

September 24, 2026

Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive (Mail Stop 8930)
Gaithersburg, Maryland 20899-8930

Submitted by email to sp800-239-comments@nist.gov

Re: Comments of ACT on AI Data Center Security Analysis: A High-Performance Computing (HPC) Driven Approach, NIST SP 800-239 ipd

The Association for Competitive Technology (ACT) appreciates the opportunity to comment on AI Data Center Security Analysis: A High-Performance Computing (HPC) Driven Approach, the initial public draft of NIST Special Publication (SP) 800-239 (the Draft), released in July 2026 by the National Institute of Standards and Technology (NIST) and prepared jointly by its Information Technology Laboratory and its Center for AI Standards and Innovation.¹

I. Introduction and Statement of Interest

ACT is a not-for-profit trade association representing the global small business technology developer community. Our members are entrepreneurs, independent developers, and small firms that build software, connected devices, and data-driven services across nearly every sector of the economy. They take part in the data center economy as suppliers, building the monitoring, observability, and security software that operators use to run and protect these facilities, and as customers, since they reach cloud and AI capacity largely by renting it from data center operators rather than by owning infrastructure. A small firm that builds its product on a cloud AI service is a tenant in the sense this draft uses the word, and it usually has no dedicated security or compliance staff and little ability to negotiate the terms on which it buys capacity.

The Draft states that nongovernmental organizations may use it on a voluntary basis, and it describes itself as a threat and security gap analysis offering basic recommendations. We ask NIST to keep that framing in the final publication and in the security guidelines it

¹ National Institute of Standards and Technology, AI Data Center Security Analysis: A High-Performance Computing (HPC) Driven Approach, NIST SP 800-239 ipd (July 2026), <https://doi.org/10.6028/NIST.SP.800-239.ipd>. Section references in these comments are to that draft.

intends to develop from this analysis by drawing on NIST SP 800-223 and NIST SP 800-234. NIST security publications are regularly incorporated into federal procurement requirements and into the security questionnaires that operators and their customers exchange, so the recommendations in this Draft are likely to reach small firms as contract terms even though the Draft itself imposes no obligation.

ACT takes a risk-based approach to artificial intelligence, under which responsibility rests with the participant in the value chain best able to understand and address a given risk.² ACT's AI Roles and Interdependency Framework identifies the stakeholder groups across the AI value chain and the responsibilities of each,³ and our comments below apply the same approach to the three parties the Draft identifies in an AI data center, namely operators, model owners, and tenants.

II. ACT's Specific Comments on the Draft

As NIST develops follow-on guidance, ACT encourages it to apply three principles throughout: assign responsibility to the entity that controls the relevant risk; state security objectives in technology-neutral, outcome-based terms; and calibrate monitoring, documentation, and assurance requirements to the risk presented.

1. Allocating recommendations among operators, model owners, and tenants

Section 3.8 describes a three-way "trust dilemma" among tenants and users, model owners, and AI data center operators. Tenants need their applications and data to remain private and secure, model owners need to protect their weights and algorithms, and operators need to protect the facility from malicious code and tenant-driven attacks. That description matches how AI data centers are used in practice, but the recommendations in Section 4 do not carry it forward. Section 4.1 recommends logging user queries and model responses, Section 4.2 recommends monitoring tenants' usage and access patterns, and Section 4.4 recommends human approval before models are promoted to production, and none of these say whether it is addressed to the operator, the model owner, or the tenant.

We ask NIST to carry the three-party model from Section 3.8 into Section 4 and into the guidelines that follow, so that each recommendation identifies the party to which it is addressed. Each party should be responsible for the controls over what it builds and runs: an operator secures the facility, the shared hardware and firmware, the orchestration layer, and the access zone it operates; a model owner

² ACT, Policy Principles for AI, <https://actonline.org/post/ai-policy-principles/>.

³ ACT AI Roles and Interdependency Framework (May 2024), <https://actonline.org/post/ai-roles-responsibilities-in-the-united-kingdom/>. The Framework was also appended to ACT's comments to the Office of Science and Technology Policy on the development of an Artificial Intelligence Action Plan (14 Mar. 2025).

secures its weights and its release process; and a tenant secures its own application, its data, and the accounts through which its staff and customers reach its service. This is the shared responsibility model that cloud providers already publish for their customers, and aligning the guidelines with it would make them easier to apply.

A small tenant can meet its obligations in full for the layers it reasonably controls, but it cannot inspect the firmware of the accelerators it rents or test the isolation between its workloads and those of other tenants, and it should not be asked to attest to either. Where a tenant's own obligations depend on controls that an operator runs, the tenant needs reliable information about those controls, and Section 4.6 already points in this direction by describing continuous monitoring and compliance results that give users transparency into the security and compliance posture of the facility. We ask NIST to identify in the follow-on guidelines the security information a tenant needs from an operator to meet its own obligations, so that the tenant can rely on the operator's statement about the controls the operator runs. We also ask NIST to distinguish tenants from users, whom the draft treats as a single party, because a firm that rents capacity and the public who use that firm's application stand in different relationships to the operator.

