CENTER FOR CYBERSECURITY POLICY AND LAW

Cybersecurity Coalition

WHITEPAPER

COORDINATING VULNERABILITY RESPONSE IN THE AGE OF AI

How Vulnerability Clearinghouses Can Scale Coordination for AI-Enabled Discovery

Ari Schwartz, Caitlin Clarke, and Timothy McGiff

SEPTEMBER 2026

Contents

I. Introduction 3

II. The Coordination Challenge 4

From Discovery to Action 5

Areas Where Proprietary Software Requires More Coordination 5

Areas Where Open Source Requires More Coordination 6

Core Coordination Functions 7

III. Private Sector Vulnerability Clearinghouses 8

Open-Source Vulnerability Clearinghouses 8

IV. Gold Eagle: A Broader Government Clearinghouse 11

Role and Activities 12

Participation and Output 13

V. Building a Clearinghouse Ecosystem 14

VI. Illustrative Division of Labor 15

VII. Principles for Interoperability 16

VIII. Next Steps and Policy Recommendations 17

I. Introduction

Artificial intelligence (AI) is changing vulnerability discovery from a time consuming and comparatively skilled, expert-driven activity into one that can occur at far greater speed and scale. AI-enabled systems could already identify vulnerabilities, determine exploitations to take advantage of them, and support mitigation. Frontier models, and increasingly, open weight models, are merely accelerating those systems' improvement at those tasks. Recent testing has demonstrated that models can autonomously convert published security patches into working exploits, compressing a process that historically required more specialized expertise and time.[1]

That shift exacerbates coordination challenges across the software supply chain. The important question is no longer simply whether more vulnerabilities can be found, it is whether the organizations that discover or receive, validate, prioritize, aggregate, mitigate, and disclose those vulnerabilities can scale their activities and keep pace with the compressed timelines needed to mitigate exploitation.[2]

For proprietary software, the basic vulnerability disclosure/coordination path is standardized under traditional international standards regimes:[3] a vulnerability is reported to the vendor or product owner, the vendor validates the issue, develops and tests a patch or other mitigation, and distributes the fix to customers. While the process is not always fast or effective, it is typically common for there to be an identifiable entity responsible for receiving the report and deciding how the product will be remediated, and mature software vendors build upon this basic pathway by operating coordinated vulnerability disclosure programs, vulnerability response centers, and bug bounty programs.

While the effects of AI will be felt everywhere, they may be especially pronounced within open-source software's vulnerability disclosure processes. Open-source software is widely adopted across critical infrastructure and within commercial products, yet vulnerability reporting and tracking is less developed, and coordination challenges are more complex. A widely used open-source component may not be consistently maintained or resourced, and open-source components do not come with the same enterprise-grade security guarantees and service level agreements (SLAs) that proprietary software does. Even where there are identifiable maintainers, they may have no relationship with or

[1] https://www.anthropic.com/research/n-days

[2] The difficulty to maintain pace with the discovery of new vulnerabilities is exemplified by the growing backlog of CVEs awaiting enrichment by the National Vulnerability Database (NVD). The National Institute of Standards and Technology (NIST) revised their CVE enrichment processes in 2026 to focus on “the most critical CVEs” due in part to the widening delta between the total number of CVE submissions received and NIST's capacity to enrich them: https://www.nist.gov/news-events/news/2026/04/nist-updates-nvd-operations-address-record-cve-growth

[3] https://www.iso.org/standard/92945.html and https://www.iso.org/standard/81807.html

knowledge of downstream users, and organizations may not even know that a vulnerable library is embedded in their products or systems. Responsibility for receiving a report, developing a fix, identifying affected products, notifying users, and distributing a mitigation is often divided among stakeholders across millions of projects that lack the knowledge of other relevant participants and established channels for working with them.

This paper examines how vulnerability clearinghouses can help to address both the scale and speed issues created by AI by providing a framework for increased, formalized, and systematized cooperation and coordination in both the proprietary and open-source ecosystems. It first describes the coordination challenge and the functions needed to handle mass intake from AI-enabled discovery, then considers why open-source software is a particularly important coordination case, evaluates emerging private-sector clearinghouse efforts, and examines Gold Eagle as the U.S. Government's broader approach to AI-driven vulnerability coordination.

