PUBLIC COMMENT OF THE BOX COMMONS

Docket No. ATR-2026-0001

Request for Information: Antitrust Guidelines for Collaborations Among Competitors

Submitted to: U.S. Department of Justice, Antitrust Division Federal Trade Commission

Submitted by: The Box Commons Brice Love, Acting Executive Director info@theboxcommons.org

Date: April 23, 2026

TABLE OF CONTENTS

I. Introduction and Interest of The Box Commons II. The AI Agent Credentialing Gap III. Response to Question 1: Topics Requiring Additional Guidance IV. Response to Question 2: New Technologies and Business Models Requiring Guidance V. Response to Question 3: Significant Legal, Economic, and Technological Developments VI. Proposed Safe Harbor Framework for AI Credentialing Standards VII. Conclusion

---

I. INTRODUCTION AND INTEREST OF THE BOX COMMONS

The Box Commons respectfully submits this comment in response to the joint request for information issued by the Department of Justice and Federal Trade Commission on February 23, 2026, seeking public input on the development of updated antitrust guidance for collaborations among competitors.

The Box Commons is a Wyoming-incorporated 501(c)(6) business league organized as a standards development organization for artificial intelligence agent credentialing, safety, and governance. The organization's mission is to develop consensus-based credentialing standards that enable AI agents to operate as accountable participants in commercial activity — supported by verifiable credentials, transparent evaluation criteria, and insurance-linked certification.

The Box Commons has engaged directly with the National Institute of Standards and Technology on AI agent governance, submitting a public comment on the NIST Artificial Intelligence 800-2 draft standard and contributing to NIST's Center for AI Standards and Innovation (CAISI) stakeholder process. The organization's founding industrial members hold a portfolio of patents covering AI agent credentialing methodologies, insurance modeling, and governance simulation, which will be made available to the standards body on royalty-free terms to ensure that the credentialing infrastructure is accessible to all market participants.

The Box Commons' interest in this proceeding is direct and urgent. As the standards body that will develop and administer AI agent credentials, the organization faces the full weight of every antitrust ambiguity in the current regulatory environment. If the revised guidelines do not provide adequate clarity and protection for AI credentialing organizations, the entire emerging ecosystem of AI agent governance — credentialing, insurance, and commercial deployment — will be chilled by legal uncertainty that only the largest technology incumbents can afford to absorb.

This comment addresses all three questions posed in the RFI and proposes a specific safe harbor framework for AI credentialing standard-setting.

II. THE AI AGENT CREDENTIALING GAP

Before addressing the specific questions, it is necessary to describe the market context that makes this comment urgent.

Artificial intelligence agents are no longer theoretical. They are autonomous software systems capable of executing multi-step tasks, interfacing with application programming interfaces, managing communications, triggering financial transactions, and operating across enterprise systems with minimal human supervision. The transition from passive AI models — which merely respond to human prompts — to active AI agents — which take independent actions within delegated authority — represents the most significant shift in commercial technology since the internet.

This shift creates an accountability gap. When a human employee makes a decision, liability follows established lines — employment law, fiduciary duties, professional licensing, insurance. When an AI agent makes a decision, none of these frameworks apply. There is no professional license for an AI agent. There is no malpractice insurance. There is no credentialing body that certifies whether a particular AI agent meets minimum competency and safety standards.

The market is beginning to address this gap through private ordering. Cyber insurance carriers are introducing AI security riders that require documented risk assessments and alignment with recognized AI risk management frameworks as prerequisites for coverage. Enterprise procurement teams are demanding evidence of model-level testing and adversarial red-teaming before deploying third-party AI agents. The National Institute of Standards and Technology has launched the AI Agent Standards Initiative to develop identity and authorization protocols for autonomous systems.

What is missing is the credentialing layer. The private insurance market needs standardized, independently verified credentials before it can systematically underwrite AI agent risk. The federal government needs private credentialing bodies to complement public standards like the NIST AI Risk Management Framework. Businesses need a reliable signal — comparable to a UL certification mark or HITRUST certification — indicating that an AI agent meets defined safety and performance criteria.

