

U.S. Nuclear Regulatory Commission Fiscal Years 2023-2027 Draft Artificial Intelligence Strategic Plan Comment Resolution Matrix

The comments included in the table below reflect the comments received in response to a Federal Register Notice (87 FR 39874) published on July 5, 2022, and a comment-gathering public meeting that took place on August 3, 2022. The full meeting transcript is available at ADAMS Accession No. ML22229A500. Quoted comments are indicated by italicized text.

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---|---------------------|--|--|--|---|
| 1 | X-energy | Ian Davis | ML22175A206 , section 4.2, paragraph 2, PDF page 25; section 4.5, paragraph 2, PDF page 27 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 29, line 15 – p. 25, line 21 | <i>My questions are related to the level of industry involvement with these activities, and specifically in two areas. One, there's a mention of an AI community of practice, and I was curious if this was an internal community or if this was going to be open to the public.</i> | The NRC agrees with this comment that from the original language used in the AI Strategic Plan it was unclear if the AI Community of Practice (AICoP) was an internal organization or open to the public. The NRC has modified the AI Strategic Plan (section 4.2) to indicate that the AI Steering Committee and AICoP are internal agency activities. |
| 2 | X-energy | Ian Davis | ML22175A206 , section 4.5, paragraph 2, PDF page 27 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p.29, line 22 – p. 30, line 2 | <i>And another section talks about pilot projects and proof-of-concept projects, and I was wondering if that was also going to be entirely internal to the NRC, or if that would involve partnerships with industry, and if so, what would that look like.</i> | The NRC agrees with this comment that the original language used in the AI Strategic Plan did not indicate if the pilots and proofs of concept would be internal or external agency activities. The NRC has modified the AI Strategic Plan (section 4.5) to clarify that the NRC plans to engage with stakeholders in public forums about potential pilot studies and proofs of concept. The NRC expects that the nuclear industry, members of the public, and the NRC will benefit from the information developed as the result of pursuing pilot studies and proofs of concept. Specific details regarding potential pilot projects and proofs of concept are not available at this time. The NRC staff plans to develop these details during the implementation of the plan. |
| 3 | Nuclear Safety and Regulatory Research Division, Idaho National Lab (INL) | Vaibhav Yadav | ML22175A206 , section 2, section 4 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 33, line 19 – p. 34, line 6 | <i>[D]oes NRC have a structure or a mechanism to enable . . . , industry engagement? Is there a -- you know, we're talking about 2027. Are there any tangible objectives and goals and timelines defined, or does NRC ha[ve] a vision to . . . have a like a competitive engagement released in next year or so that will kick off that type of industry engagement have a research development or demonstration or use case demonstrated to achieve those objectives over the course of next three to four years?</i> | The NRC agrees that the original language did not explicitly address in detail the types of industry engagement or the potential use of competition to spur additional engagement. However, the NRC did not plan to have this level of detail in the AI Strategic Plan. The NRC will consider and develop specific plans and more detailed implementation of further industry engagement and processes during the development of the AI framework envisioned in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," Strategic Goal #3, "Strengthen and Expand AI Partnerships," and Strategic Goal #5, "Pursue Use Cases to Build and AI Foundation Across the NRC." Nevertheless, NRC has added clarification to the Executive Summary and Introduction sections indicating that the NRC has developed this Strategic Plan as a tool to plan for future readiness and be prepared for new technologies such as AI. |
| 4 | Defense News | Catherine Buchaniec | ML22175A206 , (All), section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 36, line 20 – p. 37, line 6 | <i>And so my question is in the draft framework, it said that the NRC is committed to ensuring that the use of new technologies is safe and secure, and increasing AI with like any organization or mechanism can lead to increased risks. And so I'm just wondering what coordination, if any, is being done with like the country's national security agencies or departments to address any security concerns with the strategic plan or creating kind of like an across-the-board approach framework for ethical AI. Just, what's going on with cooperation or coordination?</i> | See NRC response to Comment 77. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|--------------------------|----------------|--|---|---|--|
| 5 | Member of the Public | Dan Solitz | N/A, (section 4.1, section 4.5) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 39, line 17 – p. 40, line 14 | <i>The purpose of my question is try to get a better understanding of where you're going with this AI. And let me go back to the particle board plant. It was analog system. The press had 14 openings and a cold place where coolant was -- solder (inaudible) were pushed in and then pressed then pulled out. And before they were pressed, the operator made sure the push arm was clear and all the cold plates had been cleared out. And then the operator pushed the press button. They figured out that they could save a -- they could probably do two to three more loads a day if they just relied on the relays that operated all these equipment to go ahead and close the press. And they didn't do it in my plant, they did do it in another plant, and I watched it -- watched it happen. It made me very nervous. So moving on to that to a nuclear power plant, are you considering something like calculating rod worth and fuel depletion and then automatically withdrawing the rods to -- to compensate for that? Is that the kind of use of artificial intelligence you're thinking about?</i> | Public workshops, such as those hosted by the NRC in 2021 (https://www.nrc.gov/public-involve/conference-symposia/data-science-ai-reg-workshops.html) provided some indication of where industry is looking to use AI in nuclear applications; however, the NRC is preparing for and increasing readiness for potential future applications using AI. The NRC has added clarification to the Executive Summary and Introduction sections that the NRC has developed the AI Strategic Plan to plan and prepare for new technologies involving AI. Implementation of the AI Strategic Plan will develop use cases through the framework envisioned in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," and applications to be planned in Strategic Goal #5, "Pursue Use Cases to Build and AI Foundation Across the NRC." |
| 6 | Member of the Public | Dan Solitz | N/A, (section 4.1, section 4.5) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 41, line 13 – p. 41, line 15 | <i>Well, if you could give me a range of things you're considering [where industry wants to use AI], it would be helpful.</i> | Public workshops, such as those hosted by the NRC in 2021 (https://www.nrc.gov/public-involve/conference-symposia/data-science-ai-reg-workshops.html) provided some indication of where industry is looking to use AI in nuclear applications; however, the NRC is preparing for and increasing readiness for potential future applications using AI. The NRC has added clarification to the Executive Summary and Introduction sections that the NRC has developed the AI Strategic Plan to plan and prepare for new technologies involving AI. Implementation of the AI Strategic Plan will develop use cases through the framework envisioned in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," and applications to be planned in Strategic Goal #5, "Pursue Use Cases to Build and AI Foundation Across the NRC." |
| 7 | Curtiss-Wright Scientech | Marty Murphy | ML22213A273 , (Presentation, Potential Challenges, Item 4); section 4.1; section 4.2 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 51, line 2 – p. 51, line 15 | <i>So, I think this pertains to item number four for potential challenges [the NRC should be aware of when preparing to review potential use of AI in nuclear applications]. And my question really goes to exactly how is the NRC going to consider the generation of new staff positions as it develops, and works through the implementation of the strategic plan? Especially for those applications that don't require specific NRC approval via amendment, or specific review? And what I'm talking about is inspection procedures, I didn't exactly see within the strategic plan where that was considered, or the use of say CRGR [Committee to Review Generic Requirements] was going to be applied to products that the NRC produces that don't typically go through the committee to review generic requirements.</i> | The NRC disagrees that this level of detail is appropriate for the AI Strategic Plan; however, as part of the implementation of the plan, Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," will develop the framework for overseeing AI technologies and establish the requirements and guidance for the regulatory oversight of AI in NRC-regulated activities. The commenter questioned how the NRC would manage the development of new regulatory positions outside the context of specific licensing actions. The NRC will continue to follow the applicable requirements and internal policies and procedures for developing regulations and interpretive rules. The NRC, in Management Directive 8.4, has specific processes and policies to address potential backfit or forward fit concerns. It provides processes for licensees, applicants, and the agency to raise concerns with potential new agency positions. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|--------------|----------------|---|--|---|--|
| 8 | Westinghouse | Brian Golchert | ML22175A206 , section 4.2, section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 53, line 11 – p. 54, line 1 | <i>I have some concerns about a long range for the NRC regulation. It's a broad topic [Artificial Intelligence/Machine Learning], whatever you want to call it, it covers a lot of things. And from a practical point of view, I'm worried about any future regulations will, I'll use the term what if us to death, that it will be cost ineffective for us to use these tools if you say do this model, check that model, do these uncertainty studies. I was just wondering for the strategic vision, if there had been thought to have a more intimate relationship with the potential applicants to have discussions on future regulations, future guidance to say you can't ask us to do this, it'll kill us cost effectively. Has that been considered as part of the strategic plan?</i> | The NRC understands the concern raised in the comment that requirements to implement AI uses could result in licensees and applicants forgoing AI use cases for simpler and more cost-efficient solutions. The five goals of the AI Strategic Plan will help the NRC develop appropriate future regulations for AI use. The NRC considers public involvement in, and information about, our activities to be a cornerstone of good regulation. The NRC plans to engage industry and other external stakeholders throughout the agency's work on AI as mentioned in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," and Strategic Goal #3, "Strengthen and Expand AI Partnerships." The NRC has added clarification to the Executive Summary and Introduction sections that the NRC has developed the AI Strategic Plan to plan and prepare for new technologies involving AI. No changes were made to the AI Strategic Plan in response to this comment. |
| 9 | NuclearN.ai | Bradley Fox | ML22175A206 , section 1, Table 1, PDF page 17 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 57, line 20 – p. 58, line 4 and p. 59, line 4 – p. 59, line 20 | <i>My only comments here would be to be cautious of the verbiage could impact. I think it would be good to clarify that a little further. Kind of looking back to my years in engineering, when we had that word could, there was a wide variety of what that could mean, and how that was interpreted by the inspectors at the plants, or various individuals, and that led to evaluations, and discussions. So, I think looking at how that would be interpreted up front would be great.</i> <i>I think it would be the difference between your level one, and your level two definitions in the table you showed a couple slides earlier. I think we have product running in level one right now, and you could theoretically make a case that a lot of systems that plants have could affect safety, from everything from search algorithms that might exist today, to corrective action program AI, things like that. And saying things like could affect plant safety can be carried out, or interpreted very differently by whomever is looking at that, or later evaluating that definition, and that classification so to speak. So, clarity around there is always really good. If there's guidance, in engineering world there's NUREGs, and things that would help us understand what that means.</i> | The NRC agrees with the comment that clear discussions and guidance will improve outcomes and reduce unneeded evaluations and verifications. At this time, the NRC cannot make an assertion about what specific AI application may require regulatory approval or oversight and what aspects may affect plant safety/security and control. As a result, the NRC chose to emphasize the potential to impact safety. The AI Strategic Plan is not intended to be guidance to applicants or licensees. Rather, it is a high-level document that provides a guide for the agency to plan, implement, and monitor activities to ensure the NRC staff remains prepared to review requested uses of AI. The NRC also agrees with the comment in that implementation of Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," is intended to delineate the level of NRC involvement in licensee uses of AI corresponding to the notional levels discussed in Table 1 of section 1. The table is a high-level characterization, focused on the attribute of autonomy, of various levels of AI integration in commercial nuclear activities. Through the implementation of the AI Strategic Plan, the notional levels of autonomy presented in Table 1 of section 1 may evolve. The NRC has modified section 1 to say, "As such, the NRC will treat these differences with the appropriate level of regulatory scrutiny and consider the multiple criteria necessary to determine the appropriate regulatory involvement for each level." |
| 10 | NuclearN.ai | Bradley Fox | ML22175A206 , section 1, Table 1, PDF page 17 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 58, line 4 – p. 58, line 11 | <i>I'd also love if you could look into impacts of AI on software quality assurance. And the different ways that might interact with the requirements for SQA at plants. Because we do get a lot of SQA impact questions from the operating facilities, and the folks involved in SQA there.</i> | The NRC agrees in part with the comment in that multiple factors exist in assessing AI properties such as considering how software quality assurance will be considered with AI systems. As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include model maintenance and life cycle management in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|--------------|----------------|--|--|---|--|
| 11 | Oklo | Ross Moore | ML22175A206 , section 1, Table 1, PDF page 17 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 60, line 4 – p. 60, line 19 | <i>I just want to make a quick comment slash question I guess with regard to the table [Table 1]. When I look at the table, and the levels of notional AI, and autonomy levels, level four, which is described as machine decision making with no human intervention, which I think potentially blends automation with AI. Where you can have an automated process that controls machine operability without actually integrating, or incorporating AI characteristics. And I'm curious whether that's the intent of the strategic plan, or if perhaps the strategic plan -- or if there is a potential way to delineate within the strategic plan, the difference between AI, and automation for which the NRC already has a regulatory framework surrounding.</i> | The NRC partly agrees with this comment in that it is agreed that automation and AI are two different concepts. However, Level 4 discussed in Table 1 specifically describes AI as being used in autonomy. Further, as described in Footnote 3 from page 1-3, the NRC recognizes that there are differences between automation and autonomy. In short, automation is considered to be a system that automatically acts on a specific task according to pre-defined, prescriptive rules, whereas in an autonomous AI system both the point at which action is taken and the action that is taken are the result of training an algorithm on data collected about the system. The fourth level of Table 1 is specifically focusing on AI integration to enable fully autonomous systems. No changes were made to the AI Strategic Plan in response to this comment. |