II. The Coordination Challenge

AI increases both the volume of potential cybersecurity vulnerability findings and the speed at which that vulnerability information can become operationally relevant. Traditional disclosure practices were built around a smaller number of reports moving through relatively discrete channels. AI systems can continuously scan for vulnerabilities present in configurations, deployment variance, code bases, and more. They can also reproduce findings and generate exploit-relevant information quickly. Discovery is increasingly becoming less of a constraint than remediation.

AI may significantly increase the volume of findings, but the greater challenge is ensuring sufficient capacity to validate, prioritize, remediate, and deploy mitigations before vulnerabilities can be operationalized by adversaries. AI's compression of vulnerability discovery and exploitation timelines makes it increasingly important to balance rapid public disclosure, which can accelerate ecosystem-wide remediation, with coordinated embargoes that provide critical infrastructure operators time to patch and deploy mitigations before the vulnerability details become broadly available.

From Discovery to Action

The central problem is how to transform a growing volume of vulnerability findings into trusted, contextualized, and actionable intelligence. That requires determining whether a finding is valid, which products and systems are affected, what specific lines of code are impacted, how likely exploitation is, what operational consequences may follow, whether a mitigation already exists, and who is in a position to act. Existing approaches have begun to acknowledge this, for example, CISA's Stakeholder-Specific Vulnerability Categorization (SSVC) already recognizes that prioritization requires context beyond a single severity score, including exploitation status, technical impact, safety consequences, and prevalence.[4]

Areas Where Proprietary Software Requires More Coordination

In proprietary software, vulnerability response generally has a defined center of gravity: the vendor. Researchers typically have a specific organization to contact, and the vendor is positioned to validate the report, determine which supported products are affected, develop and test a patch, coordinate disclosure, and distribute updates. Often, vendors commit to SLAs in fixing vulnerabilities in their software once reported. Customers may still face patching delays, legacy systems, or incomplete inventories, but ownership of the remediation decision is comparatively clear.

Even with the relatively well established and understood coordination mechanisms, there are issues and inefficiencies with proprietary software coordination. Foundational mechanisms like the Common Vulnerabilities and Exposures (CVE) systems and National Vulnerability Database have run into capacity limitations and funding issues, while commonly used scoring processes, like Common Vulnerability Scoring System (CVSS), often fail to provide contextualized risks. AI-scale discovery puts further stress on this model. Vendors will need better deduplication, automated intake, prioritization, and coordination with outside researchers or infrastructure providers.

[4] https://www.cisa.gov/resources-tools/resources/stakeholder-specific-vulnerability-categorization-ssvc

Areas Where Open Source Requires More Coordination

Open-source software often lacks the clear intake and remediation path of proprietary software. It does not come with an SLA. Organizations may not know which open-source components they use, where those components are deployed, or whether a vulnerable library is buried within layers of thousands of downstream dependencies. Adding to the complexity, as open-source is often a component of larger software efforts, when new open-source versions are released that mitigate vulnerabilities, there is not always a clear notification and pathway to update flawed versions. These issues are well documented, and CISA has identified visibility into open-source usage and dependencies as a core security priority. Its guidance on software bills of materials emphasizes the need for entities to identify open-source components as a prerequisite for assessing vulnerability of products and communicating risk.[5]

Responsibility for each step along the remediation path can also be ill-defined and fragmented. A maintainer may have no relationship with downstream vendors or operators, and a downstream supplier may not realize that a vulnerable component is present. A report that would go directly to a proprietary vendor may, in the open-source context, require coordination among a maintainer, foundation, package ecosystem, software supplier, security company, operator, or government body.

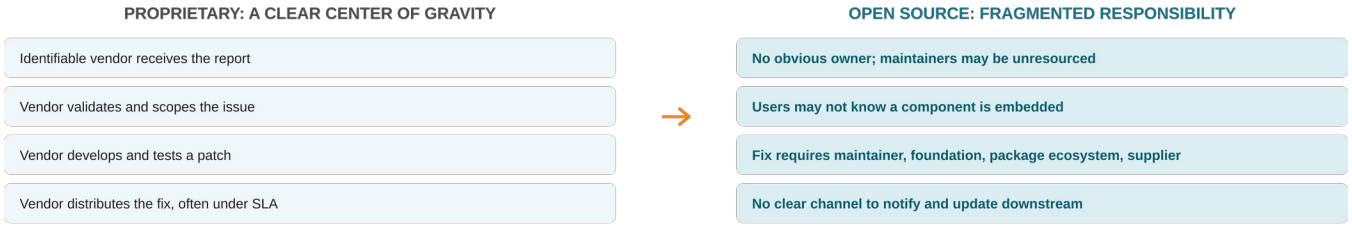

This is not to say that software supply chain security could or should strive to become more centralized, only that AI-scale discovery creates a practical need for interoperable channels that can route findings to the right parties, reduce duplicative work, and support under-resourced projects when no obvious remediation owner exists.

