


DeepSeek FASCSA Order: Model Provenance Becomes a Procurement Requirement

Why federal supply chain law now asks which model is inside the product

2026-10-22

 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 published Exclusion and Removal Order #2026-01 in the Federal Register on October 20, 2026. Issued under the Federal Acquisition Supply Chain Security Act of 2018 (FASCSA), it bars products and services from Hangzhou DeepSeek Artificial Intelligence Co., Ltd. and its affiliates from federal civilian procurement. It also requires civilian agencies to remove DeepSeek products, along with third-party products "developed using DeepSeek implementations," from their information technology systems [1][2]. Crowell & Moring reports that the Department of War (DoW) and the Director of National Intelligence (DNI) issued companion orders, posted in the SAM.gov FASCSA order listing. Together the three orders cover the civilian, defense, and intelligence parts of the federal enterprise [3].

The legal mechanics have been covered elsewhere, including CSA's research note of October 21 on how to read the "developed using" clause [4]. This note looks at the operational consequence: **model provenance has become a practical federal procurement concern**. Until this order, a contractor could meet most supply chain diligence obligations by identifying the companies it bought from. The FAR's reasonable-inquiry standard does not itself require contractors to establish model lineage. However, a contractor that cannot say which model weights sit inside a product, where those weights came from, and, in the civilian context, whether a covered model was used to build it now carries compliance risk that it cannot measure.

Several consequences follow from that shift. FASCSA orders now target a **model family**, not just a vendor. The Acronis order in 2025 named a company [5], and the 2026 DoW action against Anthropic targeted a supplier [6], but the DeepSeek orders reach open-weight artifacts that anyone can download, rehost, fine-tune, or distill. The three orders also appear to **differ in scope**. Crowell & Moring reads the DoW and DNI orders as covering applications and services that implement a DeepSeek model wherever it is hosted, including locally run open weights, while the DHS order adds products "developed using" DeepSeek implementations [3].

The compliance floor, meanwhile, is low. The FAR **"reasonable inquiry" standard** covers only information already in the contractor's possession, and it expressly does not require an audit [7]. Few organizations appear to maintain model lineage records today, so the standard currently sets a low bar, and one the procurement environment is likely to raise. Agencies and contractors should therefore build

model-level inventories, verifiable weight identity, and supplier lineage attestations now. Given three FASCSA actions in roughly a year, two of them involving AI model providers, the DeepSeek orders are unlikely to be the last FASCSA actions aimed at AI models.

Background

FASCSA and Its Implementing Clauses

FASCSA created the Federal Acquisition Security Council (FASC), an interagency body that assesses supply chain risk and recommends exclusion and removal actions. Three officials can issue binding orders. The Secretary of Homeland Security issues them for civilian agencies, the Secretary of Defense (now styled Secretary of War) for the defense enterprise and national security systems, and the Director of National Intelligence for the intelligence community and sensitive compartmented information systems [1] [3]. The FAR Council put the statute into effect through an interim rule effective December 4, 2023. That rule added clauses 52.204-28, -29, and -30 to solicitations and contracts [8].

For contractors, the working clause is FAR 52.204-30. It prohibits providing or using a covered article in contract performance unless the issuing official grants a waiver. It requires contractors to check SAM.gov for FASCSA orders at least once every three months and to conduct a reasonable inquiry when a new order could affect their supply chain. A covered article that is found must be reported to the contracting officer within three business days, followed by mitigation information within ten business days. The clause must also flow down to subcontracts [7]. The FAR's definition of reasonable inquiry is narrower than many practitioners assume. It is an inquiry designed to uncover information "in the entity's possession" about covered articles, and it expressly excludes the need for an internal or third-party audit [7].

