



INL/RPT-25-87369  
Revision 1  
DATE: 03/18/2026

# Evaluation of AI-Enabled Digital Documented Safety Analysis

March 2026

---

*Proposed Methodology for Evaluating  
Microsoft's Generative AI for Permitting  
Solution Accelerator in Nuclear Safety  
Basis Document Generation within  
NRIC's Digital Ecosystem*

Sonali Sinha Roy, Angelica Petrovic, Harleen Kaur Sandhu, Kevin  
O'Rear, Peter Suyderhoud, Jacob Rymer

*Idaho National Laboratory*





## DISCLAIMER

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

## REVISION LOG

<b>Revision No.</b>	<b>Date</b>	<b>Affected Pages</b>	<b>Description</b>
0	09/10/2025	All	New issue.
1	03/18/2026	All	Revision for comment incorporation from DOE.

## **ABSTRACT, SUMMARY, AND FOREWORD**

The National Reactor Innovation Center (NRIC) is leading a transformative initiative to accelerate advanced reactor deployment by fundamentally reimagining how nuclear safety basis documentation is developed, reviewed, and maintained. Traditional Documented Safety Analysis (DSA) processes for DOE-authorized facilities rely on static, document-centric workflows that consume significant time and resources, exemplified by recent major licensing efforts requiring hundreds of thousands of staff hours and millions of pages of documentation review. These conventional approaches create barriers to the rapid, cost-effective deployment of advanced reactors that America's future energy needs demand.

NRIC's DOE Authorization Digital Transformation Project addresses these challenges through an innovative framework that integrates artificial intelligence (AI), digital engineering, and systems-based data management into a cohesive digital ecosystem. This white paper presents NRIC's methodology for evaluating AI-enabled document generation capabilities within this broader digital infrastructure, using the Demonstration of Microreactor Experiments (DOME) facility as a pilot case study. The evaluation will assess an AI tool's ability to generate a Preliminary Documented Safety Analysis (PDSA) through progressive integration stages—from standalone document processing to full digital thread connectivity—while maintaining rigorous verification, validation, and regulatory acceptance standards.

By establishing dynamic, traceable connections between design data and safety documentation, NRIC's approach has the potential to reduce both document development time and regulatory review cycles by as much as 50%, while simultaneously improving accuracy, consistency, and traceability. This initiative represents a critical step toward establishing reusable digital infrastructure that reactor developers can leverage to accelerate their path from concept to commercial operation, directly supporting NRIC's mission to demonstrate and deploy advanced nuclear energy technologies.

## ACKNOWLEDGEMENTS

This document was sponsored by the National Reactor Innovation Center (NRIC). NRIC is a national program funded by U. S. Department of Energy's Office of Nuclear Energy and is dedicated to the demonstration and deployment of advanced nuclear energy. Neither the U.S. Government nor any agency thereof makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe on privately owned rights. References herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise do not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof. The authors would like to acknowledge the Microsoft team, including Nelli Babayan, Conor Kelly, and Henning Kilset, for facilitating access to their tool, *Generative AI for Permitting Solution Accelerator*, which will be used as part of the proposed case study.

## CONTENTS

ABSTRACT, SUMMARY, and FOREWORD .....	ii
ACKNOWLEDGEMENTS .....	iii
ACRONYMS .....	v
1. OVERVIEW OF DIGITAL DOCUMENTED SAFETY ANALYSIS .....	1
2. PROPOSED EVALUATION METHODOLOGY.....	2
2.1 Tool Review .....	3
2.2 Case Study Setup and Document Generation.....	4
2.2.1 PDSA Generation Using Standard RAG and User Interface .....	4
2.2.2 PDSA Generation Using a Connected Database .....	4
2.2.3 PDSA Generation through Digital Thread Integration .....	5
2.2.4 PDSA Generation Using Alternative Templates.....	6
2.3 Verification and Validation .....	6
2.3.1 Assessment of AI-Generated Results .....	6
2.3.2 Subject Matter Expert Assessment.....	6
2.3.3 Regulatory Acceptance Test.....	7
3. FUTURE ROADMAP FOR HOLISTIC EVALUATION .....	7
3.1 Risk Identification and Mitigation .....	7
3.2 Cost versus Benefits Assessment .....	7
3.3 Cross-Agency Review of AI Regulatory Frameworks .....	8
3.4 Evaluation Report and Improvement Roadmap.....	8
4. REFERENCES.....	9

