CSAI Foundation | Cloud Security Alliance

AI-Native CNAs Arrive as ENISA Expands Its CVE Root

AISLE's Numbering Authority Status Signals a Structural Shift in Vulnerability Governance

2026-08-07

AI-assisted Rapid Research

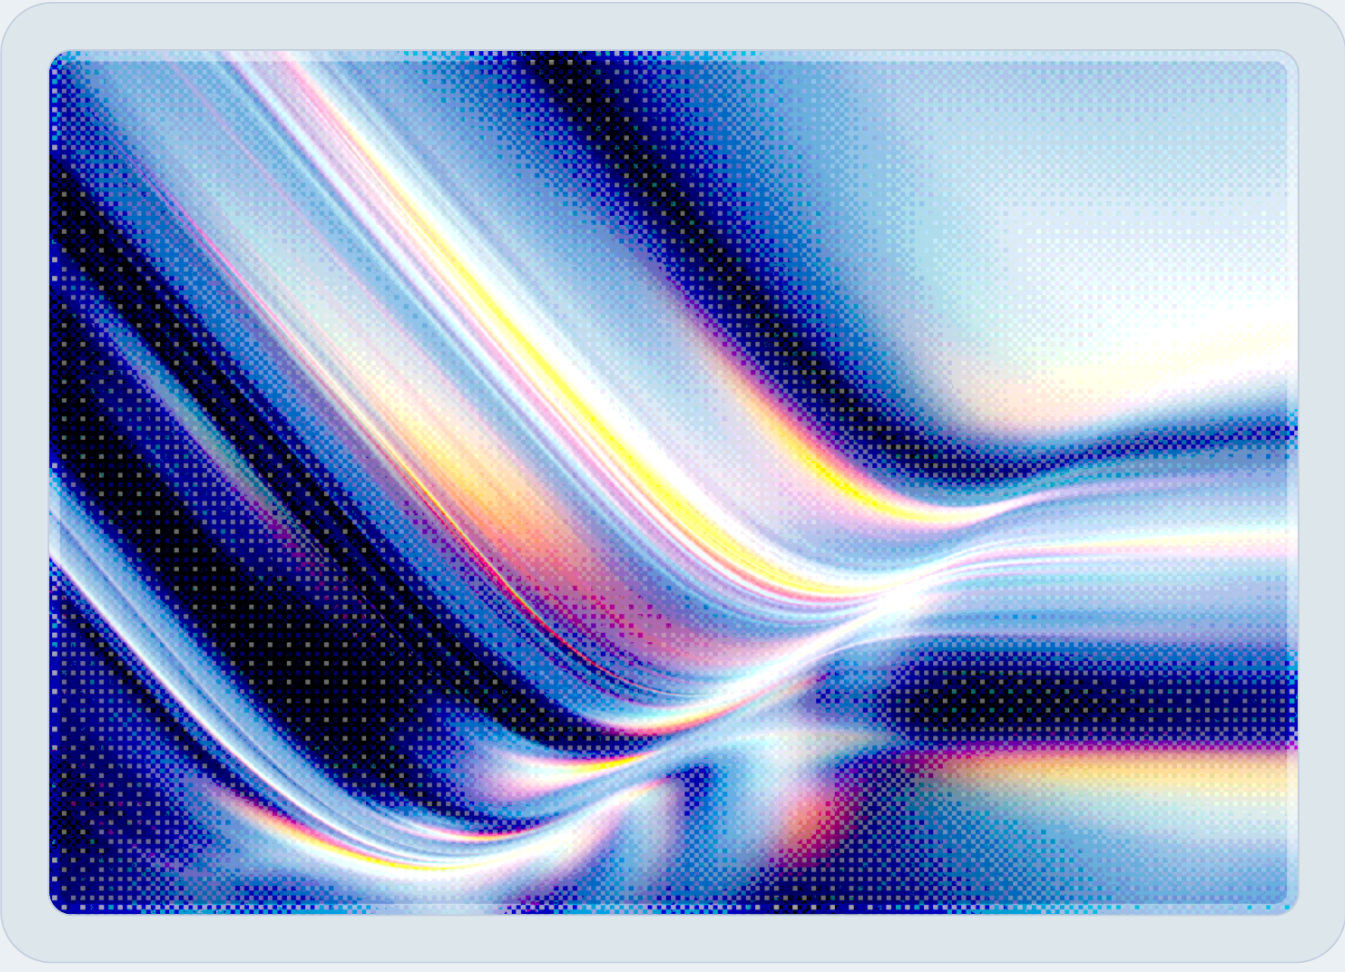

CSAI | CSA cloud security alliance®

Key Takeaways

On August 6, 2026, ENISA announced that the NATO Communications and Information Agency and AISLE, an AI-native vulnerability lifecycle management company, had joined the CVE Program as CVE Numbering Authorities under the ENISA Root [1]. The addition brings the number of CNAs operating under ENISA's Root to twenty, split between organizations onboarded directly and CNAs transferred over from the MITRE Root [1]. AISLE's inclusion is limited in scope — it covers only vulnerabilities discovered in AISLE's own products — but it signals something broader: a company whose core business is autonomous, AI-driven vulnerability discovery has been formally admitted into the governance layer that assigns and publishes CVE identifiers [2]. The proximity to another development is suggestive. Four days before AISLE's designation, Palo Alto Networks' Unit 42 published findings from its NOVA discovery system showing 14,090 confirmed vulnerabilities across 3,915 open-source projects in a two-month autonomous campaign, with 99.4 percent previously unreported and 28.6 percent rated high or critical under CVSS 3.1 (39.7 percent under CVSS 4.0) [3]. Together, these developments illustrate a CVE ecosystem attempting to scale its coordination capacity at the same moment that AI-driven discovery is scaling the volume of vulnerabilities it must coordinate.

Background

The CVE Program has historically operated under a single global Root managed by MITRE, with CVE Numbering Authorities — the organizations authorized to assign CVE identifiers within a defined scope, typically their own products or a specific research community — reporting up through that structure [4]. On November 20, 2025, ENISA became the second organization to hold Root status in the CVE Program, elevating from its prior role as a single CNA to become the governance layer for CNAs operating within the EU [4] [5]. As a Root, ENISA now recruits, onboards, trains, and supports CNAs within its scope, and it has stated its ambition to obtain Top-Level Root status — the tier immediately below the Program's overall governance board — sometime in 2026 or early 2027 [5].

That elevation has proceeded in stages. Four organizations joined as CNAs under the ENISA Root in May 2026, with seven CNAs transferring from the MITRE Root at the same time [10], and CSA's AI Safety Initiative documented the initial legal and compliance implications for multinational organizations shortly after each milestone [6] [7]. The August 6 announcement adding NATO's Communications and Information Agency and AISLE brought the ENISA Root's total to twenty CNAs, of which twelve were

onboarded directly by ENISA and eight transferred from the MITRE Root [1]. ENISA has framed the expansion explicitly in AI terms, citing "recent developments" — a reference to frontier AI models' impact on vulnerability discovery and exploitation — as the rationale for building out stronger vulnerability management infrastructure [1].

AISLE itself is a security company built around AI-native vulnerability lifecycle management. According to AISLE, its platform is designed to move from detection to verified, human-approved fixes without requiring a human analyst at each step, and it operates across cloud, on-premises, and air-gapped environments [2]. The company's researchers have previously disclosed hundreds of vulnerabilities in major open-source projects, including OpenSSL, the Linux kernel, and Apache software, routing those disclosures through the CNAs responsible for each project [2]. CNA status changes that relationship for AISLE's own products: the company can now assign CVE identifiers directly to vulnerabilities found in its own software rather than waiting on a third-party CNA, a change AISLE's CISO, COO, and co-founder Jaya Baloo tied explicitly to the standard of coordinated disclosure the company expects of others [2].

Security Analysis

Three dynamics are worth separating here, because they point toward different risks and different remediation paths.

The first is a capacity expansion whose motivation is best read from ENISA's own statements rather than assumed. ENISA's stated rationale for the Root expansion cites "recent developments" in AI-driven vulnerability discovery and exploitation as the driver for building out stronger coordination infrastructure [1]. Read alongside that framing, the practical effect of adding a second Root is to distribute CNA management across a Root in Europe and a Root at MITRE, with more Roots plausible over time — a structural response to rising disclosure volume rather than a change driven by the raw count of CNAs onboarded.

The second dynamic is what AISLE's admission represents for who gets to hold CNA authority. CNA scope is intentionally narrow — AISLE can only assign identifiers for vulnerabilities in its own products, the same restriction that applies to any vendor CNA — so this is not, by itself, a case of an AI system gaining authority over the broader vulnerability corpus. What it does establish is a precedent: a company whose primary product line is autonomous vulnerability discovery and remediation, rather than a conventional software vendor with an incidental security function, is now inside the CVE Program's formal governance structure rather than exclusively routing findings through others' CNAs. As more AI-native security firms seek CNA status for their own tooling, the Program will need clear criteria for what "own products" means when a company's core offering is finding vulnerabilities in other organizations'

code. AISLE's scope is limited to its own software, which gives it little incentive to inflate findings about products it controls, but the Program will still need to define how conflicts of interest are managed for future entrants whose CNA scope might extend closer to the products they are commercially incentivized to find flaws in.

The third dynamic, and the one the Unit 42 findings make concrete, is the volume mismatch between AI-accelerated discovery and CNA-mediated disclosure. NOVA's two-month campaign against open-source projects found 14,090 confirmed vulnerabilities, 28.6 percent of them rated high or critical under CVSS 3.1 (39.7 percent under CVSS 4.0), and 92 percent were semantic or logic flaws such as access-control gaps, path traversal, and injection vulnerabilities rather than the memory-corruption bugs traditional fuzzing typically surfaces [3]. That single autonomous campaign, run by one research team against a sample of open-source ecosystems, found more confirmed vulnerabilities in two months than the CNA structure now in place was designed to absorb at a comparable pace. Governance capacity has grown incrementally — a handful of CNAs added per Root expansion, bringing the total to twenty organizations under the ENISA Root plus more than ninety European organizations reported as eligible to join as of CSA's May 2026 analysis, alongside MITRE's existing global network [7] — while discovery volume, driven by autonomous systems like NOVA and by AI-native CNAs such as AISLE, appears to be growing far faster. That CNA structure was built around a cadence of human-paced discovery and human-mediated triage, and the mismatch between the two growth rates, rather than either rate in isolation, is the structural risk worth tracking.

A related question is who is positioned to intermediate that gap. AISLE's disclosures into OpenSSL, Linux, and Apache before it held CNA status went through the CNAs governing those projects [2], and there is no indication those CNAs' capacity expanded to match the additional disclosure volume flowing to them from AI-accelerated discovery generally. If AI-native discovery firms increasingly obtain their own CNA status, one effect may be to relieve pressure on downstream CNAs for the discoverer's own products, while doing nothing to address the disclosure bottleneck for the vast body of third-party open-source software that AI-driven tools like NOVA are now surfacing large numbers of vulnerabilities in. CSA's Project Glasswing research, published in June 2026, documented a closely related version of this gap: an AI-driven discovery initiative that found more than 10,000 critical open-source vulnerabilities in its first month, of which only 97 of 1,596 disclosed vulnerabilities had been patched by late May 2026 [8].

Recommendations

Immediate Actions

Security teams operating in or serving the EU should confirm which CVE Root — ENISA or MITRE — governs the CNAs relevant to their supply chain, since the practical effect of ENISA's Root expansion is that disclosure timelines, metadata quality, and CNA responsiveness may now differ depending on which Root a given vendor's CNA sits under. Organizations that consume vulnerability intelligence from AI-native discovery platforms, including AISLE and comparable AI-accelerated scanning tools, should verify whether findings are being disclosed through the platform's own CNA status or through the CNA of the affected open-source project, since disclosure pathway affects both timeline and the completeness of the resulting CVE record.

Short-Term Mitigations

Given that a single autonomous discovery campaign can now surface tens of thousands of vulnerabilities in open-source dependencies over a matter of weeks [3], security teams should not assume CVE assignment and publication will keep pace with discovery, and should build triage processes that can act on vendor advisories and private disclosures before a CVE record is finalized. Teams with significant exposure to the ecosystems named in the Unit 42 findings — Go, JavaScript and TypeScript, PHP, C and C++, and Java [3] — should treat semantic and logic-flaw classes such as access control and injection vulnerabilities as a growing share of what AI-driven discovery surfaces, rather than assuming memory-safety issues remain the dominant category.

Strategic Considerations

Organizations should expect the CVE Program's Root structure to continue evolving, both through ENISA's pursuit of Top-Level Root status and through additional AI-native firms seeking CNA status for their own products, and should track how the Program updates its CNA scope rules and conflict-of-interest guidance in response. Enterprise security leaders should also monitor whether the volume of AI-discovered vulnerabilities begins to outpace the capacity of the CNAs governing the open-source projects they depend on, since that gap — rather than the Root-level governance structure itself — is where disclosure delays and unpatched exposure windows are most likely to appear.

CSA Resource Alignment

This research note extends two prior CSA AI Safety Initiative publications that tracked ENISA's CVE Root elevation as it happened. "ENISA Designated as EU CVE Root: Implications for NIS2 Compliance and Cross-Border Vulnerability Disclosure" examined the initial governance shift and its interaction with NIS2's coordinated disclosure requirements for national CSIRTs [6]. "ENISA CVE Root: Dual Vulnerability Governance for Multinationals" followed with an analysis of the compliance calendar organizations now face across ENISA's Single Reporting Platform and the Cyber Resilience Act's 24-hour and 72-hour notification deadlines, alongside a count of CNAs transferring to the ENISA Root [7]. The developments described in this note — AISLE's admission as an AI-native CNA and ENISA's growth to twenty CNAs — are best read as the next installment of the trend both notes identified: an EU vulnerability governance structure that is still actively being built out even as the volume of vulnerabilities it must process accelerates.

CSA's Project Glasswing research is directly relevant to the volume mismatch discussed above: it documented an AI-driven open-source vulnerability discovery initiative outpacing maintainers' patching capacity, and it is precisely the kind of evidence the third Security Analysis dynamic in this note describes — discovery volume growing faster than the human-paced coordination structures built to process it [8].

Finally, organizations assessing how AI-native vendors like AISLE fit into their vulnerability management program should map that assessment against the Threat Vulnerability Management and AI Security domains of CSA's AI Controls Matrix (AICM) v1.1, which provides control objectives for evaluating AI-driven security tooling and the vendors that supply it [9].

References

[1] ENISA. "ENISA scales up its role in the CVE Program." ENISA, August 6, 2026.

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

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

[4] ENISA. "Stepping Up Our Role in Vulnerability Management: ENISA Becomes CVE Root." ENISA, November 20, 2025.

[5] Kevin Poireault. "ENISA Seeks Top-Tier Status in CVE Program." Infosecurity Magazine, April 15, 2026.

[6] Cloud Security Alliance AI Safety Initiative. "ENISA Designated as EU CVE Root: Implications for NIS 2 Compliance and Cross-Border Vulnerability Disclosure." Cloud Security Alliance, March 10, 2026.

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

[8] Cloud Security Alliance. "Project Glasswing: AI Discovery Outpaces Open Source Patching Capacity." Cloud Security Alliance, June 7, 2026.

[9] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.

[10] ENISA. "New CVE Numbering Authorities Under ENISA Root." ENISA, May 6, 2026.