2. Logging of prompts and model responses

Section 4.1 describes logging users' queries and API calls "along with the corresponding responses from AI models" as a key initial step in detecting suspicious activity, and Section 3.3 calls for "automated, real-time logging at every data touchpoint." The Draft does not address what should be logged, who holds the logs, how long the logs are kept, or who may read the logs. Logging of this kind serves a real security purpose, particularly in detecting model distillation and prompt injection, but full logging of content also creates a second store of sensitive data. The prompts and responses that pass through our members' applications often contain customer records, health information, source code, and confidential business information, and a copy retained in the access zone becomes a target in its own right.

Section 3.3 cites the privacy controls of the GDPR and the CCPA and the confidentiality requirements of HIPAA as reasons for rigorous data isolation, and Section 4.3 describes confidential computing as keeping data isolated even from the data center operator, so a default under which the full text of every prompt and response is retained in the access zone would work against both. Retaining that content for every request also adds storage and processing costs that are passed on to customers, and those costs weigh most on the smallest tenants, who buy capacity at published prices and cannot negotiate them.

We ask NIST to state that logging for security purposes should be proportionate to risk and consistent with data minimization. In practice, this means logging metadata by default, such as timestamps, account identifiers, request volumes, and the results of guardrail checks, which supports detection of many abuse patterns, including the high-volume querying associated with distillation. Full content should be retained only where a specific risk warrants it, for a defined period, and with access limited to security purposes. Where the content belongs to a tenant, the tenant should receive meaningful transparency and, where technically and operationally feasible, configurable controls over whether and how content is retained, subject to narrowly tailored security and legal requirements.

3. Human approval before models reach production

Section 4.4 recommends that an AI data center incorporate a mechanism that “requires final approval for critical actions,” with human oversight to ensure that permissions, version control, and approvals are in place before models are promoted to production. The draft identifies a real concern in the accidental deployment of an incorrect or unapproved model, but a fixed human approval step for every promotion to production would weigh most on small developers, who release often because speed of release is one of their few advantages over larger competitors, and who often have too few staff to separate the roles of developer and approver.

We ask NIST to frame this recommendation around its purpose, which is that only intended and authorized models reach production, and to recognize that automated controls can achieve that purpose. Signed model artifacts, version pinning, automated policy checks in the deployment pipeline, and fast rollback all guard against the accidental deployment the draft describes, and human approval can then be reserved for changes whose impact warrants it, as judged by the model owner. The guidelines should also make clear that the approval decision belongs to the organization that owns the model, while the operator’s role is to provide the permission and versioning capabilities through which the model owner exercises that control.

4. Outcome-based recommendations and small security suppliers

Many of the capabilities the Draft recommends, including AI gateway protections and guardrails that pre-evaluate prompts in Section 4.1, anomaly detection in Section 4.2, and continuous, real-time compliance monitoring in Section 4.6, are supplied to operators by specialized software firms, many of them small. When NIST develops the guidelines that follow from this analysis, we ask that it state its recommendations as security outcomes and not as particular architectures or product categories, so that operators can meet them with their own tools or with tools from third-party suppliers of any size. Recommendations tied to a particular

implementation tend to become procurement checklists that only established suppliers can satisfy, and they slow the adoption of better tools as the AI-powered attacks described in Section 3.11 continue to change.

The same concern applies to supply chain assurance, which the draft leaves to a correlative initiative. Section 3.7 rightly identifies the risk of compromised software dependencies, and we ask that assurance requirements for software suppliers be scaled to the risk a component presents, so that a small firm supplying a monitoring or security tool can demonstrate the integrity of its product without the compliance program of a large vendor. A broad base of domestic software suppliers, including small firms, is itself part of a secure and resilient data center supply chain.

III. Conclusion

ACT supports NIST's work to build a shared foundation for AI data center security through voluntary guidance, and the Draft's analysis of threats across the access, management, computing, and storage zones will be useful to operators and their customers alike. We ask for these changes so that the guidelines NIST develops from this analysis assign each recommendation to the party able to carry it out, protect tenants' data as well as the facility, and leave room for small firms both as customers of AI data centers and as suppliers of the tools that secure them.

We welcome the opportunity to discuss these comments with NIST.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

Chapin Gregor
Global Policy Counsel

Kedharnath Sankararaman
Policy Counsel

Association for Competitive Technology
1401 K St NW, Suite 501
Washington, DC 20005