CVE Program Virtual Forum: AI-Enabled Vulnerability Discovery

July 30, 2026, 10:00AM – 2:00PM EDT

Key Takeaways

Below are high-level takeaways from the forum that provide an overview of the key themes, challenges, and participant perspectives discussed. These topics are described in greater detail throughout the meeting notes.

- **AI Is Increasing Pressure Across the Vulnerability Lifecycle:** Higher report volumes are increasing workload across triage, validation, remediation, CVE assignment, and downstream vulnerability management.
- **Prioritization Matters More as Volume Grows:** Participants emphasized that more vulnerabilities do not automatically mean greater risk. Exploitability, exposure, and organizational context are increasingly important for determining which vulnerabilities require action.
- **Quality and Validation Must Keep Pace:** False positives, duplicate findings, hallucinations, and uneven report quality are creating additional work. AI can help scale triage and record development, but human verification remains important.
- **Remediation and Automation Need Greater Investment:** Participants stressed that discovery must be matched by the ability to fix vulnerabilities at scale. Automation, stronger remediation workflows, and better tooling were viewed as necessary responses to increasing volume.
- **Open-Source Maintainers Need More Support:** Resource-constrained maintainers are absorbing significant additional workload. Participants called for greater funding, expertise, patches, and tooling from organizations that depend on open-source software.
- **CVE Data and Processes Need Continued Improvement:** Participants identified gaps in product mapping, identifiers, enrichment, assignment, and interoperability that limit automation and create downstream confusion. Discussion supported further work on data standards and CVE processes but did not produce clear consensus on the scope of broader program changes.
- **Education and Engagement:** As an undercurrent to the primary discussion themes, many participants among both speakers and listeners pointed to a sustained need for increased engagement throughout the CVE community to share knowledge and enhance collaboration, particularly through CVE working groups, blog posts, and further forums.

Welcome and Overview

Alec Summers, MITRE, delivered opening remarks for the CVE Program's forum, AI-Enabled Vulnerability Discovery, on July 30, 2026. Alec prefaced the conference by sharing the strong interest from across the CVE community in leveraging artificial intelligence (AI) for vulnerability discovery, noting that more than 700 people from 70 countries registered for the discussion. Alec shared that the community has expressed interest in and concern over how AI tools, platforms, and services are increasing the volume of reported vulnerabilities, changing the workflows used to process CVEs, and creating a need for collaboration across the community.

- **Purpose:** The event opened as a CVE Program discussion on AI-enabled vulnerability discovery. It was presented as both a listening forum and a working forum, with the explicit goal of surfacing pressure points, practical responses, and areas where follow-up work may be needed rather than announcing a predetermined policy outcome.
- **Participation:** Organizers noted that more than 700 people from 70 countries had registered. The audience was described as broad and cross-functional, with representation from individual researchers, software maintainers, suppliers, CNAs, government organizations, security vendors, and enterprise consumers.
- **Impact:** The opening remarks emphasized that AI-enabled discovery is changing the speed, scale, and variability of vulnerability information across the full lifecycle. The discussion was framed around the different pressure points experienced by researchers, maintainers, suppliers, CNAs, security vendors, and downstream consumers.
- **Structure:** The forum was organized around three questions: what changes when vulnerability information moves faster and at greater scale; where existing workflows, responsibilities, or assumptions are under pressure; and where process or technology improvements could make the biggest difference.
- **Logistics:** Participants were encouraged to engage in the forum via the chat, Q&A, and outreach to the CVE Program Secretariat (cve-prog-secretariat@mitre.org). The Program shared the session recordings to the CVE Program YouTube channel.

Interactive Community Listening Session

Jay Jacobs, of Empirical Security, and Bob Lord led an interactive session that gathered live input from hundreds of participants on their roles, experiences, and perspectives regarding AI-driven changes, workload increases, and actionable findings.

- **Live Polling Approach**: Bob Lord and Jay Jacobs facilitated a survey-based, collaborative listening session to gather live input from participants on their roles, experiences, and priorities. Participants were asked to identify where they work in the vulnerability ecosystem and then answer follow-on questions from that perspective.
- **Location Warm-up and Global Reach:** The first prompt asked participants where they were joining from. The responses showed broad international participation, with a large concentration in North America, reinforcing that the event had attracted a geographically diverse audience.
- **Role-based Participation:** Participants spanned the vulnerability lifecycle, led by report intake and validation (49%) and CVE assignment or publication (44%). The largest primary-role group (31%) focused on evaluating reports, assigning CVEs, or coordinating disclosure, followed by organizational vulnerability management and policy, governance, risk, or research at 17% each.
- **Workload and Sustainability:** Responses indicated that vulnerability-related workloads have increased sharply, with more effort required to separate actionable findings from noise. Participants also indicated that current response demands are difficult to sustain, while automation has provided limited relief.
- **AI-related Change Patterns:** Participants described AI-related change in two ways: an increase in volume and a change in process. The discussion distinguished between AI simply

adding work and AI changing how work is actually done, with many responses clustering in the middle rather than at the extremes.

- **Observed AI-Related Impacts:** Free-text comments highlighted duplicate reports, false positives, hallucinations, low-quality submissions, and the need for better triage tools. Participants also described positive changes, including faster CVE description preparation, better report drafting, and the potential for more efficient use of context and evidence.
- **Interpretation of the Input:** Bob and Jay repeatedly asked participants to answer from direct experience rather than from social-media impressions or speculation. They also noted that the data would be analyzed later in the day and incorporated into a post-event summary.

Lightning Talks Session I

Jerry Gamblin and Ben Edwards of Empirical Security discussed how rising vulnerability and CVE volumes should be interpreted in an AI-enabled environment, with Jerry cautioning against panic-driven narratives around catalog growth while Ben examined whether AI is measurably accelerating CNA publication rates. They emphasized that higher volume does not automatically mean greater risk, and that the ecosystem should focus on prioritization, quality, and workflow adaptation rather than alarmist framing.

The Vulnpocalypse is Cancelled

Jerry Gamblin argued that rising CVE counts should not be framed as a catastrophe, emphasizing that CVE is a naming program and that the ecosystem is getting better at finding and documenting vulnerabilities rather than simply becoming less manageable.

- **CVE as a Naming System:** The CVE Program exists to solve a naming problem, and growth in the catalog should not automatically be interpreted as evidence that vulnerability management is becoming unmanageable.
- **Panic Framing Creates Harm:** Terms such as “apocalypse” or “tsunami can distort how vulnerability growth is understood and drive unproductive responses from leadership, vendors, and defenders.
- **More Reporting Reflects Broader Participation:** Growth in the CNA ecosystem has increased the number of organizations identifying and reporting vulnerabilities, improving visibility into issues that may previously have gone undocumented.
- **Only a Small Fraction Drives Urgent Action:** Only a limited portion of published CVEs become high-priority downstream concerns, reinforcing the need to distinguish catalog growth from vulnerabilities requiring immediate action.
- **Better Metrics Matter More Than Bigger Numbers:** Greater emphasis should be placed on exploitation and remediation outcomes rather than using total CVE volume as the primary measure of risk.