[5] https://www.cisa.gov/resources-tools/resources/cisa-open-source-software-security-roadmap

Core Coordination Functions

Addressing these coordination challenges requires a formalized and systematized process that can better connect six foundational functions:

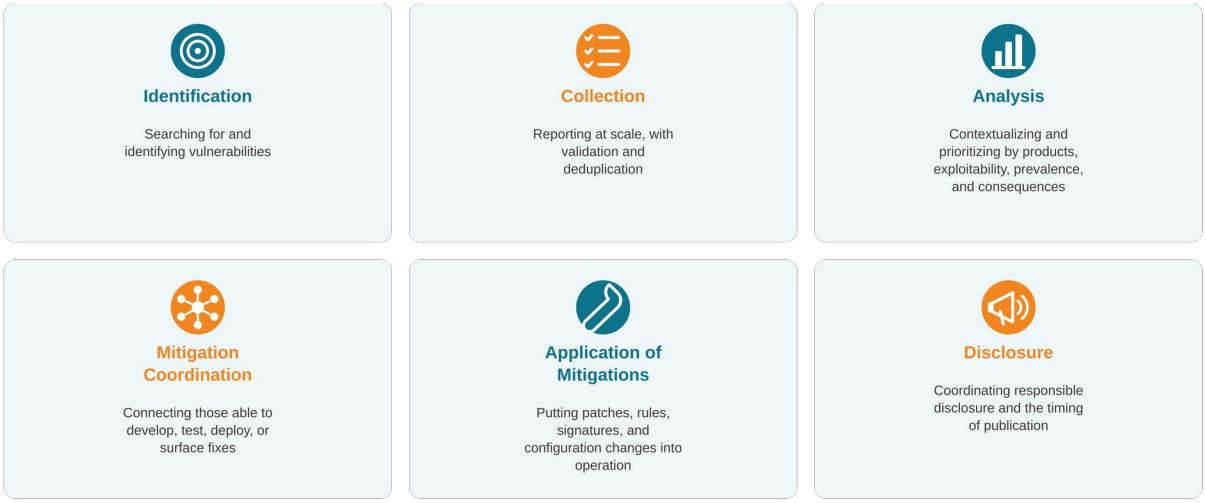

- **Identification:** searching for and identifying vulnerabilities
- **Collection:** facilitating and coordinating vulnerability reporting at scale, including validation and deduplication.
- **Analysis:** contextualizing and prioritizing vulnerability reports based on affected products, exploitability, prevalence, and operational consequences.
- **Mitigation coordination:** connecting the parties able to develop, test, deploy, or surface fixes and compensating controls across diverse use cases.
- **Application of mitigations:** enabling participating organizations to put patches, rules, signatures, configuration changes, or other protections into operation where appropriate.
- **Disclosure:** coordinating responsible disclosure and the timing of broader publication consistent with applicable best practices.

These functions need not be performed by a single entity. Effective vulnerability ecosystems are often federated, with different organizations contributing expertise at different points in the lifecycle. However, they must interoperate if AI-generated findings are to become useful guidance rather than unmanageable report volume.

III. Private Sector Vulnerability Clearinghouses

Vulnerability clearinghouses are one method to facilitate these coordination functions. For the purposes of this paper, a vulnerability clearinghouse is defined as a mechanism that coordinates or facilitates several or all of the six Core Coordination Functions. The clearinghouse model is best understood as an ecosystem: different initiatives may specialize in different stages, sectors, or software communities and still be effective if they can exchange information and hand work off cleanly. A clearinghouse model is generally most effective when attempting to connect and augment existing vulnerability management processes, filling the gaps between them rather than attempting to supersede them. The success of a clearinghouse should be measured by how efficiently it routes information to the parties best positioned to act, rather than by the volume of information it collects.

