Nuclear Security ISAC

Using AI for NRC Submittals: What is Allowed and What is Required

Establishing Human Oversight and Accountability Under Existing Regulatory Frameworks

Date: 06-05-2026

Table of Contents

Using AI for NRC Submittals: What is Allowed and What is Required

Establishing Human Oversight and Accountability Under Existing Regulatory Frameworks

Acceptability of AI-Assisted NRC Submittals

This paper defines the conditions under which AI-assisted drafting is acceptable within NRC regulatory frameworks and clarifies the responsibilities that remain with licensees. It establishes a governance model to ensure that AI use aligns with regulatory expectations for accuracy, accountability, and human oversight.

As the commercial nuclear industry integrates artificial intelligence (AI) and large language models (LLMs) to improve efficiency, questions have emerged regarding the regulatory acceptability of using AI to draft compliance documentation. The U.S. Nuclear Regulatory Commission (NRC) does not regulate AI as a standalone technology. Instead, it evaluates how AI-enabled tools are used within its existing regulatory authority. This outcome-based approach creates a gap between industry practice and explicit published guidance.

As of June 2026, the NRC has not issued comprehensive guidance addressing AI use in regulatory document preparation. However, recent regulatory gap analyses conclude that the current framework is generally sufficient to accommodate AI. In practice, AI-assisted drafting is being treated as acceptable under existing regulatory structures. This creates a divergence between formal guidance and practical implementation. While the NRC continues to emphasize that licensees are responsible for the accuracy and adequacy of their submissions, the absence of explicit boundaries introduces ambiguity. Licensees must therefore interpret how these tools fit within established regulatory constructs, increasing the importance of disciplined internal governance and qualified human oversight.

AI-assisted drafting is acceptable under current NRC regulatory frameworks because these tools are administrative in nature and do not control safety-related structures, systems, or component (SSC) functions. They do not trigger software verification and validation requirements or cybersecurity protections under 10 CFR 73.54. The NRC evaluates the accuracy and regulatory compliance of the final submission, leaving the choice of drafting tools to the licensee.

Key Policy Implications

The use of AI in NRC submittals does not reduce or transfer regulatory responsibility. Under 10 CFR 50.9, licensees are required to ensure that information submitted to the NRC is complete and accurate in all material respects or face penalties, sanctions or even criminal prosecution. This obligation applies regardless of the tools used to develop the submission. While AI tools may generate draft content, the legal and regulatory burden remains entirely with the licensee.

At the same time, the NRC has not issued explicit guidance defining acceptable use of AI in regulatory document preparation. This creates operational ambiguity, requiring licensees to interpret how AI-assisted drafting aligns with existing requirements, including quality assurance provisions such as 10 CFR 50 Appendix B.

This gap between regulatory sufficiency and practical implementation introduces risk. AI systems can produce technically sophisticated and highly confident outputs that are incorrect or unsupported. Without strong internal controls, these errors may not be detected before submission.

As a result, the effectiveness of human oversight becomes the defining factor in risk mitigation. Not all reviews are equivalent. Superficial or administrative review is insufficient to identify subtle technical errors. Effective oversight requires qualified subject matter experts with the experience necessary to validate technical accuracy, regulatory alignment, and underlying assumptions.

Executive Summary

Artificial intelligence (AI) is rapidly being adopted across the nuclear industry to support the development of regulatory documentation. The U.S. Nuclear Regulatory Commission (NRC) evaluates the adequacy and accuracy of submitted information rather than the specific tools used to produce it. As a result, AI-assisted drafting is permissible under existing regulatory frameworks.

However, the absence of explicit guidance defining acceptable AI use creates ambiguity for licensees. While the current framework is sufficient, it requires organizations to interpret how AI-assisted processes align with existing regulatory obligations and quality assurance requirements.

The primary risk associated with AI-assisted drafting is not tool use itself, but the potential for technically sophisticated yet incorrect outputs. This risk cannot be mitigated through superficial review. Effective oversight requires qualified subject matter experts capable of validating technical accuracy and regulatory alignment.