| 12 | X-energy | Ian Davis | ML22213A273 , (Presentation, Potential Challenges, Item 4) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 66, line 17 – p. 67, line 8 | <i>I also wanted to address potential challenges in line item number four there. INPO regularly releases reports talking about operator induced events, and I can foresee a conflict coming down the road where AI may have a goal of reducing operator induced events further by removing the human from the loop to some degree. And that -- to do so, we must figure out how we can create trustworthy AI. So, there's going to be a conflict of whether, or not we trust the human more, which we know the human is also [not] infallible, or we trust the AI more, and I think focusing on how to quantify, and evaluate the trustworthiness of both the human, and the AI is going to be really important to make a decision about whether, or not that application is acceptable, thank you.</i> | During the implementation of the AI Strategic Plan, as mentioned in section 4.1, the NRC will consider trustworthiness of AI in the development of the framework considered in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making." No changes were made to the AI Strategic Plan in response to this comment. |
| 13 | Westinghouse | Brian Golchert | ML22175A206 , section 1, Table 1, PDF page 17; section 4.2 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 67, line 15 – p. 68, line 4 and p. 69, line 13 – p. 69, line 20 | <i>And actually to follow up on the trust issue, looking at level one, we're looking at the possibility of using AIML optimization algorithms to design new products, or improve existing products. If we go back to the old fashioned method of using -- sorry. If we use the new method of using AIML to design it, but then go back to the old method of testing the final product, are we going to be expected to submit these AIML? I realize this is nothing in the strategic plan, but it is something to consider in the long run for regulation. There's a lot of stuff that AIML can do behind the scenes, which may, or may not be needed to be submitted, and a little curious as to how the NRC plans to fold that in.</i> <i>Well, actually I understand that, the point was if I go back, and test the new design, it doesn't matter how I got to that point. In the old days you would just have here's my drawing, here's my test, here's my results submitted to the NRC. Now, if I use AIML to get to the here's my drawing, and then do the tests, do I need to submit the AI?</i> | The NRC agrees with the comment that clarification was needed in the AI Strategic Plan that NRC staff will consider AI used for both operation or design in any NRC-regulated activity. Therefore, the italicized statement has been added to the following sentence in section 4.1: The NRC recognizes that the nuclear industry is likely to use AI in applications (e.g., notional AI and autonomy adoption levels in Table 1) <u>for the design and operation of nuclear facilities</u> that may require regulatory approval or oversight. During implementation of the AI Strategic Plan, the NRC will develop the regulatory framework in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," where the NRC staff will consider the requirements for AI used for both operation or design in any NRC-regulated activity, including whether the AI would need to be submitted. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|---|---|--|--|
| 14 | INL | Vaibhav Yadav | ML22175A206 , section 4.2 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 70, line 25 – p. 71, line 4 | <i>Talking about meeting with current regulatory framework, I was wondering if you are specifically considering Part 53, or intersection of AI based approaches to meet Part 53 requirements when it comes up?</i> | The NRC agrees with the comment in that there is recognition that AI impacts a variety of regulatory activities both ongoing and future, and the AI Strategic Plan and its implementation will support those other activities. Therefore, section 1 was revised to include the following statement: “The AI Strategic Plan was developed considering a variety of ongoing and future regulatory actions, and this strategy and its implementation will support those other activities.” The draft proposed rule for Part 53, “Risk-Informed, Technology Inclusive Regulatory Framework for Commercial Power Plant,” is currently being considered by the Commission. However, the draft proposed rule was not considered in the development of the AI Strategic Plan. |
| 15 | INL | Jason Remer | ML22175A206 , (section 4.1) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 78, line 1 – p. 78, line 21 | <i>In the past we've been working a lot in promoting digital systems as replacements for analog systems in our plants. And we worked through a lot of issues with the NRC, and I think we've got a good plan. But I just want to make sure that whatever we do in the AI area, we also just kind of incorporate all the good things we've learned about how we're going to deal with common caused failure, how redundancy, and diversity affect things. And so that whatever we come up with AI to ML, that it helps us move forward as a whole. We've found sometimes we think we've solved all the problems, and we look around, and something suddenly pops up. So, just a note of consideration, and appreciate the NRC working on this issue, I think it's very important. But I want to make sure that we can -- we kind of need to crawl before we can run, and AIML is a little bit running, and replacing analog with digital is kind of crawling. So, just a note, and again, appreciate you having this.</i> | The NRC agrees with the commenter in that there are relevant lessons-learned from prior new technology applications which can benefit the implementation of the AI Strategic Plan. Therefore, section 4.1 was revised to add that “...the NRC will leverage lessons learned from previous new technology applications in NRC-regulated activities to inform development of the AI framework.” |
| 16 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>There are 2 major issues with the NRC's strategic plan that it does adequately addresses:</i></p> <ol style="list-style-type: none"> <i>1. The maintenance requirements for models incorporated over the lifetime of the plant also known as technical debt.</i> <i>2. The overarching premise that baseline data collected in the development of said predictive models are representative of the problem scope for the entirety of operation.</i> <p><i>These two problems are even more significant than any technical development issues of predictive models.</i></p> | <p>The NRC agrees with the comment in that there are multiple factors exist in assessing AI properties. The commenter identified two potential challenges with using some implementations of AI systems: model maintenance and life cycle management. The NRC did not intend to address this level of detail in the AI Strategic Plan; however, during implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, “Ensure NRC Readiness for Regulatory Decision-Making.”</p> <p>The NRC has added Table 2, “Potential AI Technical Considerations for Regulatory Decision-Making” in section 4.1 to capture the commenters recommended addition of model maintenance and life cycle management to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|-----------------------------|-----------------------------|--|--|
| 17 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>In NUREG/CR-7294, the NRC has explored existing state of technology of AI/ML models. Repeated throughout the report is the issue of data quality, quantity, applicability, and uncertainty. These are significant non-trivial problems that currently plague all industries that utilize AI/ML. However, the difference is that in nuclear energy, the consequence of model failure is significant compared to other industries.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The commenter identified that the "...data quality, quantity, applicability, and uncertainty" are aspects to consider when evaluating the use of AI in an industry. The NRC agrees those aspects are important to consider and partly agrees that the consequence of model failure can be significant. However, the comment infers unstated assumptions regarding the use of AI and the associated level of risk. The AI Strategic Plan does not make any assumptions about potential uses of AI, as the NRC must be prepared to evaluate future applications regardless of the level of risk of the potential usage. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>The NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to capture the commenters recommended addition of risk analysis and <i>data quality, quantity, applicability, and uncertainty</i> to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 18 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>The design of the model is less relevant as nearly all models suffer from the two problems above.</i></p> | <p>The NRC agrees that the observed challenges may occur in models, but the commenter is making assumptions about the AI model alone and not about the implementation of the model in a system, or the architecture of a potential applications. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. The NRC does not make any assumptions about future AI applications, system architectures, or the quality of future submissions to the NRC. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 19 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>Take the following statement from NUREG/CR-7294 as an example:</i></p> <p><i>However, due to AI/ML uncertainty, the insufficiency of data quality and quantity, and lack of cognition about how to efficiently incorporate knowledge and data, challenges of adapting AI/ML techniques still exist. New perspectives and advanced frameworks should be proposed for different purposes in nuclear engineering. Particularly, the "black box" nature of ML/AI brings challenges with respect to the trustworthiness and transparency of the results in nuclear industry. This challenge makes the deployment of ML/AI-guided applications difficult to satisfy the regulatory requirements of NRC.</i></p> <p><i>This leads to the first problem that the strategic plan has not yet addressed.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC recognizes the statements made in NUREG/CR-7294, but the quoted text is from a subsection of the report discussing AI/ML in Nuclear Safety and Risk Analysis. This section of the report is looking at a particular application of AI/ML and not a detailed look at NRC regulatory requirements. That is to say, the quoted section is looking at the application of AI/ML to nuclear safety and risk analysis of nuclear plants, not the safety and risk of AI/ML systems and how such systems might meet regulatory requirements. The NRC will assess whether any regulatory guidance or inspection procedures need to be updated or created to clarify the process and procedure for the licensing and oversight of AI in NRC regulated activities in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>The NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include a list of potential characteristics to consider such as trustworthiness. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|-----------------------------|-----------------------------|--|---|
| 20 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>Suppose for the sake of argument that sufficient quality data exists during the software development process of the ML/AI model to develop highly accurate predictive models for an arbitrary safety critical variable (e.g., fuel centerline temperature). The model is dispatched to a plant to assist operator decisions. At first glance, this seems to be the optimal 'goal' of the ML/AI project, a dispatchable model that can improve safety; however, significant issues will arise through the lifetime of the plant. The primary problem is that the assumption that sufficient quality data is available is fundamentally flawed. The operating physics of nuclear power plants are constantly changing from beginning of life (BOL) to end of life (EOL). The data collected can only represent a subset snapshot of the reactor physics and will never represent all possible states. Therefore, it is highly probable that the model will slowly (or suddenly) become irrelevant at some changing point in the reactor lifecycle. Models developed based on BOL or snapshot reactor physics data will lose predictive accuracy overtime and be detrimental to operators rather than useful.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The commenter has illustrated examples of challenges related to the implementation of AI systems in the nuclear industry regarding available data and changes to the plant operating state, but the comment also makes assumptions about the use of AI in the nuclear industry. The NRC is aware of these potential challenges. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. During implementation of the AI strategic plan, the NRC may consider topics such as fielded performance degradation for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include fielded performance degradation in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 21 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>On the other hand, sudden changes in the reactor state (known as context shifts) can disrupt temporal predictive models (e.g., RNN). Anticipated sudden changes such as replacing fuel, or a single sensor degrading are examples of context shifts. More serious context shifts could be power transients could cause serious predictive issues. This data problem cannot be easily resolved by including all possible states of concern in the data set.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC also agrees that temporal changes in reactor state, including the indicated examples, could affect the performance of predictive models and recognizes the challenge regarding data coverage to address this issue. The AI Strategic Plan was not intended to go into this level of detail about particular technologies or applications. During implementation of the AI strategic plan, the NRC may consider topics such as changes in distribution of data and fielded performance degradation for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include data drift and fielded performance degradation in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 22 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>Even limiting the data set to the subset including anticipated operational occurrences (AOO) or design basis accidents (DBA), would rely on strong assumptions about how systems fail that may be physically infeasible. In any regard, the dataset for anticipated possible states would be intractable, expensive, and would most likely rely on multiple models developed for different context scenarios. The last point is due to the over generalization problem of neural network-based ML/AI models. In essence, as the problem scope increases to include more scenarios, the accuracy of the model steadily decreases.</i></p> | <p>See NRC Response to Comment 21.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|-----------------------------|-----------------------------|--|---|
| 23 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <i>In general, there is no acceptable solution to the latter problem [of over generalization of neural network-based ML/AI models] as there are too many highly speculated (and unsupported) scenarios to consider to develop a comprehensive model. Therefore, if the former approach is adopted, it may be possible to update the ML/AI model routinely to reflect changing states of the NPP.</i> | The NRC does not intend to speculate on solutions to the challenges that have been presented by the commenter. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making." No changes were made to the AI Strategic Plan in response to this comment. |