## FIGURES

Figure 1. Proposed evaluation methodology workflow. ....	3
Figure 2. Example of a knowledge graph. ....	5
Figure 3. Digital thread integration.....	5

## ACRONYMS

AI	Artificial Intelligence
DSA	Documented Safety Analysis
DOE	Department of Energy
DOME	Demonstration of Microreactor Experiments
EBR	Experimental Breeder Reactor
FAA	Federal Aviation Administration
FDA	Food and Drug Administration
INL	Idaho National Laboratory
LLM	Large Language Model
NHTSA	National Highway Traffic Safety Administration
OCR	Optical Character Recognition
OSHA	Occupational Safety and Health Administration
NLP	Natural Language Processing
NRC	Nuclear Regulatory Commission
NRIC	National Reactor Innovation Center
PDSA	Preliminary Documented Safety Analysis
RAG	Retrieval-Augmented Generation
SME	Subject Matter Expert
U.S.	United States

# Evaluation of AI-Enabled Digital Documented Safety Analysis

## 1. OVERVIEW OF DIGITAL DOCUMENTED SAFETY ANALYSIS

The National Reactor Innovation Center (NRIC) is leading an effort to accelerate nuclear regulatory processes by utilizing generative AI for safety basis documentation. The development of safety basis and licensing documentation for nuclear reactors remains a major challenge in the deployment of advanced nuclear technologies. Traditional Documented Safety Analysis (DSA) processes can be improved by dynamic documents with better traceability and new safeguards to reduce the likelihood of human error. The regulatory review and engagement processes are similarly constrained, often requiring significant effort and extensive manual verification. For instance, the United States (U.S.) Nuclear Regulatory Commission (NRC) review of the NuScale application required over 250,000 staff hours and the evaluation of approximately two million pages of documentation [Bergman 2021]. To accelerate the commercial deployment of advanced reactors, the National Reactor Innovation Center (NRIC) and Idaho National Laboratory (INL) propose a digital DSA framework that integrates safety analysis and licensing documentation into a broader digital engineering ecosystem [Lund et al. 2024]. This approach would create a dynamic, data-driven safety basis that evolves with the reactor design. As a continuation of the effort to build this digital ecosystem, NRIC aims to utilize AI-enabled document generation to further accelerate the delivery of advanced nuclear technologies.

The integration of digital engineering and AI technologies has demonstrated substantial benefits in reducing project risk and improving efficiency. For instance, digital twin technologies have been shown to reduce schedule delays by approximately 20% [Ritter and Rhodes 2023], and GE Vernova reported \$1.6 billion in cost savings through digital engineering initiatives [GE Vernova]. AI-powered document processing tools complement these advances, outperforming traditional optical character recognition (OCR) by understanding context and relationships within text and enabling intelligent automation of complex documentation tasks [Patel 2025]. Mandelli et al. [2023] demonstrated that natural language processing (NLP) can extract structured data from unstructured equipment reliability reports to support nuclear safety analysis.

Companies like Microsoft are actively developing commercial solutions for automated documentation generation in nuclear safety analysis and licensing. A methodology is proposed in this paper for evaluating the feasibility and performance of a tool implementing the digital DSA framework through a targeted application case study. The case study will evaluate the feasibility of using a generative AI tool to create safety documentation from existing data within the NRIC digital ecosystem. The Microsoft Generative AI for Permitting Solution Accelerator has been selected as the digital tool under evaluation.