To address this, this paper defines the conditions under which AI-assisted drafting is acceptable within NRC regulatory frameworks and establishes a structured governance approach based on three layers: AI-assisted generation, human validation, and human authorization. This model ensures that AI remains an administrative tool while preserving full human accountability for all regulatory submissions.

This paper represents the first in a series addressing the use of artificial intelligence in nuclear cybersecurity and regulatory applications. It focuses specifically on AI-assisted drafting of regulatory documentation, primarily leveraging large language models, and does not address broader policy considerations related to other AI modalities, including visual or operational AI systems.

Background and Emerging Practices

Recent advances in machine learning, particularly generative AI technologies, including large language models (LLMs) and other emerging modalities such as visual AI, have driven the nuclear industry to explore automating repetitive administrative tasks. Industry stakeholders are using AI architectures and Retrieval-Augmented Generation (RAG) to assist in generating incident reports, corrective action program (CAP) evaluations, environmental impact statements, and 10 CFR 50.59 applicability screenings (Blue Wave AI Labs, 2024; Nuclearn, 2025).

AI systems can ingest large volumes of regulatory precedent, licensing basis documents, and historical plant records to rapidly generate draft content and extract relevant compliance data (U.S. Nuclear Regulatory Commission, 2024; Suman, 2021). This capability offers significant efficiency gains in document development and review processes.

The use of AI in highly regulated environments introduces questions related to quality assurance, compliance, and the role of human expertise. The NRC evaluates the adequacy and accuracy of submissions rather than the tools used to produce them (Pensado et al., 2024). However, AI-generated outputs may appear technically sound while containing incorrect or unsupported conclusions, creating a risk that is not immediately apparent (NIST, 2026).

Regulatory Basis for AI-Assisted Drafting

The NRC does not regulate artificial intelligence as a standalone technology. Instead, it evaluates how AI-enabled tools are used within activities that fall under its existing regulatory authority (NRC, 2023).

The agency's Artificial Intelligence Strategic Plan (NUREG-2261) and subsequent AI Regulatory Gap Analysis (AIRGA) conclude that the current regulatory framework is sufficient to accommodate the use of AI in nuclear applications (Pensado et al., 2024). With respect to regulatory documentation, the analysis confirms that existing requirements do not prescribe the tools used to prepare and produce required reports and submissions (Pensado et al., 2024).

This establishes the core regulatory principle: the NRC evaluates the adequacy of the submission, not the method of its preparation. This approach aligns with federal guidance from the Office of Management and Budget (OMB), which directs agencies to regulate AI applications based on outcomes and statutory authority rather than the underlying technology (OMB, 2020).

Policy Determination: AI Drafting as an Administrative Function

AI-assisted drafting tools are classified as administrative functions within the NRC regulatory framework. The NRC evaluates the adequacy and accuracy of submitted outcomes, not the tools used to generate them. Because AI systems used for document preparation do not perform safety functions, control plant systems, or execute licensing calculations, they do not fall within the scope of safety-related structures, systems, and components.

As a result, these tools do not trigger the regulatory requirements associated with safety-related software, including formal software verification and validation processes or cybersecurity protections applied to critical digital assets. AI-assisted drafting remains outside the boundary of operational and and safety and important-to-safety systems.

This classification establishes the basis for their permissibility. AI may be used to support the development of regulatory documentation, provided that the final submission meets all applicable requirements for accuracy, completeness, and regulatory compliance. Responsibility for those outcomes remains entirely with the licensee. This determination is consistent with 10 CFR 50.9, which requires that all information submitted to the NRC be complete and accurate, regardless of how it is generated.

Required Human Oversight Model

AI-assisted drafting requires structured and qualified human oversight. While these tools are already being used across the industry, often through informal or ad hoc processes, their use in regulatory workflows demands clearly defined roles and accountability.

Not all human review is equal. The primary risk of AI-assisted drafting is the generation of technically sophisticated and highly confident outputs that may be incorrect or unsupported. Superficial or administrative review is insufficient to identify these errors. Effective oversight requires qualified subject matter experts with the experience necessary to validate technical accuracy, regulatory alignment, and underlying assumptions.