| 24 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <i>Generally, there are three approaches for implementing ML/AI models in continuously evolving environments: (1) locked models, (2) updating locked models, and (3) continuously updating models.</i> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC agrees that the three approaches mentioned are potential methods to handle models operating in a continuously evolving environment. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to capture the commenters recommended addition of model maintenance and life cycle management to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 25 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <i>In the most basic approach, a locked model is a model with weights and parameters that cannot change regardless of the state of the plant. The benefits to this approach are that the model can be reviewed and controlled for quality without concern of any deviation from function. The outputs from the model can be anticipated based on the training data and determined if they are valid or not. However, the drawback is that the model is guaranteed to become irrelevant over time in any continuously evolving environment. It also removes the primary benefit of ML/AI models which is the ability to adapt and continuously 'learn' the operational environment.</i> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC agrees that one potential approach to creating an AI system in an evolving environment is to lock the model, and recognizes a potential drawback is the model performance could potentially degrade in time; however, the NRC does not agree that locking the model always results in a model becoming irrelevant with time. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to capture the commenters recommended addition of model maintenance and life cycle management to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|-----------------------------|-----------------------------|---|---|
| 26 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>The second approach is the updating locked model approach. In this approach, locked models (with updated weights) are periodically dispatched as software updates whenever a sufficient change in the operational environment is detected. Each of the locked models are verified against a benchmark to ensure that they meet performance and safety requirements. In this approach, the ML/AI models can still update and 'learn' however under the strict oversight of developers. The drawback to this approach, however, is that given multiple versions of the same model will be developed, which model should be trusted and used. If an earlier version of the model contradicts later versions of the same model, should operators not trust the later version of the model? Version contradiction and management is the chief problem with the updating locked model approach.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC agrees that one potential approach to creating an AI system in an evolving environment is the "...updating locked model approach..." and recognizes a potential drawback is "...given multiple versions of the same model will be developed, which model should be trusted and used." The NRC will not speculate in the AI Strategic Plan if "Version contradiction and management is the chief problem with updating the locked model approach." Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to capture the commenters recommended addition of model maintenance and life cycle management to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 27 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>The last approach is the conventional approach to ML/AI models, that is to develop a continuous learning model with build-in restrictions and limitations and to hope that they ML/AI model will always perform accordingly. The model will adapt to any minute changes in the operational environment from BOL to EOL and to theoretically provide always accurate results. However, this approach is incredibly dangerous.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC agrees a potential approach to creating an AI system in an evolving environment is "...develop a continuous learning model with build-in restrictions and limitations." The NRC has not conducted a review of an AI application which would operate in a continuous learning environment which may have this concern. The NRC does not speculate on system architecture functionality and cannot conduct a review of an application until it is submitted for regulatory review. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to capture the commenters recommended addition of model maintenance and life cycle management to the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|----------------|-----------------------------|-----------------------------|---|---|
| 28 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>Routinely we have seen optimal models that have been developed in the ideal laboratory conditions only to be corrupted, cause unanticipated side effects, or abuse constraints (in the reward function) and result in loss to stakeholders. It is also incredibly difficult to design comprehensive constraints such that the model will always perform as expected. It is difficult in conventional systems (which is why software failure routinely occurs) and it is nearly impossible for a continuously learning system.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC agrees the observations discussed in this comment can be a challenge for continuously learning systems. The NRC does not speculate on system architecture functionality and cannot judge an application until it is submitted to the NRC for regulatory review. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail. During implementation of the AI strategic plan, the NRC may consider topics such as this for Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include risk analysis in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 29 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>This leads to the second major problem not covered by the strategic plan, that is the data control for risk applications.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC does not intend to develop a specific strategic goal which focuses on data control for risk applications, but instead will consider data control and risk analysis as AI technical considerations in the development of the framework in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include risk analysis, data drift, and data applicability in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 30 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>A new emerging topic in ML/AI is out-of-distribution (OOD) detection. The premise being that ML/AI models are only highly accurate whenever the input data is within the training data subset (i.e., interpolation). However, ML/AI models routinely fail at extrapolation task. OOD detection therefore is 'sensing' or calculating how 'far' away the input data is from known training data. The Mahalanobis distance is one example metric used to gauge distance of an input to training data distributions. Training data will always be a subset of the operational environment (for the reasons discussed above).</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. The NRC does not intend to develop a specific strategic goal which focuses on out-of-distribution detection, but instead will consider such topics as AI technical considerations in the development of the framework in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include data drift and fielded performance degradation in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|---------------------|-----------------------------|-----------------------------|---|---|
| 31 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>Furthermore, training data is typically developed in a highly augmented environment (i.e., without data noise, normalized, trimmed, synthesized, etc.). This means that while training performance can be incredible high (>90%), when applied to real world applications, the models typically experience a 30-40% decrease in predictive accuracy. Therefore, a way to measure OOD is absolutely required to trust the predictions made by an ML/AI model. The significance of OOD detection is not mentioned in the strategic plan nor in NUREG/CR-7294.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. However, the NRC cannot speculate on system architecture functionality and cannot conduct a review of an application until it is submitted for regulatory review. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI and does not speculate on what those approaches may entail.</p> <p>The NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include test, evaluation, verification and validation (TEV&V) and model maintenance in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 32 | Member of the Public | Edward Chen | ML22175A206 | ML22208A014 | <p><i>In summary, the current strategic plan should incorporate:</i></p> <ol style="list-style-type: none"> <i>1. The NRC's anticipated plan on post ML/AI deployment and continued maintenance. Specifically, how will developers maintain their models to keep them relevant.</i> <ol style="list-style-type: none"> <i>a. Will ML/AI models be locked models or continuously updating models and what type of framework is sufficient to ensure models remain relevant over time?</i> <i>2. Out-of-distribution is a key area of research when it comes to the trustworthiness and reliability of ML/AI models. Regardless of the type of model developed, a framework to ensure model relevancy (via OOD detection methods) must be in place to ensure models are relevant.</i> | <p>The NRC agrees in part with the comment in that multiple factors exist in assessing AI properties. The NRC does not agree that the AI Strategic Plan should include specific plans at this time on how to address post-AI/ML deployment, maintenance, model degradation, or model relevancy using out-of-distribution detection as these topics require significant regulatory research to determine how to best approach such areas. The NRC did not intend to address this level of detail in the AI Strategic Plan; however, during implementation of the AI Strategic Plan, the NRC will consider topics such as this as part of the AI framework developed in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making."</p> <p>However, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include the commenter's suggested topical areas for future consideration such as risk analysis, model maintenance, data drift, fielded performance degradation, data drift, life cycle management, and data considerations. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 33 | Member of the Public | Rob De La Espriella | ML22175A206 | ML22208A016 | <p><i>The AI Strategic Intelligence Plan is on target if not a little late. As a former decorated NRC Resident Inspector, we could have used AI algorithms to help inspectors at the nuclear plants to carve through the mountains of reports and documents, to help us proactively identify negative performance trends before they could escalate to bigger consequences.</i></p> | <p>The NRC agrees there could be potential in using AI algorithms to assist staff in performing their duties and leverage institutional knowledge. The AI Strategic Plan is an outwardly facing examination to ensure readiness to review potential industry uses of artificial intelligence. Appendix B of the AI Strategic Plan acknowledges there could be the potential for internal uses of AI at the NRC and insights gained through the implementation of the strategic plan. No changes were made to the AI Strategic Plan in response to this comment.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|----------------------|---------------------|-----------------------------|-----------------------------|--|---|
| 34 | Member of the Public | Rob De La Espriella | ML22175A206 | ML22208A016 | <p><i>My comment is not on the much-needed emphases that this Plan brings about, but on the flip side of the equation; there are key human skills that are extremely important to the NRC and its inspectors in the field, and these skills cannot be replaced by machines. I'm concerned the NRC has no plan to strengthen these human skills on par with their AI strategic plans. There is a false sense that AI will be able to solve complex problems, and that is not the case. AI can provide answers, but humans still need to ask the right questions.</i></p> <p><i>Each year, CEOs and executives from corporations all over the world meet at the World Economic Forum (WEC). One of their products is a future of jobs report, and since its inception in 2016, these executives have called out the human skills that are increasing in demand.[1] Consistently at the top of the list since 2016 are "critical thinking" and "complex problem solving." [2] However, these skill sets are not give the appropriate level of emphasis by academia and corporate institutions.[3]</i></p> <p><i>Here are the top 10 human skills needed to thrive in the age of AI, according to the WEC...is there an equally ambitious plan to strengthen all of these skills for NRC employees?</i></p> <ol style="list-style-type: none"> 1. Complex Problem Solving 2. Critical Thinking 3. Creativity 4. People Management 5. Coordinating with Others 6. Emotional Intelligence 7. Judgement/Decision Making 8. Service Orientation 9. Negotiation 10. Cognitive Flexibility | <p>The NRC agrees there could be potential benefit in using AI algorithms to assist staff in performing their duties and leverage institutional knowledge, but the intent would not be to replace human skills by algorithms. The NRC's Strategic Plan, NUREG-1614, Volume 8 (ML22067A170), Goal 2, "Continue to Foster a Healthy Organization" addresses the agency's overarching goals related to employee development. The NRC will provide future opportunities for stakeholder engagement in development of the AI framework. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 35 | Member of the Public | Rob De La Espriella | ML22175A206 | ML22208A016 | <p><i>Bottom line: critical thinking is the foundation for complex problem solving and the practical application of these skills are not being given proper emphasis by mainstream academia. (Their curriculum is largely theoretical and not practical). I encourage the NRC to evaluate their own training programs to ensure these valuable human skills are not eroded by over-reliance on AI and machine learning.</i></p> | <p>The NRC agrees there could be potential benefit in using AI algorithms to assist staff in performing their duties and leverage institutional knowledge, but the intent would not be to replace human skills by algorithms. The NRC's Strategic Plan, NUREG-1614, Volume 8 (ML22067A170), Goal 2, "Continue to Foster a Healthy Organization" addresses the agency's overarching goals related to employee development. The NRC will provide future opportunities for stakeholder engagement in development of the AI framework. No changes were made to the AI Strategic Plan in response to this comment.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---|----------------|---|--|---|---|
| 36 | Virginia Tech National Security Institute (VTNSI) | Tyler Cody | ML22175A206 , section 4.1, paragraph 3, PDF page 24 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 25, line 14 – p. 28, line 2 | <p><i>I had a comment on the question one here about specific recommendations or improvements. On -- (inaudible) pretty specific, but on line 34, page 4-2 in the strategic goal 1 section, like 4.1, the statement reads, The NRC will undertake research to develop an AI framework to determine the approach to assess areas such as, but not limited to explainability, trustworthiness, bias, robustness, ethics, security risks, and technical readiness of AI.</i></p> <p><i>So while these topics are definitely top of mind in the computer science community and like the people who are very close to artificial intelligence, from a systems engineering perspective, I think that these various topics are akin to properties we would like an AI solution to have. And that there are other concerns which are fundamental to engineering generally, but may take on different forms in AI that it's important to consider.</i></p> <p><i>So I just want to mention too, basic and related ones. And the first is test and evaluation, which is you need to just even see if these properties are actually possessed by the AI. So it's one thing to say, oh, we have a method for explainability, but how do we even know that method's working? What kind of scenarios will we test that method in?</i></p> <p><i>So under test and evaluation, there are two sort of issues I think about and I just want to highlight them for the sake of being a little thought- provoking. But the first is that there's this prevailing focus on using held-out but identically distributed test sets of data.</i></p> <p><i>And identically distributed testing checks that the algorithm's working, but like the learning algorithm or the AI algorithm, but it's not actually testing if, say, the function approximation that it produces, right. Because the AI produces some function that tells us, okay, give me these inputs, I'll give you those outputs.</i></p> <p><i>It's not actually telling us if that's going to work as intended in the variety of scenarios we're going to face during operation. So the first part is this identically distributed testing. But the second is sort of built on that, because while the AI solutions themselves will be like these input-output components for most part, they also have these system-level influences. So, and they create system-level effects, and you know, system- level outcomes that we look at.</i></p> <p><i>And so just the input-output component level testing might lack the scope to properly identify operating envelopes of AI solutions in terms of their system or the environment they're in, like what the operating conditions are, etc.</i></p> <p><i>So backing out, I think testing evaluation has these dual challenges where the - - right, currently the component-level tests are insufficient.</i></p> <p><i>But also, the component-level tests won't be all that we need to make sure that, you know, we're producing the outcomes we want from the AI.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. As pointed out in the comment, the NRC appreciates the recommendation to consider identically distributed testing and system-level influences as part of the test and evaluation of AI models. As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include the higher-level recommendations of test, evaluation, verification and validation (TEV&V) and life cycle management in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|-----------------------------|-----------------------------|---|---------------------------------|
