Congress of the United States
House of Representatives
Washington, DC 20515

April 29, 2026

Mr. Michael Truell
Co-Founder and Chief Executive Officer
Anysphere, Inc.
2261 Market Street STE 86466
San Francisco, CA 94114

Dear Mr. Truell:

The House Committee on Homeland Security and the House Select Committee on the Chinese Communist Party (the Committees) are conducting a joint investigation into the growing national security risks created by the integration of People's Republic of China (PRC)-developed artificial intelligence (AI) models into the software tools and development environments on which American enterprise, government, critical infrastructure, and national defense increasingly depend.

As part of this investigation, the Committees are examining a pattern of conduct by PRC-based AI laboratories involving the large-scale theft of proprietary capabilities from American frontier AI systems through adversarial distillation, the redistribution of those stolen capabilities as open-weight models available for global download, and the incorporation of PRC-origin models into products used daily by hundreds of thousands of American developers and engineers. Alarmingly, this conduct is part of a broader PRC campaign to accelerate its AI capabilities through the exploitation of American innovation, including through espionage, intellectual property theft, and other unlawful or deceptive means.

In February 2026, multiple leading American frontier AI laboratories disclosed, both to Congress and publicly, that three PRC-based AI companies, specifically DeepSeek, Moonshot AI, and MiniMax, had carried out coordinated campaigns to extract advanced capabilities from American AI systems through adversarial distillation.[1] According to those disclosures, the three PRC laboratories collectively generated more than 16 million exchanges with American AI systems through approximately 24,000 fraudulent accounts, deliberately circumventing regional access restrictions and terms-of-service prohibitions.[2] The targeted capabilities included some of the most commercially valuable and security-sensitive functions American systems offered, including agentic reasoning, autonomous tool use, and advanced software development.[3] The PRC laboratories also reportedly routed their operations through commercial proxy networks to

[1] *Memo from OpenAI to the House Select Committee on the Chinese Communist Party* (Feb. 12, 2026); *see also* Anthropic, *Detecting and Preventing Distillation Attacks*, ANTHROPIC BLOG (Feb. 23, 2026).
[2] *Id.*
[3] *Id.*

evade geographic restrictions, reconstituted terminated accounts within hours, and, in at least one documented case, redirected half their extraction traffic to a newly released American model within a single day of its launch.[4]

The billions of dollars American companies invest in foundational research, compute infrastructure, and security engineering is being undercut by a sustained extraction campaign conducted at a fraction of the cost of independent development. This threat is not limited to commercial harm. American frontier AI laboratories invest heavily in security testing and in building guardrails designed to prevent their models from being used to develop weapons, automate software vulnerability discovery and exploitation, generate tailored disinformation, or assist in the synthesis of dangerous chemical or biological agents.

When capabilities are stripped out through distillation and repackaged without equivalent safeguards, the resulting models may become available to hostile state actors, terrorist organizations, and criminal enterprises. Once released as open-weight, those capabilities cannot readily be recalled and may proliferate beyond the reach of any single company's contractual controls or any single government's enforcement authorities. The Trump Administration reached the same conclusion on April 23, 2026, when the White House Office of Science and Technology Policy issued a memorandum formally characterizing these campaigns as deliberate, industrial-scale operations and warning that distillation-derived models allow PRC actors to "deliberately strip security protocols from the resulting models and undo mechanisms that ensure those AI models are ideologically neutral and truth-seeking."[5]

On March 19, 2026, Anysphere, Inc., through its Cursor product, released a new model called Composer 2 and marketed it as offering "frontier-level coding intelligence," with benchmark results that reportedly exceeded those of leading American models at roughly one-tenth the inference cost.[6] Anysphere described the model as the product of "continued pre-training of a base model combined with reinforcement learning," but did not identify the underlying base model at launch.[7] Within hours, an independent developer examining Cursor's Application Programming Interface (API) traffic discovered an internal model identifier indicating that Composer 2 was built on Kimi K2.5, an open-weight model released in January 2026 by Moonshot AI, a Beijing-based company backed by Alibaba Group.[8] Anysphere confirmed the model's origin only after that independent discovery.[9] Co-founder Aman Sanger subsequently acknowledged that it was "a miss to not mention the Kimi base."[10] Moonshot AI is

---

[4] *Id.*
[5] Memorandum from Michael J. Kratsios, Assistant to the President for Science and Technology, to the Heads of Executive Departments and Agencies, NSTM-4, Adversarial Distillation of American AI Models (Apr. 23, 2026), https://whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf.
[6] Cursor, *Introducing Composer 2,* CURSOR BLOG (Mar. 19, 2026).
[7] Sumit Pandey, *A $29B Startup Got Caught. A Developer, an API Call, and 24 Hours.*, MEDIUM (Mar. 26, 2026).
[8] *Id.*
[9] Mannat Dora, *Cursor Founder Clears Air on Kimi Model Use in Composer 2: Here's All You Need to Know*, THE ECONOMIC TIMES (Mar. 21, 2026).
[10] *Id.*

one of the three PRC-based laboratories publicly implicated in the industrial-scale distillation campaigns described above.[11]