Measuring AI Acceleration at the CNA Level

Ben Edwards examined whether AI is measurably accelerating vulnerability publication, emphasizing that increased discovery does not translate directly into CVE growth because publication still depends on human review, validation, and coordination.

- **Measuring the AI Effect:** Edwards examined CNA-level publication data for evidence of an AI-driven increase, noting signs of recent growth while cautioning against drawing conclusions from anecdotal reports alone.
- **Vulnerabilities Are Not Automatically CVEs:** AI may increase vulnerability discovery, but findings still must move through human review and coordination before becoming published CVE records.
- **Prioritization Will Become More Important:** Growing volumes will require CNAs to make stronger decisions about which findings should be processed first rather than treating every submission equally.
- **Record Quality Must Be Preserved:** Increased volume should not come at the expense of CVE quality. Validation and quality controls will remain essential as AI-assisted reporting expands.
- **Tooling Is Necessary to Scale:** Manual processes alone will not keep pace with higher volumes. AI and other tooling can support triage and record preparation, but outputs still require verification.
- **Consumers Face the Same Scaling Problem:** Downstream users will also need stronger prioritization and tooling to determine which vulnerabilities warrant attention as the volume of available information grows.

AI-Enabled Discovery: Ecosystem Perspectives

Alec Summers moderated a panel discussion on how AI-enabled vulnerability discovery is increasing pressure across the ecosystem through higher report volume, false positives, duplicate findings, mapping ambiguity, and coordination strain. Panelists included Chris Robinson from the Linux Foundation, Stig Palmquist from the Comprehensive PURL Archive Network, Deana O'Meara from NVIDIA, Ken Williams from Broadcom, Aubrey Olandt from Red Hat, and Daniel Larson from CISA. The panel emphasized that the main challenge is not just finding more issues, but helping the ecosystem validate, fix, map, and communicate them effectively.

Upstream Open-Source Pressure and Support Gaps

Christopher Robinson focused on the growing burden AI-enabled vulnerability discovery places on upstream open-source projects, particularly projects with limited funding, staffing, and maintainer support.

- **Upstream Maintainer Pressure:** Increased vulnerability reporting is adding significant triage and remediation work for maintainers who already operate with limited resources.
- **Support Must Extend Beyond Discovery:** Organizations that depend on open source should contribute expertise, patches, feedback, and tooling rather than placing the full burden of newly discovered vulnerabilities on upstream maintainers.
- **Better Submission Practices:** Reports should provide actionable evidence, including proof of concept, tests, and proposed patches where possible, giving maintainers a stronger starting point for validation and remediation.
- **Embedded Security Support:** Security expertise embedded within open-source communities or foundations can help bridge gaps between vulnerability researchers and maintainers and accelerate response.

False Positives, Maintainer Burden, and Report Quality

Stig Palmquist focused on the burden that high-volume, uneven-quality AI-generated reports place on open-source maintainers and emphasized that vulnerability reporting should help maintainers resolve issues rather than create additional work.

- **Report Volume Is Increasing Maintainer Burden:** Maintainers are receiving more reports, including duplicates and false positives, while the length of AI-generated submissions can make legitimate findings harder to assess.
- **Reports Should Be Concise and Actionable:** Effective reports should provide the minimum necessary detail to understand and reproduce the issue rather than overwhelming maintainers with generated text.
- **Evidence Improves Report Quality:** Proof of concept, tests, and proposed patches can help validate findings and give maintainers a clearer path toward remediation.
- **Discovery Should Support Remediation:** Finding a vulnerability is only part of the process. Reporting practices should make it easier for maintainers to validate and fix the issue instead of transferring the researcher's workload upstream.
- **Incentives Should Favor Fixes:** The ecosystem should place greater value on remediation and consider stronger incentives for fixing vulnerabilities, not only discovering and reporting them.

Remediation Investment, Harness Design, and Internal Coordination

Deana O'Meara described how the declining cost of vulnerability discovery is shifting the operational challenge from finding vulnerabilities to validating, coordinating, and fixing them at scale.

- **Remediation Investment Must Increase:** As AI makes vulnerabilities easier to discover, organizations need corresponding investment in remediation capacity rather than focusing primarily on additional discovery.
- **Harness Design Matters:** The effectiveness of AI-assisted security work depends heavily on the surrounding harness, context, and workflow, not simply the underlying model.
- **Context Can Reduce False Positives:** Project documentation and resources such as security.md files can give AI systems better context and improve the quality of generated findings.
- **Internal Coordination Is a Scaling Challenge:** Identifying the correct product or engineering team for a finding can be difficult at higher volumes, creating a need for better attribution and routing.
- **AI Can Support the Response:** AI-assisted skills, prompts, portals, and workflows can help organizations coordinate and remediate vulnerabilities in bulk.

Operation Scaling Challenges and Human Oversight

Ken Williams focused on the operational challenges of scaling AI-enabled vulnerability discovery, including uneven tool access, variable output quality, duplicates, and the continued need for human oversight.

- **Volume Is Increasing Rapidly:** AI tools are producing substantially more findings across both open and closed software, including vulnerabilities in pre-production code.

- **Access and Cost Remain Constraints:** Frontier cybersecurity models require financial resources, token budgets, and infrastructure that are not equally available across organizations or researchers.
- **Quality Varies Widely:** False positives and hallucinations remain significant concerns, with results heavily influenced by the quality of the harness and supporting skills.
- **Duplicate Findings Add Work:** Organizations may discover vulnerabilities internally only to receive the same findings externally months later, creating additional coordination and triage burden.
- **AI-Generated Fixes Require Review:** AI can identify legitimate vulnerabilities while still producing weak remediation recommendations, reinforcing the need for humans throughout triage and remediation.

Downstream Mapping, Standardization, and Practical Fixes

Aubrey Olandt focused on the downstream challenge of accurately mapping upstream vulnerabilities to affected products and determining whether findings apply to specific customer environments.

- **Product Mapping Remains Difficult:** Multiple downstream components, inconsistent naming, and incomplete identifiers make it difficult to connect upstream vulnerabilities to affected products.
- **Poor Mapping Creates False Positives:** Weak or incomplete metadata drives manual triage and can create confusion for customers trying to determine whether they are affected.
- **Standardization and Enrichment Are Needed:** Stronger baseline CVE data and better enrichment would improve downstream correlation and reduce reliance on supplemental vulnerability databases.
- **PURLs Provide a Practical Approach:** Embedding upstream source references as PURLs in SBOMs and querying those references directly can improve connections between upstream and downstream information.
- **Upstream Support Remains Essential:** Better downstream data alone will not resolve the problem without additional resources and support for upstream maintainers.