| 37 | VTNSI | Tyler Cody | ML22175A206 | ML22230A047 | <p>1. Are there any specific recommendations or improvements to consider in the development of the AI Strategic Plan?</p> <p><i>On line 34 of page 4-2 in section 4.1, "Strategic Goal 1":</i></p> <p><i>The document states:</i></p> <p><i>"The NRC will undertake research to develop an AI framework to determine the approach to assess areas such as, but not limited to, explainability, trustworthiness, bias, robustness, ethics, security, risks, and technical readiness of AI."</i></p> <p><i>While these topics are certainly top-of-mind for the computer science community, from a systems engineering perspective, they are akin to "properties" one would like AI solutions to have. There are other concerns which are fundamental to engineering generally but take on different forms for AI, and need to be addressed in whatever AI framework is developed. I will mention two interrelated basic issues. The first is test and evaluation (T&E), which is needed to see if these properties are actually possessed by the AI. The second, related issue is life cycle.</i></p> <p><i>As stated in the abstract, "AI are machine-based systems that go beyond defined results and scenarios". Accordingly, there are two issues with T&E: the first is the prevailing focus on "held-out-but-identically-distributed" test sets of data. Identically distributed testing checks that the algorithm is working—not that the function approximations it produces will work as intended in the variety of scenarios it will face during operation. The second is that AI solutions are often an input-output component, but AI solutions are influenced by the systems they operate within, they create systems-level effects and and systems-level outcomes. Component-level testing lacks the scope to properly identify the operating envelopes of AI solutions in terms of the context of their system and environment. Therefore, T&E for AI has dual challenges of (1) insufficient component-level tests and (2) component-level tests only being a piece of the picture.</i></p> | See NRC response to Comment 36. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|---|--|---|
| 38 | VTNSI | Tyler Cody | ML22175A206 , section 4.1, paragraph 3, PDF page 24 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 28, line 3 – p. 29, line 1 | <p><i>And I know I'm going on here, but briefly, the second related issues is on the life cycle, and it's very closely related, right. So AI solutions do have life cycles. But that's kind of an understudied topic. So whatever framework that's proposed should involve some plans for test and evaluation over the life cycle to sort of monitor the health of the AI, if you will.</i></p> <p><i>And also address the concepts of, like, system maintenance that we have generally, right. So for the AI, what does it do to retrain it or recalibrate it, etc. As well as retirement, you know, what is the process for retiring an AI model.</i></p> <p><i>So I will concede that, you know, this is really relevant in dynamic settings. So you know, in -- in a nuclear setting, things are very constrained. So it's important to say okay, how much variance can we expect, where will that variance come from. And that's where life cycle and test and evaluation I think should be focused on.</i></p> <p><i>So in short, my specific recommendation is to consider an emphasis on test and evaluation and life cycle management, in addition to these pillars or these properties that are currently listed on line 34.</i></p> | The NRC agrees with the comment in that multiple factors exist in assessing AI properties. As pointed out in the comment, the NRC appreciates the recommendation to consider test and evaluation over the life cycle of AI models. As such, the NRC has added Table 2, "Potential AI Technical Attributes for Consideration in Regulatory Framework" in section 4.1 to include test, evaluation, verification & validation, model maintenance and life cycle management in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties. |
| 39 | VTNSI | Tyler Cody | ML22175A206 | ML22230A047 | <p><i>AI solutions have life cycles, but they are understudied. Whatever framework that is proposed should involve plans for T&E over the life cycle, to monitor the "health" of AI solutions, as well as address system maintenance (retraining, recalibration, etc. of AI solutions) and retirement. In dynamic settings, which I am not sure exist widely in nuclear applications, continuous T&E and continuous re-engineering are important concepts. So, in short, the specific recommendation is to consider an emphasis on T&E and on life cycle management in addition to those "properties" that are currently listed on line 34 of page 4-2. While not all applications of AI may merit all "properties" currently listed, all applications will merit a T&E and life cycle management plan.</i></p> | See NRC Response to Comment 38. |
| 40 | VTNSI | Tyler Cody | ML22175A206 | ML22230A047 | <p>2. What goals, objectives, or strategies within the NRC's current strategic plan should be added, enhanced, or modified in the AI Strategic Plan?</p> <p><i>The response to Question 1 applies to Question 2. Additionally, I think there should be a strategy which raises assurance in AI by focusing on engineering processes for training, deploying, maintaining, and retiring models. The emphasis on processes is meant to distinguish assurance approaches which focus on individual models or solution methods. There is a tendency in AI to focus on validating and accrediting trained models, as opposed to the processes that train, deploy, maintain, and retire them.</i></p> | The NRC partly agrees with the comment. The NRC does not intend to develop a specific strategic goal which focuses on assurance in AI, but instead will consider AI assurance in engineering processes in the development of the framework in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making." While the list in section 4.1 is not intended to be all inclusive, the NRC has added Table 2, "Potential AI Technical Attributes for Consideration in Regulatory Framework" in section 4.1 to include assurance processes in the list of potential characteristics to consider. The NRC will provide future opportunities for stakeholder engagement in development of the framework and communication of the timeline envisioned for those activities. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---------------------------------------|---|---|--|
| 41 | VTNSI | Tyler Cody | ML22175A206, section 4.2, section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 47, line 12 – p. 49, line 5 | <p><i>I had a comment for the third question here about potential near-term and far-term AI activities. I just want to bring awareness to this group that the systems engineering research community has a growing and active interest in systems engineering methods and best practices for AI and AI- enabled systems.</i></p> <p><i>And that these lines of research are directly concerned with the engineering process. So like all the way from need analysis through decomposition to components and recomposition to a system, deployment maintenance, retirement, etc. And so in the near term, I just, there is this ongoing work that I think will address a gap in the literature as well as the research community in increasingly like the workforce.</i></p> <p><i>Because currently there's a lot of focus on the AI algorithms or on those properties like explainability, trustworthiness, privacy, etc. But just not as much research into like what's the engineering process that it followed, will produce the AI solution, you know, that fits our requirements and our needs.</i></p> <p><i>So I think that's a near-term area that is being worked on elsewhere and is very relevant to this community. And in the farther term, the same community is looking at the use of digital models as part of digital engineering and model-based systems engineering activities.</i></p> <p><i>And those activities will expand to AI- enabled systems. Which basically means that increasingly there'll be digital processes for verification and validation and accreditation of AI solutions using things like digital twins, and model- based systems engineering to, you know, to make use of those digital twins to test the performance, to test against requirements, you know, to vary the conditions in those virtual environments.</i></p> <p><i>So I just, yeah, I wanted to mention the Department of Defense is funding this research, and they have a growing portfolio of this research. And so with respect to strategic goals 2 and 3, I'd really suggest trying to form some connections with the systems engineering research community.</i></p> | <p>Thank you for your comment. The NRC appreciates the recommendations on both near- and far-term activities with respect to AI research in the systems engineering community. During implementation of the AI Strategic Plan, the NRC will consider engaging with the systems engineering research community in implementing Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," Strategic Goal #2, "Establish an Organizational Framework to Review AI Applications," and Strategic Goal #3, "Strengthen and Expand AI Partnerships." The NRC will also provide future opportunities for stakeholder engagement in development of the AI framework. No changes were made to the AI Strategic Plan in response to this comment.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|-----------------------------|-----------------------------|---|---------------------------------|
| 42 | VTNSI | Tyler Cody | ML22175A206 | ML22230A047 | <p>3. What are potential near-term, or far-term, AI activities that the NRC should be aware of when finalizing and prioritizing the AI Strategic Plan, or associated supporting research?</p> <p><i>The systems engineering research community has a growing, active interest in systems engineering methods and best practices for AI and AI-enabled systems. These lines of research are directly concerned with the engineering process, from need analysis, through decomposition into components, recomposition into a system, deployment, maintenance, and retirement.</i></p> <p><i>In the near term, this work addresses a gap in the literature and research community. There is heavy research into AI algorithms, and into secondary concerns like those properties listed under Strategic Goal 1, but there is not as much research into engineering processes, which, if followed, will produce AI solutions that meet stakeholder needs and requirements.</i></p> <p><i>In the far-term, the use of digital models as part of digital engineering and model-based systems engineering activities will expand to AI-enabled systems. This means digital processes for verification, validation, and accreditation of AI solutions using digital twins and model-based systems engineering to evaluate test performance and to test against requirements in (virtually) varied conditions. The Department of Defense has a growing portfolio of this research, and so, especially for Strategic Goals 2 and 3, it is suggested that the NRC connect with the systems engineering research community.</i></p> | See NRC response to Comment 41. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|--|---|---|---|
| 43 | VTNSI | Tyler Cody | ML22213A273, (Presentation, Potential Challenges, Item 4); section 1, Table 1, PDF page 17 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 64, line 9 – p. 66, line 2 | <p><i>This time I have a comment on the fourth item about potential challenges, and it relates to the question of where to place regulations, I guess. It's a bit academic, but it's pretty important. Regarding operations, there are many no free lunch theorems of statistical learning theory that suggests no single model can be optimized for all conditions at once.</i></p> <p><i>So, the concept of test, and evaluation as a pre-deployment activity, or accreditation as an exclusively pre-deployment activity is in pretty strong conflict with the first principles behind machine learning solutions. So, basically as conditions change, for example between operations, or as platforms degrade, or with changes in use, if there is a material difference in the data that flows through the model, then the performance of the model is expected to change.</i></p> <p><i>And if it was well fit originally, then by definition it has to be worse fit now, because it was well fit to conditions that are no longer occurring. So, this all goes to suggest that domain adaptation, like the fact that we have to update our models is the rule, not the exception. And that so called universal models, for example general purpose vision models are the exception, not the rule.</i></p> <p><i>So, I think when you think about that table, and the criterion for what level of autonomy are we at, there's other things to consider like is this thing going to need to adapt over time? Is this a stationary environment, or not? Meaning is the distribution of the data changing over time, is this thing doing online learning? And all those little check boxes sort of back you into the quadrant of how worried do I need to be about unexpected changes in the behavior of my model?</i></p> <p><i>Because some technician tightened something somewhere, does that need to be folded in to how -- do you need to incorporate the policies for how you change the physical system with the policies for the AI? And do those things need to be coupled, et cetera, I think come down to those kinds of analysis. Again, just a comment, but I hope thought provoking. Thank you.</i></p> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. As such, the NRC has added Table 2, "Potential AI Technical Attributes for Consideration in Regulatory Framework" in section 4.1 to include domain adaptation and fielded performance degradation in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> <p>The NRC interprets that the commenter's use of the term operator means the end user of the technology and is not intended to mean the term as defined in 10 CFR 55.4. The NRC will provide future opportunities for stakeholder engagement in development of the AI framework.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---------------------------|----------------|-----------------------------|-----------------------------|--|--|
| 44 | VTNSI | Tyler Cody | ML22175A206 | ML22230A047 | <p>4. What are potential challenges the NRC should be aware of when preparing to review potential use of AI in nuclear applications?</p> <p><i>Regarding operations, the no free lunch theorems of statistical learning theory suggest that no single model can be optimized for all conditions at once. And so, the concept of T&E as a pre-deployment activity is in conflict with the first principles behind ML solutions. That is, as conditions change, e.g., between operations, or as platforms degrade, or with changes in use, if there is a material difference in the data that flows through the model, then the performance of the model is expected to change. This suggests that domain adaptation is the rule, not the exception. Conversely, so-called universal models, e.g., general purpose vision models, are the exception, not the rule. Thus, operators looking for accreditation ought to have in place procedures for mitigating risks due to changes in the distribution of data. Important to such procedures is an acknowledgement of catalysts; what aspects of the system wherein the AI model will operate can change and can those changes create a material difference in the data that flows through the AI model? Maintenance of physical parts and machinery, changes in use, and variation in environmental factors are examples of catalysts that exist outside the scope of the AI model itself.</i></p> | See NRC response to Comment 43. |