## 2. PROPOSED EVALUATION METHODOLOGY

The AI tool developed by Microsoft, called Generative AI for Permitting Solution Accelerator, is designed to support automated document generation for nuclear regulatory applications. The NRIC team proposes a phase-wise methodology for evaluating the efficacy of Microsoft's AI tool, as illustrated in Figure 1. The methodology not only tests the tool's technical capabilities but also assesses its suitability for nuclear licensing workflows, ensuring alignment with DOE safety documentation standards and regulatory expectations. This methodology will be implemented for a structured case study in which the Microsoft Generative AI for Permitting Solution Accelerator tool will generate the Preliminary Documented Safety Analysis (PDSA) for the Demonstration of Microreactor Experiments (DOME) test facility, which is being developed and constructed by NRIC at INL. The NRIC-DOME test bed is repurposed from the containment structure of the Experimental Breeder Reactor-II (EBR-II) and is designed to host experimental reactors up to 20 megawatts thermal, which will enable industry testing of advanced reactor concepts.

A PDSA is the preliminary safety basis for DOE facilities associated with the as-designed configuration. It describes the facility, its safety structures, systems, and components, provides hazard and accident analyses, and details the safety management plan. A PDSA is followed by a more comprehensive DSA prior to turning over the as-constructed facility to operations.

The purpose of this evaluation is to utilize Microsoft's AI tool to generate the PDSA for the NRIC-DOME test bed. Initially, the official DOE-approved DOME PDSA will serve as the benchmarking reference against which the AI-generated version will be evaluated during the verification and validation phase. If the tool is able to satisfactorily generate documents in known, predefined formats, the next step would be to generate documents from the same data in multiple formats or templates.