From Vendors to Models

The statute went unused for nearly two years after the interim rule. The first order was published in September 2025, when ODNI excluded Swiss data protection firm Acronis AG from intelligence community procurement [5][9]. In early 2026 the Department of War designated Anthropic a supply chain risk under both 10 U.S.C. § 3252 and FASCSA. That designation reached all Anthropic affiliates and products [6], and on September 25, 2026 the D.C. Circuit upheld the FASCSA designation [10]. Both actions targeted a company. A contractor could comply by checking vendor names, corporate affiliations, and product SKUs against the order.

The DeepSeek orders change that. DHS frames the risk around DeepSeek's AI models and applications. Crowell & Moring reports that the civilian order describes DeepSeek's freely available models as presenting multiple threat vectors: known technical vulnerabilities, an obligation to provide the government of the People's Republic of China certain access, and the possibility that information collected through DeepSeek apps and models could yield "actionable intelligence" about the U.S. government [3]. The technical-vulnerability element of that rationale is consistent with findings from NIST's Center for AI Standards and Innovation (CAISI). In September 2025, CAISI reported that DeepSeek models lagged U.S. frontier models on performance and security and showed shortcomings that may pose risks to application developers, consumers, and national security [11]. CAISI did not, however, assess the PRC legal access obligations that form the second element of the DHS rationale.

The technical distribution of DeepSeek models is what makes this order different from the vendor-centric orders before it. DeepSeek released the R1 weights under the MIT License and explicitly permitted "any modifications and derivative works, including, but not limited to, distillation for training other LLMs." It also published six distilled models built on Qwen2.5 and Llama 3.x base models [12]. Major U.S. cloud platforms made R1 available as well. Amazon Bedrock offered it through Bedrock Marketplace and Custom Model Import, and from March 2025 as a fully managed serverless model [13]. In practice, a covered DeepSeek model can show up in a federal supply chain under a U.S. cloud provider's name, inside a third-party product with no DeepSeek branding, or as weights carrying a Llama or Qwen label.

Security Analysis

Three Orders, Three Scope Definitions

Reporting so far suggests that the three orders do not define their covered articles the same way. Only the DHS order has been published with reasoning. The DoW and DNI orders appear in the SAM.gov FASCSA order data without explanation [3]. The table below summarizes the scope differences reported so far. Organizations should confirm these readings against the order texts on SAM.gov and Regulations.gov (docket DHS-2026-0628) [1] before relying on them.

Issuing Official	Systems in Scope	Covered Articles (as reported)	Provenance Question Raised
Secretary of Homeland Security	Federal civilian executive branch IT, excluding	DeepSeek and affiliate products and services, plus third-party products and services "developed using	Does the product contain DeepSeek weights, and was it

Issuing Official	Systems in Scope	Covered Articles (as reported)	Provenance Question Raised
	national security and SCI systems	DeepSeek implementations" [1] [3]	built with a DeepSeek model or tool?
Secretary of War	Defense agencies and national security systems	Applications and services that implement a DeepSeek AI model, regardless of hosting location, including locally integrated open weights [3]	Does any component run inference on DeepSeek weights, anywhere?
Director of National Intelligence	Intelligence community and SCI systems	Applications and services that implement a DeepSeek AI model, as with DoW [3]	Same as DoW, applied to IC procurements

A contractor that serves both civilian and defense customers may therefore face two different tests for the same product. One test is about runtime: is a DeepSeek model executing inside the product? The other also reaches development: was a DeepSeek model used to build the product? The October 21 CSA note examines the interpretive range of the development test, including distillation, synthetic training data, and AI-assisted coding [4]. The practical point here is simpler. Both tests require evidence about model identity, and the government's standard procurement identifiers (vendor name, Unique Entity ID, product SKU, FedRAMP authorization boundary) record none of it.