| 45 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <p><i>How does this strategic plan influence the already ongoing procurements (SOL_31310022R0041, SOL_31310022Q0064, and SOL_31310022R0037) around AI concepts?</i></p> | The AI Strategic Plan does not impact planned or ongoing procurements around AI concepts as the AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities. The AI Strategic Plan will not address existing acquisition strategies or procurement processes. No changes were made to the AI Strategic Plan in response to this comment. |
| 46 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <p><i>We recommend that NRC engage all parts of the organization to establish a unified approach to obtain industry support and improve the procurement processes.</i></p> | The AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities. The AI Strategic Plan will not address existing acquisition strategies or procurement processes. However, the NRC does agree that organizational communication and collaboration are critical to ensure a unified regulatory approach as discussed in Strategic Goal #2, "Establish and Organizational Framework to Review AI Applications." The NRC will establish an internal Artificial Intelligence Steering Committee (AISC) to facilitate effective coordination and collaboration across the NRC to ensure readiness for reviewing the use of AI in NRC-regulated activities. No changes were made to the AI Strategic Plan in response to this comment. |
| 47 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <p><i>We ask that NRC consider pursuing a larger contract or vehicle which would support AI/ML efforts over a longer period of time and better engage Industry, Technology, Research and Standards Organizations.</i></p> | The AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities. The AI Strategic Plan will not address existing acquisition strategies or procurement processes. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---------------------------|----------------|---|-----------------------------|---|--|
| 48 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>Would the NRC please elaborate on how this Strategic Plan for AI will be applied and implemented in a contractual sense across the various communities of interest across NRC?</i> | The AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities. The AI Strategic Plan will not address existing acquisition strategies or procurement processes. The NRC did not intend to address this level of detail or type of information in the AI Strategic Plan. No changes were made to the AI Strategic Plan in response to this comment. |
| 49 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>How is this standardized and normalized against FCEB standards and policy guidance from all Government organizations responsible for defining these standards?</i> | The AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities and activities undertaken with respect to each strategic goal may also benefit the agency in preparing for potential future internal agency AI usage. The NRC will gain knowledge and expertise in a wide range of skills and capabilities such as artificial intelligence and data literacy that could support expanding the use of data for decision-making in the agency. However, the AI Strategic Plan does not attempt to focus on addressing using AI for internal agency purposes or considering standards or policy guidance from other government organization with respect to agency AI implementation. No changes were made to the AI Strategic Plan in response to this comment. |
| 50 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>There will be elements of technology specifications and interfaces that should be common and standardized, there will be institutional and operational guidance that should be standard internal to the organization, and externally to organizations that interface with NRC and its communities of interest.</i> | The NRC recognizes that AI is a data-driven technology and as such that common technology specifications and interfaces benefit both internal and external users of the technology. However, the AI Strategic Plan was developed to support agency preparation for external AI usage in NRC-regulated activities. As part of Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," and Strategic Goal #3, "Strengthen and Expand AI Partnerships," the NRC will participate with standards development organizations and work with Federal agencies and the international regulatory community to offer critical expertise and perspectives to inform the drafting and revision of AI standards and guidance documents which may be used to support review and evaluation of AI uses in NRC-regulated activities. No changes were made to the AI Strategic Plan in response to this comment. |
| 51 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 27 | ML22231A174 | <i>Would NRR please specify the specific NRC organizations that will be involved in developing strategy, approach, and governance to this Artificial Intelligence Strategy (e.g. ADM, OCIO, NSIR, etc.).</i> | A diverse group of organizations across the agency will be represented on the AI Steering Committee, including but not limited to, the Office of the Chief Information Officer, the Office of Nuclear Reactor Regulation, the Office of Nuclear Regulatory Research, the Office of Nuclear Security and Incident Response, and the Regional Offices. The NRC will provide future opportunities for stakeholder engagement in development of the AI framework discussed in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making" and during the implementation of Strategic Goal #3, "Strengthen and Expand AI Partnerships." No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---------------------------|----------------|---|-----------------------------|--|---|
| 52 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 27 | ML22231A174 | <i>Has the NRC considered how this AI Strategic Plan will be developed in coordination with Licensees and other stakeholders across the community?</i> | <p>The NRC recognizes the importance of stakeholder engagement and transparent public outreach. There will be opportunity for stakeholder engagement during the development of the AI framework as part of Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making", and any potential changes to regulations. No changes were made to the AI Strategic Plan in response to this comment.</p> <p>Consistent with NRC practice, notification of all public meetings will be published 10 days in advance, including engagement through AI-specific public workshops, Advisory Committee on Reactor Safeguards public meetings, the NRC Regulatory Information Conference, and any future public meetings to engage stakeholders on regulatory research plans or further information related to implementation of the strategic goals. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 53 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 27 | ML22231A174 | <i>NRC will certainly need to support an internal community of interest and external organizations when instituting an AI Strategic Plan. Identifying the structures and mechanisms in this Strategic Plan to support governance, standards, organizational responsibilities and management oversight would all help to establish this importance and commitment of NRC to this AI Strategic Plan.</i> | <p>The NRC agrees that successful implementation of the AI Strategic Plan requires effective coordination and collaboration across the NRC, at both the management and staff levels. The NRC will establish an internal Artificial Intelligence Steering Committee (AISC) to provide cross-office coordination and direction to ensure readiness for regulatory decision-making and develop AI governance. The AISC will include senior management with responsibility for AI technology across the agency. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 54 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 35 | ML22231A174 | <i>Would NRR please specify the importance of "explainability and uncertainty"?</i> | <p>The NRC notes that in order to assess the safety and security of AI applications in NRC-regulated activities, it is important to develop a regulatory framework to assess the explainability and uncertainty of AI models. The NRC recognizes that multiple attributes may be used to quantify AI explainability and uncertainty. As such, the NRC plans to undertake research to develop an AI framework which considers methods to ensure AI explainability and assess uncertainty. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 55 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 35 | ML22231A174 | <i>We believe data collection, analysis, summarization and visualization are essential tools when considering "explainability and uncertainty" as it applies to design, implementation, decision making and near-real time operation of complex models and solutions.</i> | <p>The NRC agrees with the comment in that multiple factors exist in assessing AI properties. As such, the NRC has added Table 2, "Potential AI Technical Considerations for Regulatory Decision-Making" in section 4.1 to include data quality, quantity, applicability, and uncertainty in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---------------------------|----------------|---|-----------------------------|---|---|
| 56 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 24 of 40, line 35 | ML22231A174 | <i>We highly recommend avoiding black box models in this environment. It is critical to conduct post-mortems when failures occur. Overly complex models that cannot readily be summarized present risks to operations and overall safety.</i> | The NRC agrees there are potential shortcomings surrounding AI implementations which rely on "...black box models..." and methods to communicate risks and model confidence are important to consider and partly agrees that the consequence of model failure can be significant. However, the comment infers unstated assumptions regarding the use of AI and the associated level of risk. The AI Strategic Plan does not make any assumptions about potential uses of AI as the NRC must be prepared to evaluate potential AI applications regardless of the risk level of the potential usage. Nevertheless, the NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. No changes were made to the AI Strategic Plan in response to this comment. |
| 57 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 25 of 40, lines 15-16 | ML22231A174 | <i>Would the Government consider establishing an Advisory Board consisting of leaders from Industry, Technology, Research and Standards Organizations to provide NRC with a mechanism to balance advances in AI with the operational, business and organizational needs of the Government?</i> <i>Relevant areas of interest across the participants could include best practices and centralized resources for areas such as Standards, Governance, Security, Education, Safety, Incident Response, Emergency Management, and periodic investigations supporting NRC strategies and future plans.</i> | The NRC appreciates the importance of considering external stakeholder perspectives on a variety of topics including AI and there will be continued opportunities for stakeholder engagement throughout the implementation of the AI Strategic Plan. The creation and activities of Federal Advisory Committees is governed by the Federal Advisory Committee Act (FACA), and the NRC does not currently plan to create an AI Advisory Committee under FACA. No changes were made to the AI Strategic Plan in response to this comment. |
| 58 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 25 of 40, lines 15-16 | ML22231A174 | <i>We recommend that the NRC define the organization modeling from other Government Advisory Boards and solicit participation of Industry, Academia Technology, Research and other Government Agencies to serve the best interests of NRC and the American public over an extended period of time. This advisory board would be best positioned to help define requirements, validate the application of technology to satisfy mission and operational requirements, and continuously provide insights across the spectrum of advances in the AI field and technologies over an extended period of time.</i> | See NRC response to Comment 57. |
| 59 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 25 of 40, lines 15-16 | ML22231A174 | <i>How will the NRC Strategic Plan for AI be applied across the various communities of interest both within NRC and external to NRC?</i> | The AI Strategic Plan is a high-level document that serves as a guide for the staff to plan, implement, and monitor work associated with AI activities. Granular schedule details are not appropriate for the AI Strategic Plan. The AI Strategic Plan details the AI-related goals and objectives on which the NRC will focus for FY'23 through FY'27. The NRC plans to engage external stakeholders, including the public, throughout the agency's implementation of the AI Strategic Plan. As mentioned in the public meeting, early NRC activities include developing an organizational framework, steering committee and working groups, and public workshops. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|---------------------------|----------------|---|-----------------------------|--|--|
| 60 | Gunnison Consulting Group | Sean Fromm | ML22175A206 page 25 of 40, lines 15-16 | ML22231A174 | <i>What governance model will NRC employ to best deliver the promise of AI to this community? There will be elements of technology specifications and interfaces that should be common and standardized, there will be institutional and operational guidance that should be standard internal to the organization, and externally to organizations that interface with NRC and its communities of interest.</i> | The NRC did not intend to address this level of detail in the AI Strategic Plan; however, any future governance model for potential internal agency use of AI will be developed in coordination of the agency evidence-building plan (ML22066B056), as required by the Evidence Act and the NRC's Information Technology/Information Management Strategic Plan (ML19323D858). The main purpose of the AI Strategic Plan is to enable continued effective review and evaluation of the use of AI in NRC-regulated activities. No changes were made to the AI Strategic Plan in response to this comment. |
| 61 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>Would NRC please provide [or] specify where to locate their respective AI Use Cases in support of this Strategic Plan or where they will be provided when they are developed?</i> | The NRC did not intend to provide this level of detail in the AI Strategic Plan; however, the NRC can clarify that Executive Order 13960, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," does not apply to independent regulatory agencies like the NRC. No changes were made to the AI Strategic Plan in response to this comment. |
| 62 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>How will these Use Cases mature, evolve and be maintained with advances in technology, and improvements from operational lessons learned?</i> | Strategic Goal #5, "Pursue Use Cases to Build and AI Foundation Across the NRC," focuses on developing and pursuing use cases, consistent with priority question two of the agency evidence-building plan (ML22066B056), to build technical expertise for reviewing the use of AI in NRC-regulated activities. To better understand how AI algorithms, models, and claims are validated and tested, the NRC needs to undertake research to develop use cases with data from various sources and in multiple forms. These use cases will help the staff gain AI expertise that could be used in performing regulatory reviews or assessments for a wide range of potential AI applications but are simply for research and development purposes and not for operation or deployment. No changes were made to the AI Strategic Plan in response to this comment. |
