Atlantic Council
CYBER STATECRAFT INITIATIVE

ai x ac

March 2026

Issue brief Securing cloud infrastructure for AI

Written by Sara Ann Brackett

With AI raising the stakes of cloud security and key cybersecurity institutions weakened or dissolved, this brief outlines needed policy steps to promote transparency and accountability across the cloud ecosystem.

Executive summary

Securing artificial intelligence (AI) infrastructure requires ensuring the security of the cloud ecosystem. The cloud infrastructure that implements and executes AI workloads presents an opening for adversaries that existing vulnerability management institutions were not designed to cover. This brief examines the mechanisms through which vulnerabilities in cloud infrastructure are discovered, disclosed, communicated, and remediated, and finds them to be inadequate to meet the security demands of an ecosystem in which AI has a growing impact.

Nation-state actors continue to target cloud environments, compressing vulnerability discovery and exploitation timelines. At the same time, public vulnerability data, anchored by the Common Vulnerabilities and Exposures (CVE) ID system and the linked National Vulnerability Database (NVD), faces severe strain. The policy institutions tasked with addressing cloud security face leadership vacuums, funding uncertainty, and competing priorities.

Community and industry driven efforts to respond to these challenges remain fragmented and voluntary, while providers operate without public accountability. Making progress on these urgent challenges requires policy mechanisms to incentivize and mandate clarity and transparency in the cloud ecosystem.

Background

In the United States, several essential cybersecurity authorities and institutions face simultaneous disruption. The Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA), passed in 2022, proposed to establish the country's first mandatory incident reporting regime for critical infrastructure sectors, has seen the publication of its final rule delayed to May 2026. The Cybersecurity Information Sharing Act of 2015 (CISA 2015), which provides liability and antitrust protections for companies sharing threat indicators with the federal government and each other, lapsed on September 30, 2025 and received only a temporary extension through September 2026, with bipartisan reauthorization efforts stalled in Congress. The Cybersecurity and Infrastructure Security Agency (CISA) continues to operate without a confirmed director, after the Senate failed to act on the administration's nominee, and the agency has seen workforce reductions that have diminished its operational readiness. The Cyber Safety Review Board (CSRB), which conducted a landmark investigation of the 2023 Microsoft Exchange Online compromise, was dissolved in early 2025.

In Europe, a set of untested regulatory instruments are taking effect. The European Union (EU) NIS2 Directive expanded cybersecurity obligations for 18 critical sectors. The Cyber Resilience Act (CRA), adopted in 2024 with enforcement beginning in 2027, will require digital product manufacturers to build in cybersecurity capabilities and provide vulnerability disclosure mechanisms. The United Kingdom (UK) Cyber Security and Resilience Bill is progressing through Parliament with similar objectives.

The US, UK, and EU have also adopted policy approaches and established agencies specific to AI. In the US, the Trump administration's AI Action Plan outlined goals of exporting the US AI stack abroad. The EU AI Act imposed obligations on AI developers based on the level of risk posed by specific AI models. The US and UK also established AI-specific testing and research organizations, the Center for AI Standards and Innovation (CAISI) and the AI Security Institute (AISI) respectively.

Cloud computing

Cloud computing describes a model of access to computing resources, where customers specify workloads, or defined sets of tasks, which cloud providers implement and execute. This model of access is an important part of AI development and deployment, and frontier AI companies have partnered with cloud providers to ensure access to cutting-edge compute resources.

Compute and virtualization services allocate processing power and include the orchestration platforms that oversee AI training and inference workloads. Data and storage services comprise the managed databases and object storage for training datasets, model weights, and inference outputs. Observability and logging services collect the telemetry essential to detecting anomalies and investigating incidents. Identity and access management services control who and what can interact with cloud resources.

Layered on top of these foundational services are AI-specific runtimes and serving frameworks (the managed environments in which models are loaded and scaled) as well as the web and API gateways through which users interact with AI systems. Each of these categories presents distinct vulnerabilities. A flaw in a container escape mechanism raises different remediation questions than a misconfiguration in a logging pipeline, yet both can impact the confidentiality, integrity, and availability of an AI workload.

Fraying public vulnerability infrastructure

The NVD has served for nearly two decades as an authoritative source for enriched vulnerability data, powering compliance frameworks, automated scanning tools, and risk assessments across both the public and private sectors. Budget constraints and rising submission volumes have degraded its reliability as an operational resource. The National Institute of Standards and Technology (NIST) acknowledged in early 2025 that a 32 percent increase in CVE submissions during 2024 meant the backlog was still growing.

CISA maintains the Known Exploited Vulnerabilities (KEV) catalog, which allows the agency to publicly announce vulnerabilities that have been exploited in the wild. As of March 2026, the catalog contains 1,551 vulnerabilities, making it a useful smaller-scale prioritization signal, especially in comparison to the NVD's 339,010 vulnerabilities. CISA's ability to update and maintain the KEV database is likely affected by the ongoing partial shutdown of CISA's parent agency, the Department of Homeland Security, and the mass layoffs of CISA employees since January 2025.