Why Vendor Identity No Longer Answers the Question

Software supply chain diligence has matured around the idea that a component is identified by its publisher and version. Software bills of materials (SBOMs), package registries, and code signing all rest on that idea. Model weights break it in at least three ways. First, weights can be renamed or rehosted freely. A model served from a U.S. hyperscaler endpoint under a marketplace listing has, from the buyer's point of view, a U.S. supplier. Under the DoW and DNI readings reported by Crowell & Moring, however, hosting location does not take a DeepSeek model out of scope [3]. Second, derivatives keep the lineage without keeping the name. A fine-tune of a DeepSeek checkpoint, or one of DeepSeek's own published distills on a Llama base [12], may not carry "DeepSeek" in any field a procurement system would see. Third, in CSA's observation, many AI products route requests through model gateways, aggregators, or agent frameworks that pick a model at runtime. In those architectures, which model actually served a given request is a configuration fact, not a contract fact, and it can change without any procurement event.

This is the gap the reasonable inquiry standard exposes. Because the FAR limits the inquiry to information already in the contractor's possession and disclaims any audit requirement [7], a contractor whose suppliers never disclosed model lineage could, on a narrow reading, satisfy the clause while unknowingly delivering a covered article. That reading may be legally defensible, but if covered articles turn up after the fact, contracting officers and agency CISOs are likely to respond by writing explicit model disclosure requirements into solicitations. Contractors who build lineage records now will be better placed than those who later have to reconstruct them under a ten-business-day mitigation deadline.

What Counts as Provenance Evidence

If model provenance is going to carry compliance weight, organizations need to separate types of evidence by how much assurance they actually give. The hierarchy below is CSA's assessment, not a government standard. It is meant to help organizations decide which evidence to collect, and from whom.

Evidence Type	What It Establishes	Limitations
Supplier attestation of model lineage	The supplier's representation of base model, fine-tuning sources, and distillation inputs	Only as reliable as the supplier's own records; weak for deep or multi-tier supply chains
Model card, configuration, and tokenizer metadata	Declared architecture, base model, and license, as published with the weights	Easily edited or stripped; derivatives may inherit or omit fields
Cryptographic hash of deployed weights against a known-artifact registry	Exact identity of an unmodified checkpoint	Any fine-tune or quantization changes the hash; requires a maintained reference set
Architectural and tokenizer fingerprinting	Probable family membership even after renaming	Probabilistic; families with shared architectures may be hard to tell apart
Behavioral fingerprinting via probe prompts	Likely lineage of an API-only model whose weights cannot be inspected	Research-grade; results can be contested, and gateways can change the backing model

Evidence Type	What It Establishes	Limitations
Build-time training and data lineage records	Whether a covered model contributed synthetic data, distillation signal, or code to a product	In CSA's observation, rarely kept today; must be recorded at build time because it cannot be reliably reconstructed afterward

Two observations follow from the table. Only the last row provides first-party evidence on the DHS "developed using" test, and it is the evidence organizations are least likely to have. Supplier attestations in the first row can address that test indirectly, but only as reliably as the supplier's own records. The second and third rows are cheap to collect for self-hosted models and should be standard practice in any organization that runs open weights. Organizations that consume models only through APIs will rely mostly on the first row. That makes contract language requiring model disclosure and change notification the most practical control available to them.

Risk Beyond the Compliance Boundary

The orders apply only to federal agencies and their contractors. Private companies and the public remain outside their legal reach [2]. Even so, the security rationale DHS gives does not stop at the federal boundary. If its concerns about data exposure, state access obligations, and technical vulnerabilities hold for a federal agency, they deserve at least a risk assessment in any regulated enterprise, and CAISI's evaluation offers a public, government-produced basis for one [11]. It is also worth separating two risk types that the orders bundle together. Data-flow risk, where prompts and outputs travel to DeepSeek-operated infrastructure, is largely removed when weights are self-hosted or consumed through a U.S. cloud provider such as Amazon Bedrock [13]. Model-integrity risk is not. The DHS order itself states that the models "may contain malware or other security vulnerabilities" [1], and behavioral weaknesses, susceptibility to malicious instructions, and, as a general property of model weights rather than a documented DeepSeek finding, any trained-in behavior go wherever the weights go. The DoW and DNI decision to reach locally integrated open weights [3] is consistent with treating model-integrity risk, not just data-flow risk, as material, although those orders were issued without published reasoning. Private organizations that have treated self-hosting as a complete mitigation should revisit that assumption.