Open-Source Vulnerability Clearinghouses

The unique challenges associated with coordinating open-source software vulnerabilities has led to the establishment of several private-sector vulnerability clearinghouse initiatives. These initiatives should not be viewed as competing attempts to create a single universal authority for open-source. They are better understood as complimentary efforts that contribute different capabilities and coverage across the broader vulnerability-management ecosystem.

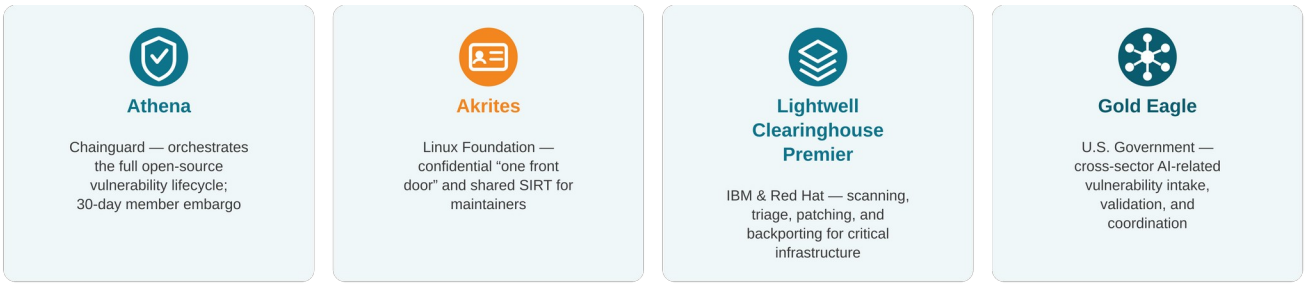

Athena[6]

COORDINATED BY CHAINGUARD

Description: Developed and coordinated by Chainguard, Athena is a centralized clearinghouse dedicated to addressing vulnerability management in the open-source ecosystem. Athena orchestrates stakeholders through the full vulnerability lifecycle to protect both upstream maintainers and downstream consumers of open-source software.

Related Functions: Collection, Analysis, Mitigation Coordination, Application of Mitigations, Disclosure

Role in the Ecosystem: Athena acts as a conduit and coordinating mechanism for stakeholders across the vulnerability lifecycle. It pools findings from member companies using frontier models through initiatives like Project Glasswing and Daybreak. It assesses, validates, deduplicates, and enriches those findings, remediates entire classes of discovered vulnerabilities across project versions and language ecosystems, informs and enables platform, network, and infrastructure mitigations, integrates with cybersecurity partners, and drives coordinated disclosure upstream. Athena members have a 30-day embargo to access patches and get them implemented in product environments before they are publicly disclosed by Chainguard or through Akrites, depending on whether the vulnerability is exploitable at the project's latest version.

Participation: Participation in Athena is assessed primarily on the basis of capabilities. Athena broadly categorizes participating members as either vulnerability-submitting members or mitigation partners. Vulnerability-submitting members must be using frontier AI models to scan sandboxed applications so they can submit those findings for triage, deduplication, enrichment, and remediation. Mitigation partners can implement traffic, endpoint, or network controls or scan customer environments for affected dependencies. Athena's membership is actively continuing to expand across several critical industries and global entities, not just those in the United States. Organizations interested in joining Athena may reach out via the Athena website.

Output: Submitting members receive access to patched artifacts during a 30-day embargo period, while network and cyber partners can use Athena-derived information to create protective rules or alert customers to exposed open-source dependencies. Chainguard also publishes OSV records for “silent” vulnerabilities, or vulnerabilities that exist in older versions but are fixed in the project's latest version. These vulnerabilities were never assigned a CVE tag, causing conventional scanning tools to miss them because there was no record that they exist. For vulnerabilities that exist at the latest version of a project, Chainguard works with the Akrites Initiative to inform the relevant maintainers and get the vulnerability and patch information applied as a fix in a future version.

[6] https://www.chainguard.dev/athena

Akrites[7]

COORDINATED BY THE LINUX FOUNDATION

Description: Developed and coordinated by the Linux Foundation, Akrites provides a centralized, confidential, and structured clearinghouse and shared Security Incident Response Team (SIRT) to handle vulnerability remediation across critical open-source software dependencies. It is designed as a "one front door" for maintainers by validating and deduplicating reports, coordinating fixes, and managing disclosure.