Recent reporting further underscores that even Anysphere appears to recognize that AI assisted software development can magnify software supply chain risk if underlying dependencies are not tightly controlled.[12] On April 21, 2026, Cursor announced a partnership with Chainguard, an open-source security company, to steer AI generated code toward vetted open-source components and reduce the risk that developers unknowingly pull vulnerable or malicious libraries and container images into production environments.[13] That development is notable because it reflects an apparent acknowledgment by Cursor that agentic and "vibe coded" development can cause dependency selection and package inclusion decisions to occur at a scale and speed that outpaces ordinary human review, and because it highlights that the security of an AI coding environment depends not only on the model itself, but also on the provenance and integrity of the packages, libraries, and images the system recommends, retrieves, or incorporates into downstream software.[14] In environments handling sensitive government, defense-industrial, or critical infrastructure code, those software supply chain risks carry obvious national security implications.

As you know, Cursor is one of the most widely adopted AI-powered software development environments in the United States, with more than one million daily active users, more than 50,000 business customers, and adoption by more than half of the Fortune 500.[15] An AI-powered coding tool, by its nature, may operate with access to categories of proprietary and security-sensitive information that few other software products routinely process. When a developer uses Cursor, the tool may ingest the code being written, surrounding project files, conversation history, and indexed portions of a broader codebase. Depending on the use case, that information may include security architectures, cryptographic implementations, authentication and access-control logic, vulnerability remediation code, trade secrets, and other proprietary business logic.

According to Anysphere's public security materials, relevant Cursor systems are covered by SOC 2 Type II certification.[16] At the same time, Anysphere's public materials do not indicate that Cursor has obtained FedRAMP authorization or that it meets the requirements that may apply in environments handling controlled unclassified information, export-controlled technical data, or other similarly sensitive government-related information, including environments subject to NIST SP 800-171 requirements, CMMC-related obligations, or ITAR- and EAR-related restrictions.[17] Given the nature of the code and technical data that AI-powered software development tools may access, that limited public compliance posture raises questions about the

---

[11] *Supra* note 1.
[12] Sam Sabin, *Exclusive: Cursor Taps New Security Partner in Push to Secure Vibe Coding*, AXIOS (Apr. 2026).
[13] *Id*.
[14] *Id*.
[15] Rachel Metz, *AI Coding Assistant Cursor Draws a Million Users Without Even Trying*, BLOOMBERG (Apr. 7, 2026).
[16] Cursor, *Security*, CURSOR, https://cursor.com/security (last visited Mar. 9, 2026).
[17] *Id.*

circumstances in which Cursor is being marketed, adopted, and used, particularly in certain government, defense, and critical infrastructure-related environments.

These issues are unfolding against a broader and accelerating strategic shift. PRC-developed open-weight AI models have experienced rapid global adoption. In late 2024, PRC models reportedly accounted for an estimated one percent of global AI workloads.[18] By the end of 2025, that share had reportedly grown to an estimated 30 percent.[19] PRC-developed models also have been reported to exhibit censorship aligned with Chinese Communist Party (CCP) positions on politically sensitive topics, and federal testing has found that leading PRC models echoed CCP-approved narratives at rates far exceeding comparable American systems.[20] Beijing has been explicit in its view that the global distribution of open-weight AI serves PRC strategic interests.[21] What is at issue, therefore, is not simply market competition, but the growing risk that software systems used across the American economy, government, and defense industrial base will come to depend on models developed by PRC-linked laboratories and shaped by PRC strategic objectives.

Accordingly, to inform the Committees' joint investigation into the national security risks associated with the theft of American AI capabilities, the downstream deployment of PRC-developed open-weight models, and the integration of such models into software tools used across the American economy, including in connection with government, defense-industrial, and critical infrastructure-related work, the Committees request that Anysphere, Inc. provide the following documents and information no later than May 13, 2026:

1. All documents and communications relating to any direct or indirect relationship, partnership, licensing arrangement, technical collaboration, data-sharing agreement, joint development effort, or financial transaction between Anysphere and any PRC-based entity, or any entity owned, controlled, or materially funded by a PRC-based entity, including but not limited to Moonshot AI, DeepSeek, MiniMax, Alibaba Group (including Alibaba Cloud), Zhipu AI, ByteDance, Tencent, Baidu, and each of their subsidiaries, affiliates, investors, and agents.

2. All documents and communications relating to the selection, evaluation, and integration of Moonshot AI's Kimi K2.5 model as the base for Composer 2, including the following:

    (a) All models evaluated as alternatives, and the criteria used to compare them.
    (b) All internal analyses, memoranda, or presentations comparing PRC-origin and non-PRC-origin models on the basis of performance, cost, safety, or any other dimension.
    (c) All assessments, formal or informal, of the legal, reputational, national security, or supply chain risks associated with building on a base model of PRC origin.

