


DeepSeek FASCSA Order: What "Developed Using" Means for AI Supply Chains

Security Implications and Guidance for Federal Agencies, Contractors, and AI Product Vendors

2026-10-21

 AI-assisted Rapid Research



CSAI

cloud
security
alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

The Department of Homeland Security (DHS) published a Federal Acquisition Supply Chain Security Act (FASCSA) exclusion and removal order on October 20, 2026. It directs federal civilian executive branch agencies to stop procuring products and services from Hangzhou DeepSeek Artificial Intelligence and its affiliates. It also requires agencies to remove those products, together with "third-party products and services developed using DeepSeek implementations," from agency information technology [1][2]. It is the first FASCSA order issued by DHS and the first to cover all federal civilian executive branch agencies [14]. The two earlier FASCSA actions were the September 2025 ODNI order against Acronis and the Department of War's March 2026 designation of Anthropic. Unlike both, this order reaches explicitly past the named source into third-party products that other vendors built on its technology [8][14][23].

The derivative clause is the operative risk for industry. The order never defines "DeepSeek implementations," and the phrase could plausibly reach self-hosted DeepSeek open weights, fine-tunes and quantizations of those weights, and possibly models distilled from DeepSeek outputs. In practice, vendors will likely field the questions about whether their products fall inside the clause. FAR 52.204-30 limits contractors' formal duty to a "reasonable inquiry" of information already in their possession, but agency customers may ask for broader attestations than the regulation requires [6]. The timeline compounds the uncertainty. Secretary Markwayne Mullin signed the order on July 24, 2026, and agency exception requests were due within 60 days of issuance. That deadline passed about four weeks before the order appeared in the Federal Register [1][14]. Vendors have no exception route under the order itself, although contractors may disclose and request a waiver under FAR 52.204-30 when the order is added to a contract by modification [6].

Detection is uneven across the supply chain. Model inventories and AI bills of materials can show whether DeepSeek weights or endpoints are running in a product. Few organizations keep records that can show whether DeepSeek outputs were used for distillation, synthetic training data, or evaluation [14]. Where FAR 52.204-30 is in a contract, it requires contractors to search SAM.gov for applicable FASCSA orders, report covered articles they find within three business days, and flow the clause down to subcontractors [6]. Commercial organizations outside the order's legal scope should still act. Much of the order's reasoning, particularly the data-flow risk of PRC-hosted services, applies equally to regulated enterprises, and its "developed using" standard may become a template for future AI supply chain controls.

Background

The FASCSA Authority

Congress enacted FASCSA in December 2018. The law created the Federal Acquisition Security Council (FASC), an interagency body that assesses information and communications technology supply chain risk. The FASC can recommend that a product or source be excluded from federal procurement or removed from federal systems. Under 41 U.S.C. § 1323, three officials may issue orders on those recommendations, each for a different part of government. The Secretary of Homeland Security issues orders for civilian agencies, the Secretary of Defense for defense and most national security systems, and the Director of National Intelligence for the intelligence community and sensitive compartmented information systems [3]. The FASC finalized its implementing rule at 41 CFR Part 201-1 in August 2021 [4][5]. The Federal Acquisition Regulation then added contractor-facing clauses at FAR 52.204-29 and 52.204-30 through an interim rule that took effect in December 2023 [6][7].

No FASCSA order was issued until September 2025. That month the Office of the Director of National Intelligence (ODNI) issued the first, excluding Acronis AG and its affiliates from intelligence community procurements and requiring removal of covered Acronis articles from IC systems [8]. The statute was next used against an AI developer. On March 5, 2026, the Department of War issued a FASCSA supply chain risk designation of Anthropic covering only Department of War contracts [14]. On September 25, 2026, the U.S. Court of Appeals for the D.C. Circuit upheld that determination under 41 U.S.C. § 4713. The court read "supply chain risk" broadly: a supplier need not have hostile or malicious intent, and a model developer's training methodology could constrain agency operations even without post-deployment access [23]. A separate action against Anthropic under 10 U.S.C. § 3252 was set aside in the Northern District of California and should not be conflated with the FASCSA determination [23]. Against that backdrop, the DeepSeek order is the first issued by DHS, which makes it the first FASCSA order that applies across the federal civilian enterprise [14].