Related Functions: Collection, Analysis, Mitigation Coordination, Application of Mitigations, Disclosure

Role in the Ecosystem: Akrites focuses on disclosure and remediation coordination rather than bug discovery itself. It accepts reports from external discovery efforts and finders and coordinates the downstream work needed to get a validated issue fixed and disclosed. Akrites also complements existing OpenSSF and Linux Foundation efforts and collaborates with open-source foundations like the Rust Foundation and CNCF. All of these efforts help address capacity limitations among maintainers and strengthen existing initiatives.

Participation: Participants must be current Linux Foundation members or invited open-source foundations or projects under participation and confidentiality terms. There are varying membership tiers available for organizations depending on capability and entity type. Organizations are invited to inquire about membership on the Akrites website.

Output: Akrites provides upstream maintainers with validated, deduplicated reports and engineering assistance to prepare and test patches. Tested security patches are published back to the originating project, while participating operators and consumers benefit from coordinated embargoes and synchronized disclosure. The model also contemplates acting as a "maintainer of last resort" for critical projects that lack active maintainers.

[7] https://akrites.org/

Lightwell Clearinghouse Premier[8]

COORDINATED BY IBM AND RED HAT

Description: Jointly developed and coordinated by IBM and Red Hat, Lightwell Clearinghouse Premier is a project within the broader Lightwell effort to secure the open-source software supply chain. The Lightwell Clearinghouse Premier serves as a trusted clearinghouse for vulnerability management, coordinating scanning, triage, patching, backporting, and upstream collaboration.

Related Functions: Collection, Analysis, Mitigation Coordination, Application of Mitigations, Disclosure

Role in the Ecosystem: Lightwell Clearinghouse Premier currently focuses its effort on pre-selected critical US infrastructure entities. These entities receive the services provided by the Lightwell Network, as well as "vulnerability reporting and remediation for member-specific package versions, novel vulnerability verification and disclosure handling, anonymized access to other member requests, and Technical Account Manager (TAM) services."[9]

Participation: The Lightwell Clearinghouse is initially focused on the financial services sector and is expected to expand into additional critical infrastructure sectors.

Output: Mitigations are provided back to originating open-source projects, while Lightwell Network members receive digitally signed binaries, source code, and compliance artifacts, including software bills of materials, for integration into existing pipelines.

IV. Gold Eagle: A Broader Government Clearinghouse

The United States federal government has also pursued the development of a vulnerability clearinghouse solution that appears to envision encompassing both proprietary and open-source software. The Trump administration outlined its approach and intent in the Promoting Advanced Artificial Intelligence Innovation and Security Executive Order that was published on June 2, 2026.[10] That Executive Order ordered the development of an "AI cybersecurity clearinghouse" built "in

[8] https://www.redhat.com/en/lightwell

[9] The Lightwell Network provides annual, consolidated access to Lightwell security remediations and mitigations for eligible open-source vulnerabilities, supporting remediation across the application ecosystem as coverage expands. Customers access the Lightwell repositories through their existing software delivery workflows.

[10] https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/

voluntary collaboration with the AI Industry and operators of critical infrastructure."[11] The effort was tasked with coordinating and deconflicting scanning and identification of software vulnerabilities, validating vulnerability discoveries, and coordinating the prioritization of remediation efforts and the distribution of patches.[12]

The White House formally announced the creation of their vulnerability clearinghouse, designated "Gold Eagle," in a press release on July 14.[13] The press release explained that Gold Eagle "[represented] a new operational model for cyber defense," and stated that it had "already begun to intake and prioritize Identified cybersecurity vulnerabilities from across industries and sectors, coordinate scanning verifications, and ultimately ensure the security of our nation's software and networks."[14] Further operational details of Gold Eagle were revealed on August 14 by CISA.[15]