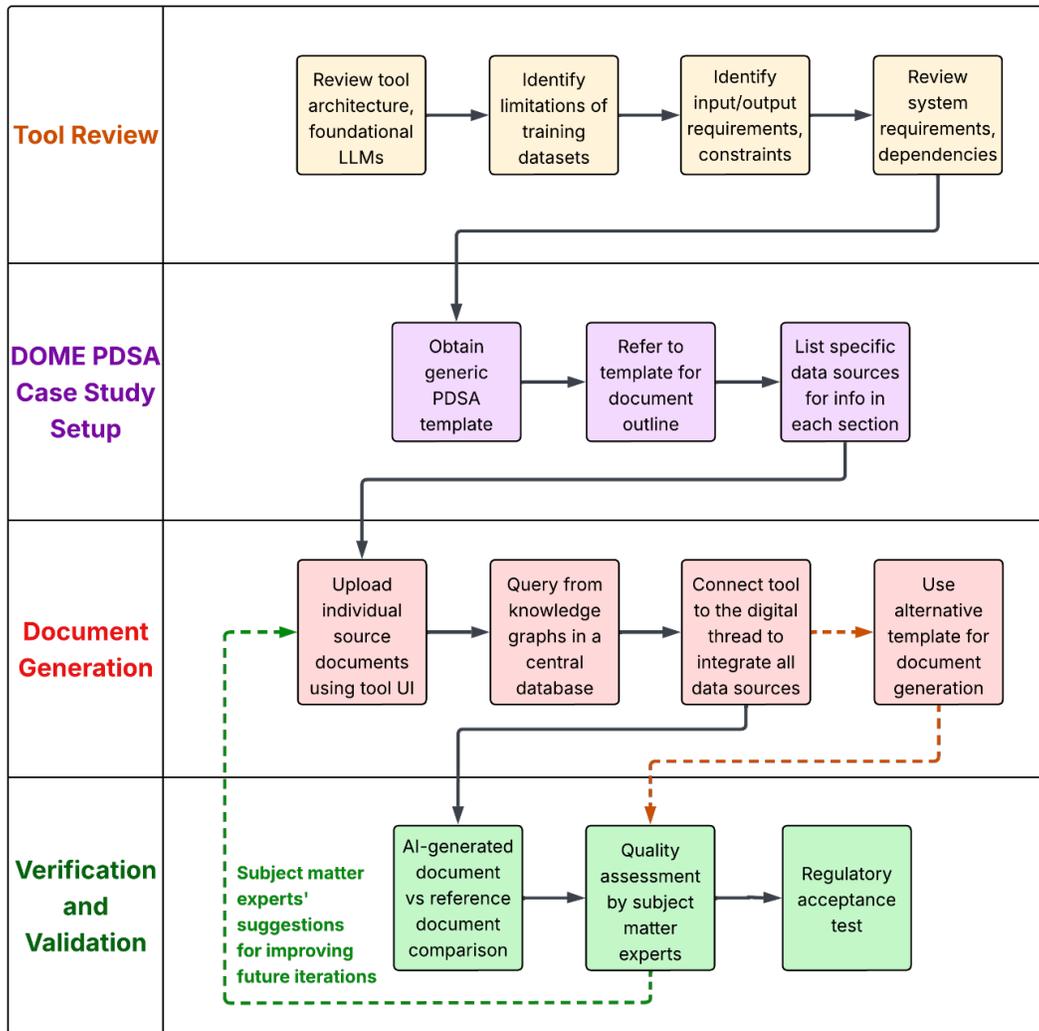


Figure 1. Proposed evaluation methodology workflow.

## 2.1 TOOL REVIEW

The evaluation begins with a comprehensive review of the Microsoft Generative AI for Permitting Solution Accelerator tool. The Microsoft tool comes with commercially available large language models such as gpt-4o for text generation, embedding, and reasoning. Models may be interchangeable from available options on the Azure OpenAI Service/AI Foundry. To effectively use the Microsoft tool, the evaluation team must be aware of any system requirements or dependencies, prompt-formatting requirements, retrieval-augmented generation (RAG) document constraints, alternative input interfaces (e.g., JSON, XML), and output structure constraints.

## **2.2 CASE STUDY SETUP AND DOCUMENT GENERATION**

The evaluation team will use a generic PDSA template to create the document outline in the Microsoft Generative AI for Permitting Solution Accelerator tool. This generic template will guide the prompts used to generate content for each section of the DOME PDSA. Data sources for each section will be identified using the DOE-approved DOME PDSA and through consultations with the nuclear safety analysts and design engineers who created it.

A progressive approach is proposed to integrate Microsoft's tool into the NRIC digital ecosystem. The evaluation approach involves using a constant set of project data sources for each PDSA section. However, these data sources are housed in different storage locations in different stages of document generation maturity. Sections 2.2.1 through 2.2.3 identify the increasing connectivity of successive stages.

### **2.2.1 PDSA GENERATION USING STANDARD RAG AND USER INTERFACE**

In the first stage, the documents identified as data sources (e.g., requirements, design documents, engineering calculation and analysis reports) for the PDSA will be uploaded through the Microsoft tool's user interface to an object storage location within INL's Microsoft Azure GovCloud tenant. The tool vectorizes the documents to create a knowledge base for RAG, which is then utilized to generate sections of the PDSA in response to user prompts. These prompts help the tool identify the desired content for the section and the data source to use for generating its response.

### **2.2.2 PDSA GENERATION USING A CONNECTED DATABASE**

The second stage of document generation will connect the Microsoft AI tool to INL's DeepLynx Nexus data warehouse. This database will act as a centralized repository for the data sources which were previously individually uploaded to the Microsoft tool. DeepLynx Nexus will create a knowledge graph of the information contained within these sources and the relationships between various entities, as illustrated by Figure 2. Each PDSA section will be generated by pulling information from this external database via query syntax within the section prompts. The knowledge graphs created within DeepLynx Nexus utilize predefined ontologies and formal semantics, providing better context for RAG.

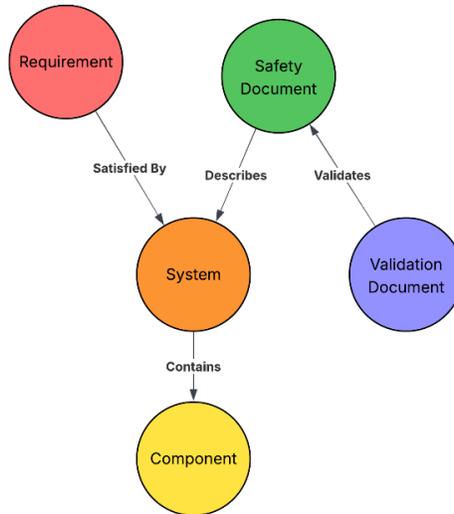


Figure 2. Example of a knowledge graph.

### 2.2.3 PDSA GENERATION THROUGH DIGITAL THREAD INTEGRATION

The third and most advanced stage of document generation will focus on connecting the Microsoft Generative AI for Permitting Solution Accelerator tool to an integrated digital thread solution. Instead of just acting as a repository for source files, DeepLynx Nexus will be connected to independent project databases where design information is stored and linked. DeepLynx will dynamically pull information from these project data sources for use in the generation of documents, as shown in Figure 3. Knowledge graphs will be generated using the information and relationships contained within the project source-of-truth databases. Therefore, DeepLynx Nexus will be used as a data broker between the Microsoft tool and the original data source, thereby maintaining a dynamic and traceable digital thread and ensuring that up-to-date information is used for document generation.

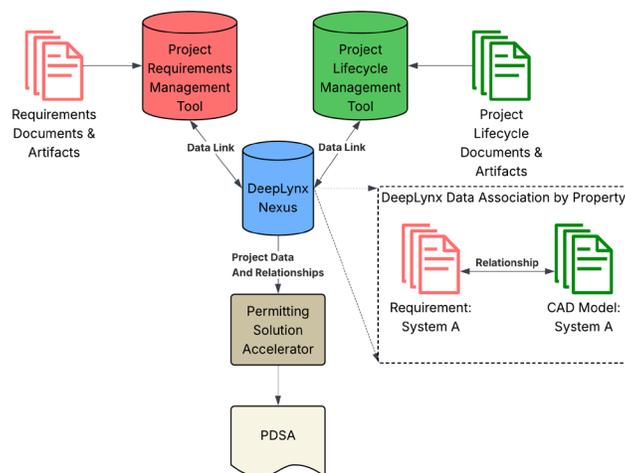


Figure 3. Digital thread integration.

The tool will then generate the PDSA sections from information contained within the disparate data sources related and accessed through DeepLynx Nexus. The goal for full integration of the digital thread is to utilize these individual pieces of project data in lieu of querying complete documents as in the previous stages of document generation.

## **2.2.4 PDSA GENERATION USING ALTERNATIVE TEMPLATES**

In the first phase of the evaluation, the document template will be based on the DOE-approved DOME PDSA. This will help the evaluation team conduct a section-by-section comparison of the AI-generated PDSA against the DOE-approved manually generated benchmark PDSA. However, the ultimate goal is to use the Microsoft tool to generate safety documents in a variety of formats and templates. Therefore, once the first cycle of verification and validation is complete, the team will generate a new PDSA using the same test case (i.e., NRIC DOME) and the same data sources, but in a different document template.

## **2.3 VERIFICATION AND VALIDATION**

The PDSA generated by Microsoft's AI tool must be rigorously evaluated to identify any errors or omissions in the document that may point to limitations in the tool's performance. To achieve this, a three-fold assessment process is proposed, as detailed below.

### **2.3.1 ASSESSMENT OF AI-GENERATED RESULTS**

The first step is a quality assessment of the AI-generated results. AI tools are susceptible to issues like hallucination and context misinterpretation, sometimes even with RAG. Therefore, the AI-generated DOME PDSA must be thoroughly examined by the evaluation team to ensure completeness, comprehensiveness, and conformance with the DOE PDSA format and guidelines. The information must be verified for factual correctness to ensure that the AI-generated content is free from biases and hallucination. It is also important to check for the correct use of technical terminology and consistency in units and systems of measurements. For the tool to be reliable, it must be able to demonstrate a reasonable degree of reproducibility (i.e., while the exact words and sentences may be different each time a section is generated, the general content should remain the same). The evaluation team must also test the tool's ability to handle edge cases, such as prompts referencing information beyond its training data or RAG knowledge base.

### **2.3.2 SUBJECT MATTER EXPERT ASSESSMENT**

The technical quality of the AI-generated results is best evaluated by subject matter experts (SMEs), i.e., nuclear safety analysts. The core evaluation team will, therefore, work with an SME to create a quality "checklist" for the DOME PDSA consisting of metrics such as accuracy, completeness, consistency, time-to-completion, and so on. The SMEs will review, evaluate, and score the AI-generated PDSA based on the checklist. In the first cycle where the DOE-approved PDSA is used as the benchmark, the SMEs will compare the AI-generated PDSA with the benchmark document. To minimize biases, the SMEs evaluating the AI-generated document must be completely independent from individuals associated with writing the original DOE-approved benchmark PDSA. In succeeding phases where an alternative document template will be used, the SMEs will evaluate the AI-generated PDSA as a standalone document.

### **2.3.3 REGULATORY ACCEPTANCE TEST**

The final step in the verification and validation process for the AI-generated DOME PDSA is a regulatory acceptance test. Once the final version of the AI-generated PDSA (using the alternative template) is ready, the evaluation team will identify DOE staff who could conduct a “mock review” of the document. These personnel will evaluate the AI-generated PDSA against acceptability criteria used for real PDSAs submitted to DOE. If the AI-generated document meets the DOE acceptance criteria, the case study will be deemed a success.

## **3. FUTURE ROADMAP FOR HOLISTIC EVALUATION**

The proposed case study of the Microsoft Generative AI for Permitting Solution Accelerator tool to generate the DOME PDSA is the first step in the holistic evaluation of AI-augmented digital DSA procedures. A few considerations for future steps are outlined in the following subsections.

### **3.1 RISK IDENTIFICATION AND MITIGATION**

To ensure the reliability and credibility of AI-generated safety documents, a risk mitigation strategy must be put in place. Potential risks include inaccuracies in AI-generated text, gaps in source document coverage, or misinterpretation of technical content. For the proposed case study, these risks will be mitigated through a multi-layer review process that includes cross-checking AI outputs against the DOE-approved DOME PDSA, consultation with nuclear safety analysts and design engineers, and validation against regulatory requirements. Additional safeguards include transparent traceability of AI-generated content to source materials, and incorporation of human oversight to identify and correct errors. This integrated approach will minimize errors throughout the document to ensure alignment with DOE safety standards. A similar risk mitigation strategy can be devised for future use cases.

### **3.2 COST VERSUS BENEFITS ASSESSMENT**

The use of Microsoft’s AI tool to generate documents for regulatory approvals presents both cost considerations and measurable benefits. On the cost side, initial expenses include licensing and maintaining the AI tool, configuring digital thread integration, uploading data sources, and dedicating SMEs to oversee validation. Additional effort will be required for training personnel, establishing traceability, and performing human-in-the-loop reviews.

AI significantly reduces the time required to draft and assemble content, accelerating document preparation by automating repetitive tasks such as system descriptions and accident analyses. This enables nuclear safety analysts to focus on high value oversight rather than manual drafting. The integration of the digital thread ensures improved traceability, reduced duplication of effort, and streamlined future updates to safety documentation. Over time, efficiencies gained from automation lower overall labor costs, shorten review cycles, and enhance consistency across deliverables.

While this case study evaluates the tool’s ability to generate DOE documentation, the Microsoft tool is agnostic to both generated document format and ingested data set. The process utilized for this case study can be repeated for other regulations and document types, as needed. Thus, the potential exists to extend the cost and efficiency benefits to multiple institutions.

### **3.3 CROSS-AGENCY REVIEW OF AI REGULATORY FRAMEWORKS**

It is essential to examine how other federal regulatory agencies are approaching artificial intelligence governance to guide the development of AI tools for nuclear licensing. The Federal Aviation Administration (FAA), for example, has published a roadmap for AI safety assurance, which outlines guiding principles for integrating AI into aviation systems [FAA]. Similarly, agencies such as the U.S. Food and Drug Administration (FDA), Occupational Safety and Health Administration (OSHA), and the National Highway Traffic Safety Administration (NHTSA), have begun developing frameworks or issuing guidance on the responsible use of AI in their respective domains. Including these agencies in a cross-agency review can provide valuable insights into best practices for transparency, validation, human oversight, and regulatory acceptance. Many of these principles may be directly applicable to the nuclear industry's adoption of AI-enabled tools.

### **3.4 EVALUATION REPORT AND IMPROVEMENT ROADMAP**

To support continuous improvement, a comprehensive evaluation report will be written upon the conclusion of this case study. This report will document key findings, lessons learned, and actionable recommendations for Microsoft, focusing on enhancing the tool's performance, usability, and regulatory alignment. It will incorporate feedback from SMEs and regulatory reviewers and include a regulatory feedback mechanism to capture reviewer insights. Additionally, the report will propose a dashboard framework to track real-time performance metrics, such as accuracy, completeness, and reproducibility, enabling transparent monitoring and iterative refinement of the Microsoft Generative AI for Permitting Solution Accelerator tool as it finds new nuclear energy applications and use cases across various organizations.

## 4. REFERENCES

- [1] ADAMS U.S. NRC. (2025). *Agencywide Documents Access and Management System (ADAMS)*. <https://adams-search.nrc.gov/home> [Accessed Aug 26, 2025].
- [2] Bergman, T. (2021). *Lessons-Learned from the Design Certification Review of the NuScale Power, LLC Small Modular Reactor*. (No. LO-0221-74035). NuScale Power, Rockville, MD. <https://www.nrc.gov/docs/ML2105/ML21050A431.pdf>
- [3] Darrington, J. W., J. M. Browning, and C. S. Ritter. (2020). *Deep Lynx: Digital Engineering Integration Hub*. Idaho National Laboratory. <https://github.com/idaholab/Deep-Lynx>
- [4] FAA. (2024). *FAA Roadmap for Artificial Intelligence Safety Assurance, Version I*. US Department of Transportation. [https://www.faa.gov/aircraft/air\\_cert/step/roadmap\\_for\\_AI\\_safety\\_assurance](https://www.faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance) [Accessed Aug 26, 2025].
- [5] Swartzendruber, Tracy. (2025). "What is a Digital Twin?" *GE Vernova Electrification Software* (blog). August 1, 2025. <https://www.gevernova.com/software/blog/what-digital-twin> [Accessed Aug 26, 2025].
- [6] Lund, M. L., K. M. O'Rear, J. S. Rymer, K. J. Terrill, and P. A. Suyderhoud. (2024). *Framework for a digital documented safety analysis* (No. INL/RPT-24-78848-Rev000). Idaho National Laboratory (INL), Idaho Falls, ID. [https://inldigitallibrary.inl.gov/sites/STI/STI/Sort\\_112773.pdf](https://inldigitallibrary.inl.gov/sites/STI/STI/Sort_112773.pdf)
- [7] Mandelli, D., C. Wang, J. Cogliati, and V. Agarwal. (2023). "Data Fusion of Numerical and Textual Equipment Reliability Data: A Knowledge-Graph-based Approach." In *Proceedings of the Probabilistic Safety Assessment and Management (PSAM) 2023 Topical Conference*. <https://www.ideals.illinois.edu/items/128993>
- [8] Patel, D. (2025). "Comparing Traditional OCR with Generative AI-Assisted OCR: Advancements and Applications." *International Journal of Science and Research (IJSR)*, 347–351. <https://doi.org/10.21275/sr25603211507>
- [9] Ritter, C., M. Rhodes. (2023). "Incorporating Digital Twins in Early Research and Development of Megaprojects to Reduce Cost and Schedule Risk." *INCOSE Insight.*, vol. 26, issue 3. <https://doi.org/10.1002/inst.12457>
- [10] U.S. DOE. (2016). *Review and Approval of Nuclear Facility Safety Basis and Safety Design Basis Documents*. (No. DOE-STD-1104-2016). Retrieved from <https://www.standards.doe.gov/standards-documents/1100/1104-astd-2016/@@images/file> [Accessed Aug 26, 2025].