What the Order Says

The signed order, designated Exclusion and Removal Order #2026-01, appears as an appendix to the Federal Register notice [1]. It applies to federal civilian executive branch agencies "except for their national security systems and secure compartmented information systems." (The statute and the order's own footnotes use "sensitive compartmented information systems" [3].) It directs two separate actions. First, agencies must exclude DeepSeek, DeepSeek-affiliated entities, and their products from procurements, effective immediately. That exclusion extends to agency-approved product lists, service provider equipment lists, e-commerce marketplaces, and similar procurement channels. Second, agencies

must remove "products provided by DeepSeek, DeepSeek-affiliated entities, and third-party products and services developed using DeepSeek implementations" from civilian agency information technology systems. Removal is required immediately after issuance unless the agency seeks an exception. A footnote in the order says the exclusion reaches "third party implementations of DeepSeek used to provide products and services to civilian agencies" [1].

The order's stated rationale rests on several threat vectors. DeepSeek's servers are located in the People's Republic of China (PRC), and PRC law gives government cyber and intelligence functions broad authority to access private company infrastructure. DeepSeek's hosted services collect personally identifiable information and may fingerprint users. Prompts that federal employees or vendors submit could therefore transfer sensitive data to the PRC, where it "could produce actionable intelligence on U.S. persons and/or government programs" [1]. The order also states that "the models themselves may contain malware or other security vulnerabilities." It supports that claim with NowSecure's February 2025 finding of multiple critical vulnerabilities in the DeepSeek iOS application [1][10]. That evidence concerns the mobile app, however, not the model weights. In its footnotes the order also cites the House Select Committee on the CCP's April 2025 report on DeepSeek [11].

The exception process follows 41 CFR 201-1.304(b). Agencies may request that the order not apply to them or to certain actions, or ask for a delay in compliance. Requests had to be submitted in writing within 60 days of issuance with a "compelling justification," such as impact on mission-critical functions, along with any alternative mitigations [1]. During the period before full removal, the order encourages agencies to take interim measures. These include monitoring traffic, enhancing logging, applying data exfiltration controls, and flagging anomalous activity. CISA may provide technical assistance to agencies without the internal capability to comply, and the Secretary must review the order within one year and annually after that [1].

Why the Technical Context Matters

DeepSeek differs from the hardware and software vendors that supply chain exclusions have usually targeted, because much of its technology is distributed as open model weights. DeepSeek publishes models such as DeepSeek-R1 under the MIT license. That license allows commercial use, modification, and derivative works, and the R1 model card names distillation for training other models as a permitted use [12]. DeepSeek also released six dense "distilled" models: Qwen 2.5 and Llama 3 base models fine-tuned on reasoning data generated by DeepSeek-R1 [12]. As of October 21, 2026, the deepseek-ai organization on Hugging Face lists more than 100 public model repositories [13]. Hugging Face's model-tree metadata also records more than 500 community models that declare DeepSeek-R1 alone as their base

model, including fine-tunes, adapters, quantizations, and merges [12]. NIST's Center for AI Standards and Innovation (CAISI) reported in September 2025 that downloads of DeepSeek models on model-sharing platforms had risen nearly 1,000% since January 2025 [9].

Security Analysis

The Interpretive Problem: What Counts as a "DeepSeek Implementation"

The order's removal clause depends on a phrase it never defines. "DeepSeek implementations" plausibly covers DeepSeek's hosted application and API, and it almost certainly covers DeepSeek weights that a third party runs on its own infrastructure. Beyond those cases, the answer is less clear. Commentators have identified at least four ways a vendor might have "used" DeepSeek in development: direct runtime use, distillation, synthetic data generation, and fine-tuning. They have also observed that CISA had published no implementing guidance on the "developed using" standard when the order became public [14][15]. The table below sets out CSA's reading of how each category of AI supply chain dependency may relate to the order's text. It is an interpretation for risk-planning purposes, not legal advice, and it may change when CISA issues guidance.