The order's timing also has a procedural consequence. The signed order is dated July 24, 2026, but it was not filed for public inspection until October 19 [1][2]. Because the order requires agency exception requests "within 60 days of the issuance" of the order, that window had already closed by the time of publication [1]. Contractor obligations under FAR 52.204-30 attach through SAM.gov posting and

incorporation into contracts, not through Federal Register publication [7]. Contractors should therefore check SAM.gov and their contract modifications rather than treat October 20 as the start date of their obligations.

Recommendations

Immediate Actions

Federal contractors and civilian agencies should check SAM.gov for the DeepSeek FASCSA orders, using GSA's FASCSA order guidance [14], and determine which orders their contracts incorporate. Each then needs a model-level inventory, not just a vendor-level one. That inventory should cover every self-hosted model with its declared base model and a weight hash, every model API with the provider and the specific model identifier in use, and every gateway or agent framework with the full set of models it can route to. Where a DeepSeek model or a known derivative is found in contract performance, contractors should start the FAR 52.204-30 three-business-day reporting clock [7] and send questions to CISA's FASCSA Order Implementation Team at FASCSA.Orders@cisa.dhs.gov [1]. Organizations should also send short written lineage questionnaires to AI suppliers now. The answers become "information in the entity's possession" and turn the reasonable inquiry from a passive exercise into a documented one.

Short-Term Mitigations (30 to 90 Days)

Within the next quarter, organizations should move from inventory to verification. For self-hosted models, that means keeping a reference set of hashes for known DeepSeek checkpoints and published distills [12] and checking deployed weights against it as part of model onboarding. For gateways and routers, it means pinning allowed model identifiers for federal workloads and logging which model served each request, so that a report under the ten-business-day mitigation requirement can be backed by evidence. Procurement and legal teams should add model disclosure, lineage attestation, and change notification clauses to AI supplier agreements, covering base models, fine-tuning sources, and any use of a covered model for distillation or synthetic data generation. Organizations serving both civilian and defense customers should document which scope reading they apply under each order, and why. The DHS and DoW scope tests appear to differ [3], and a documented position is easier to defend than an improvised one.

Strategic Considerations

The broader lesson is that AI model identity is becoming a compliance-relevant attribute, and supply chain programs should treat it that way. Organizations should adopt a machine-readable AI bill of materials that records model lineage, training-data sources, and build-time use of other models alongside conventional software components. The build-time record matters most for "developed using" tests, because it is the one piece of evidence that cannot be rebuilt reliably after the fact. Architectures that abstract model choice behind a gateway with an enforceable allowlist will let organizations respond to future orders by changing configuration rather than re-engineering products. The progression from Acronis [5] to Anthropic [6] to DeepSeek [1] shows FASCSA moving from a dormant authority to an active one, and the last two actions suggest it is becoming a tool of AI policy. Organizations should therefore plan on more orders, possibly naming other model families, and should build provenance capabilities that generalize instead of single-purpose DeepSeek screening.

CSA Resource Alignment

CSA's October 21, 2026 research note, *DeepSeek FASCSA Order: What "Developed Using" Means for AI Supply Chains* [4], is the companion to this document. It maps seven dependency patterns, from hosted API use to AI-assisted coding, against the civilian order's derivative clause and rates how detectable each one is. This note builds on that analysis by asking what evidence would show whether each pattern is present, and by placing the civilian order beside the DoW and DNI orders and their reported narrower runtime focus.