CNA-LR Assignment Pressure and Coordination Complexity

Daniel Larson described how higher AI-generated vulnerability volume is increasing pressure on CNA-LRs, particularly when findings are incomplete, difficult to route, or involve projects without an obvious responsible CNA.

- **Assignment Decisions Are Becoming More Complex:** CNA-LRs may need to make assignment decisions using incomplete or immature technical information when publicly disclosed vulnerabilities fall outside established CNA scopes.
- **Supplier Validation Can Become a Bottleneck:** CNA-LRs depend on suppliers for product expertise, but overloaded suppliers may struggle to validate findings within expected disclosure timelines.
- **Routing Is Increasingly Difficult:** AI tools are identifying vulnerabilities in dependencies and abandoned projects where responsibility for CVE assignment may be unclear.

- **Coordination Across Roots Is Required:** Difficult cases may require coordination among CNA-LRs to determine the appropriate assignment path and transfer records when the responsible owner is later identified.
- **Backlogs Increase Timeline Pressure:** Growing assignment queues and compressed disclosure timelines can create bottlenecks, increasing the importance of efficient handoffs while preserving the context and authority behind CVE assignments.

Vulnerability Data: Consumer Perspectives

Chris Folk, MITRE, moderated a panel that discussed the necessity of prioritization, automation, and quality improvements in vulnerability management. Speakers included Chris Nims of Capital One, Ray Carney of Tenable, Saeed Abbasi of Qualys, and Robert Hansen of Rude Evidence. The panel highlighted the divergence between disclosed and exploited vulnerabilities while advocating for structural changes to CVE enrichment and software development practices.

Automation and Continuous Deployment

Chris Nims focused on using automation and frequent software delivery to shift vulnerability management away from prolonged prioritization debates and toward faster remediation.

- **Automate Remediation**: AI is being used to generate vulnerability fixes, reducing the amount of manual engineering work required to address findings.
- **Fix Rather Than Debate**: Engineering teams can spend more time determining whether a vulnerability is exploitable than it would take to remediate the underlying issue.
- **Continuous Deployment Supports Scale**: Frequent release cycles allow fixes to move into production quickly rather than accumulating against traditional remediation timelines.
- **Human Review Remains Important:** Approximately 80% of AI-generated fixes discussed required no human modification but fixes still receive human review.
- **Improve Security at the Source:** Greater automation should be paired with stronger code quality and development practices that address vulnerabilities earlier in the software lifecycle.

Separating Vulnerability Volume from Real-World Risk

Robert Hansen focused on the widening gap between the number of vulnerabilities disclosed and the much smaller subset that creates meaningful risk for defenders.

- **Disclosure Volume Does Not Equal Threat:** Growth in published vulnerabilities is significantly outpacing growth in vulnerabilities actively used by adversaries.
- **Prioritization Requires Better Signal:** Defenders need to distinguish vulnerabilities that materially affect their environments from the much larger volume of disclosed findings.
- **Exploitation Should Inform Decisions:** Real-world attacker behavior and exploitation data provide stronger prioritization signals than vulnerability volume alone.
- **Perfect Remediation Is Not Realistic:** Vulnerability management should account for practical risk and operational constraints rather than assuming every disclosed issue can receive equal attention.

Validating Vulnerabilities and Understanding Exploit Chains

Saeed Abbasi focused on validation and vulnerability chaining, emphasizing that individual vulnerabilities may not provide enough information to understand an actual attack path.

- **Validation Is Essential:** Vulnerability findings need confirmation before organizations can reliably use them for prioritization and remediation decisions.
- **Vulnerabilities Must Be Viewed in Context:** Individual weaknesses can become more significant when combined with other vulnerabilities or environmental conditions.
- **Chaining Needs Better Representation:** CVE information should better capture exploit paths and relationships between vulnerabilities when those relationships affect real-world risk.
- **Actionable Context Improves Prioritization:** Additional information about how a vulnerability can be reached or combined with other weaknesses can help downstream users make better decisions.

Turning CVE Data into Actionable Risk Decisions

Ray Carney focused on the distinction between CVE as a vulnerability identification system and the additional information consumers need to make operational decisions.

- **CVE Identifies but Does Not Prioritize:** A CVE record establishes that a vulnerability exists but does not by itself tell an organization whether that vulnerability should be addressed first.
- **Disclosed and Exploited Vulnerabilities Are Diverging:** Vulnerability disclosure continues to grow, while only a narrower subset of those vulnerabilities becomes weaponized.
- **Consumers Need Operational Context:** Effective prioritization requires information such as active exploitation, environmental reachability, relevant threat actors, campaigns, and organizational exposure.
- **AI Can Connect Vulnerability Data:** AI may provide greater downstream value by rapidly connecting CVEs with threat, exploitation, exposure, and organizational context rather than simply finding more vulnerabilities.

Lightning Talks Session II

The second round of lightning talks included discussions on variance in the quality of AI-assisted vulnerability submissions, improving how vulnerability identification connects to CWE root-cause classification, and a maintainer-informed approach to improving early-stage triage for vulnerability reporting.

Waiting for the Fever to Break

Erick Galinkin of NVIDIA focused on the uneven quality of AI-assisted vulnerability submissions and the need to prevent low-quality findings from increasing maintainer and PSIRT workload.

- **Report Quality Remains Uneven:** AI-generated reports range from useful findings to submissions containing fabricated evidence, incorrect reasoning, or bugs that are not meaningfully exploitable.
- **Validation Must Come Before Submission:** Researchers should confirm findings and provide credible proof of concept rather than automatically submitting model-generated results.

- **Severity Can Be Overstated:** AI-assisted reports can inflate vulnerability severity, particularly when models are asked to generate scores without systematically evaluating the underlying scoring criteria.
- **Automated Triage Has Limits:** Models can misclassify findings, fail without clear warning, or be manipulated, making human oversight important for vulnerability assessment and triage.
- **Structured Scoring Performs Better:** Models may produce more useful severity assessments when evaluating individual CVSS attributes against defined criteria rather than directly generating a numerical score.

From Binary Vulnerability Detection to CWE Classification

Obinna Okeke of Bowling Green State University presented research examining how well LLMs move beyond identifying vulnerable code to assigning the specific CWE needed for root-cause classification.

- **Detection Outperforms Classification:** Models are more effective at determining whether code is vulnerable than identifying the precise CWE associated with the weakness.
- **Specific CWE Selection Remains Difficult:** Models may recognize the correct general weakness family while selecting the wrong specific classification.
- **AI Can Narrow Candidate Categories:** Current models may be useful for reducing the set of possible CWEs, but the results do not support relying on them as authoritative classifiers.
- **Classification Quality Matters for CVE Workflows:** Accurate root-cause classification requires greater precision than simply determining that a vulnerability exists.
- **Dataset Limitations Were Acknowledged:** Obinna's research relied on inherited CVE-to-CWE mappings from SecVulEval, which sourced metadata from NVD, and acknowledged that those mappings may contain noise.