Building this credentialing infrastructure requires competitor collaboration. Technology companies, cybersecurity firms, AI developers, insurers, and patent holders must work together within a standards development organization to define what "safe" and "competent" mean for an AI agent. They must pool technical expertise, agree on interoperability protocols, and establish certification criteria. This is precisely the type of collaboration that the withdrawn 2000 guidelines were designed to evaluate — and precisely the type that currently operates in a regulatory vacuum.

III. RESPONSE TO QUESTION 1: TOPICS REQUIRING ADDITIONAL GUIDANCE

The agencies ask what topics would benefit from additional guidance, including "joint licensing arrangements" and "agreements adopting conditions for dealing with competitors." Two topics are critical to the AI credentialing ecosystem.

A. Credentialing Decisions by Standards Development Organizations

The most urgent topic requiring guidance is the antitrust treatment of credentialing decisions — specifically, the decision to grant, deny, suspend, or revoke a credential.

Under current law, the line between legitimate standard-setting and unlawful group boycott is dangerously unclear for credentialing organizations. When a standards body denies or revokes an AI agent's credential based on published technical criteria, the denied party may characterize that decision as a "concerted refusal to deal" under Section 1 of the Sherman Act. If insurers condition coverage on that credential, the denied party's argument gains force: the credentialing body has effectively excluded a competitor from the market.

The Supreme Court's framework in *Allied Tube & Conduit Corp. v. Indian Head, Inc.*, 486 U.S. 492 (1988), provides some structure but insufficient clarity. *Allied Tube* established that private standard-setting enjoys no blanket antitrust immunity, and that economically interested parties who subvert the process face liability under the rule of reason. The Court recognized that standard-setting can be procompetitive when conducted with adequate procedural safeguards — open participation, conflict-of-interest disclosure, balanced representation, and documented appeals processes.

What *Allied Tube* did not address is the specific context of credentialing — where the standard-setting organization does not merely publish recommendations but actively evaluates individual applicants against those standards and renders pass/fail judgments. Credentialing adds an enforcement mechanism to standard-setting, and that enforcement mechanism is where the antitrust risk concentrates.

The revised guidelines should explicitly address credentialing and provide that a credential denial or revocation by a standards development organization does not constitute a group boycott or concerted refusal to deal when the following conditions are met:

1. The credentialing criteria were developed through a consensus process consistent with the principles of OMB Circular A-119;
2. The criteria are published, technology-neutral, and available for public review;
3. The applicant received documented notice of the specific deficiency;
4. A documented appeals process is available and was made available to the applicant; and
5. The credentialing decision was based on the published criteria, not on the competitive interests of any individual participant in the standard-setting process.

This framework draws directly from *Allied Tube*'s emphasis on procedural safeguards while providing the operational specificity that a credentialing body needs to function without perpetual litigation risk.

B. Insurance-Linked Credentialing

The second topic requiring explicit guidance is the relationship between private credentialing and insurance underwriting.

The most powerful market mechanism for AI agent governance is insurance. When cyber insurers require a recognized credential as a condition of coverage — analogous to how property insurers require UL certification for electrical components — the credential becomes a de facto market requirement. This mechanism is faster, more adaptive, and more enforceable than government regulation.

However, the convergence of credentialing and insurance creates a specific antitrust concern. If a standards body's credential becomes the effective gateway to insurance coverage, and insurance coverage becomes the effective gateway to commercial deployment, then the standards body wields enormous market power — the power to exclude.

The revised guidelines should address this directly. When an insurer independently adopts a standards body's credential as an underwriting requirement, that adoption does not transform the credentialing body into a market gatekeeper for antitrust purposes, provided:

1. The standards body maintains structural independence from the insurers who adopt its credentials — no insurer holds a controlling governance position;
2. The credential criteria are not designed to favor or disfavor specific commercial products or technologies;
3. Multiple pathways to demonstrating compliance exist (the standards body does not monopolize testing or certification); and
4. The insurer's adoption of the credential was an independent commercial decision, not the result of coordination with the standards body to exclude competitors.

This framework is essential because the alternative — leaving insurance-linked credentialing unaddressed — creates a paradox. The market needs credentials to enable insurance. Insurance needs credentials to assess risk. But the combination of the two, without antitrust guidance, creates legal exposure that deters the very organizations willing to build the credentialing infrastructure.