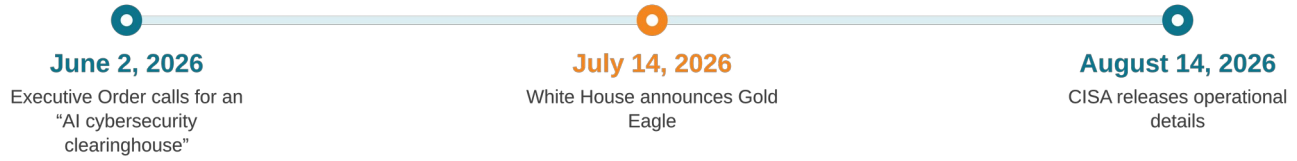

Role and Activities

Official announcements and publications have outlined that Gold Eagle seeks to operate as a coordination mechanism for vulnerabilities identified at scale, including through AI-enabled discovery. Its activities reportedly include receiving vulnerability findings, coordinating scanning and verification, validating findings, prioritizing vulnerabilities and remediation activities, and coordinating the distribution of vulnerability information and patches. Unlike the private sector clearinghouse initiatives, the White House describes Gold Eagle's work in broad terms of addressing vulnerabilities across industries and sectors and the security of software, networks, and critical infrastructure.

Treasury, DHS/CISA, the Department of War, and the Office of the National Cyber Director are participating in the initiative, alongside companies involved in AI, cybersecurity, open-source software, and critical infrastructure. The White House describes this coordination as a means of collectively

[11] Ibid.
[12] Ibid.
[13] https://www.whitehouse.gov/releases/2026/07/white-house-launches-gold-eagle-initiative-for-unprecedented-cybersecurity-vulnerability-coordination/
[14] https://www.whitehouse.gov/releases/2026/07/white-house-launches-gold-eagle-initiative-for-unprecedented-cybersecurity-vulnerability-coordination/
[15] https://www.cisa.gov/resources-tools/resources/gold-eagle-advancing-ai-driven-vulnerability-reporting

identifying risks, prioritizing action, and providing actionable threat and remediation information to defenders.[16]

While open-source software appears to be one component of Gold Eagle's listed activities, and the White House's launch announcement specifically identifies open-source software partners as participants in the coordinated system for receiving and patching vulnerabilities, it is notable that the language used suggests it is structured around collaboration between federal agencies and private-sector participants, and it does not clearly articulate a role for open-source maintainers.[17]

This potentially places Gold Eagle alongside the specialized open-source clearinghouse initiatives discussed above but leaves ambiguity as to how it may ultimately fit in the vulnerability clearinghouse ecosystem. The currently available government materials do not establish a detailed division of responsibilities between Gold Eagle and efforts such as Akrites, Athena, or Lightwell. Accordingly, the extent to which Gold Eagle routes open-source findings to those initiatives, performs overlapping functions, or establishes other forms of operational coordination remains an area where additional implementation detail would be useful.

Participation and Output

CISA describes Gold Eagle as a software capability through which organizations can report, validate, and track AI-related vulnerabilities at scale.[18] Organizations across industries and sectors using AI to identify vulnerabilities at scale are encouraged to participate. As described in the existing draft, reports submitted through Gold Eagle that satisfy coordinated vulnerability disclosure requirements initiate cases in the Vulnerability Information and Coordinated Environment (VINCE), the web-based collaboration platform developed by Carnegie Mellon University's CERT Coordination Center and sponsored by CISA.

[16] https://www.whitehouse.gov/releases/2026/07/white-house-launches-gold-eagle-initiative-for-unprecedented-cybersecurity-vulnerability-coordination/
[17] Terms commonly used include organization, industry, industry partners, private partners, and private sector.
[18] https://www.cisa.gov/resources-tools/resources/gold-eagle-advancing-ai-driven-vulnerability-reporting

V. Building a Clearinghouse Ecosystem

The emerging vulnerability clearinghouse landscape is best understood as an ecosystem of complementing efforts with unique specializations, focuses, relationships, and capabilities. This emerging ecosystem comprises all the varied capabilities and stakeholders necessary to mitigate the AI specific challenges of increased vulnerability volume and the compressed response timelines, while also strengthening longstanding vulnerability management practices, including the coordination of disclosure embargoes to provide affected organizations sufficient time to develop, test, and deploy patches.

Akrites, Athena, and Lightwell are examples of specialized private-sector efforts focused heavily on open-source coordination, while Gold Eagle seeks to have a broader cross-sector role in AI-related vulnerability intake and coordination. These newer efforts also operate alongside and augment established programs — including CVE, VINCE, JCDC, OpenSSF initiatives, vendor security teams, and sector-specific organizations — that already perform adjacent or complementary functions.[19]

Software ecosystems are global and therefore interoperability should extend beyond U.S.-based participants. Effective coordination mechanisms should facilitate international participation where appropriate. Given this range of actors, the central question is therefore not "which effort should perform which function," but "how can these initiatives work together effectively?" The policy objective should be to make these efforts interoperable enough that information can move to the actor best positioned to use it. A finding should not need to be independently rediscovered, revalidated, and re-triaged by every organization it touches. Similarly, participation in one clearinghouse should not prevent an issue from reaching another coordinator with more relevant technical or sector-specific expertise.