Protecting CVE Quality While Reducing Maintainer Burden

Jessy Ayala of the University of California, Irvine, focused on the challenges volunteer and resource-constrained open-source maintainers face when evaluating vulnerability reports and described a maintainer-informed approach to improving early-stage triage.

- **Existing Processes Can Burden OSS Maintainers:** CVE workflows designed around organizations with dedicated security teams may be difficult for volunteer maintainers with limited security resources.
- **Low-Quality Reports Multiply Triage Work:** AI-assisted submissions can include inflated severity, weak evidence, or references to files and functions that do not exist in the affected project.
- **Project Context Improves Assessment:** Vulnerability assessment should account for project-specific security policies, scope, and definitions of severity rather than evaluating findings without maintainer context.
- **MINT Provides an Early Quality Gate:** Jessy's team is developing MINT, a maintainer-informed triage framework that checks whether referenced artifacts exist, evaluates report quality and severity in the context of the affected project, and produces maintainer-facing summaries for human validation. It is designed to serve as an accessible, transparent quality gate early in the CVE lifecycle without requiring major changes to existing standards.

- **Human Validation Remains Central:** The framework is intended to reduce maintainer workload and improve report quality without replacing human judgment or requiring major changes to existing CVE standards.

Interactive Community Listening Session Takeaways

Jay Jacobs, Empirical Security, and Bob Lord joined fellow panelist Lisa Olson of Microsoft to discuss participants' survey responses (see page 12 for detailed survey responses), revealing increased workload, mixed views on automation, and differing priorities between consumers, CNAs, and researchers. Alec Summers facilitated the discussion.

- **Workload and Automation Results:** Survey responses strongly agreed that AI has increased vulnerability-related workload for many participants, while the effect of automation on reducing manual work was mixed. Open-source maintainers in particular tended to report little reduction in manual effort.
- **Role-based Differences:** Researchers generally reported more positive effects from automation, while consumers and maintainers emphasized increased workload, prioritization pressure, and the need for better tools. The discussion reinforced that the impact of AI depends heavily on where a participant sits in the ecosystem.
- **Program Operations and Data Access:** Participants ranked CNA operations, the CVE record format, public CVE data access, and federation-related work as major areas for improvement. Several comments also highlighted that many upstream developers do not interact directly with the CVE Program and therefore need more practical support.
- **Data Quality and Prioritization Themes:** Participants also emphasized the key themes of data quality, interoperability, assignment, grouping, and the difficulty of changing record formats or workflows quickly. Some participants wanted stronger prioritization mechanisms, while others emphasized the need for stability given limited resources.
- **Interpretation of the Results:** The facilitators treated the polling as a beta test and said that the data would be analyzed further. The discussion made clear that the next steps would include a deeper review of the comments, broader survey work, and synthesis for future working-group efforts.

Follow-Up Actions

- **CVE Program Blog Sharing (Secretariat):** Share the CVE Program discussion blog link in the chat for participants who have not read it.
- **CVE Record Enrichment Feedback (Lisa Olson):** Follow up with participants who expressed concerns about CVE record enrichment to gather more detailed feedback for potential improvements.
- **Panelist Q&A Follow-Up (All panelists):** Review and respond to questions directed to panelists in the Q&A section after the session.
- **Survey Follow-Up (Secretariat):** Send a short survey to all attendees to collect additional feedback and inform future event planning.
- **Survey Expansion (Alec Summers, Jay Jacobs, TWG):** Plan and execute a broader survey to gather input from a wider audience, using this event as a beta test.

- **Blog Post and Data Summary (Secretariat):** Prepare and publish a blog post summarizing meeting feedback and survey data, including potential projects for working groups.
- **Working Group Project Identification (Lisa Olson, TWG):** Review survey and meeting feedback to identify specific projects for working groups to address.
- **Free Text Survey Analysis (Jay Jacobs/Bob Lord):** Perform detailed analysis of free text survey responses, including sentiment and extraction of specific quotes for future reference.