IV. RESPONSE TO QUESTION 2: NEW TECHNOLOGIES AND BUSINESS MODELS REQUIRING GUIDANCE

The agencies ask what new technologies and business models would benefit from additional guidance, including "algorithmic pricing" and "information and data sharing." Two developments in the AI agent ecosystem require attention.

A. AI Agents as Autonomous Economic Actors

The 2000 guidelines were drafted in a market where economic actors were humans and corporations. AI agents are neither. They are autonomous software systems that make decisions, execute transactions, and interact with other agents and humans in commercial contexts. The revised guidelines must account for a market in which some of the "competitors" are not human and some of the "collaborations" involve defining the rules under which these non-human actors may participate.

Standard-setting for AI agents is categorically different from traditional product standard-setting. When the National Fire Protection Association sets standards for electrical conduit, it is evaluating a static physical product. When an AI credentialing body sets standards for an AI agent, it is evaluating a dynamic, learning system whose behavior may change over time. The credentialing process must therefore include ongoing monitoring, periodic re-evaluation, and the ability to suspend or revoke credentials — activities that look, from an antitrust perspective, like ongoing coordination among competitors.

The guidelines should recognize that continuous credentialing — including monitoring, re-evaluation, and revocation — for AI agents is a structural feature of the technology, not evidence of ongoing anticompetitive coordination. The dynamic nature of AI systems requires standards bodies to maintain active oversight in ways that static product certification does not.

B. Royalty-Free Patent Licensing to Standards Bodies

The relationship between patent holders and standards development organizations is a well-established topic in antitrust law, but the AI credentialing context introduces new dynamics.

The Box Commons' founding industrial members hold patents covering AI agent credentialing methodologies, insurance modeling for autonomous systems, and governance simulation. These patents will be made available to the standards body on royalty-free terms, ensuring that the credentialing infrastructure can be adopted without proprietary barriers. This arrangement is explicitly procompetitive: it makes patented technology available to an industry consortium on open terms, rather than locking it behind licensing walls.

The guidelines should confirm that a royalty-free or FRAND licensing commitment by a patent holder to a standards body is a procompetitive act that weighs in favor of the licensor under rule-of-reason analysis. Specifically:

1. A FRAND or royalty-free commitment reduces the risk that a patent holder will use standard-essential patents to hold up the standard-setting process;
2. A FRAND or royalty-free commitment demonstrates that the patent holder's participation in the standard-setting process is motivated by the standard's success, not by exclusionary leverage;
3. The existence of a FRAND or royalty-free commitment should be treated as evidence of pro-competitive intent when evaluating the patent holder's conduct in standard-setting activities.

This is particularly important for small patent holders. A large technology company can absorb the litigation risk of participating in standard-setting without explicit safe harbors. A small patent holder cannot. Without clarity on the procompetitive treatment of FRAND commitments, small innovators are deterred from contributing their patented technologies to collaborative standards — the exact opposite of the outcome the antitrust laws are designed to promote.

C. Collective AI Threat Intelligence Sharing

The development of AI credentialing standards requires the sharing of sensitive technical information about AI vulnerabilities, failure modes, and adversarial attack vectors. Participants in a credentialing consortium must exchange data about how AI agents can be exploited, deceived, or caused to malfunction — information that, outside the standards context, could constitute commercially sensitive competitive intelligence.

The agencies should provide explicit guidance that the sharing of AI vulnerability data, adversarial testing results, and failure-mode analyses through a recognized standards development organization for the purpose of developing safety and credentialing standards does not constitute an unlawful exchange of competitively sensitive information, provided:

1. The information shared is limited to what is necessary for the standard-setting purpose;
2. The information is shared through a structured process with documented confidentiality protections;
3. Participants do not use the forum to exchange pricing, output, or customer allocation information; and
4. The resulting standards are available for adoption by any market participant, not limited to consortium members.

This carve-out is analogous to existing protections for information sharing in cybersecurity contexts under the Cybersecurity Information Sharing Act of 2015 (CISA), 6 U.S.C. § 1501 et seq., and extends the principle to the AI-specific context of agent credentialing.