Dependency Pattern	Example	Plausible Scope Reading	Detectability
DeepSeek-hosted app or API	Employees using the DeepSeek chat app; a SaaS product calling DeepSeek's API	Directly covered as a DeepSeek product or service	High: network egress, app inventories, API keys
Self-hosted DeepSeek weights	A vendor serving DeepSeek-V4 or R1 on U.S. cloud GPUs	Very likely covered as a third-party implementation	High if model inventory exists; weight hashes and model cards identify origin
Derivative weights	Fine-tunes, LoRA adapters, quantizations, or merges of DeepSeek checkpoints	Likely covered; the weights inherit DeepSeek parameters	Moderate: depends on <code>base_model</code> metadata and supplier disclosure

Dependency Pattern	Example	Plausible Scope Reading	Detectability
DeepSeek-published distills	DeepSeek-R1-Distill-Qwen-32B, R1-Distill-Llama-70B	Likely covered; published by DeepSeek despite non-DeepSeek base architecture	Moderate: model names often reveal lineage, but renamed copies may not
Third-party distillation or synthetic data	A vendor's in-house model trained partly on DeepSeek-generated outputs	Uncertain; turns on whether "developed using" reaches training inputs	Low: rarely recorded and not reliably recoverable from weights
Architectural techniques	Reimplementing published methods such as mixture-of-experts routing or GRPO	Unlikely to be covered; ideas from published research, not implementations	Not meaningfully detectable
AI-assisted code generation	Product code written with help from a DeepSeek-backed coding assistant	Uncertain; a literal reading of "developed using" could reach it	Very low unless development tooling is logged

The first four rows are primarily an inventory problem, one that existing tooling can address where model inventories and supplier disclosures exist. The last three are an evidentiary problem that current tooling does not solve. Runtime AI inventories and SBOM-style tools can identify a model that is deployed in a product. They cannot reconstruct which models generated training data, evaluation sets, or code months earlier [14]. A vendor asked to certify that a product was not "developed using DeepSeek implementations" may therefore be unable to answer with confidence. The difficulty grows if the vendor's own upstream model suppliers did not track that lineage. The D.C. Circuit's treatment of training methodology as relevant to supply chain risk in the Anthropic case suggests that development-time questions of this kind are unlikely to be dismissed as out of scope [23].

The order's text also contains internal inconsistencies that may complicate scoping. The Federal Register notice names "Hangzhou DeepSeek Artificial Intelligence Co., Ltd." The signed order names "Hangzhou DeepSeek Artificial Intelligence Basic Technology Research Co., Ltd.," and its affiliate footnotes refer to "DeepSeek AG" [1][2]. The affiliate definition borrowed from 48 CFR 2.101 depends on control

relationships, which may be hard to establish for a privately held PRC company. These details are unlikely to narrow the order's practical effect. They do suggest that contracting officers and vendors should expect clarifying guidance and should document their own interpretive choices in the meantime.

Risk Analysis Beyond the Compliance Question

DHS's threat rationale combines two distinct risk classes, and they call for different controls. The first is data-flow risk. Prompts, files, and telemetry sent to DeepSeek-operated infrastructure fall under PRC jurisdiction. This risk applies to the hosted app and API, and network egress controls, data loss prevention, and endpoint application control are the standard mitigations. The risk largely disappears when DeepSeek weights run entirely on infrastructure under U.S. control, which the order's derivative clause nonetheless reaches. That choice suggests DHS is concerned with more than data exfiltration.