Listening Session detailed Survey Responses

1. Please rate the following:

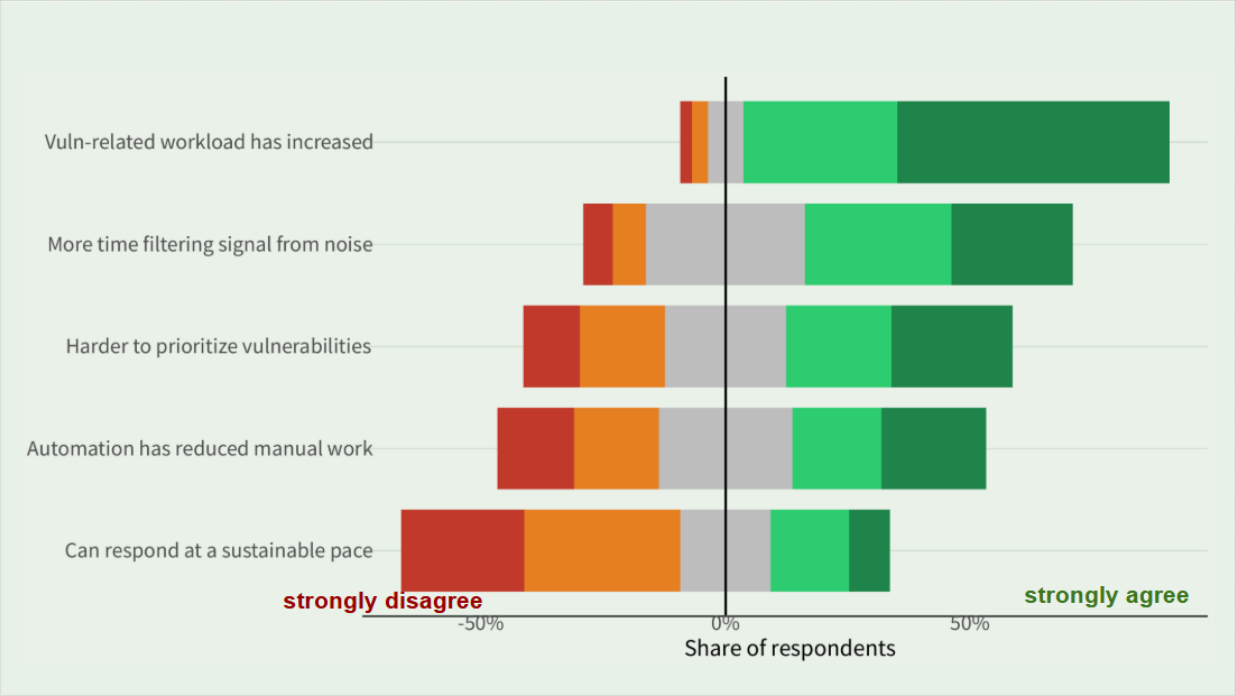

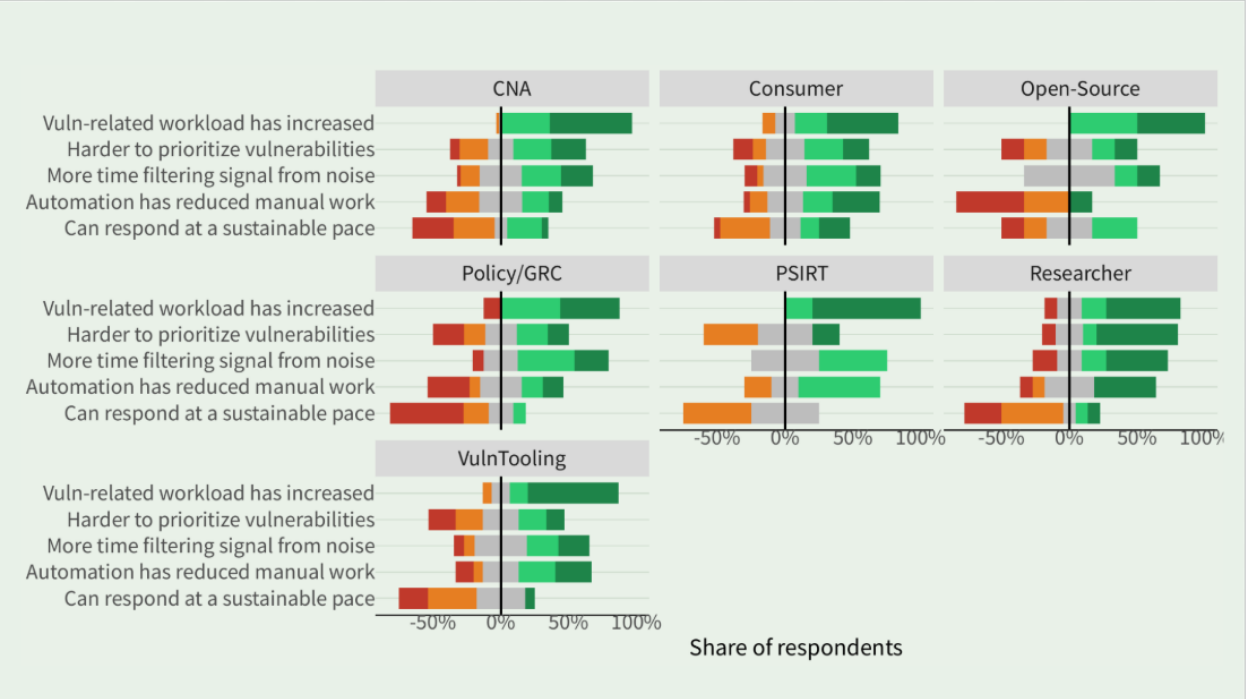

2. How would you characterize AI-related changes?

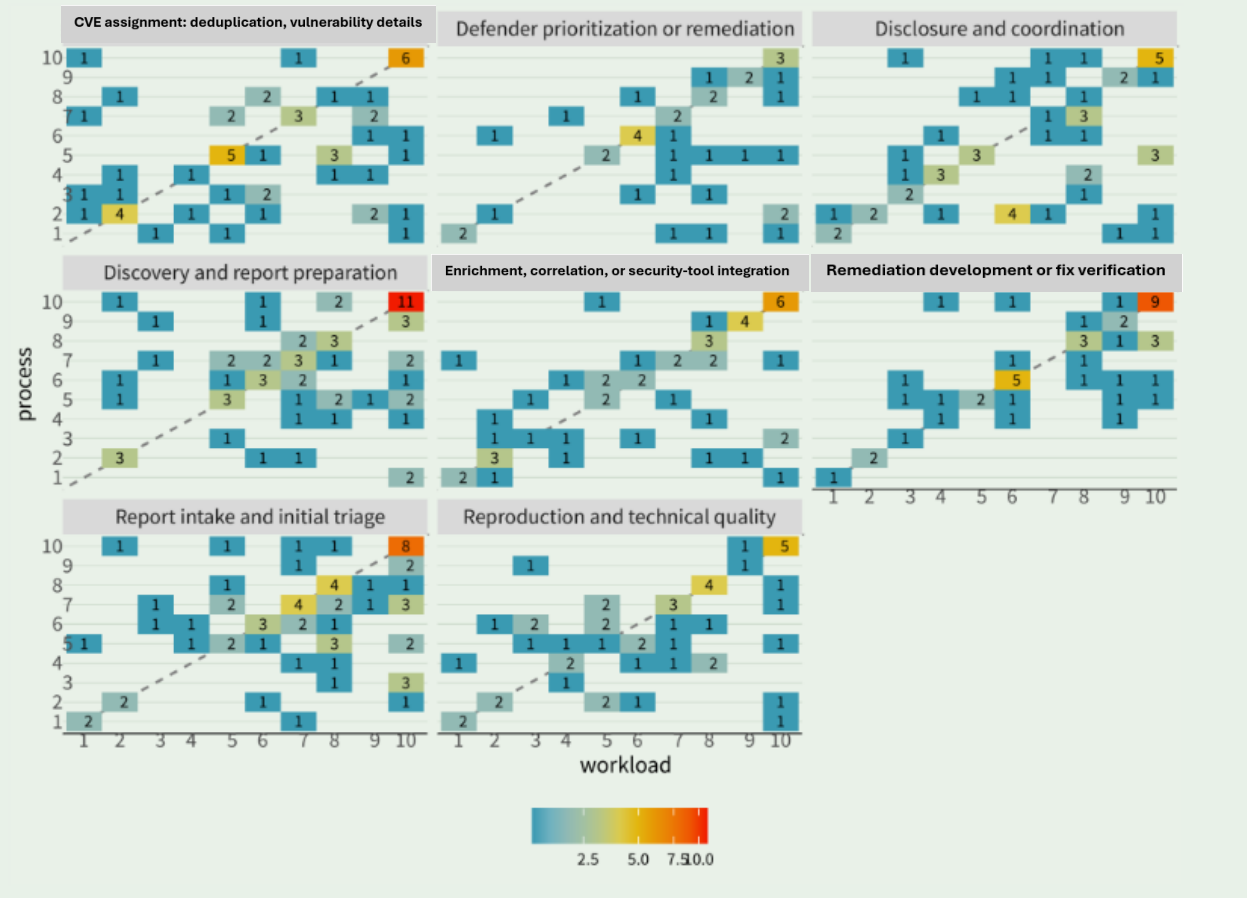

3. **Where is the most CVE Program work needed?**

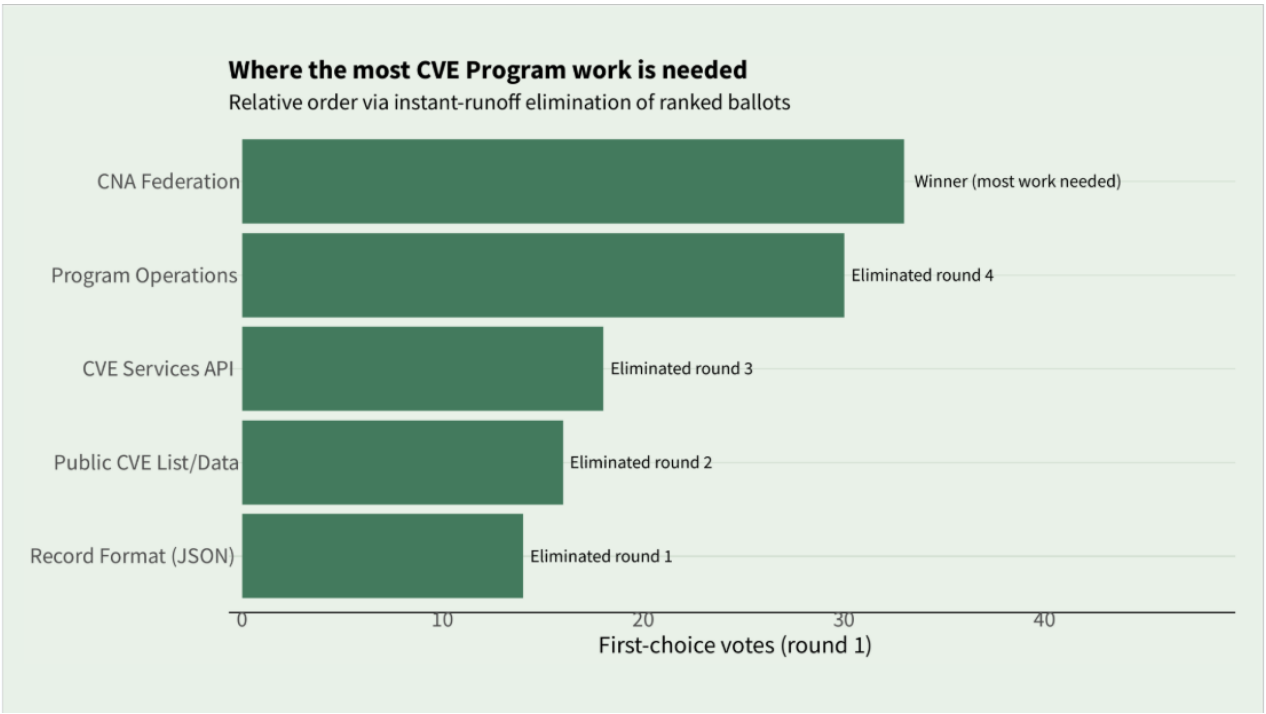

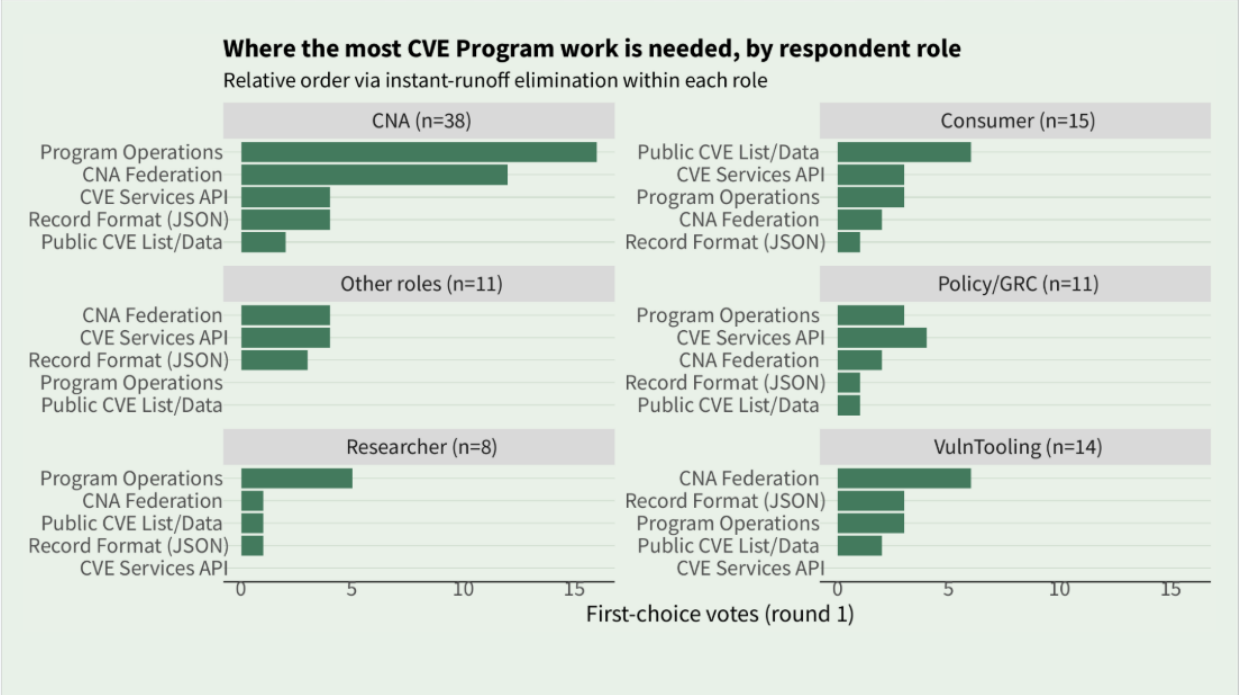

4. How much does each capability need to change?

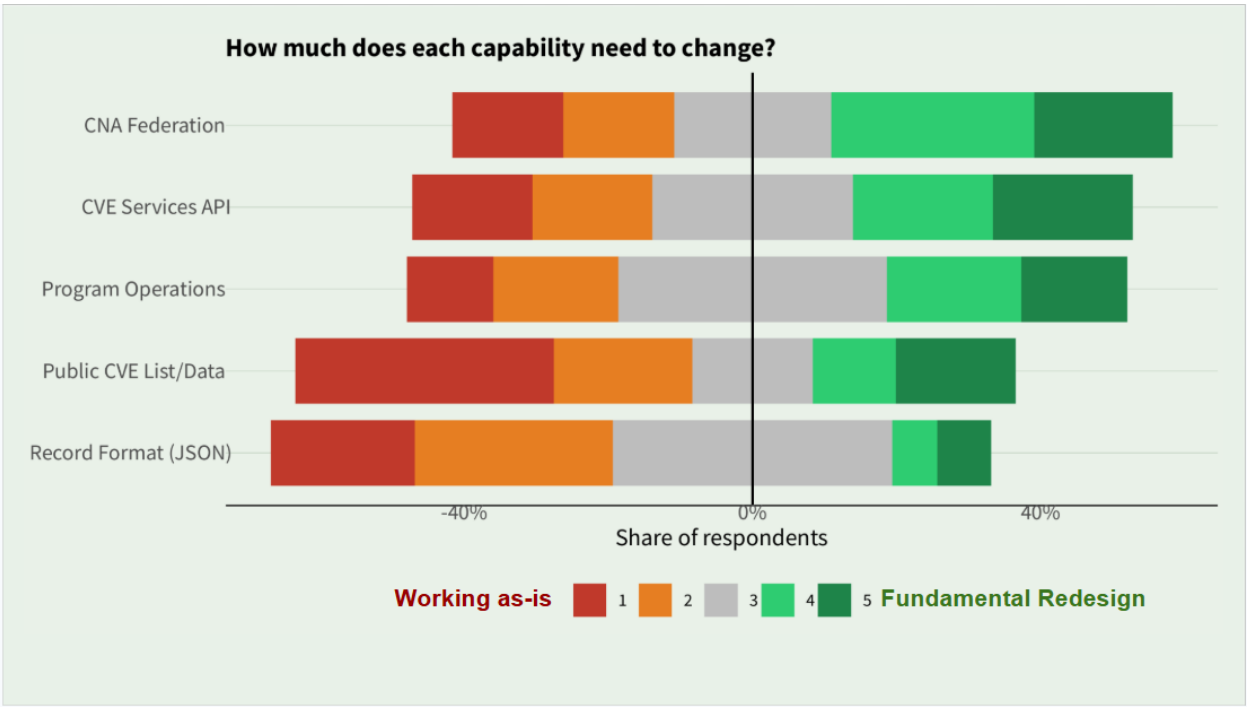

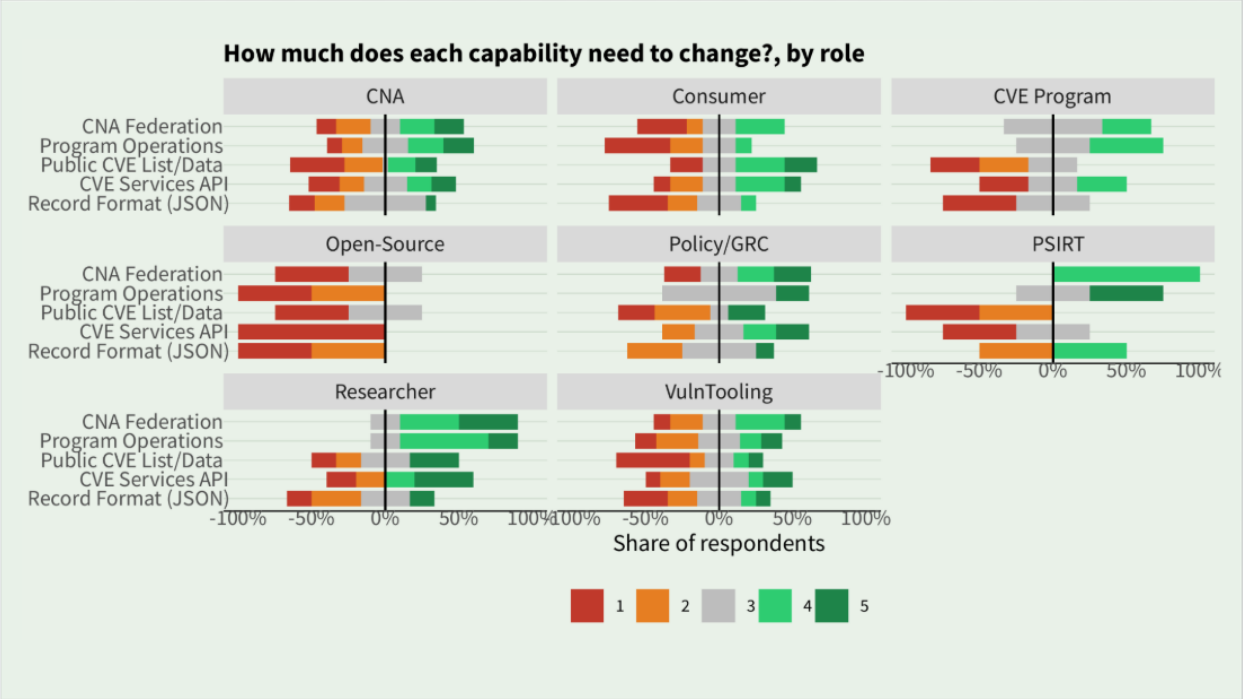

5. Describe any concrete AI-related changes you have observed (+ or -): what changed, what work it caused or removed, and what happened as a result.

Overall summary

The dominant story is that AI has sharply increased the volume of vulnerability discovery and reporting, but not in a cleanly manageable way. Most participants described a large rise in incoming findings, CVE requests, and