V. RESPONSE TO QUESTION 3: SIGNIFICANT LEGAL, ECONOMIC, AND TECHNOLOGICAL DEVELOPMENTS

The agencies ask what legal, economic, or technological developments should be considered in revising the guidelines. Three developments are dispositive.

A. The Standards Development Organization Advancement Act and AI Standards

The Standards Development Organization Advancement Act of 2004 (SDOAA), 15 U.S.C. §§ 4301–4306, extended the protections of the National Cooperative Research and Production Act to standards development organizations. Under the SDOAA, an SDO that files a notification with the Department of Justice is entitled to have its standard-setting activities evaluated under the rule of reason rather than per se analysis, and successful antitrust plaintiffs are limited to actual damages rather than the treble damages ordinarily available under the Sherman Act.

These protections are significant but insufficient for the AI credentialing context. The SDOAA provides a damages limitation, not a safe harbor. An SDO that files an SDOAA notification still faces the full cost of antitrust litigation — discovery, expert testimony, trial — before reaching the question of damages. For a large industry consortium with institutional legal resources, this risk is manageable. For a small, newly formed standards body focused on an emerging technology, the litigation cost alone is prohibitive, regardless of the damages cap.

The revised guidelines should address this gap by providing that an SDO that has filed an SDOAA notification and maintains procedural safeguards consistent with *Allied Tube* and OMB Circular A-119 is presumptively procompetitive. This presumption would not immunize the SDO from antitrust scrutiny but would signal to the agencies — and to private plaintiffs — that the burden of demonstrating anticompetitive harm is substantial. The practical effect would be to reduce the deterrent that litigation risk poses to new entrants in the standards ecosystem.

B. The Asymmetric Threat to Small Standards Bodies

The withdrawal of the 2000 guidelines in December 2024 created a regulatory vacuum that disproportionately affects small organizations.

Google, Microsoft, Amazon, and OpenAI can participate in collaborative standard-setting with relative impunity. They maintain in-house antitrust counsel, have standing outside relationships with antitrust litigators, and can absorb the cost of defending against antitrust claims as a routine business expense. An adverse verdict would be financially significant but not existential.

A small standards body — a newly formed 501(c)(6) with limited initial funding, a volunteer board, and no institutional legal reserves — faces a fundamentally different risk calculus. A single antitrust lawsuit, even one that ultimately fails, can consume the organization's resources and effectively terminate its operations. The chilling effect is not theoretical. It is the predictable consequence of removing regulatory guidance without providing a replacement.

The revised guidelines should acknowledge this asymmetry and provide that the agencies will exercise prosecutorial discretion in favor of small standards bodies that demonstrate good-faith compliance with recognized procedural safeguards. This acknowledgment would not change the substantive law. It would provide the predictability that small organizations need to participate in collaborative standard-setting — which is the stated purpose of the guidelines.

C. The Convergence of Federal AI Standards and Private Credentialing

The National Institute of Standards and Technology has launched the AI Agent Standards Initiative through its Center for AI Standards and Innovation (CAISI). This initiative is developing identity and authorization protocols for autonomous AI systems. NIST has explicitly recognized that private-sector credentialing will complement federal standards — the federal government does not intend to become the sole credentialing authority for AI agents.

This creates a policy imperative for antitrust guidance. If the federal government is relying on private standards bodies to credentialize AI agents, the antitrust framework must enable those standards bodies to function. The revised guidelines should explicitly reference the complementary relationship between federal AI standards (NIST AI RMF, CAISI Agent Standards Initiative) and private credentialing, and should provide that private credentialing activity that operates within the NIST framework is entitled to a favorable inference under rule-of-reason analysis.

This is not a request for government endorsement of any particular credentialing body. It is a request for the agencies to acknowledge what NIST has already recognized: that private credentialing infrastructure is a critical component of AI governance, and that antitrust law should enable rather than impede its development.