The second class is model-integrity risk: the possibility that the weights themselves encode behavior adverse to the deployer. The order's statement that the models "may contain malware or other security vulnerabilities" is a hedged claim, and the order cites no public evidence of malicious payloads in DeepSeek weights [1]. Public evidence does show behavioral weaknesses that matter for deployers. CAISI's evaluation found that agents built on DeepSeek's most secure tested model, R1-0528, were on average 12 times more likely than evaluated U.S. frontier models to follow malicious instructions that diverted them from user tasks. Under a common jailbreaking technique, R1-0528 complied with 94% of overtly malicious requests, compared with 8% for U.S. reference models. CAISI also found that DeepSeek models echoed inaccurate or misleading CCP narratives four times as often as U.S. reference models [9]. Distillation and fine-tuning may carry these behavioral properties into derivative models, which offers a plausible security basis for following the model into downstream products. Public examples also show that post-training can change some inherited behaviors. Perplexity's R1-1776, for example, was post-trained from DeepSeek-R1 specifically to remove CCP-aligned censorship [24]. The extent of inheritance in any given derivative therefore remains an open empirical question that has not been systematically characterized.

The order also exposes a broader structural weakness in AI supply chains. Software supply chain programs matured around package manifests, signed artifacts, and build provenance. Model supply chains have equivalents in emerging formats such as CycloneDX ML-BOM and the SPDX 3.0 AI profile, and in Hugging Face's `base_model` metadata [16][17][18]. These formats are relatively new, and in CSA's assessment adoption remains uneven. Training-data lineage, including whether a model learned from another model's outputs, appears to be the provenance layer least often captured. The DeepSeek order makes that gap a procurement exposure, not only a governance concern.

Who Is Affected and How

The order applies only to federal civilian executive branch agencies [1][3], and contractors are reached through FAR 52.204-30 and its companion clauses [6][7]. Under FAR 52.204-30, contractors must search SAM.gov for applicable FASCSA orders and review it at least every three months. If they identify a covered article in contract performance, they must report it to the contracting officer within three business days and follow up on mitigation within ten business days. The clause flows down to subcontracts. The contractor's duty is defined as a "reasonable inquiry," meaning an inquiry into information already in the contractor's possession, and it expressly excludes the need for an internal or third-party audit [6]. An order issued after a solicitation becomes a contractual prohibition only once a solicitation amendment or contract modification adds it. The clause's notice and reporting provisions, however, already require contractors to check SAM.gov for new orders during performance and to report identified covered articles within three business days [6]. The period before a modification is therefore a time for discovery and reporting, not a grace period, and CSA's judgment is that it is best used for discovery. Private companies, state governments, and individuals are outside the order's legal scope. Regulated commercial entities may still find that customers, auditors, and insurers begin asking the same question the order asks.

Recommendations

Immediate Actions

Federal civilian agencies and their contractors should build a model-level inventory of every AI capability in production and in the development pipeline. The inventory should cover agent frameworks, embedded features in commercial SaaS, and coding assistants, and should record the specific model family, version, and hosting location for each. It should go beyond vendor-level questionnaires, because AI-enabled products increasingly route among several underlying models. Network teams should block or alert on egress to DeepSeek-operated endpoints and application domains, and endpoint teams should confirm that DeepSeek mobile and desktop applications are removed from managed devices. Contractors should search SAM.gov for the order, determine which contracts include FAR 52.204-30, and prepare to meet the three-business-day reporting obligation if a covered article is found [6]. Agencies that believe they need relief should contact FASCSA.Orders@cisa.dhs.gov promptly, even though the formal 60-day exception window measured from issuance appears to have passed [1][14]. The order does not provide for late requests, but CISA's implementation team is the designated point of contact for "additional information, assistance, and reporting" [1].