third-party vulnerability intake, often paired with lower signal-to-noise, more duplication, more false positives, and more pressure on already-constrained triage, remediation, and coordination teams.

At the same time, many respondents also said AI is genuinely useful for analysis, correlation, CVE drafting, enrichment, reproduction, exploitability testing, translation, and workflow automation. In short:

- AI is creating more work by increasing discovery volume and report noise.
- AI is also helping absorb some of that work through automation and faster analysis.
- The net effect for many teams is still backlog, stress, and process redesign.

A smaller group reported no meaningful change yet.

What changed

The most commonly observed changes were:

1. Big increase in vulnerability/report volume

Participants described:

- 10x or greater increases in report volume
- Orders-of-magnitude increases in internal findings
- Huge spikes in CVE assignment/publication requests
- More third-party component CVEs to evaluate

- More proof-of-concept-driven reports
- More “chunky” report waves and vendor dumps with hundreds of issues

This was the clearest and most repeated theme.

2. Lower signal-to-noise in many reports

Many respondents said AI has made it easier for less-skilled reporters to submit plausible-looking but weak findings:

- More duplicates/collisions of findings
- More “slop submissions”
- More bug reports framed as vulnerabilities
- More reports with no demonstrated security impact
- More findings that ignore product trust or threat models
- Long, polished, AI-written disclosures that still aren’t valid vulnerabilities
- False positives from LLM-driven discovery

Several noted that obvious AI slop is disappearing; reports now often look professional even when incorrect.