Effective human validation requires traceability of source material used to generate and support AI-assisted content. Reviewers must be able to identify, verify, and trace all technical statements to authoritative sources, including licensing basis documents, regulatory guidance, and approved analyses. AI-generated content that cannot be traced to verifiable sources should not be accepted without independent validation.

This oversight is implemented through two complementary roles: human-in-the-loop (HITL) and human-in-command (HIC).

Human-in-the-loop (HITL)

Human-in-the-loop (HITL) involves active validation during document development. The reviewer functions as a technical filter, applying established human performance tools such as self-checks, peer reviews, and independent verification to identify errors and ensure accuracy before the document advances.

Human-in-command (HIC) and the NRC Autonomy Model

Human-in-command (HIC) and the NRC Autonomy Model represents the final authority responsible for the submission. Within the NRC's autonomy framework, AI-assisted drafting operates at the lowest levels of machine independence, where all outputs must be vetted and executed by a human decision-maker. The HIC assumes full legal and technical responsibility for the content of the submission.

AI may assist in generating and organizing information, but it does not replace human judgment. Responsibility for validation, decision-making, and submission remains entirely with qualified personnel.

Recommended Governance Framework for AI-Assisted Drafting

AI-assisted drafting requires a structured governance model that defines the boundaries between machine assistance and human accountability. This model aligns with the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0), particularly the GOVERN and MAP functions, which emphasize defined human roles and operator proficiency (NIST, 2023).

Implementation is achieved through a three-layer framework:

1. **AI-Assisted Generation**
   The AI tool acts purely as an administrative engine. Driven by specific, narrow user prompts, the AI agent searches vetted, localized datasets to extract requirements, formulate outlines, and generate a preliminary text draft. The AI's role ends at text generation. It possesses no authority to finalize or submit the document.

2. **Human Validation**
   The generated draft enters the HITL phase. A qualified engineer or designated subject matter expert assumes control of the document to conduct a substantive review. This includes verifying all technical claims against primary sources, ensuring traceability to authoritative references, confirming alignment with current regulatory requirements, and ensuring that no unsupported or fabricated information is included (NIST, 2026).

3. **Human Authorization**
   The finalized document proceeds to the HIC. This is the accountable manager or executive who holds the legal authority to submit documentation to the NRC. By signing the document, the HIC formally adopts the contents of the submission as their own.

   This framework ensures that AI remains an administrative tool while enforcing rigorous human validation, source verification, and final accountability at every stage of the process.

Clarification on Regulatory Boundaries (V&V and Cybersecurity)

AI-assisted drafting tools do not fall within the scope of software verification and validation (V&V) or cybersecurity requirements applied to safety-related systems. These tools are administrative in nature and are not connected to plant operational or and safety and important-to-safety functions.

A common misconception is that all software used by a licensee must meet the same requirements as systems supporting plant operations. Under 10 CFR 50 Appendix B, quality assurance criteria apply to structures, systems, and components that prevent or mitigate the consequences of postulated accidents, as well as activities affecting those systems. AI tools used for document generation do not perform safety

functions, control instrumentation and control systems, or execute licensing calculations. As a result, they are not subject to formal V&V requirements for safety-related software.

Similarly, AI-assisted drafting tools do not meet the definition of Critical Digital Assets under 10 CFR 73.54. Cybersecurity requirements focus on protecting systems that directly affect safety, security, or emergency preparedness. AI drafting tools operate within enterprise information technology environments and do not interface with operational technology or plant control systems. They cannot manipulate plant equipment or alter safety parameters.

Because these tools remain outside the boundary of safety and operational systems, their use does not trigger the cybersecurity controls required for protected assets under NRC regulations. This includes maintaining traceability of source material to support auditability and regulatory defensibility.