The **AI Controls Matrix (AICM) v1.1** [15] provides the control framework for turning provenance into practice. Its supply chain management, transparency, and accountability controls set expectations for supplier disclosure and third-party risk assessment. Its model security and data lineage controls cover the inventory, integrity verification, and training-data tracking recommended above. AICM's mapping to STAR for AI also gives AI suppliers a structured way to publish the lineage attestations that federal contractors will increasingly ask for.

CSA's **AI Model Risk Management Framework** [16] describes model cards, data sheets, and risk cards as the core documentation for model governance. In light of the DeepSeek orders, those artifacts move from good practice to compliance evidence. A model card that records base model and derivation history, and a data sheet that records synthetic data sources, are directly relevant to answering a reasonable inquiry.

Finally, CSA's research note on **sovereign AI vendor dependency and the Pentagon-Anthropic designation** [17] examines the same FASCSA authority applied to a U.S. frontier model provider. Read together with the DeepSeek orders, it supports this note's view that supply chain risk designations can now reach model providers of any national origin. Model-agnostic architecture and documented provenance are therefore general resilience measures, not responses to a single country or vendor.

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.](#)" Federal Register, 91 FR 66346, October 20, 2026.
- [2] Swati Khandelwal. "[U.S. Orders DeepSeek Removed From Federal Civilian Systems, Bans New Purchases.](#)" The Hacker News, October 21, 2026.
- [3] Crowell & Moring LLP. "[FASCSA Orders Require Exclusion and Removal of DeepSeek AI.](#)" Crowell & Moring Client Alert, October 21, 2026.
- [4] Cloud Security Alliance AI Safety Initiative. "[DeepSeek FASCSA Order: What 'Developed Using' Means for AI Supply Chains.](#)" CSA Labs, October 21, 2026.
- [5] Covington & Burling LLP. "[First Order Issued Under the Federal Acquisition Supply Chain Security Act, Triggering Immediate Requirements on Contractors.](#)" Inside Government Contracts, September 2025.
- [6] Mayer Brown LLP. "[Anthropic Supply Chain Risk Designation Takes Effect: Latest Developments and Next Steps for Government Contractors.](#)" Mayer Brown, March 10, 2026.
- [7] FAR Council. "[FAR 52.204-30 Federal Acquisition Supply Chain Security Act Orders—Prohibition.](#)" Acquisition.gov, accessed October 22, 2026.
- [8] Crowell & Moring LLP. "[Coming December 4: Do You Know Where Your Supply Chain Risks Are? FAR Council Issues Interim Rule Requiring Contractor Diligence for FASC Exclusion and Removal Orders.](#)" Crowell & Moring Client Alert, October 2023.
- [9] Morgan Lewis. "[FASC Issues First FASCSA Exclusion Order: Implications for Federal Contractors and ICTS Supply Chains.](#)" Morgan Lewis, October 1, 2025.
- [10] Holland & Knight LLP. "[D.C. Circuit Upholds Exclusion of Anthropic from DoW Supply Chain.](#)" Holland & Knight Alert, September 29, 2026.
- [11] NIST Center for AI Standards and Innovation. "[CAISI Evaluation of DeepSeek AI Models.](#)" National Institute of Standards and Technology, September 2025.
- [12] DeepSeek-AI. "[DeepSeek-R1 Model Card.](#)" Hugging Face, January 2025.

- [13] Amazon Web Services. "[DeepSeek-R1 Is Available Fully Managed in Amazon Bedrock](#)." AWS What's New, March 10, 2025.
- [14] General Services Administration. "[FASCSA Order Check](#)." GSA SmartPay, accessed October 22, 2026.
- [15] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." Cloud Security Alliance, 2025.
- [16] Cloud Security Alliance AI Technology and Risk Working Group. "[AI Model Risk Management Framework](#)." Cloud Security Alliance, 2024.
- [17] Cloud Security Alliance AI Safety Initiative. "[Sovereign AI Dependency: The Pentagon-Anthropic Concentration Trap](#)." CSA Labs, April 13, 2026.