3. Better offensive and defensive capability

AI is changing both discovery and response:

- Easier vulnerability discovery and source review
- AI-assisted pen-testing and code analysis
- Faster exploit and PoC development
- Better large-scale pattern recognition
- Faster CVE drafting and enrichment
- Better ability to test exploitability in labs

Some participants reported that AI-assisted exploits are becoming more effective and harder to detect.

4. New AI-specific risks

Some respondents pointed to risks introduced by AI adoption itself:

- Prompt attacks
- Data leakage
- Non-human identity / IAM challenges
- Tech debt from agentic IAM
- Supply-chain issues from immature AI libraries
- Increased AI-related findings and disclosures

What work AI caused

The main additional work created was:

1. Much more triage and validation

Teams now spend more time:

- Sorting signal from noise
- Checking whether findings are real vulnerabilities
- Determining CVE qualification
- Deduplicating repeated reports
- Validating AI-generated PoCs, exploit claims, and patches
- Arguing against machine-generated but incorrect reports

This was often described as the biggest burden.

2. More remediation and patching pressure

Participants reported:

- More issues to fix, often without added headcount
- Difficulty prioritizing because volume is too high
- Remediation windows exceeded
- Need to evaluate compensating controls when patches lag
- More work on sloppy or incomplete fixes
- Increased patch/disclosure coordination effort

Several said the volume is forcing teams to fix first and evaluate later.

3. More coordination and process overhead

AI has created operational strain across the vulnerability pipeline:

- Backlogs in CVE publishing and enrichment
- Delays in downstream review
- Scope conflicts across vendors/CNAs
- More difficult disclosure coordination
- Need for separate repos/queues for AI-submitted reports
- Need to coordinate shared disclosure across ecosystems

4. More management, customer, and policy communication

Respondents mentioned increased work responding to:

- Executive fear and confusion
- Customer questions about AI and vulnerability handling
- Social media noise and branded-flaw hype
- Pressure from CISOs and senior leadership
- Policy uncertainty around AI in secure development and vuln management

What work AI removed or reduced

Despite the strain, participants also described real efficiency gains.

Common areas where AI reduced effort:

- CVE record drafting and description writing
- CVE enrichment at scale

- Correlation and remediation analysis
- Reproduction of issues
- Initial analysis of third-party CVEs
- Pattern recognition across large datasets
- Security/risk compliance baseline reviews
- Translation and contact discovery during coordination
- Faster advisory writing, notifications, and documentation

Some teams said AI shifted analysts away from repetitive intake work and toward:

- Risk assessment
- Exception management
- Remediation coordination
- Strategic decision-making

What happened as a result

The most common outcomes were:

Negative outcomes

- Large backlogs in triage, publishing, and enrichment
- Burnout and fatigue across the pipeline
- Strained relations with researchers, maintainers, and open source communities
- Lower trust in reports and AI-generated fixes
- Delayed CVE issuance/publication
- Reduced ability to give each report proper attention
- More management anxiety and escalation
- Difficulty staffing and budgeting appropriately

Positive outcomes

- Faster analysis and issue reproduction
- Better exploitability testing
- More scalable automation
- Faster document generation and coordination support
- Better handling of large CVE datasets
- In some cases, improved report quality and better PoCs than earlier AI-era submissions

Key themes

1. Volume explosion is the defining change

This was the strongest consensus theme. Nearly every major downstream problem traces back to sharply increased discovery and reporting volume.

2. AI increased both capability and noise

Participants repeatedly described a dual effect:

- AI finds more real issues
- AI also generates more junk, duplication, and uncertainty

3. Triage is the main bottleneck

The biggest pain point is no longer just discovery. It is validating, prioritizing, deduplicating, and deciding what actually matters.

4. Remediation capacity is not keeping up

Even when findings are real, organizations often cannot fix them fast enough. AI is accelerating discovery faster than remediation teams can respond.

5. Automation is no longer optional

Many respondents described creating new repos, workflows, or AI-assisted pipelines specifically to cope with AI-driven scale. Several implied that "AI to assess AI reports" is becoming inevitable.

6. Trust has become harder

There is growing distrust of:

- AI-generated reports
- AI-generated fixes
- AI-produced severity/CVSS scoring
- AI use by novice researchers
- AI use in open-source coordination

At the same time, polished presentation no longer signals quality.

7. Human burnout is widespread

Burnout, fatigue, frustration, and loss of morale were mentioned directly and indirectly many times. This is not just a tooling issue; it is a workforce issue.

8. Executive attention is up, but understanding is uneven

Leaders are hearing about AI-related vulnerability trends, often through hype or second-hand narratives, which is increasing pressure on teams without always improving decision quality.

9. The ecosystem is unevenly prepared

Several comments pointed to gaps in:

- tooling
- staffing
- standards/categorization
- reporting consistency
- AI risk frameworks
- funding access to frontier models
- CNA coordination

10. Some quality is improving, but the net burden remains high

A notable minority said current AI-assisted reports are getting better:

- better language

- better PoCs
- less obvious slop
- more verifiable findings

But even these respondents usually still reported higher workload due to sheer volume.

Important tensions in the responses

There were a few clear contradictions worth noting:

Report quality: worse or better?

Both were reported.

- Many said quality has dropped sharply due to false positives, verbosity, and low-context reports.
- Others said quality has recently improved, especially PoCs and readability.

Best interpretation: early obvious slop is evolving into more polished submissions, which may be better written but still costly to validate.

Signal-to-noise: lower or higher?

Mostly lower, but not universally.

- Most respondents said noise increased.
- A few said signal-to-noise improved, even while backlog still grew.

Best interpretation: teams with stronger filters or more mature intake may be seeing better yields than those exposed to open submission channels.

Has anything changed?

A minority said “no meaningful change.”

Best interpretation: impact is uneven across roles, sectors, and maturity levels. Frontline triage/CVE coordination functions seem to feel the change most strongly.

Short takeaway

The core pattern is:

- AI has massively increased vulnerability discovery and reporting.
- That has created more triage, deduplication, coordination, remediation, and publishing work than most teams can absorb.
- AI also helps automate analysis and documentation, but not enough to offset the growth in volume.
- The result is a system under strain: more findings, more backlog, more burnout, more executive pressure, and a growing need for automated yet trustworthy workflows.

Details regarding the participants and submissions

- 102 submissions from 67 participants
- 20 participants submitted multiple responses
- A participant counts only once per theme, even across several submissions

Theme	Participants	Share
Data quality and interoperability	20	29.9%
Volume, speed, and operational capacity	17	25.4%
CNA governance and accountability	16	23.9%
Risk, actionability, and enrichment	16	23.9%
Scope of AI-related vulnerabilities	13	19.4%
Automation and tooling	13	19.4%
CVE assignment, reservation, and grouping	10	14.9%
Messaging, education, and AI skepticism	8	11.9%
Community guidance and coordination	8	11.9%
Human oversight and AI limitations	7	10.4%