Implications for Industry and Regulators

The integration of AI into regulatory workflows introduces significant efficiency gains, but it does not change the underlying responsibilities of licensees or the regulatory expectations of the NRC.
The current framework provides flexibility for innovation, but that flexibility is coupled with strict accountability. Licensees retain full responsibility for the accuracy and completeness of all submitted information, regardless of the tools used to produce it.

If a flawed assumption bypasses a weak human review, the resulting regulatory violation belongs to the licensee, not the software provider.

For industry, this means that adopting AI requires disciplined governance and expert-driven validation. For regulators, it reinforces that existing outcome-based frameworks remain effective, provided oversight continues to focus on the adequacy and accuracy of submitted information.

Conclusion

AI-assisted drafting is acceptable under current NRC regulatory frameworks. Because these tools are administrative in nature and do not perform and safety and important-to-safety functions, they do not trigger software verification and validation or cybersecurity requirements associated with protected systems. The NRC continues to evaluate the adequacy and regulatory compliance of the final submission, leaving the choice of drafting tools to the licensee.

However, the use of AI does not reduce responsibility. The primary risk lies in the potential for technically credible but incorrect outputs. Human presence alone is not sufficient. Effective oversight requires qualified expertise capable of validating technical accuracy and regulatory intent.

By implementing structured governance through human-in-the-loop validation and human-in-command authorization, organizations can realize the efficiency benefits of AI while preserving the accountability and rigor that underpin nuclear safety.

Disclosure

This document was developed by human authors and may have utilized AI (AI) tools to assist in drafting, editing, or organizing content. Consistent with the governance principles outlined in this paper, all AI-assisted content was subject to human validation and review by qualified subject matter experts.

The authors retain full responsibility for the accuracy, completeness, and regulatory alignment of the information presented in this document.

References

1. Blue Wave AI Labs. (2024). Blue Wave AI Labs 50.59 screener [Presentation] [ML24263A261].
2. Cybersecurity and Infrastructure Security Agency, Australian Signals Directorate, National Security Agency, et al. (2025). Principles for the secure integration of artificial intelligence in operational technology.
3. National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1
4. National Institute of Standards and Technology. (2026). Practices for automated benchmark evaluations of language models (NIST AI 800-2 ipd). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.800-2.ipd
5. Nuclearn. (2025). 50.59 – Prior NRC approval? Integrating AI into design processes [ML25273A301].
6. Office of Management and Budget. (2020). Guidance for regulation of artificial intelligence applications (M-21-06). Executive Office of the President.
7. Pensado, O., LaPlante, P., Hartnett, M., & Holladay, K. (2024). Regulatory framework gap assessment for the use of artificial intelligence in nuclear applications [ML24290A059]. Southwest Research Institute; Prepared for the U.S. Nuclear Regulatory Commission.
8. Suman, S. (2021). Artificial intelligence in nuclear industry: Chimera or solution? Journal of Cleaner Production, 278, 124022. https://doi.org/10.1016/j.jclepro.2020.124022
9. U.S. Nuclear Regulatory Commission. (2024). Applying a natural language processing model to analyze regulatory documents [ML24345A023].
10. U.S. Nuclear Regulatory Commission. (2023). Artificial intelligence strategic plan: Fiscal years 2023–2027 (NUREG-2261) [ML23132A305].
11. U.S. Nuclear Regulatory Commission. (2025). Nuclear AI insights [Presentation] [ML25273A290].
12. U.S. Nuclear Regulatory Commission. (2013). Regulatory Guide 1.168: Verification, validation, reviews, and audits for digital computer software used in safety systems of nuclear power plants (Rev. 2) [ML13073A210].
13. U.S. Nuclear Regulatory Commission. (n.d.). 10 CFR 50.9 – Completeness and accuracy of information. Code of Federal Regulations. https://www.ecfr.gov/current/title-10/part-50/section-50.9
14. U.S. Nuclear Regulatory Commission. (n.d.). 10 CFR Part 50, Appendix B – Quality assurance criteria for nuclear power plants and fuel reprocessing plants. Code of Federal Regulations. https://www.ecfr.gov/current/title-10/part-50/appendix-Appendix%20B%20to%20Part%2050