| 63 | Gunnison Consulting Group | Sean Fromm | ML22175A206 | ML22231A174 | <i>For reference: EO 13960, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government", requires all FCEB agencies to publish their AI use cases. FCIO Council led coordination on this and provided further direction that all should be published on the respective agency websites at [https://[agency].gov/data/AI_Inventory]. An advisory board, as previously mentioned, would be an excellent partner in establishing these AI Use Cases.</i> | The NRC did not intend to provide this level of detail in the AI Strategic Plan; however, the NRC can clarify that Executive Order 13960, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," does not apply to independent regulatory agencies like the NRC. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|-----------------------------|-----------------------------|--|--|
| 64 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <p><i>DataRobot has chosen to comment on several aspects of the NRC's AI Strategic Plan. We agree with the overall goal of this strategic plan "to ensure continued staff readiness to review and evaluate AI applications effectively and efficiently," but we have made comments on specific definitions and additional considerations under the strategic goals that we believe will make a stronger, and more expansive plan.</i></p> <p><i>The Nuclear Regulatory Commission (NRC) defines AI as a "a machine-based system that can go beyond defined results and scenarios and has the ability to emulate human-like perception, cognition, planning, learning, communication, or physical action." There are a variety of domains and categories that fall into this definition, including natural language processing, machine learning, and deep learning that were already mentioned in the document, but it is limited in scope. There are two additional categories that should be considered: time series and unsupervised learning. Time series is the forecasting of events based on historical data for either one or multiple series, and could be valuable for the NRC to consider especially with forecasting when their equipment might fail, or forecasting staffing needs for specific programs. There is tremendous value in using time series forecasting problems and utilizing these algorithms would expand the number and complexity of use cases that the NRC would consider under their AI strategy.</i></p> | <p>The NRC believes the definition of AI in the AI Strategic Plan is sufficiently broad to encompass unsupervised and time-series learning problems. The definition of AI in the strategic plan is applicable only to the AI Strategic Plan; it is not intended to be a binding regulatory definition. The definition of AI may be refined during implementation of the AI Strategic Plan to reflect experience and stakeholder feedback. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 65 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <p><i>Additionally, in the introduction page of the document, the AI strategy "considers an evolving landscape where computers use data and unseen behavior to construct the underlying algorithmic model, draw inferences, and define the rules to achieve a task." While this is generally true for supervised algorithms where the[ir] data has been labeled to make inference, this leaves out unsupervised methods of modeling where there is no labeled data. These use cases are great for clustering and anomaly detection where the patterns are hard for analysts to discern. Including unsupervised learning into the definition also expands the use cases that could be considered.</i></p> | <p>See NRC response to Comment 64.</p> |
| 66 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <p><i>We would also like to comment on Figure 1, a diagram that identifies specific categories under the umbrella term of Artificial Intelligence. This diagram categorizes natural language processing, deep learning, machine learning, and robotic process automation (RPA). While we generally agree with the first three mentioned, as they use probabilistic methods to compute outcomes, we believe that robotic process automation doesn't necessarily fall within that category. RPA does automate certain repetitive tasks, but it uses a rules-based system that a human has to manually define first. RPA does not use historical data and algorithms to produce inference like the others do. It could be used to act on inference or information from the statistical outputs from NLP, ML, and deep learning, but it doesn't produce predictions.</i></p> | <p>The NRC agrees with the comment and as a result has removed robotic process automation from Figure 1 in section 1 of the AI Strategic Plan.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|-----------------------------|-----------------------------|---|--|
| 67 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <i>Under Strategic Goal 1: Ensure NRC Readiness for Regulatory Decision-making, there is an urgency that “focuses on developing the regulatory guidance and tools to prepare the staff to assess 20 AI as part of NRC regulatory activities.” The deployment of these AI models and being able to serve them in a production environment is as crucial as developing them, and we believe that Machine Learning Operations (MLOps) must be considered in the strategic plan. The NRC must consider where they want to host these models, and what environment they can do so securely (whether that be on-premises, or on the cloud). Having an infrastructure to host and deploy these models is key to getting value out of AI.</i> | <p>The NRC agrees the principles in machine learning operations (MLOps) are very useful to ensure that deployment considerations are not neglected at the expense of up-front model development. However, the overall goal of the AI Strategic Plan is to ensure the staff’s readiness to effectively and efficiently review and evaluate the use of AI in NRC-regulated activities, not develop and deploy internal agency AI applications.</p> <p>The NRC appreciates the comment to consider MLOps as an important aspect of the AI lifecycle and agrees that multiple factors exist in assessing AI properties. As such, the NRC has added Table 2, “Potential AI Technical Considerations for Regulatory Decision-Making” in section 4.1 to include model maintenance and life cycle management in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 68 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <i>Additionally, there are monitoring considerations in a deployed environment. Even the best models degrade over time due to differences in real-world data and the data that the models were trained on, and being able to track, identify, and retrain models on the newest set of data is imperative for any use case. In addition to the monitoring of these models, governance should play a role in whatever tools the NRC eventually uses. There should be user roles that limit who can create models, who can deploy them, and who can monitor them.</i> | <p>The NRC agrees the appropriate model monitoring, governance, and oversight are important considerations to ensure deployed AI models do not degrade over time. However, the overall goal of the AI Strategic Plan is to ensure the staff’s readiness to effectively and efficiently review and evaluate the use of AI in NRC-regulated activities, not develop and deploy internal agency AI applications.</p> <p>The NRC appreciates the comment to consider model degradation as an important aspect of the AI lifecycle and agrees that multiple factors exist in assessing AI properties. As such, the NRC has added Table 2, “Potential AI Technical Considerations for Regulatory Decision-Making” in section 4.1 to include fielded performance degradation in the list of potential characteristics to consider. The NRC further recognizes that this may not be an exhaustive list of factors for final consideration and as the NRC implements the AI Strategic Plan, it may include other factors in assessing AI properties.</p> |
| 69 | DataRobot | Chetin Durak | ML22175A206 | ML22231B047 | <i>These models should also be extensible so they can support a variety of prediction methods including, but not limited to, batch predictions, real-time streaming, and scheduled jobs from a variety of existing databases.</i> | <p>The NRC agrees that there is practical benefit for AI model developers to include flexibility in their models so they can benefit from a variety of prediction methods and communication from a variety of information repositories. However, the overall goal of the AI Strategic Plan is to ensure the staff’s readiness to effectively and efficiently review and evaluate the use of AI in NRC-regulated activities, not develop and deploy internal agency AI applications. Therefore, the attributes discussed in the comment are not appropriate for the AI Strategic Plan as the NRC must be prepared to evaluate any implementation of potential AI applications and must be prepared to make an evaluation of the model implementation and application architecture that is submitted for regulatory review. The NRC does not make any assumptions about future AI applications, system architectures, or the quality of future submissions to the NRC.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|--------------------------------|----------------|---|--|---|---|
| 70 | Nuclear Energy Institute (NEI) | James Slider | ML22175A206 , section 4 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p.23, line 11 – p. 24, line 2 | <i>Just a couple of questions to get the conversation started. Matt, the timeframe given for this plan is fiscal years '23 through '27, and I'm curious, you're going to be halfway through fiscal '23 before you have the final plan out. And I'm wondering, there's nothing in this document that indicates what's going to be happening over those succeeding couple of years in the scope of this document. Can you describe or give any sort of overview of what we can expect the NRC to focus on in the latter half of fiscal '23 out through '27?</i> | The AI Strategic Plan is a high-level document that serves as a guide for the staff to plan, implement, and monitor work associated with AI activities. Granular schedule details are not appropriate for the AI Strategic Plan. The AI Strategic Plan details the AI-related goals and objectives on which the NRC will focus for FY'23 through FY'27. The NRC plans to engage external stakeholders, including the public, throughout the agency's implementation of the AI Strategic Plan. As mentioned in the public meeting, early NRC activities include developing an organizational framework, steering committee and working groups, and public workshops. No changes were made to the AI Strategic Plan in response to this comment. |
| 71 | NEI | James Slider | ML22175A206 Executive Summary (page ix) | ML22231B053 | Timeframe of this Plan – <i>In line 17, as in the title of the report, the text says this strategic plan covers fiscal years 2023-2027. Despite that, the remainder of the strategic plan gives no indication of the timing or sequence of any elements of the work that would be subject to this plan. What are the order and time frames in which major elements of the strategic plan would be performed?</i> | See NRC response to Comment 70. |
| 72 | NEI | James Slider | ML22175A206 , section 1, Table 1, PDF page 17 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p.42, line 24 – p. 43, line 19 | <p><i>On the question about levels of autonomy, I don't know if this goes to Matt or Luis, but I'm curious, the table that you showed implies that there's one criterion that goes into defining those levels.</i></p> <p><i>And I think that the answer is more complicated than that, and I'm just wondering if you have any thoughts at this point about how -- how you might be distinguishing, say, levels two and three in more ways than just that description of -- of autonomy versus independence and so forth, as you showed on your chart.</i></p> <p><i>And if you don't have an answer to that question, I'm just curious where in the timeframe of this plan you would expect to flesh out those criteria that distinguish, say, AI applications at level two that you might require some regulatory review from, say, levels -- AI level three applications where it might not be -- might not fall under NRC's aegis.</i></p> <p><i>I'm just curious if you have any thoughts on that or when that would be defined.</i></p> | <p>The NRC agrees with the comment in that implementation of Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making," is intended to delineate the level of NRC involvement in licensee uses of AI corresponding to the notional levels discussed in Table 1 of section 1. The table is a high-level characterization, focused on the attribute of autonomy, of various levels of AI integration in commercial nuclear activities. The NRC notes that the AI Strategic Plan is not guidance to licensees or applicants.</p> <p>Through the implementation of the AI Strategic Plan, the notional levels of autonomy presented in Table 1 of section 1 may evolve. The NRC has modified section 1 to say, "As such, the NRC will treat these differences with the appropriate level of regulatory scrutiny and consider the multiple criteria necessary to determine the appropriate regulatory involvement for each level." The NRC has added clarification to the Executive Summary and Introduction sections that the NRC has developed the AI Strategic Plan to plan and prepare for new technologies involving AI.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|--|--|---|
| 73 | NEI | James Slider | ML22175A206 section 1, "Introduction" (pages 1-1 to 1-3) | ML22231B053 | <i>Levels of AI. On page 1-3, between lines 3 and 4, Table 1 describes four levels of AI. It implies, but does not explicitly state, which of these levels might warrant an NRC regulatory decision. The NRC should strive to delineate the level of NRC involvement in licensee uses of AI corresponding to each of the four levels. For example, the description of Level 1 implies that NRC would have little or no need to weigh in on such AI uses. The NRC might have interest in such AI uses just for awareness. By contrast, for AI Level 4, it seems clear that NRC could have a more substantial regulatory role. Other situations might be less clear. For example, an AI use at one of the lower levels that involves forward-looking modeling could conceivably invite greater regulatory scrutiny than is apparent in the Table 1 descriptions. In other words, while Table 1 suggests the level of NRC involvement in an AI use might be determined by a single criterion (e.g., the level of autonomy), in fact multiple criteria might be necessary to determine the level of NRC involvement appropriate to each level of AI use.</i> | See NRC response to Comment 72. |
| 74 | NEI | James Slider | ML22213A273 , (Presentation, Slide 10, Bullet 2) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 45, line 13 – p. 45, line 17 | <i>If I could ask a follow-up question on your slide 10, on the second bullet, where you talked about the interdisciplinary team of AI subject matter experts across the Agency. I'm curious what -- what other disciplines besides expertise in AI are important to the success of that team?</i> | The NRC agrees with this commenter that a range of expertise is needed to evaluate AI applications. The implementation of the plan will include consideration of staffing needs and what domain expertise is needed to review AI applications. To further clarify this point, the NRC has modified the AI Strategic Plan introduction (section 4) to include the statement that "...the successful implementation of the goals outlined in this strategic plan will involve NRC staff with varied expertise." However, the NRC has not determined what additional disciplines are important to the success of the team. |
| 75 | NEI | James Slider | ML22175A206 , section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 55, line 3 – p. 55, line 17 | <i>I just want to exhort you, and Matt, and Terry to please engage industry early, and often as this work unfolds. I think what I hear in Brian's comment, and others is a high level of interest in working with the NRC collaboratively to find an appropriate level of regulatory scrutiny that assures public confidence in these tools. But at the same time, doesn't kill them in the crib. Because they do offer the prospect of improving safety in our operating plants, and we very much want to do those things that enhance safety, but it's important that the level of regulation be appropriate, and fitting for the technological challenges that these represent.</i> | See NRC Response to Comment 8. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|---|--|---|