Meanwhile, despite a clear recommendation from the CSRB's review of the 2023 Microsoft Exchange incident, cloud providers do not comprehensively disclose security vulnerabilities or flaws within their cloud services that do not require customer action to fix. In 2024, both Microsoft and Google announced that they would issue CVEs for critical vulnerabilities, which are only a subset of the overall vulnerability landscape. Vulnerability scoring and severity evaluations are complex and involve judgement calls, so allowing providers to determine which vulnerabilities they disclose distorts publicly available data on cloud security issues.

Failing to issue a CVE identifier for a security flaw also precludes the vulnerability from being included in the KEV database, limiting the ability of US government agencies to publicly communicate evidence of exploitation. Companies can refuse to acknowledge security incidents or transparently communicate with customers in the absence of policy obligations, as Oracle's communications around an incident in May 2025 exemplified.

Hyperscale cloud providers also operate vulnerability reward programs (VRPs) that incentivize external researchers to report flaws. According to a program website, Google's Cloud VRP has issued $3,574,399 in awards over the past year. These programs are voluntary, variable in scope and payout, depend on the communications channels offered by the cloud provider, and are not subject to public reporting obligations.

Provider programs are siloed. No mechanism exists for identifying shared flaws across cloud platforms or generating a system-wide view of collective vulnerability data. This limitation is consequential considering research demonstrating that

independently developed cloud services can harbor similar security flaws due to shared open-source dependencies or common architectural patterns. The absence of cross-provider coordination means that when a researcher identifies a vulnerability pattern in one cloud platform, there is no systematic process for evaluating whether the same pattern exists in others.

AI services are not immune from these systemic challenges. A July 2025 container escape vulnerability in the NVIDIA Container Toolkit, discovered by Wiz researchers, highlighted that security issues in popular libraries impact customers regardless of their cloud provider.

Community-driven projects have attempted to address the lack of standardized tracking mechanism for cloud security issues. The Wiz-backed Open Cloud Vulnerability and Security Issue Database catalogs publicly known cloud vulnerabilities and flaws, providing researchers and practitioners with a centralized reference for flaws that might otherwise be scattered across notification methods. The ONUG Cloud Security Notification Framework addresses the lack of a common data model for security notifications across providers. While these efforts are valuable, neither possesses the institutional backing to compel provider participation or to generate the kind of systematic accounting of vulnerabilities in cloud platforms which could form the basis of further policy action.

AI changes the risk landscape

As a target, AI infrastructure concentrates extraordinarily valuable intellectual property within cloud environments: model weights, proprietary training data, novel research methods, and fine-tuning configurations, all of which are only as secure as the weakest component of their infrastructure. The scarcity of compute resources specific to AI may lead organizations to deprioritize security requirements in favor of rapid access to processing power. The emergence of AI-focused cloud providers, newer entrants that may lack the mature security operations and vulnerability management programs of established hyperscale cloud providers, creates additional points of systemic risk.

As a tool, AI is reshaping the vulnerability landscape on both the offensive and defensive sides, rapidly accelerating the pace of vulnerability discovery and exploit development. Google's Project Zero reported 20 vulnerabilities in popular open-source packages, each of which was discovered and reproduced by an AI agent without human intervention. A similar collaboration between Mozilla and Anthropic discovered 22 vulnerabilities in Firefox and crafted partial exploits for each of them. Wiz's first-ever cloud hacking competition surfaced over 11 vulnerabilities in open-source code comprising foundational layers of cloud infrastructure. Open-source maintainers and operators of bug bounty programs have raised alarm about the increasing volume of AI-generated bug reports, which are of varying quality and require significant effort on the part of developers and maintainers to evaluate.

Recommendations

Government agencies must respond to the changing cloud vulnerability landscape. As experts have warned, failing to keep pace with the rapid rate of developments in offensive cyber risks of AI will have security consequences. Managing the risks posed by AI for vulnerability discovery and exploitation requires recommitting to known best practices, which can serve as a foundation for future policy experimentation and adaptation.

1. Follow through on lapsed and languishing cybersecurity efforts

Congress should reauthorize the Cybersecurity Information Sharing Act of 2015, which remains the foundational legal framework enabling voluntary cyber threat intelligence sharing between the private sector and federal government. Temporary extensions do not provide sufficient assurance and protection to organizations committing to information sharing, and even brief lapses disrupt long-standing collaborations.

Consistent with the AI Action Plan, the federal government should establish a dedicated Artificial Intelligence Information Sharing and Analysis Center (AI-ISAC) to centralize threat intelligence specific to AI systems, model vulnerabilities, and adversarial exploitation techniques. This body could facilitate real-time coordination across industry, academia, and government, led by DHS in collaboration with CAISI and the Office of the National Cyber Director (ONCD).