[19] A wide range of well-established and nascent initiatives in both open-source and proprietary software may become important contributors to the wider vulnerability ecosystem. This may include collaborative open-source efforts such as the OpenAI, Trail of Bits, and HackerOne involved Patch the Planet, which connects security engineers with open-source projects; Project Quiltworks, a Crowdstrike led coalition dedicated to closing the exploit gap created by frontier AI in a coordinated manner; and the Google Fairwind Program, which provides trusted partners with early access to advanced AI models and tools to find, verify, and fix vulnerabilities.

KEY TAKEAWAY

The success of a clearinghouse should be measured by how efficiently it routes information to the parties best positioned to act, rather than by the volume of information it collects. The relevant test is not whether one initiative is comprehensive, but whether a report can reach the right actor quickly and retain the context needed for action.

VI. Illustrative Division of Labor

Interoperability does not require every clearinghouse or vulnerability-management program to perform the same functions. The initiatives described above have different participants, capabilities, areas of expertise, and existing relationships, and those differences allow them to play complementary roles within the broader vulnerability-management ecosystem. Clarifying where those roles overlap and where information may need to move between organizations can help identify both potential gaps and unnecessary duplication.

The following illustrative division of labor provides one way to consider how responsibilities could be distributed across the ecosystem. It is not intended to prescribe fixed roles for individual initiatives, but rather to demonstrate how different actors could contribute at different stages of the vulnerability lifecycle while preserving existing programs and areas of expertise.

Role	Proprietary Vendor	Specialized Open-Source Effort	Gold Eagle / Government Coordination
Receive / validate reports	Primary owner for its products	Provide a clear front door where maintainers or ecosystems lack one	Accept large-scale reports; route and deconflict across programs
Develop software patch	Vendor engineering team	Maintainers, foundations, or clearinghouse engineering support	Generally coordinate rather than become the patch developer
Prioritize risk	Product and customer context	Component, dependency, and maintainer context	Cross-sector, critical infrastructure, and national-risk context
Deploy compensating controls	Vendor / customers / security providers	Partners may create network, endpoint, or artifact mitigations	Connect affected sectors and support broad defensive coordination
Open-source cases	As downstream user or distributor	Primary specialized coordination path	Use existing efforts where possible; fill gaps and connect to critical infrastructure

VII. Principles for Interoperability

As the vulnerability clearinghouse ecosystem continues to evolve, it would be beneficial for stakeholders to come to agreement on core principles for interoperability to guide development. For example:

- Vulnerability clearinghouses and their associated mechanisms should endeavor to use common identifiers and machine-readable formats, and ensure bi-directional interoperability with the CVE and NVD ecosystems so findings and remediation information can be exchanged without repeated manual translation.

- Maintain provenance, attribution, and version history for vulnerability information so updates, corrections, enrichment activities, and remediation status can be tracked transparently over time.
- Standardizing and explicit handoff rules should be developed that allow a clearinghouse to route an issue to a vendor, open-source coordinator, sector partner, or government program without losing provenance or context.
- Ensuring deduplication across discovery sources so mass scanning does not create multiple parallel cases for the same vulnerability.
- Leverage existing identifiers and standards, including CVE and related vulnerability data formats, wherever possible to avoid fragmentation and duplicative reporting requirements.
- Ensuring vulnerability remediation measures (i.e., patches) address the vulnerability without introducing a breaking change in the package's functionality or a new vulnerability.
- Developing deduplication guidance or version guidance for organizations where multiple patched or otherwise remediated versions of a package exist simultaneously.
- Implementation of need-to-know controls for sensitive pre-disclosure information, combined with clear rules for access, retention, correction, and eventual disclosure.
- Implementation of mechanisms for feeding mitigation outcomes and recurring vulnerability patterns back into AI discovery, secure development, and defensive controls.
- Establishing clear boundaries around government use of submitted vulnerability information, with strict parameters for exclusive support for remediation and responsible disclosure, except where it is specifically provided or authorized for other legitimate government purposes.