Short-Term Mitigations

Over the next 30 to 90 days, organizations should extend the inventory into model lineage. For every self-hosted or vendor-supplied model, they should record the base model, any fine-tuning or adapter lineage, and quantization source [18]. Where weights are claimed to be unmodified, teams should confirm that file hashes match the published checkpoint. Hashes will not match for fine-tunes, adapters, and quantizations, so for those categories teams should rely on supplier attestation and, where the stakes justify it, weight-similarity or architecture-fingerprint analysis against candidate base models. Suppliers should receive written attestation requests that ask specifically about distillation, synthetic data generation, and evaluation use of DeepSeek outputs, not only about runtime deployment. Organizations should also document the interpretive position they take for the uncertain categories in the table above. Wherever DeepSeek-derived components are found, teams should apply the order's interim cybersecurity measures (traffic monitoring, enhanced logging, and exfiltration controls) until removal or replacement is complete [1]. Teams that rely on DeepSeek-derived models for cost or performance reasons should test replacements before any contract modification forces a cutover.

Strategic Considerations

The DeepSeek order is unlikely to be the last action of its kind, and its "developed using" formulation could be reused for other model developers or jurisdictions. Organizations should therefore treat model provenance as a standing supply chain control, not a one-time response to this order. In practice, this means generating machine-readable AI bills of materials in formats such as CycloneDX ML-BOM or SPDX 3.0 for internally developed models and AI-enabled products [16][17]. It also means recording training-data lineage, including synthetic data sources, at the time models are built, because that information cannot be reliably reconstructed later. Procurement contracts should require suppliers to disclose model lineage and to notify customers when a regulatory action affects a component. Finally, architectures should keep model substitution practical through model-agnostic gateways and tested fallback paths, so that runtime exposure to a future exclusion order can be addressed through configuration instead of re-engineering.

CSA Resource Alignment

CSA's July 2026 whitepaper *Sovereign AI Risk: When Your AI Vendor Gets Export-Controlled* addresses the mirror image of this order: a government action that suddenly cuts off access to a U.S. model [19]. Its core recommendations apply directly here. They include an urgent inventory of model integrations that pins model versions in code and configuration, documentation of workflows without a tested fallback,

and a model-agnostic gateway with at least two tested alternatives for each critical workflow. Together, the two events show that regulatory action can remove a model from an enterprise's supply chain in either direction, and that resilience planning should not assume model availability is solely a commercial question.

CSA's May 2026 whitepaper *AI Model Repository Monoculture* examines how concentration on a few model hosting platforms turns individual supply chain weaknesses into ecosystem-scale events [20]. Its recommendations on model inventory and dependency mapping, signing and pinning, and the "verification gap" in public repositories correspond to the derivative-weights rows of the scoping table in this note. Organizations that already apply those controls will be better placed to answer whether DeepSeek parameters are present in their deployed models.

CSA's August 2026 research note *AISI: Open-Weight Models Close Cyber Capability Gap* reported UK AI Security Institute findings that DeepSeek V4-Pro was among the open-weight models approaching recent frontier closed-model performance on offensive cyber tasks at much lower cost [21]. The capability-to-cost ratio it documents is consistent with the broad commercial uptake reflected in CAISI's download figures [9], and that uptake helps explain why the derivative clause has such broad reach. The note also warns that open-weight models can be self-hosted and stripped of vendor safeguards, which reinforces this note's distinction between data-flow risk and model-integrity risk.

The AI Controls Matrix (AICM) v1.1 provides the control vocabulary for operationalizing these recommendations [22]. Its supply chain management, transparency, and accountability controls address supplier attestation and model provenance. Its data security and privacy lifecycle controls address prompt and telemetry flows to foreign-hosted services, and its model security controls address the integrity of third-party weights. Organizations can map their DeepSeek inventory, attestation, and replacement activities to AICM controls to produce evidence for auditors. The same mapping will support assurance requests from federal customers or STAR-for-AI assessments.