Congress should ensure CISA and NIST, as the stewards of the KEV catalog and NVD, receive sustained, adequate resourcing to fulfill their role in the vulnerability management ecosystem.

Congress should re-establish the Cyber Safety Review Board, resolving issues with the original board's investigation by giving the board both subpoena power and sufficient staff to support critical investigations. Congress should also clarify criteria for incidents reviewable by the board. As recent analysis in Lawfare argued, a review board specific to AI could investigate the role of AI in cyberattacks. An AI-specific body should also be scoped to include cloud security incidents, reflecting the cloud's critical role as AI infrastructure.

2. Incentivize and disclose high-quality public vulnerability data for cloud computing

ONCD should lead on establishing a comprehensive, government-backed information and data sharing solution to drive more effective vulnerability management across the cloud ecosystem. Policy design in cloud cybersecurity suffers from a lack of high-quality public data on critical and non-critical vulnerabilities; patterns of misconfiguration; and trends in exploitation techniques by threat actors.

ONCD's leadership on this challenge, as a component of the National Cybersecurity Strategy's goal to shape adversary behavior, could improve collaboration with the private sector and cloud providers, while avoiding diverting CISA from its core mission of protecting government and critical infrastructure systems.

Greater data transparency and disclosure of vulnerabilities across the ecosystem could serve as the foundation of prioritization processes across the cloud ecosystem, increasing pressure on providers to tackle vulnerabilities and classes of vulnerabilities that have security consequences for government entities and companies worldwide. That prioritization should privilege shared vulnerabilities, architectural flaws, and common weaknesses present across multiple hyperscale providers, which no current mechanism systematically or publicly identifies and addresses.

3. Lead on international coordination

The US government should pursue alignment with the European Union Agency for Cybersecurity and allied governments on cloud vulnerability disclosure norms. The US government's support of vulnerability databases and coordination efforts creates benefits for other countries. Questions about the stability of that support spur divergent efforts, such as the European Union Agency for Cybersecurity establishing its own database for cataloging cyber vulnerabilities.

Ensuring that AI safety institutions and cybersecurity agencies share information and coordinate on vulnerability management, rather than operating in parallel silos, should be an explicit element of efforts to mitigate the risks of emerging AI for cloud computing security.

The United States and the United Kingdom should begin by harmonizing their own practices and then extend that alignment to the EU. The forthcoming CIRCIA rule and the UK's Cyber Security and Resilience Bill offer opportunities to embed cloud-specific vulnerability and incident reporting requirements that can serve as reference points for international coordination. Encouraging international allies to adopt the same approach to disclosing vulnerabilities in compute infrastructure can contribute to changing the incentive structure of the AI compute industry, shifting it towards greater transparency from cloud providers.

Conclusion

Trust in cloud computing cannot be sustained without visibility. The physical location of a data center does not determine its vulnerability to misconfigured access controls, unpatched container runtimes, or supply chain compromises. The cloud infrastructure that underpins AI development and deployment is subject to a vulnerability management regime designed for a different era of computing. Cloud-specific security flaws fall between existing institutional mandates, and the organizations building the most consequential AI systems lack the ability to demand transparency from the infrastructure they depend on.

The building blocks for a better approach exist, but the policy architecture to connect them is missing. The United States and its allies and partners possess both the responsibility and the capacity to design an approach to cloud vulnerability management that matches the scale and complexity of the systems it is meant to protect. The question is whether they will do so before the gap between the complexity of cloud infrastructure and the maturity of the institutions overseeing it becomes the defining vulnerability of the AI era.

About the author

Sara Ann Brackett is an associate director with the Cyber Statecraft Initiative, part of the Atlantic Council Tech Programs. She focuses her work on open-source software security, cloud computing, and software supply-chain risk management within the Cyber Statecraft Initiative's cybersecurity and policy portfolio.

Acknowledgments

The author would like to thank the Cyber Statecraft Initiative and Atlantic Council Tech Programs teams for their support and guidance throughout this project. Thank you to Trey Herr and Tess deBlanc-Knowles, who provided thoughtful feedback, and to Safa Shahwan Edwards and Jen Roberts, who were instrumental in planning and executing a workshop that informed this paper. Thank you to Nikita Shah, whose feedback shaped earlier iterations of this brief and its accompanying visualizations. The author would also like to thank the workshop participants who shared their expertise and perspectives under Chatham House Rule.

About the center

The **Cyber Statecraft Initiative** works at the nexus of geopolitics, technology, and security to craft strategies to help shape the conduct of statecraft and to better inform and secure users. This work extends through the competition of state and non-state actors, the security of the internet and computing systems, the safety of operational technology and physical systems, and the communities of cyberspace. The Initiative convenes a diverse network of passionate and knowledgeable contributors, bridging the gap among technical, policy, and user communities.

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

Atlantic Council
1400 L Street NW, 11th Floor
Washington, DC 20005
(202) 778-4952
www.AtlanticCouncil.org