| 76 | NEI | James Slider | ML22175A206 , section 2, section 3, section 4.1, section 4.2 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 55, line 25 – p. 56, line 20 | <i>A follow up question on the five year horizon that's posited at the beginning of this plan. On the one hand I would urge the NRC to be as prompt as possible in getting down the road, and developing the regulatory guidance. Because the feedback I've gotten from my industry stakeholders is that in the absence of clear regulatory guidance, that they're unlikely to tackle AI applications that require NRC review. So, you'll see -- but on the other hand, that's also a choice to help industry build experience with AI applications on business processes, for example. Like AI applied to the corrective action program, and so forth. But we would, as quickly as we can, to begin to extend the reach of these AI tools for the benefit of safety, and it's going to be important for the NRC to get its regulatory guidance out in a timely way that helps industry to understand what the challenges are -- the regulatory challenges are, and move ahead with those more ambitious applications.</i> | The AI Strategic Plan is a high-level document that serves as a guide for the staff to plan, implement, and monitor work associated with AI activities. As part of implementing the AI Strategic Plan, the NRC will prepare timelines and objectives relevant to, among other things, development of regulations and guidance. The NRC plans to engage external stakeholders, including the public, throughout the agency's implementation of the AI Strategic Plan. The NRC recognizes the importance of stakeholder engagement and transparent public outreach. There will be opportunity for stakeholder engagement during the development of the AI framework, and any potential changes to regulations. No changes were made to the AI Strategic Plan in response to this comment. |
| 77 | NEI | James Slider | ML22175A206 , section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 62, line 7 – p. 62, line 20 | <i>But whatever the agency can do to accelerate the work that is called for in this plan is, I think important, and beneficial to NRC, and the industry. In that spirit, I would like to ask, in section four the plan mentions working with other federal agencies in the draft to inform the drafting of AI standards, and guidance documents. And I'm curious what other federal regulatory agencies have experience with AI that is relevant to the NRC's mission. I wonder if you might be able to comment on that.</i> | The NRC is aware that there are other federal agencies that have similar regulatory structures to regulate safety-critical applications. The NRC is aware of AI research and experience at other federal agencies such as National Institute of Standards and Technology (NIST) (https://www.nist.gov/artificial-intelligence), Food and Drug Administration (FDA) (https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-artificial-intelligence), Federal Aviation Administration (FAA) (https://www.transportation.gov/AI), and Department of Energy (DOE) (https://www.energy.gov/ai/artificial-intelligence-technology-office). The NRC has an existing MoU with the U.S. Department of Energy on "Operating Experience and Applications of Data Analytics" (ML21069A328). In addition, the NRC has an MOU with a non-Federal entity, the Electric Power Research Institute on "Collaborative Nuclear Safety Research" (ML21263A196). During implementation of the AI Strategic Plan, other federal agency's publicly available plans, regulations, and guidance with respect to AI will be surveyed to gain lessons-learned or insights on existing potential regulatory frameworks that will inform the objectives of Strategic Goal #1, "Ensuring NRC Readiness for Regulatory Decision-Making." No changes were made to the AI Strategic Plan in response to this comment. |
| 78 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Relevant Experience. <i>On page 4-2, in lines 34-42, the text indicates that NRC will work with other federal agencies to inform the drafting of AI standards and guidance documents. What other federal regulatory agencies have experience with AI that is relevant to the NRC mission?</i> | See NRC response to Comment 77. |
| 79 | NEI | James Slider | ML22175A206 , section 4.3 | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 63, line 10 – p. 63, line 17 | <i>...if the NIST group, or NIST publishes a standard on say explainability, you would see the NRC as -- the NRC would take that guidance on what explainability means, and how you demonstrate it, et cetera, et cetera, and interpret that in the NRC's domain of safety regulation, is that how that would flow down?</i> | The NRC will consider consensus codes and standards as appropriate to meet its statutory obligations. As the NRC considers codes and standards during the implementation of the AI Strategic Plan, it will engage stakeholders in transparent and public processes. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|--|--|---|
| 80 | NEI | James Slider | ML22175A206 , (All) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 72, line 21 – p. 73, line 7 | <i>Just a final question on process. I'm just wondering Matt, if you could just tell us what we can expect to see on this project over the next say six months, between now, and when the final strategic plan is ready for release next spring. You mentioned in your previous slide that you've got an ACRS meeting planned for November. Comments on this draft are due August 19th. And I'm just wondering if you can give us some idea of what we might be seeing in terms of public communications, or public releases of information between now, and next spring.</i> | Consistent with NRC practice, notification of all public meetings will be published 10 days in advance, including engagement through AI-specific public workshops, Advisory Committee on Reactor Safeguards public meetings, the NRC Regulatory Information Conference, and any future public meetings to engage stakeholders on regulatory research plans or further information related to implementation of the strategic goals. No changes were made to the AI Strategic Plan in response to this comment. |
| 81 | NEI | James Slider | ML22175A206 , (All) | Public Meeting Transcript (ADAMS Accession No. ML22229A500) p. 76, line 15 – p. 76, line 16 | <i>My ears perked up when you said the word white papers, can you elaborate on that?</i> | If a prospective applicant chooses to engage in preapplication activities, it may choose to submit white papers as part of its pre-application engagement. Additionally, the NRC is aware that industry groups have also voluntarily submitted white papers related to advanced reactor topics for the NRC's consideration. Such white papers are not required for the NRC staff to implement the AI Strategic Plan. No changes were made to the AI Strategic Plan in response to this comment. |
| 82 | NEI | James Slider | ML22175A206 Executive Summary (page ix) | ML22231B053 | Short/Mid/Long-Term Actions – <i>In lines 28-31, the text speaks of "actions recommended for the short-term (<1 year), mid-term (1-3 years), and long-term (3-5 years) time horizons..." These time horizons are not elucidated anywhere else in the document. (Note that the same three time-horizons are mentioned again in section 4, on page 4-2, in lines 1-3, but not explained there either.) What is the relevance of these three time-horizons to the strategic plan? What specific actions are contemplated in each of the three time-horizons?</i> | The NRC agrees that the time horizons were not specifically addressed further in the plan. The NRC staff will address such level of detail on short and long-term actions in the implementation of the AI Strategic Plan. As such, the timelines referenced on page ix and page 4-2 have been removed. |
| 83 | NEI | James Slider | ML22175A206 Executive Summary (page ix) | ML22231B053 | <i>What are the short-term actions that must be completed within one year?</i> | See NRC response to Comment 82. |
| 84 | NEI | James Slider | ML22175A206 Executive Summary (page ix) | ML22231B053 | <i>During the August 3, 2022, public meeting, the NRC staff said that the strategic plan is to be finalized in Spring 2023 That would be halfway through the first year in the scope of the strategic plan. When does the clock start on actions that should be completed within the one-year "short-term" window?</i> | See NRC response to Comment 82. |
| 85 | NEI | James Slider | ML22175A206 section 1, "Introduction" (pages 1-1 to 1-3) | ML22231B053 | Definition of AI. <i>Mentioned in the Executive Summary, the Introduction, and the Glossary is a definition of AI. Establishing an agency-wide definition of AI is important to this strategic plan and the NRC's next steps on preparing for reviewing AI applications. We would encourage the NRC to begin by defining key terms and concepts so that discussions and future applications proceed with a common language. We would welcome the opportunity to engage with the NRC on the development of this common language to avoid future misunderstandings on nomenclature.</i> | The NRC agrees that terms and concepts should be clearly defined. As stated throughout the strategic plan (e.g., executive summary, introduction), the definition of AI in the strategic plan is applicable only to the strategic plan; it is not intended to be a binding regulatory definition and may be refined to reflect experience and stakeholder feedback. The NRC plans to engage stakeholders, including the public, throughout the agency's implementation of the AI Strategic Plan. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|--|-----------------------------|---|---|
| 86 | NEI | James Slider | ML22175A206 section 1, "Introduction" (pages 1-1 to 1-3) | ML22231B053 | Opportunities to Improve Safety – On page 1-1, lines 23-29, the text says, "AI provides new opportunities for organizations to enhance safety and security, improve processes, leverage historical and current data, identify research needs, and even explore autonomous control and operation...The NRC is committed to enabling the safe and secure use of new technologies, especially those that can enhance the safety and security of nuclear facilities." We agree that AI offers great potential to improve safety and security. We encourage the NRC to use its work on AI as a catalyst that supports efficient NRC decision-making on all technological innovations that benefit safety and security. Timely decision-making will be especially important in the rapidly evolving area of AI. | The NRC agrees with the comment that there may be opportunities to gain lessons-learned in the implementation of the AI Strategic Plan that could potentially improve the efficiency of future decision-making on technological innovations in the nuclear sector. Regarding potential uses of AI to increase NRC efficiency in decision-making, any future governance model for potential internal agency use of AI will be developed in coordination of the agency evidence-building plan, as required by the Evidence Act and the NRC's Information Technology/Information Management Strategic Plan (ML19323D858). No changes were made to the AI Strategic Plan in response to this comment. |
| 87 | NEI | James Slider | ML22175A206 section 1, "Introduction" (pages 1-1 to 1-3) | ML22231B053 | "Could." The descriptions of AI Levels 2-4 use the word "could", i.e., "could impact safety." The NRC should be more specific in distinguishing Levels 2-4 than the all-encompassing word, "could." | The NRC disagrees with this comment. The NRC cannot, at this time, make an assertion about what specific AI application may require regulatory approval or oversight and what aspects may affect plant safety/security and control. The AI Strategic Plan is not guidance to applicants or licensees. Rather, it provides a guide for the agency to plan, implement, and monitor activities to ensure the NRC staff remains prepared to review requested uses of AI. No changes were made to the AI Strategic Plan in response to this comment. |
| 88 | NEI | James Slider | ML22175A206 section 2, "Vision" (page 2-1) section 4, "Strategic Goals" (page 4-1) | ML22231B053 | "Ensure." The stated vision is, "The NRC's vision is to continue to keep pace with technological innovations to <u>ensure</u> the safe and secure use of AI in NRC-regulated activities." [Emphasis added.] We respectfully suggest the NRC's vision should be to "enable" the safe use of AI, not to "ensure." The licensee is responsible for the safety and security of its licensed activities; thus, it is the licensee who ensures the safety and safe use of AI in its licensed facility. The NRC's role is to <u>oversee</u> the licensee's design and implementation of AI and verify that the AI meets regulatory requirements. Equally important, however, is the NRC striking the balance needed to regulate appropriately to assure adequate protection but not excessively to the point of stifling AI-based innovations that would benefit public health and safety. (This comment about the term "ensure" also applies to the description of the ultimate outcome of Strategic Goal 1 mentioned in section 4, page 4-1, lines 7-10, which repeats the NRC Vision statement.) | The NRC agrees with the intent of the comment insofar as the clarity of the vision statement could be increased. The AI Strategic Plan has been revised in section 2, and section 4, to state that "The NRC's vision is to continue to keep pace with technological innovations to allow for the safe and secure use of AI in NRC-regulated activities, when appropriate." |
| 89 | NEI | James Slider | ML22175A206 section 3, "Purpose and Drivers" (page 3-1) | ML22231B053 | Implications of the Five-Year Horizon: The NRC assumes that licensing applications that include the use of AI technologies may be submitted for NRC review and approval in the next five years. <i>i. What is significant about this five-year horizon to the NRC?</i> | The statement in the strategic plan regarding the five-year horizon reflected the NRC's understanding that licensees may submit applications for NRC review and approval within the next five years that include the use of new AI technologies. The implementation of the AI Strategic Plan is not dependent on a licensee submittal in any particular timeframe. Reference to the fact that applications may be submitted in the next five-years mentioned in section 3 and section 4.1 has been changed to "...the next few years...". The NRC notes that the FY23-27 duration of the AI Strategic Plan is meant to match the agency's four-year strategic plan duration (ML222067A170). |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|-----------------------------|--|--|
| 90 | NEI | James Slider | ML22175A206 section 3, "Purpose and Drivers" (page 3-1) | ML22231B053 | Implications of the Five-Year Horizon: <i>The NRC assumes that licensing applications that include the use of AI technologies may be submitted for NRC review and approval in the next five years.</i> ... <i>ii. How does this horizon relate to the NRC's short-term actions (<1 year), mid-term actions (1-3 years), and long-term actions (3-5 years) mentioned in the Executive Summary and section 4?</i> | The NRC will address this level of detail during the implementation of the plan. The high-level AI Strategic Plan is not intended to contain such a level of granularity; therefore, the NRC has removed the timelines referenced in the Executive Summary and section 4. |
| 91 | NEI | James Slider | ML22175A206 section 3, "Purpose and Drivers" (page 3-1) | ML22231B053 | Implications of the Five-Year Horizon: <i>The NRC assumes that licensing applications that include the use of AI technologies may be submitted for NRC review and approval in the next five years.</i> ... <i>iii. Does the five-year horizon mean that all the work identified in this strategic plan must be completed within five years? If so, what is that scope of work?</i> | The implementation of the AI Strategic Plan is not dependent on a licensing submittal within any particular timeframe but is intended to ensure that the NRC maintains its readiness to review applications involving new or unique AI technologies. The NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. Reference to the fact that applications involving the use of AI may be submitted in the next five years in section 3 and section 4.1 have been modified to "...the next few years..." as the FY23-27 duration of the AI Strategic Plan is meant to be consistent with the duration of the agency's four year strategic plan (ML22067A170). The NRC staff will provide further detail on the scope of work to be conducted under the plan in future public forums as part of the implementation of the plan. |