VI. PROPOSED SAFE HARBOR FRAMEWORK FOR AI CREDENTIALING STANDARDS

Drawing together the analysis above, The Box Commons proposes that the revised guidelines include an explicit safe harbor for AI credentialing standard-setting that incorporates the following elements:

Safe Harbor A: SDO Safe Harbor for AI Security Standards

Consensus-based standard-setting for AI agent credentialing, safety, and security shall be evaluated under the rule of reason and shall be presumptively procompetitive when conducted by an organization that:

- Has filed or intends to file an SDOAA notification under 15 U.S.C. § 4305;
- Maintains open membership or participation;
- Uses a consensus process consistent with OMB Circular A-119;
- Publishes its standards for public review and comment; and
- Maintains documented conflict-of-interest policies.

Safe Harbor B: Collective AI Threat Intelligence

The sharing of AI vulnerability data, adversarial testing results, and failure-mode analyses through a recognized SDO for the purpose of developing credentialing standards shall not constitute an unlawful exchange of competitively sensitive information, provided the information exchanged is limited to what is necessary for the standard-setting purpose and does not include pricing, output, or customer allocation data.

Safe Harbor C: Credentialing Decisions

The denial, suspension, or revocation of a credential by an AI credentialing SDO shall not constitute a group boycott or concerted refusal to deal when:

- The credentialing criteria were developed through a consensus process;
- The criteria are published, technology-neutral, and publicly available;
- The applicant received notice of the specific deficiency;
- A documented appeals process was available; and
- The decision was based on published criteria, not competitive interests.

Safe Harbor D: Insurance-Linked Credentialing

The independent adoption by an insurer of a credentialing SDO's standards as an underwriting requirement shall not transform the SDO into a market gatekeeper for antitrust purposes, provided:

- The SDO maintains structural independence from adopting insurers;
- The credential criteria are technology-neutral;
- Multiple compliance pathways exist; and
- The insurer's adoption was an independent commercial decision.

Interaction of Safe Harbors

These safe harbors are designed to operate as a system. Safe Harbor A protects the formation and governance of the standards body. Safe Harbor B enables the technical work. Safe Harbor C

protects the credentialing output. Safe Harbor D protects the market mechanism — insurance — through which the credential achieves commercial significance.

Without all four, the system fails. A standards body that can form and deliberate (A and B) but faces litigation when it issues credentials (C) or when insurers adopt those credentials (D) will not survive. The framework must protect the full lifecycle of the credentialing process.

VII. CONCLUSION

The Department of Justice and Federal Trade Commission have an opportunity to shape the antitrust framework for AI governance for the next decade. The withdrawal of the 2000 guidelines has left a vacuum that the market cannot fill without legal clarity.

The Box Commons is a newly formed standards development organization whose founding industrial members have committed their patented technologies on royalty-free terms for collaborative AI credentialing standard-setting. The organization has engaged directly with NIST's AI Agent Standards Initiative and is building the credentialing infrastructure that the private insurance market and federal agencies need to govern AI agent deployment. This record reflects a commitment to the procompetitive deployment of technology through open standards — not exclusionary leveraging.

The AI agent credentialing ecosystem requires four specific safe harbors: protection for the standards body itself, protection for the technical information sharing that enables standard development, protection for credentialing decisions, and protection for the insurance linkage that gives credentials commercial force. Without this framework, the market for AI agent governance will be dominated by organizations large enough to absorb antitrust litigation risk — precisely the incumbents whose market power the antitrust laws exist to check.

The agencies' stated goal is to provide "businesses — especially in dynamic sectors — with predictability, compliance tools, and safety zones distinguishing procompetitive conduct from unlawful coordination." AI agent credentialing is the most dynamic sector requiring this predictability. The four safe harbors proposed in this comment provide the compliance tools. The procedural requirements embedded in each safe harbor — consensus processes, published criteria, appeals mechanisms, structural independence — provide the substantive safeguards against anticompetitive abuse.

Small standards bodies cannot afford to wait for case-by-case adjudication to establish that AI credentialing is procompetitive. The revised guidelines should establish that principle now, while the credentialing ecosystem is being built — not after it has been chilled into nonexistence by legal uncertainty.

The Box Commons appreciates the opportunity to submit this comment and welcomes further engagement with the agencies on these issues.

Respectfully submitted,

Brice Love Acting Executive Director The Box Commons c/o Northwest Registered Agent LLC 30 N Gould Street, Suite N Sheridan, WY 82801 info@theboxcommons.org