---

[18] Farhan Kabir, *Open-Source AI: Why China is Winning the Global Adoption Race*, MEDIUM (Feb. 18, 2026).
[19] *Id*.
[20] NIST, *CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks*, NIST (Sep. 30, 2025).
[21] Owen J. Daniels and Hanna Dohmen, *China's Overlooked AI Strategy: Beijing Is Using Soft Power to Gain Global Dominance*, FOREIGN AFFAIRS (Jul. 25, 2025).

(d) All communications, whether internal or external, regarding PRC national intelligence law, cybersecurity law, data security law, export control law, or related legal obligations.
(e) All communications, including communications among officers, directors, employees, and board members, regarding whether and when to disclose the model's provenance to users, customers, enterprise accounts, investors, or regulators.
(f) All communications with Moonshot AI, Fireworks AI, or any other intermediary relating to the licensing, inference hosting, or commercial terms of the Kimi K2.5 integration.

3. A complete technical description of every data pathway through which user-generated content is transmitted, processed, cached, stored, logged, made available, or otherwise handled in connection with Anysphere's Cursor-branded AI coding products and services, including source code, natural-language prompts, conversation history, file metadata, codebase embeddings, and telemetry, including:

   (a) A description of the full lifecycle of such data from the point of local collection, access, or transmission through final disposition, including any routing, preprocessing, inference, storage, monitoring, safety, retention, or product- or model-improvement functions.
   (b) Identification of all entities involved in that processing, whether Anysphere, any affiliate, or any third party acting on Anysphere's behalf.
   (c) Identification, for each such entity, of its name and corporate parentage, its country of incorporation, the physical location of all servers or systems on which the relevant data is processed or stored, and whether the entity, or any corporate parent of that entity, is subject to PRC jurisdiction under any applicable law.

4. All documents relating to any security audit, penetration test, red-team evaluation, model integrity assessment, supply chain risk analysis, or backdoor detection effort conducted on any PRC-developed model integrated into Cursor's products, whether performed internally, by a third-party firm, or by any government entity. If no such assessment has been conducted with respect to any such model, provide the reasons no such assessment was undertaken.

5. Copies of all agreements between Anysphere and every model inference provider, cloud infrastructure provider, API routing service, and subprocessor that handles, or has access to, user-submitted code or data, including all data retention, data processing, and zero-data-retention agreements.

   (a) For each such agreement, identify every contractual provision, exception, or carve-out under which user code or data may be retained, logged, cached, used for model training or improvement, or made accessible to any third party, and identify whether any such agreement permits the provider to use such data for any purpose beyond real-time inference.

6. A description of all steps Anysphere has taken, or intends to take, to accomplish each of the following:

    (a) Ensure that no user code or data is transmitted to, processed on, or accessible from, any server or system located in the PRC or operated by any entity subject to PRC jurisdiction.
    (b) Provide users and enterprise customers with real-time, affirmative disclosure of the identity and country of origin of every model processing their data at the time of processing.
    (c) Independently verify the integrity of all base models for the presence of backdoors, poisoned weights, hidden data exfiltration mechanisms, or deliberately engineered vulnerabilities prior to deployment.
    (d) Determine, for each relevant product offering and customer environment, whether and to what extent the product may be used in contexts subject to federal security, defense-industrial, or export-control requirements, including FedRAMP, NIST SP 800-171, CMMC-related obligations, and ITAR- and EAR-related restrictions, and describe any steps Anysphere has taken to notify customers of the product's current compliance posture and limitations.
    (e) Disclose to current and prospective customers the full provenance of every AI model integrated into Cursor's products, including the identity and nationality of the developer and any relationship between the developer and a foreign government or state-affiliated entity.

The Committees further request that appropriate personnel from Anysphere appear for an in-person briefing on these matters, including the issues identified in this letter and Anysphere's response thereto, no later than May 20, 2026.

To arrange document production, coordinate the requested briefing, discuss the scope of these requests, or address any questions regarding this letter, please contact majority staff of the Committee on Homeland Security at (202) 226-8417 and majority staff of the Select Committee on the Chinese Communist Party at (202) 226-1541. The Committees request that Anysphere raise any questions or concerns promptly so that they may be resolved without delaying compliance.

Pursuant to Rules X and XI of the U.S. House of Representatives, the Committee on Homeland Security has jurisdiction over homeland security policy and oversight of "all Government activities relating to homeland security, including the interaction of all departments and agencies with the Department of Homeland Security." House Resolution 5 grants the Select Committee investigative jurisdiction over matters relating to "countering the economic, technological, security, and ideological threats of the Chinese Communist Party to the United States and allies and partners of the United States." Upon receipt of this letter, please preserve all hard copy and electronic documents and communications related to its subject matter.

Sincerely,

ANDREW R. GARBARINO
Chairman
Committee on Homeland Security

JOHN MOOLENAAR
Chairman
Select Committee on China

cc: The Honorable Bennie Thompson, Ranking Member
Committee on Homeland Security

The Honorable Ro Khanna, Ranking Member
Select Committee on China