References

- [1] U.S. Department of Homeland Security. "[Notice of Order Regarding Exclusion and Removal of DeepSeek Products From Federal Civilian Executive Branch Information Technology Systems and Procurements \(Exclusion and Removal Order #2026-01, Appendix\)](#)." Federal Register, Vol. 91, No. 201, pp. 66346–66350, October 20, 2026.
- [2] Federal Register. "[Notice of Order Regarding Exclusion and Removal of DeepSeek Products From Federal Civilian Executive Branch Information Technology Systems and Procurements](#)." Docket DHS-2026-0628, October 20, 2026.
- [3] Legal Information Institute. "[41 U.S. Code § 1323 – Functions of the Federal Acquisition Security Council](#)." Cornell Law School, accessed October 21, 2026.
- [4] Federal Acquisition Security Council. "[Federal Acquisition Security Council Rule](#)." Federal Register, 86 FR 47581, August 26, 2021.
- [5] Electronic Code of Federal Regulations. "[41 CFR Part 201-1 – General Regulations](#)." eCFR, accessed October 21, 2026.
- [6] Acquisition.gov. "[FAR 52.204-30, Federal Acquisition Supply Chain Security Act Orders—Prohibition](#)." General Services Administration, accessed October 21, 2026.
- [7] Acquisition.gov. "[FAR 52.204-29, Federal Acquisition Supply Chain Security Act Orders—Representation and Disclosures](#)." General Services Administration, accessed October 21, 2026.
- [8] National Law Review. "[ODNI Issues First FASCSA Order Targeting Acronis AG](#)." National Law Review, September 24, 2025.
- [9] National Institute of Standards and Technology. "[CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks](#)." NIST, September 30, 2025.
- [10] Andrew Hoog. "[NowSecure Uncovers Multiple Security and Privacy Flaws in DeepSeek iOS Mobile App](#)." NowSecure, February 6, 2025.
- [11] House Select Committee on the Strategic Competition Between the United States and the Chinese Communist Party. "[DeepSeek Unmasked: Exposing the CCP's Latest Tool for Spying, Stealing, and Subverting U.S. Export Control Restrictions](#)." U.S. House of Representatives, April 16, 2025.

- [12] DeepSeek-AI. "[DeepSeek-R1 Model Card](#)." Hugging Face, January 2025 (model tree accessed October 21, 2026).
- [13] Hugging Face. "[deepseek-ai Organization Page](#)." Hugging Face, accessed October 21, 2026.
- [14] Tech Times. "[DHS Filed DeepSeek Ban 28 Days After Agencies Lost Right to Request Exceptions](#)." Tech Times, October 20, 2026.
- [15] Forkast. "[DHS Ordered Every Federal Agency to Purge DeepSeek. The Exclusion Follows the Model In to Anything Built On It](#)." Forkast, October 20, 2026.
- [16] OWASP CycloneDX. "[Machine Learning Bill of Materials \(ML-BOM\)](#)." CycloneDX, accessed October 21, 2026.
- [17] SPDX. "[SPDX AI and Dataset Profiles](#)." Linux Foundation SPDX, accessed October 21, 2026.
- [18] Hugging Face. "[Model Cards](#)." Hugging Face Hub Documentation, accessed October 21, 2026.
- [19] Cloud Security Alliance AI Safety Initiative. "[Sovereign AI Risk: When Your AI Vendor Gets Export-Controlled](#)." Cloud Security Alliance, July 2, 2026.
- [20] Cloud Security Alliance. "[AI Model Repository Monoculture: How Ecosystem Concentration Transforms Supply Chain Vulnerabilities into Platform-Scale Events](#)." Cloud Security Alliance, May 14, 2026.
- [21] Cloud Security Alliance. "[AIS: Open-Weight Models Close Cyber Capability Gap](#)." Cloud Security Alliance, August 9, 2026.
- [22] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." Cloud Security Alliance, June 22, 2026.
- [23] Amy L. Fuentes, David S. Black, Ann E. McRitchie, Tanner N. Slaughter, and Marissa C. Serafino. "[D.C. Circuit Upholds Exclusion of Anthropic from DOW Supply Chain Under FASCSA](#)." Holland & Knight, September 29, 2026.
- [24] Maximilian Schreiner. "[Perplexity AI Removes Chinese Censorship from Deepseek R1](#)." The Decoder, February 19, 2025.