This list is not meant to be exhaustive or authoritative, but it could serve as a starting point for discussions between stakeholders.

VIII. Next Steps and Policy Recommendations

As AI increases the volume and speed of vulnerability discovery, the effectiveness of the emerging clearinghouse ecosystem will depend in part on how well new and existing efforts interact. The development of multiple initiatives creates opportunities to bring different capabilities and expertise to the vulnerability-management process, but it also raises practical questions about duplication, information sharing, routing, and responsibility.

The next phase of this work should therefore focus on how these efforts can operate together as AI-enabled vulnerability discovery continues to scale. The following areas provide a starting point for considering how policymakers, clearinghouse operators, vendors, researchers, and other participants

can strengthen coordination while maintaining the distinct roles of existing vulnerability-management programs.

1. **Define the clearinghouse problem around the scale of vulnerabilities and not the software model (open or proprietary).** Policy should begin with the operational consequences of AI-driven mass vulnerability discovery across the software ecosystem. AI should be integrated into existing vulnerability management processes and standards wherever possible rather than creating separate reporting or coordination regimes.

2. **Preserve clear vendor responsibility where it already exists.** For proprietary software, clearinghouses should improve intake, deduplication, validation, prioritization, and coordination without weakening the basic expectation that vendors remain responsible for validating and remediating vulnerabilities in their products.

3. **Use specialized open-source initiatives rather than recreating them inside government.** Gold Eagle should establish strong interfaces with existing open-source clearinghouses, foundations, and project security teams and should route findings to those actors when they are better positioned to coordinate with maintainers.

4. **Strengthen foundational vulnerability infrastructure.** Governments and industry should support modernization of vulnerability management systems through machine-readable formats, APIs, provenance tracking, version history, automation-friendly tooling, and scalable enrichment capabilities.

5. **Identify clearinghouses as serving as a primary means to provide context for vulnerability prioritization, including empirical evidence of exploitability.** This context is generally derived from highly operationally sensitive information and a class of data not normally a major component of existing vulnerability programs, which have historically dealt with academic/POC exploits around single points in an attack chain or low-complexity attack chains. At AI-driven discovery volumes, disclosure alone cannot tell a defender which vulnerabilities warrant action. That determination requires testing the vulnerability against the configurations, versions, and exposure conditions in which it actually appears. Clearinghouses should therefore treat exploitation-testing results as critical context, including negative results establishing that a vulnerability is not exploitable in common deployments, and should create channels for the practitioners who generate that evidence to contribute to it.

6. **Focus Gold Eagle on the coordination gaps only the government can fill.** Its comparative advantage lies in cross-sector deconfliction, critical infrastructure context, escalating risk visibility into the C-Suite, broad participation, and the ability to connect findings with the organizations that can act. It should avoid becoming a universal patching organization or a substitute for mature vendors and community processes.

7. **Build for bulk learning as well as case management.** Clearinghouses should analyze patterns across AI-generated findings to identify recurring vulnerability classes, dependencies, and mitigations. Outputs may include development guidance, network or endpoint protections, scanning signatures, and other defenses that reduce risk across multiple vulnerabilities at once.

8. **Invest in remediation scalability.** Efforts to accelerate vulnerability discovery should be matched by investments in validation, prioritization, patch development, mitigation deployment, and measurement of remediation outcomes.

9. **Make interoperability a design requirement.** Clearinghouse efforts should use compatible identifiers, formats, interfaces, and disclosure practices so that the ecosystem can divide work without fragmenting information. The relevant test is not whether one initiative is comprehensive, but whether a report can reach the right actor quickly and retain the context needed for action.

10. **Bound government participation with trust protections.** Government involvement can be valuable for critical infrastructure and cross-sector coordination, but vulnerability information may be highly sensitive.

As these important steps for vulnerability coordination evolve and strengthen, it is essential that policy makers and cyber industry leaders embrace a foundational principle that at its core, vulnerability disclosure is merely an alert mechanism. Knowing a bug exists without internal operational readiness is virtually useless without knowing, e.g., the software you have and especially the software you depend on for critical operations. Effective implementation of the core Cybersecurity Framework (CSF)[20] is the key to actionable defense and cyber resilience.

[20] https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf