI with resume and delete actions, replacing the separate secondary view.

Connector threads (Slack, Discord, Telegram, Linear, Google Chat, WhatsApp) now auto-recover when their bound session is gone, replaying the turn against a new session instead of requiring manual edits to `threads.json`.

Honors per-server `timeout` in `cline_mcp_settings.json` for `initialize`, `tools/list`, and `tools/call`, removing the hardcoded 5-second limit for slow MCP servers.

OpenRouter now defaults to `anthropic/claude-sonnet-5`.

Plugins can now emit telemetry through `ctx.telemetry`.

▾ 1 more release in this issue · 2026-07-31

sdk/sdk/v0.0.67 NOTES ↗ STABLE

Cline SDK v0.0.67 adds session forking, live transcript reads, plugin telemetry, and configurable MCP timeouts.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch sdk/sdk/v0.0.67
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout sdk/sdk/v0.0.67
```

└──▷ USE IT

Read the live in-memory transcript of a running session so a plan/act rebuild has full context rather than starting from an empty history.

```typescript
const messages = await clineCore.readLiveMessages();
```

Set a longer MCP server timeout for slow-starting servers instead of hitting the old hardcoded 5 s limit.

```json
// cline_mcp_settings.json
{
  "mcpServers": {
    "my-slow-server": {
      "command": "npx",
      "args": ["my-mcp-server"],
      "timeout": 120
    }
  }
}
```

Adds session forking and user-run message APIs so a host can edit an earlier prompt: fork before a selected user run, trim checkpoint history, and restore prior messages.

Adds `ClineCore.readLiveMessages` for reading a resident session's in-memory transcript, enabling plan/act rebuilds during in-flight turns to start from full history.

Enables plugins to emit telemetry through `ctx.telemetry` from both the subprocess sandbox and in-process execution.

Honors the per-server `timeout` in `cline_mcp_settings.json` for MCP `initialize`, `tools/list`, and `tools/call` calls, defaulting to 60 seconds (clamped 1–3600 s) instead of hardcoded limits.

Reasoning controls (effort, budget, on/off) are now driven by the models.dev catalog and normalized before provider encoding, with Anthropic's mandatory/impossible thinking modes handled explicitly and out-of-range budgets clamped.

+1 more

└─▷ **BREAKING ON UPGRADE**

! The Enable R1 messages format option has been removed from the OpenAI Compatible provider.

WAS THIS USEFUL? 👍 👎

OpenHands

Feedback ⓘ v1.8.0 NOTES ↗

OpenHands: AI-Driven Development

OpenHands v1.8.0 lets you toggle installed MCP servers on or off directly from their cards.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.8.0 https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout v1.8.0
```

Adds enable/disable toggle for installed MCP servers directly from the server card in the UI.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

LangGraph

Feedback ⓘ checkpointpostgres==3.1.1 NOTES ↗

Build resilient agents.

LangGraph checkpoint-postgres 3.1.1 adds opt-in omit_expired to skip expired checkpoint rows on read.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch checkpointpostgres==3.1.1
https://github.com/langchain-ai/langgraph.git
# already have the repo? check out this version:
$ git checkout checkpointpostgres==3.1.1
```

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Daytona

Feedback ⓘ Org members command and client-side HTTP timeout NOTES ↗

Daytona is a self-hosted development environment manager that automates provisioning, configuration, and management of standardized, secure dev environments across any infrastructure.

Daytona adds an org members command and client-side HTTP timeout support.

Adds an org members command for managing organization membership.

Introduces client-side HTTP timeout configuration.

WAS THIS USEFUL? 👍 👎

Promptfoo

Feedback ⓘ snapshot-20260731 NOTES ↗

Promptfoo is an open-source tool for testing, evaluating, and red-teaming LLM applications to detect vulnerabilities, improve reliability, and compare model outputs.

Promptfoo January 2026 adds adaptive rate limiting, Transformers.js local inference, telecom red team plugins, and video generation providers.

└──▷ **TRY IT**

View Promptfoo logs directly from the CLI after a long evaluation run to inspect rate limiting or provider errors.

```
$ promptfoo logs
```

Load multiple environment files (e.g., base config + environment-specific secrets) when running an evaluation.

```
$ promptfoo eval --env-file .env.base --env-file .env.prod
```

Adds Transformers.js provider to run Hugging Face models locally in Node.js or browser without external API calls.

Adds Vercel AI Gateway and Cloudflare AI Gateway providers for routing requests through managed gateways.

Adds AWS Bedrock Video (Nova Reel, Luma Ray 2) and Azure AI Foundry Video (Sora) providers for video generation.

Adds native session endpoint support to the HTTP provider for stateful multi-turn conversations.

Adds `apiBaseUrl`, `websocketUrl` override, and function call support to the xAI Voice provider.

+18 more

WAS THIS USEFUL? 👍 👎

NoLabs nono

Feedback ⓘ v0.71.0 NOTES ↗ CODE ↗

Sandbox any AI agent in seconds - zero setup, zero latency.

nono v0.71.0 adds a new `platform` subcommand for enrollment and audit delivery to a control plane.

└─▷ GET THIS VERSION ▾

```
$ git clone --branch v0.71.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.71.0
```

└─▷ BREAKING ON UPGRADE

! The `nono learn` command has been removed; it was deprecated since v0.50.1.

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ v2.5.1 NOTES ↗

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, and security risks.

SkillSpector v2.5.1 adds `SKILLSPECTOR_MAX_LLM_CONCURRENCY` to throttle async LLM batch requests for rate-limited providers.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.5.1 https://github.com/NVIDIA/SkillSpector.git
# already have the repo? check out this version:
$ git checkout v2.5.1
```

└──▷ **TRY IT**

Serialize LLM analyzer requests to avoid bursting a rate-limited provider such as Anthropic.

```
$ export SKILLSPECTOR_MAX_LLM_CONCURRENCY=1
export SKILLSPECTOR_PROVIDER=anthropic
export ANTHROPIC_API_KEY=sk-ant-...
skillspector scan ./my-skill/
```

Adds `SKILLSPECTOR_MAX_LLM_CONCURRENCY` environment variable to control the default concurrency of asynchronous LLM analyzer batches (default 10; values below 1 clamp to 1; blank or invalid values retain the default).

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

GitHub CLI

Feedback ⓘ v2.97.0 NOTES ↗

GitHub's official command line tool

gh project commands now support field/option names, plus Grok skill host and Devin agent added.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.97.0 https://github.com/cli/cli.git
# already have the repo? check out this version:
$ git checkout v2.97.0
```

└──▷ TRY IT

Update a project item's Status field by name instead of requiring numeric field/option IDs.

```
$ gh project item-edit 1 --owner monalisa --url <url> --field "Status"
--value "In Progress"
```

List project items with Status and Priority shown as named columns for quick triage.

```
$ gh project item-list 1 --owner "@me" --field "Status" --field
"Priority"
```

Adds name-based field and single-select option resolution to `gh project item-edit`, so you can set fields like `Status` to `In Progress` without knowing internal IDs.

Adds `--field` flag to `gh project item-list` to display named project fields as extra columns in output.

Adds Grok skill host support to `gh skill` agent integrations.

Replaces Windsurf with Devin in the `gh skill` agent roster.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Jul 31, 2026 CHANGELOG ↗

Cloudflare is a cloud-based platform that provides website performance optimization, DDoS protection, CDN services, and application security to safeguard and accelerate internet traffic.

Cloudflare adds Stream key rotation, AI Search SDK integrations, Browser Run human handoff, Gateway DNS TTL cap, and more.

└──▷ TRY IT

Rotate compromised broadcast credentials for a Stream live input without changing its identifier.

```
$ curl --request POST
'https://api.cloudflare.com/client/v4/accounts/{account_id}/stream/liv
e_inputs/{live_input_identifier}/rotate_keys' --header 'Authorization:
Bearer <API_TOKEN>'
```

Profile a Worker's cold-start CPU usage and bundle size before deploying, to catch heavy dependencies early.

```
$ wrangler check startup
```

Add grounded AI Search retrieval to a LangChain RAG pipeline using hybrid search against an existing knowledge base.

```python
from langchain_cloudflare import CloudflareAISearchRetriever

retriever = CloudflareAISearchRetriever(
    account_id="ACCOUNT_ID",
    api_token="API_TOKEN",
    instance_name="knowledge-base",
    retrieval_type="hybrid",
)
results = retriever.invoke("How do I configure Workers AI?")
```

Adds a `rotate_keys` API endpoint for Stream live inputs, revoking old broadcast credentials and returning fresh ones without changing the live input identifier; response now includes a `keysRotatedAt` field.

Adds `wrangler check startup` bundle size reporting and CPU startup profiling — raw/compressed sizes and sampled/active/GC/idle time summary printed directly in the terminal.

Adds AI Search integrations for Vercel AI SDK (`ai-search-provider` package), LangChain (`CloudflareAISearchRetriever` in `langchain-cloudflare`), and a new Cloudflare Agents SDK guide for stateful chat agents backed by AI Search.

Sets Node.js 24.18.0 as the default runtime for Workers Builds; Node.js 22.23.2 also preinstalled.

Cloudflare MCP servers now support the MCP 2026-07-28 Specification with stateless per-request servers; `/mcp` endpoint accepts stateless requests from 2025 Streamable HTTP clients.

+3 more

└──▷ **BREAKING ON UPGRADE**

- ! Historical `/sse` URLs on Cloudflare MCP servers no longer serve the deprecated HTTP+SSE transport — they now alias the Streamable HTTP handler; clients that force SSE transport must switch to Streamable HTTP or automatic transport detection.
- ! DoH JSON API (`application/dns-json`) DNSSEC records now emit numeric algorithm identifiers (e.g., `8` instead of `RSASHA256` , `13` instead of `ECDSAP256SHA256` , `15` instead of `ED25519`) and numeric digest types (e.g., `2` instead of `SHA-256`) — parsers expecting mnemonic strings will break.
- ! DoH JSON API (`application/dns-json`) `data` field format changes for CAA, NAPTR, RP, IPSECKEY, SVCB, HTTPS, TLSA, SSHFP, and OPENPGPKEY records: previously returned RFC 3597 generic hex (`\# <length> <hex>`), now returns standard presentation format — consumers parsing hex encoding will break.
- ! DoH JSON API (`application/dns-json`) HINFO `data` field now individually quotes each character-string (e.g., `"\"Intel Xeon\" \"Linux\""` instead of `"Intel Xeon Linux"`) — parsers expecting a single unquoted string will break.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 31, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue