TECHNET
THE VOICE OF AMERICAN INNOVATION

1420 New York Avenue NW, Suite 825
Washington, D.C. 20005
www.technet.org | @TechNetUpdate

September 16, 2026

National Institute of Standards and Technology
U.S. Department of Commerce
100 Bureau Drive
Gaithersburg, MD 20899

Re: NIST AI 300-1, Guidance and Templates for Public-Facing AI Documentation, An AI Standards "Zero Draft"

To Whom It May Concern:

TechNet appreciates the opportunity to provide comments on the National Institute of Standards and Technology's (NIST) initial public draft of *Guidance and Templates for Public-Facing AI Documentation*, NIST AI 300-1. TechNet is the national, bipartisan network of technology CEOs and senior executives that promotes the growth of the American innovation economy. Our diverse membership includes leading companies across the artificial intelligence (AI) ecosystem that develop, deploy, integrate, and rely upon AI technologies across a wide range of industries and use cases.

TechNet strongly supports NIST's continued leadership in developing voluntary, consensus-based approaches to AI standards. Well-designed documentation standards can improve communication among organizations across the AI lifecycle, support effective risk management, reduce duplicative compliance work, and make it easier for organizations to evaluate and responsibly integrate AI technologies. We particularly appreciate NIST's recognition that documentation processes should be manageable, context-sensitive, and calibrated to organizational objectives and available resources.

As NIST develops the final standards draft and advances it through the voluntary consensus standards process, TechNet encourages NIST to focus on four principles that will improve the framework's practical utility and international relevance: promoting global harmonization and interoperability with existing standards; recognizing shared responsibility across the AI lifecycle; preserving flexibility for organizations to use automated, human, or hybrid approaches to monitoring, evaluation, and risk mitigation; and facilitating the reuse of existing documentation to minimize duplicative burdens.

Promote Global Harmonization and Interoperability

TechNet strongly supports NIST's goal of developing the Zero Draft for eventual consideration through INCITS/AI and ISO/IEC JTC 1/SC 42. The value of the resulting standard will increase significantly if organizations can use it not only in the United

States, but also as part of globally interoperable AI governance, risk management, and compliance programs.

NIST has already taken important steps in this direction. The Zero Draft uses ISO/IEC drafting conventions and expressly contemplates eventual integration into the ISO/IEC standards ecosystem. It also appropriately identifies interoperability as an important characteristic of effective documentation and recognizes that standardized content and structure can improve reuse, comparison, accessibility, and automation.

TechNet encourages NIST to build on this foundation by explicitly mapping the final document's terminology, fields, profiles, and documentation objectives to relevant existing and emerging international and voluntary consensus standards. This should include relevant ISO/IEC standards, such as ISO/IEC 42001 on AI management systems, ISO/IEC 23894 on AI risk management, ISO/IEC 42005 on AI system impact assessments, the ISO/IEC 5259 series on data quality for analytics and machine learning, and ISO/IEC 23053 on AI systems using machine learning. NIST should also consider relevant standards developed through other standards development organizations, including IEEE standards addressing AI transparency, documentation, data governance, and model interoperability, as well as emerging CEN-CENELEC standards supporting implementation of the EU AI Act, particularly EN 18284 on quality and governance of datasets in AI. As these standards continue to develop, NIST should seek to minimize unnecessary divergence and enable organizations to leverage documentation developed under one recognized framework to satisfy substantially equivalent documentation objectives under another.

This interoperability is particularly important given the increasingly global nature of AI development and deployment. Organizations operating across jurisdictions should not be required to maintain parallel documentation systems simply because different standards use different terminology, structures, or formats to request substantially equivalent information. Such fragmentation increases engineering and compliance costs without necessarily improving AI risk management. NIST has an opportunity to promote greater global alignment by designing the final standard so that existing documentation and controls can be readily mapped, referenced, and reused across internationally recognized frameworks, allowing organizations to focus additional resources on areas where genuinely new information or risk-management measures are needed.

We therefore strongly encourage NIST to complete Annex C, "Relationships with Other ISO/IEC Standards on AI," as a substantive component of the final draft. Rather than providing only a general discussion of related standards, Annex C should identify relevant relationships between the Zero Draft's fields and established international AI standards and, where appropriate, indicate how existing documentation can satisfy or support the Zero Draft's objectives.

NIST should also clarify how the final standard relates to its broader portfolio of AI guidance and profiles, including the AI Risk Management Framework (AI RMF), AI Cybersecurity Framework Profile, AI for Critical Infrastructure work, and NIST

guidance on testing, evaluation, verification, and validation (TEVV). Where a field, control, or documentation practice satisfies an objective in another NIST framework, NIST should identify that relationship so practitioners can understand how the documents work together and avoid duplicative implementation efforts.
Similarly, NIST should explain how the final standard relates to emerging work on AI Bills of Materials (AIBOMs) and associated industry standards and open-source efforts, including work supported by the Linux Foundation. AIBOMs and public-facing model and dataset documentation can serve distinct but complementary purposes. Clarifying where their data elements overlap, can be referenced, or can be made machine-interoperable would help organizations integrate the final standard into existing AI supply-chain transparency and governance practices.

TechNet also recommends that Section 4.2.6, Artifact Interoperability, expressly recognize interoperability across governance and documentation frameworks, in addition to interoperability in documentation format and structure.

Suggested additional text for Section 4.2.6:

"To promote interoperability across jurisdictions and organizational governance frameworks, documentation artifacts should, where appropriate, be capable of referencing or incorporating information developed pursuant to recognized international standards, industry standards, or established organizational documentation practices. Organizations should not be expected to reproduce substantially equivalent information where an existing documentation artifact adequately addresses the relevant field or objective."

Similarly, NIST should make clear in Clause 5 and the profiles in Annex A that fields may be populated through references to existing documentation where doing so provides the information required by the relevant field. The draft already uses this approach in several contexts, including references to dataset documentation within model documentation. Applying this principle more broadly would reduce duplication and improve maintainability.

TechNet also encourages NIST to develop, or enable development of, a dedicated procurement profile. Public- and private-sector procurement presents a distinct documentation context in which purchasers need sufficient information to assess suitability, risk, security, and governance without imposing duplicative or bespoke disclosure requirements on providers. A procurement-focused profile could identify a common, risk-based set of documentation fields that purchasers can rely on while permitting providers to reference existing documentation and established assurance materials.

Finally, NIST should provide practical examples demonstrating how the templates and profiles can be implemented. Worked examples using NIST-developed AI tools or systems, such as Dioptra where appropriate, could illustrate how fields should be populated, how references to existing documentation can be used, and how organizations can apply the standard proportionately in different contexts. Such

examples would be especially useful during piloting and would help surface ambiguities before the document advances through the standards process.

Recognize Shared Responsibility Across the AI Lifecycle

TechNet encourages NIST to more explicitly recognize shared responsibility across the AI lifecycle. AI systems frequently involve multiple organizations that perform different functions and exercise control over different components of the system. These may include model developers, cloud service providers, infrastructure providers, application developers, integrators, deployers, customers, and operators.

Responsibility for identifying, managing, and mitigating risk should reflect these roles and the practical ability of each actor to control the relevant source of risk. An upstream model or infrastructure provider may be best positioned to address risks inherent to the model or service it provides, while a downstream deployer or operator may be best positioned to address risks arising from its particular system architecture, configuration, data, operational environment, or use case.

The Zero Draft already provides a foundation for this approach. It distinguishes model-level documentation from documentation of complete AI systems and recognizes that complete systems may include complex combinations of components beyond the model itself. It also describes model maintenance and monitoring in terms of both processes and responsibilities and, in the default model profile, appropriately conditions certain recommendations on whether a model provider actually deploys and operates the model.

TechNet encourages NIST to apply this principle consistently throughout the final profile and their other AI drafts. The framework should avoid language that could be interpreted as assigning an upstream provider responsibility for risks or controls that arise primarily from the decisions of downstream deployers, integrators, or operators. Conversely, downstream organizations should not be able to rely on the profile to shift responsibility for deployment-specific risks entirely to technology providers.

This principle is well established in cloud computing. Cloud service providers and their customers commonly use shared responsibility models to distinguish responsibilities controlled by the service provider from responsibilities controlled by the customer. Similar concepts can help clarify responsibility across the AI lifecycle.

NIST should explicitly recognize established shared responsibility models where appropriate and encourage documentation that helps parties understand these responsibility boundaries. This would make the Zero Draft more useful to both providers and downstream users by helping each organization identify which safeguards it can reasonably expect from another actor, and which safeguards remain within its own control.

TechNet recommends adding guidance to the Model Governance and Maintenance and Monitoring fields explaining this principle.

Suggested additional text for Model Governance, Root Field 8:

"Where multiple entities contribute to the development, provision, integration, deployment, or operation of an AI model or AI system, documentation may describe relevant allocations of responsibility among those entities. Responsibilities should reflect each entity's role, technical capabilities, access to relevant information, and ability to identify, manage, or mitigate the applicable risk. Documentation should not imply that responsibility for deployment-specific risks necessarily rests with the model provider where those risks arise from downstream configuration, integration, data, use, or operational decisions."

Suggested additional text for Maintenance and Monitoring, Root Field 7:

"Documentation should distinguish, where relevant, between monitoring and maintenance activities performed by the model provider and those expected to be performed by downstream deployers, integrators, operators, or other parties. Existing shared responsibility frameworks may be referenced where they clearly describe these respective responsibilities."

This clarification would also improve the usefulness of the profile for critical infrastructure and other highly complex environments. Critical infrastructure operators may rely on AI models, cloud services, software, proprietary systems, operational technology, and additional controls supplied by multiple entities. A documentation framework should assist these organizations in understanding responsibility boundaries, not provide a mechanism through which any participant in the lifecycle can shift all risk to another actor.

Suggested additional text for Roles and Responsibilities subfield:

TechNet recommends that NIST consider adding an optional "Roles and Responsibilities" subfield within the model profile. Such a field could identify relevant responsibilities retained by the model provider and those that depend upon downstream implementation, without attempting to comprehensively document the entire AI system.

Preserve Flexibility for Automated, Human, and Hybrid Risk Controls

TechNet supports NIST's technology-neutral approach to evaluation and encourages NIST to carry that principle consistently throughout the final document.

The default model profile appropriately recognizes that evaluation methods may be fully automated, performed by humans, or involve both approaches. NIST should apply the same principle when discussing monitoring, incident detection, risk treatment, safety controls, and other operational safeguards.

Human judgment plays an important role in many AI risk management activities. However, manual human intervention is not always technically feasible, operationally

scalable, or the most effective means of mitigating risk. In many environments, particularly those involving large-scale cloud services or high-volume AI deployments, automated controls can respond more rapidly, consistently, and effectively than manual intervention.

Appropriate safeguards may include automated monitoring, classifiers, access controls, rate limits, automated containment mechanisms, model or service-level restrictions, anomaly detection, technical policy enforcement, or other automated controls. In other circumstances, human review may be necessary. Many organizations will appropriately use a combination of both.

NIST should therefore avoid establishing, explicitly or implicitly, a hierarchy in which human intervention is presumed to be more effective or responsible than a properly designed automated control. The relevant question should instead be whether the safeguard is appropriate to the risk, effective in the relevant context, and subject to appropriate testing and validation.

This principle is particularly important for the Maintenance and Monitoring provisions of the model profile, which already contemplate technical monitoring mechanisms such as analysis of logs and classifiers.

Suggested additional text the following general guidance to the model profile:

"Risk management, monitoring, evaluation, and response mechanisms may be automated, human-led, or hybrid. Organizations should select approaches based on the nature and severity of the risk, technical feasibility, reliability, operational scale, latency requirements, and the effectiveness of the mitigation. This document should not be interpreted as establishing a general preference for manual human intervention where automated or hybrid controls can effectively achieve the relevant risk management objective."

To make this guidance operational, NIST should also provide greater clarity and consistency around the meaning of "risk" and "severity" across its AI guidance. The final standard should explain the factors organizations should consider when assessing the severity of a risk or use case and provide illustrative, non-exhaustive examples. This would help organizations calibrate documentation and controls proportionately while avoiding rigid classifications that may not translate across technologies, sectors, or deployment contexts.

Suggested clarification for the Monitoring Mechanisms subfield in the default model profile:

"Monitoring mechanisms may include automated, human, or hybrid processes. Automated mechanisms may be particularly appropriate where continuous monitoring, rapid detection, or response at scale is necessary."

This approach would preserve flexibility as AI systems and risk-management technologies continue to evolve while maintaining focus on outcomes rather than prescribing particular operational processes.

Facilitate Reuse of Existing Documentation and Minimize Duplicative Burdens

Across each of these areas, TechNet encourages NIST to allow organizations to leverage existing documentation wherever it satisfies the objectives of the relevant field, rather than requiring organizations to create duplicative materials.

NIST appropriately recognizes that documentation imposes real resource demands and that documentation processes must remain manageable. The final standard should therefore encourage organizations to reference existing materials wherever those materials provide the information needed for the relevant field.

Organizations today may already maintain model cards, system documentation, internal risk assessments, evaluation reports, certifications, cloud shared responsibility documentation, security documentation, governance records, or materials developed pursuant to international standards. Requiring organizations to translate the same information into multiple substantially overlapping artifacts could undermine the Zero Draft's stated goal of making documentation more efficient and useful.

NIST should make clear that conformity depends on whether the relevant information is available and appropriately structured or referenced, rather than whether an organization has created duplicative documentation specifically to demonstrate conformity with the final standard.

Suggested addition to Clause 5.1:

"Organizations should, where practical, reuse or reference existing documentation that satisfies the objectives and requirements of a field rather than creating duplicative documentation. References may include documentation developed pursuant to other recognized standards, governance frameworks, certifications, contractual arrangements, or established organizational practices, provided the referenced material is sufficiently accessible to the intended audience and contains the information necessary to satisfy the applicable field."

NIST could further operationalize this principle through a mapping table or implementation checklist showing how commonly used documentation tools and artifacts, such as model cards, evaluation reports, risk assessments, AI Bill Of Materials (AIBOMs), and other established documentation practices, correspond to fields in the final standard. Such a crosswalk would help organizations identify where existing materials already satisfy the standard and where additional documentation may be needed.

This approach would reinforce the draft's principles of interoperability, maintainability, and manageability while substantially improving its utility for organizations that already operate mature AI governance programs.

Maintain Flexible Default Templates

The ability to tailor documentation according to application domain, risk, technical characteristics, geographic considerations, and other contextual factors is important to avoiding a one-size-fits-all documentation regime.

As NIST finalizes the default guidance, we encourage the agency to preserve flexibility and avoid converting large numbers of recommended or optional fields into universal requirements. The information necessary to understand a model or dataset can vary substantially according to the technology, use case, intended audience, risk profile, organizational role, and availability of information.

NIST should therefore continue to reserve mandatory fields for information that is broadly necessary across contexts and use recommended, optional, or conditionally required designations for information whose relevance depends upon particular circumstances. Profiles developed for more specialized contexts can appropriately impose additional requirements where justified.

NIST should also ensure that the final standard does not require public disclosure of confidential, proprietary, security-sensitive, or otherwise protected information. Because the document is intended to govern public-facing documentation, each field and profile should be designed so organizations can satisfy the standard without exposing trade secrets, sensitive security information, personal data, or other information that would be inappropriate for public release. Where relevant information cannot be publicly disclosed, the standard should permit appropriately scoped descriptions, references, or indications that additional information exists through controlled channels.

TechNet also supports the draft's use of more precise documentation qualities such as "freshness" and "correctness." NIST should use consistent terminology across its AI standards and guidance and consider incorporating these concepts into future updates to the AI RMF and related materials where they provide greater operational clarity than broader or less precisely defined concepts. Greater terminological consistency across NIST's AI portfolio would make the documents easier for practitioners to implement together.

This flexible structure will be particularly important as the final standard enters the international consensus standards process. A baseline that is sufficiently adaptable across organizations, industries, technologies, and jurisdictions is more likely to achieve broad adoption and provide a durable foundation for additional sector-specific profiles.

Conclusion

TechNet appreciates NIST's continued leadership in developing voluntary, consensus-based AI standards and its commitment to incorporating real-world stakeholder experience into the Zero Draft process. Standardized approaches to AI documentation have the potential to improve interoperability, facilitate responsible AI adoption, and reduce unnecessary friction across the AI ecosystem.

As NIST advances this work toward a final standard, we encourage it to ensure that the resulting framework promotes global harmonization and interoperability with existing international standards; reflects shared responsibility across the AI lifecycle based on each actor's role, control, and capabilities; preserves flexibility for organizations to use automated, human, or hybrid approaches to risk management; and enables organizations to leverage existing documentation rather than creating duplicative materials. Together, these principles will help ensure that the final standard is practical, scalable, globally interoperable, and capable of supporting both responsible AI development and continued innovation. TechNet looks forward to continuing to work with NIST and other stakeholders as this effort advances through the voluntary consensus standards process.

Sincerely,

Linda Moore
President and CEO