CSAI Foundation | Cloud Security Alliance

The CVE Program Adds Its First AI-Native CNA

2026-08-08

AI-assisted Rapid Research

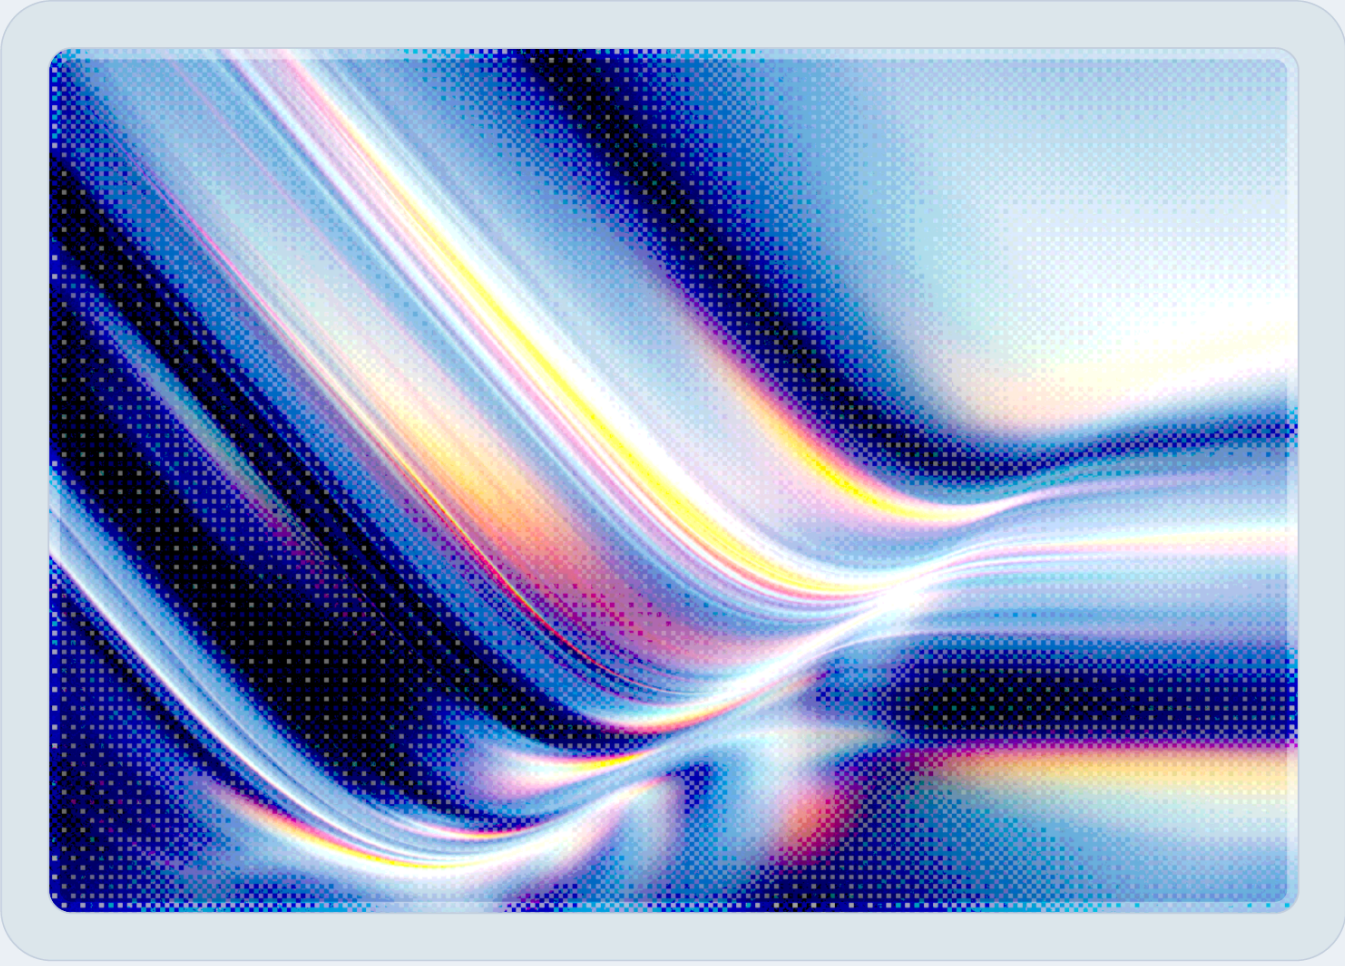

CSAI

CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

On July 22, 2026, AI-native vulnerability lifecycle company AISLE became a CVE Numbering Authority (CNA) under the CVE Program, authorized to assign its own CVE identifiers for flaws found in its own products [1][2]. Six days later, on July 28, Anthropic received the same designation as the CVE Program's 282nd U.S.-registered CNA, scoped to "software, services, and open-source projects they develop, maintain, or distribute" [3]. AISLE's designation sits under ENISA's CVE Root rather than the traditional CISA/MITRE root, one of 20 CNAs now recruited, trained, and overseen by the European Union Agency for Cybersecurity since it became a full Root in November 2025 [1]. Both designations are narrowly scoped to each company's own software; neither grants AISLE or Anthropic blanket authority to mint CVE identifiers for the third-party open-source vulnerabilities their AI models discover elsewhere, a distinction that matters because both companies have been prolific external bug-hunters. AISLE has disclosed hundreds of vulnerabilities in OpenSSL, Linux, Apache, and OpenEMR since emerging from stealth [2], while Anthropic's Project Glasswing reported more than 23,000 potential vulnerabilities across the first half of 2026, of which only 126 became published CVE records and 97 were confirmed patched — a conversion rate Forkast News calculates at roughly 6 percent, though the outlet does not specify the denominator behind that figure [3]. CISA officials have publicly acknowledged this asymmetry and confirmed a separate "AI researcher" CNA pilot allowing selected AI companies to assign identifiers for vulnerabilities their research uncovers, signaling that the AISLE and Anthropic designations are the leading edge of a broader structural change rather than isolated administrative actions [8].

Background

The CVE Program has spent the past eighteen months absorbing two simultaneous shocks: a funding and governance crisis centered on its historical dependence on a single CISA-funded MITRE contract, and a discovery-volume crisis driven by AI systems that find vulnerabilities far faster than the human-staffed pipeline that turns a report into a numbered, enriched CVE record can process them. The governance crisis peaked in 2025 when CISA's contract with MITRE approached expiration with no confirmed renewal, prompting an emergency extension and the formation of the independent CVE Foundation to explore governance models less dependent on a single national funder [6]. By early 2026 the CVE board was told there would be no funding cliff, but the episode accelerated a parallel effort already underway: internationalizing the program's root structure so no single government's budget

decision could interrupt global vulnerability identification. ENISA, the EU's cybersecurity agency, had signaled interest in becoming a top-level Root alongside CISA and MITRE, and formally achieved that status in November 2025 [1][7]. Since then it has recruited 12 CNAs directly and absorbed 8 more transferred from the MITRE root, bringing its total to 20, with the NATO Communications and Information Agency and AISLE among its most recent additions [1]. Hans de Vries, ENISA's Chief Cybersecurity and Operations Officer, tied the expansion explicitly to the AI-driven discovery surge, stating that "recent developments in global cybersecurity, coupled with Frontier AI models' impact on vulnerability discovery, have underscored the need to build strong vulnerability management infrastructure" [1].

In CSA's assessment, the discovery-volume crisis has had the more consequential effect on how the CVE Program functions day to day. Palo Alto Networks' Unit 42 recently disclosed that its own autonomous vulnerability discovery system, NOVA, analyzed 3,915 open-source projects over two months and surfaced 14,090 findings, 99.4 percent of which had not been previously reported, with roughly 40 percent scored high or critical severity under CVSS 4.0 [4]. That single research effort, run by one vendor over sixty days, approaches a meaningful fraction of the roughly 48,000 CVEs published across the entire program in all of 2025 [11]. AISLE's own public disclosures, including 12 CVEs in OpenSSL, among them a defect dating back nearly three decades [10], illustrate that this discovery capability is finding exploitable flaws in load-bearing open-source infrastructure, addressing a common criticism that AI-discovered reports are frequently low-value. Anthropic's Project Glasswing produced a similarly large discovery volume but a strikingly low conversion rate into disclosed, patched CVEs, a gap CSA has previously characterized as the shift of the program's bottleneck from discovery to verification and remediation [9]. CISA's own leadership has acknowledged the tension publicly: Lindsey Cerkovnik, chief of the agency's Vulnerability Response and Coordination Branch, said at the opening of VulnCon26 in Scottsdale, Arizona, in April 2026 that the program is "at a turning point," with some AI tools finding genuinely valid vulnerabilities while others surface issues of little practical value, and that AI companies "should be better represented" within the CNA ecosystem [5][8].

Security Analysis

The AISLE and Anthropic CNA designations should be read as two distinct but related governance signals rather than a single event. The first signal is about self-accountability: an AI-native vulnerability company and a frontier AI lab have both agreed to hold their own software to the same coordinated-disclosure standard they apply, or advocate applying, to everyone else. AISLE described the designation in these terms in its own announcement, stating that "coordinated disclosure is one of the most important tasks in cybersecurity, and that starts with holding your own products to the same standard you expect of everyone else" [2]. That framing matters because AISLE, Anthropic, and comparable AI-

native security vendors are simultaneously becoming major sources of externally reported vulnerabilities and major deployers of autonomous, agentic tooling in enterprise environments; a CNA obligation for their own products creates a public, auditable disclosure trail for security tools to which many customers are extending elevated access to their own codebases and infrastructure. It is a narrower, more limited commitment than it might first appear, however. CNA scope, for both companies, covers only vulnerabilities discovered in software the CNA itself develops or distributes; it does not authorize AISLE to assign a CVE number to an OpenSSL flaw its models find, or Anthropic to number a WordPress plugin vulnerability Claude's evaluation infrastructure surfaces. Those disclosures still route through the affected project's own CNA, a coordinating CNA, or MITRE's default assignment path, exactly as they did before either company held CNA status. Readers who assume the designation gives either firm broad minting authority over the ecosystem's growing AI-discovered vulnerability backlog would be mistaken.

The second, more structural signal concerns where governance authority for AI-era vulnerability disclosure is consolidating, and it points toward ENISA rather than exclusively toward the historically US-centered CISA/MITRE axis. AISLE's founding team draws heavily on European cybersecurity leadership, including former Avast chief executive Ondrej Vlcek and former KPN, Avast, and Rapid7 CISO Jaya Baloo [12], and its choice of ENISA as CNA root, rather than the CISA/MITRE root Anthropic used, may reflect that lineage as much as any policy preference. Whatever the specific reason, the practical effect, at least in these two early cases, is that AI-native vulnerability disclosure is not consolidating exclusively under American governance infrastructure, at the same moment ENISA's own leadership has stated an ambition to "step up" toward parity with CISA and MITRE as a top-level Root [1]. For organizations that rely on the CVE identifier system as a shared, jurisdiction-neutral reference for patch prioritization, this diversification arguably improves resilience: no single government's contracting decision can now interrupt CVE issuance for the growing set of CNAs recruited under ENISA, though it also introduces new coordination demands across roots. It also means security teams tracking vulnerability provenance for audit or supply-chain purposes should expect an increasingly multi-root, multi-jurisdiction CNA landscape rather than the near-single-root structure that prevailed through most of the CVE Program's history.

In CSA's assessment, the volume mismatch documented by Unit 42's NOVA findings and Anthropic's Glasswing conversion rate is the most consequential thread for practitioners, independent of which company holds which CNA badge. A system that finds 14,090 previously unreported issues across fewer than 4,000 projects in two months, or 23,000 across six months from a single research program, suggests a volume that a CNA and enrichment pipeline still calibrated to pre-AI submission rates is struggling to absorb — a conclusion consistent with the low conversion rate Forkast News reports for Glasswing (roughly 6 percent, though the outlet does not detail its denominator), which points toward a growing backlog of known-but-unaddressed flaws rather than simply a faster version of the old process [3][4]. Formalizing AI companies as CNAs for their own products does not resolve that bottleneck; if anything it adds an additional stream of high-volume submissions into a triage system that CSA has

previously found is already straining under earlier waves of AI-assisted discovery [9]. CISA's own leadership has framed the needed response not as tracking more vulnerabilities but as building better triage: Cerkovnik's observation that "not all vulnerabilities matter at the same level" is consistent with the risk-based prioritization approach, weighting exploitability and exposure over raw CVE count, that CISA officials have begun advocating publicly [5][8].

Recommendations

Immediate Actions

Vulnerability management teams should update their internal CNA and vendor-mapping references to reflect that AISLE and Anthropic are now authorized to self-assign CVE identifiers for their own products, and should confirm that any CVE record bearing an AISLE or Anthropic CNA prefix is being sourced and validated through normal channels rather than assumed to carry special authority over third-party findings. Security teams that consume AISLE's or Anthropic's own security advisories should verify whether a given disclosure was self-numbered under the new CNA status or routed through a different CNA, since the distinction affects how quickly and reliably a fix will be published.

Short-Term Mitigations

Organizations with a formal coordinated vulnerability disclosure (CVD) intake process should extend their provenance-labeling practices, tagging inbound reports by whether they originate from an AI-native discovery pipeline, to specifically flag reports attributable to AISLE, Anthropic's Project Glasswing, or comparable frontier-model research efforts, since these sources have demonstrated both high volume and, per CISA's own assessment, uneven per-report value. Security leaders should also monitor CISA's AI researcher CNA pilot program as it matures, since an expansion of that pilot's scope beyond each participant's own products toward broader third-party numbering authority would materially change how enterprises should route and prioritize AI-discovered vulnerability reports [8].

Strategic Considerations

CISOs and vulnerability management leads should track ENISA's continued growth as a CVE Root, now covering 20 CNAs across direct recruitment and MITRE-root transfers, as a durable feature of the CVE ecosystem rather than a transitional arrangement, and should factor multi-root governance into any long-term tooling or reporting dependency on a single national CVE infrastructure [1][7]. Organizations should also watch pending congressional efforts to codify CVE Program governance, including the

board-composition proposal from Representatives Ramirez and Whitesides that CISA has cautioned against over-specifying, since the outcome of that debate will shape how quickly and under what rules additional AI companies gain CNA status in the next twelve to eighteen months [8].

CSA Resource Alignment

CSA's ENISA CVE Root: Dual Vulnerability Governance for Multinationals is the most directly relevant prior CSA work, having already examined ENISA's elevation to CVE Program Root status and the resulting growth of its CNA ecosystem, including the compliance obligations that shift creates for multinational organizations under the EU's Cyber Resilience Act and NIS2 Directive. This note's findings extend that analysis with concrete evidence that ENISA's Root expansion continues at pace, now encompassing 20 CNAs, and that AI-native firms such as AISLE are becoming part of the population of CNAs ENISA recruits, trains, and oversees.

CSA's Project Glasswing and the AI Vulnerability Disclosure Velocity Crisis supplies the broader strategic framing for the volume mismatch documented here, having already argued, from Glasswing's earlier disclosure data, that AI-paced discovery is structurally outrunning human-paced disclosure and remediation capacity. The roughly 6 percent conversion rate Forkast News now calculates for Glasswing's first-half-2026 output is consistent with the low patch rate that earlier CSA analysis identified, and organizations that treat vulnerability remediation as a continuous, automated discipline rather than a queue-based process will be better positioned to absorb the additional report volume that AI-native CNAs are likely to generate.

More generally, the AI Controls Matrix (AICM v1.1) remains the standing reference for organizations building governance around AI-assisted vulnerability discovery and disclosure workflows, particularly its Threat and Vulnerability Management domain, which addresses the provenance-labeling and triage practices this note recommends for AI-sourced vulnerability reports.

References

[1] ENISA. "ENISA scales up its role in the CVE Program." European Union Agency for Cybersecurity, 2026.

[2] GlobeNewswire. "AISLE Named a CVE Numbering Authority, Formalizing Its Commitment to Transparent Vulnerability Disclosure." GlobeNewswire, July 22, 2026.

[3] Forkast News. "Anthropic's CNA Designation Marks the Industrialization of Vulnerability Discovery." Forkast News, July 30, 2026.

[4] Unit 42. "The Frontier AI Vulnerability Burst: Industrializing Autonomous Zero-Day Discovery in Open-Source Software." Palo Alto Networks Unit 42, 2026.

[5] Infosecurity Magazine. "AI Companies To Play Bigger Role in CVE Program, Says CISA." Infosecurity Magazine, April 15, 2026.

[6] CSO Online. "CVE program funding secured, easing fears of repeat crisis." CSO Online, 2026.

[7] Cloud Security Alliance. "ENISA CVE Root: Dual Vulnerability Governance for Multinationals." Cloud Security Alliance, May 18, 2026.

[8] Nextgov/FCW. "CISA cautions against rigid rules for future of cyber vulnerability program." Nextgov/FCW, August 7, 2026.

[9] Cloud Security Alliance. "Project Glasswing and the AI Vulnerability Disclosure Velocity Crisis." Cloud Security Alliance, May 24, 2026.

[10] Infosecurity Magazine. "Autonomous System Uncovers Long-Standing OpenSSL Flaws." Infosecurity Magazine, January 28, 2026.

[11] Gamblin, Jerry. "2025 CVE Data Review." JerryGamblin.com, January 1, 2026.

[12] GlobeNewswire. "AISLE Emerges from Stealth with New AI-Native Cyber Reasoning System to Finally Start Driving Application Vulnerability Backlogs to Zero." GlobeNewswire, October 16, 2025.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.