| 92 | NEI | James Slider | ML22175A206 section 3, "Purpose and Drivers" (page 3-1) | ML22231B053 | Operating Plant Horizon: <i>NEI discussions with a cross-section of operating plant representative suggest it is unlikely that they will be ready to submit license applications involving AI technologies in the timeframe contemplated in the strategic plan. This answer is highly uncertain due to the varying levels of interest in AI, the rapid evolution of AI technologies, and differing tolerances for business risk among the operating plant owners. Operating plant companies want to gain experience with non-safety uses of AI (for example, in Balance of Plant functions and business system functions), develop expertise and confidence in the use of AI and understanding how to leverage AI, and advance the digitization of plant systems and plant data to the degree necessary to enable effective use of AI. If the horizon for the first operating plant license amendment request involving AI is more than five years away, how would that change the NRC's plans described in NUREG-2261?</i> | The NRC strives to keep pace with new technologies and is maintaining its awareness and surveying the landscape. The NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. The implementation of the plan is not dependent on a licensing submittal in any predetermined timeframe but is intended to ensure that the NRC maintains its readiness to review applications involving new or unique technologies. No changes were made to the AI Strategic Plan in response to this comment. |
| 93 | NEI | James Slider | ML22175A206 section 3, "Purpose and Drivers" (page 3-1) | ML22231B053 | Advanced Reactors Horizon: <i>We expect the panoply of advanced reactor developers to take a variety of approaches on incorporating AI technologies in their designs and operations. Some may try to "push the envelope," looking for opportunities to leverage AI technologies to the maximum extent in their design and initial licensing. Others may take more conservative approaches that avoid the uncertainties of seeking NRC approval of an AI application during initial licensing review. NEI is seeking more information from the advanced reactors community to better understand what those plans look like. In the meantime, the NRC should look at its existing requirements and guidance on software and programs as the starting point for evaluating potential AI applications.</i> | The NRC agrees that different applicants or advanced reactor developers may take a variety of approaches on incorporating AI technologies. As described in section 4.1 of the AI Strategic Plan, "the NRC will assess whether any regulatory guidance (e.g., regulatory guides or standard review plan sections) or inspection procedures need to be updated or created to clarify the process and procedure for the licensing and oversight of AI in NRC-regulated activities." Such an assessment will involve looking at the NRC's existing guidance. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|-----------------------------|--|---|
| 94 | NEI | James Slider | ML22175A206 section 4, "Strategic Goals" (pages 4-1 to 4-5) | ML22231B053 | <p>Engaging with stakeholders. On page 4-1, in lines 16-17, the text says, "The NRC is committed to engaging the industry and relevant stakeholders to maintain awareness of industry efforts (AI Strategic Goal 3) and prepare for regulatory reviews." The wording implies this engagement is important (or most important) in Strategic Goal 3. We agree that NRC engaging with industry and other stakeholders is vital in Strategic Goal 3, but we would emphasize that engagement is just as important in the other Strategic Goals as well. We strongly encourage the NRC to approach every aspect of developing its capability to make regulatory decisions pertaining to AI with a commitment to engaging with industry and stakeholders at all stages. By engaging stakeholders through all the NRC Strategic Goals, the NRC is more likely to foster buy-in and support from the public and the regulated industry.</p> <p>During the August 3, 2022, public meeting, the NRC staff expressed an expectation that significant industry engagement will also occur through pilot studies mentioned under Strategic Goal 5 (page 4-5, line 14 of the plan). The NRC appears to be counting on pilot studies to help determine the depth of review needed in AI applications. The NRC said pilot studies would be discussed during a proposed AI workshop in 2023 (date to be determined). However helpful pilot studies may be, the timing of pilot studies is unknown and cannot be predicted with confidence. Hence, pilot studies must be considered a complement to, not a substitute for, direct engagement with industry as soon as practical and as often as necessary to advance the development of appropriate regulatory guidance.</p> | <p>The NRC agrees in part with this comment. The NRC agrees stakeholder engagement is important, however, the majority of stakeholder engagement is envisioned to occur in implementing Strategic Goal #3 as it will inform development of other strategic goals. Per Strategic Goal #3, the NRC commits to engage with stakeholders, including the public, on development of the AI framework. Other strategic goals may involve engagement with stakeholders as applicable to update or solicit feedback on implementing the strategic goal objectives.</p> <p>The NRC disagrees with the commenter's statement that the NRC appears to be counting on pilot studies. As described in section 4.5 of the AI Strategic Plan, the NRC "is planning to investigate" engaging with the nuclear industry on "potential pilot studies and proofs of concept[.]" Section 4.5 further states that pilots or proofs of concept "may help" identify challenges associated with implementing the NRC's AI framework.</p> <p>The NRC considers public involvement in, and information about, our activities to be a cornerstone of good regulation. We recognize the public's interest in the NRC's regulatory activities and strive to make timely information available to the public and provide the public with meaningful engagement opportunities. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 95 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | <p>"Consideration or Oversight". On page 4-2, in lines 13-16, the text says, "The NRC anticipates that within the next five years an existing licensee, new, or advanced nuclear technology applicant may employ AI in such a manner that it requires NRC regulatory consideration or oversight." What is significant about use of the term "regulatory consideration or oversight" here instead of saying, for example, "a regulatory decision"?</p> | <p>The goal of the AI Strategic Plan is to ensure the NRC staff continues to be positioned to effectively and efficiently review and provide oversight of the implementation of new and evolving technologies. This language is intended to capture license applications and other activities, such as regulatory inspections. No changes were made to the AI Strategic Plan in response to this comment.</p> |
| 96 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | <p>Timing. Does the NRC intend to complete all the work derived from Strategic Goal 1 within five years?</p> | <p>The NRC strives to keep pace with new technologies and is maintaining its awareness and surveying the landscape. The NRC maintains a posture to be ready for potential future applications and ensure future readiness should an applicant choose to use AI. The implementation of the plan is not dependent on a licensing submittal in any predetermined timeframe but is intended to ensure that the NRC maintains readiness to review applications involving new or unique technologies. Given the rapidly evolving pace of AI technologies, the NRC staff expects that certain aspects of the work to support Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making" (e.g., continuing to maintain awareness of technological innovations in AI) will be an ongoing effort. No changes were made to the AI Strategic Plan in response to this comment.</p> |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|-----------------------------|--|---|
| 97 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Deadline. <i>When would NRC be ready to review an application with an AI component?</i> | The NRC staff is currently capable of reviewing applications involving the use of AI. The overall objective of the AI Strategic Plan is to continue to be flexible to accommodate rapidly evolving technological innovations so that the NRC remains ready to review and evaluate applications involving AI. No changes were made to the AI Strategic Plan in response to this comment. |
| 98 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Pace of Developments. <i>Licensees are utilizing project management methods, such as AGILE, to determine more quickly whether a given approach is successful and, if it is not, quickly learn from the failure and adapt. The technologies of artificial intelligence are rapidly evolving as well. This means that the NRC's pace of development and pace of its future regulatory decisions on AI applications must be correspondingly swift. If the NRC takes one to two years to review a license application involving an AI tool, the NRC's slow pace will impair the industry's pursuit of beneficial AI tools and potentially, new reactor design development since it is expected that AI will have a significant influence on the new designs. The AI strategic plan does not indicate how the plan will enable more timely and efficient decisions on AI cases than is the norm today for "ordinary" license amendment requests.</i> | The NRC agrees in part with the comment with respect to the need to have processes in place related to AI that will improve the timeliness and efficiency of regulatory decisions commensurate with the increasingly rapid pace of AI technology development and deployment. The NRC disagrees that NRC's review will impair industry's pursuit of AI tools if the licensees and applicants submit high-quality license amendments and applications. As such, the AI Strategic Plan, in particular Strategic Goal #4, "Cultivate an AI-Proficient Workforce" and Strategic Goal #5, "Pursue Use Cases to Build an AI Foundation Across the NRC," seek to proactively prepare an AI-proficient workforce and build an AI foundation across the NRC to enable more timely and efficient decisions on AI cases. During implementation of the Strategic Plan, the NRC will consider new and updated methodologies, such as AGILE, that will continue to improve the efficiency and effectiveness of NRC reviews. No changes were made to the AI Strategic Plan in response to this comment. |
| 99 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Fees. <i>The AI strategic plan does not hint at when or how the NRC will address the impact that review fees could have on licensee business cases for AI projects. As the NRC further develops its thinking on AI, we would encourage that the review fees do not inhibit the development of AI applications that could benefit public safety and security. Consideration should be given to the use of fee waivers until the review of AI applications becomes routine and predictable. At that point, we would encourage the NRC to consider setting a fixed price or cap on review fees for AI applications.</i> | The NRC's fee regulations are primarily governed by two laws: (1) the Independent Offices Appropriation Act, 1952 (IOAA) (31 U.S.C. 9701), and (2) the Nuclear Energy Innovation and Modernization Act (NEIMA) (42 U.S.C. 2215). The IOAA authorizes and encourages Federal agencies to recover, to the fullest extent possible, costs attributable to services provided to identifiable recipients. Under NEIMA, the NRC must recover, to the maximum extent practicable, approximately 100 percent of its annual budget, less the budget authority for excluded activities. Each fiscal year, the NRC revises 10 CFR Parts 170 and 171 that set out the fees to be charged to licensees, applicants, and others requests NRC review. Fee waivers may be requested on case-by-case basis consistent with requirements in Parts 170 and 171. No changes were made to the AI Strategic Plan in response to this comment. |
| 100 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Metrics. <i>How will NRC and industry know that the NRC is ready to perform its first regulatory review? Said differently, what are the NRC's measures of success for the work implied by Strategic Goal 1? Assuming test cases will be developed, will the NRC consider doing a "dry run" of their regulatory review on a test case? This would clearly provide industry a path forward particularly if the amount of effort associated with the review of the test case is included in a report.</i> | As part of implementation of the AI Strategic Plan, actions that clearly indicate the objectives, milestones and outcomes will be developed and implemented following finalization of the AI Strategic Plan. The NRC staff will provide further detail on the scope of work to be conducted under the plan in future public forums as part of the implementation of the plan, such as development of regulations and guidance or the need for any regulatory review test cases to support Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making." The NRC recognizes the importance of stakeholder engagement and transparent public outreach. There will be opportunity for stakeholder engagement during the development of the AI framework, and any potential changes to regulations. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|---|-----------------------------|---|--|
| 101 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Relevant Experience. On page 4-2, in lines 34-42, the text indicates that NRC will work with other federal agencies to inform the drafting of AI standards and guidance documents. How will the NRC identify and apply lessons learned by those agencies in reviewing, approving, and using AI applications for their industries? | As part of the implementation of Strategic Goal #3, “Strengthen and Expand AI Partnerships,” the NRC will determine the specific types of engagement with other federal agencies. This could include means such as memoranda of understanding (MOUs), cooperative technical exchanges, conferences, seminars, or workshops like those hosted by the NRC in 2021 (https://www.nrc.gov/public-involve/conference-symposia/data-science-ai-reg-workshops.html). As discussed in section 4.2 of the AI Strategic Plan, the AI Community of Practice will facilitate the internal forum for identifying and sharing lessons learned. Specific details will be developed during the implementation of Strategic Goal #2, “Establish an Organizational Framework to Review AI Applications.” No changes were made to the AI Strategic Plan in response to this comment. |
| 102 | NEI | James Slider | ML22175A206 SG 1 – Ensure NRC Readiness for Regulatory Decision Making | ML22231B053 | Relevant Experience. On page 4-2, in lines 34-42, the text indicates that NRC will work with other federal agencies to inform the drafting of AI standards and guidance documents. We understand that the NRC is participating in the work of the American Society of Mechanical Engineers (ASME) in the development of a verification, validation and uncertainty quantification standard for computational models. In this and in any other new regulatory guidance that might apply to AI, we suggest that the NRC carefully weigh the costs and benefits to avoid creating regulatory requirements so onerous they preclude the use of AI. | The NRC will consider consensus codes and standards as appropriate to meet its statutory obligations. As the NRC considers codes and standards and the development of guidance during the implementation of the AI Strategic Plan, it will engage stakeholders in transparent and public processes. No changes were made to the AI Strategic Plan in response to this comment. |
| 103 | X-Energy | Clevin Canales | ML22175A206 section 4, “Strategic Goals” (pages 4-1 to 4-5) | ML22339A038 | (Page 4-2) Will the NRC make public the information they discover about “available AI standards”? In the event that there are no available AI standards, as may be the case, what will the NRC’s course of action be? | One of the NRC values is openness. It states that “Nuclear regulation is the public’s business, and it must be transacted publicly and candidly. The public must be informed about and have the opportunity to participate in the regulatory processes as required by law. Open channels of communication must be maintained with Congress, other government agencies, licensees, and the public, as well as with the international nuclear community.” In fulfilling its values, the NRC will consider consensus codes and standards as appropriate. As the NRC considers codes and standards during the implementation of the AI Strategic Plan, it will engage stakeholders in transparent and public processes. In accordance with Strategic Goal #1, “Ensure NRC Readiness for Regulatory Decision-Making,” and Strategic Goal #3, “Strengthen and Expand AI Partnerships,” as described in Section 4.1 of the AI Strategic Plan, the NRC staff will undertake research to develop an AI framework, which may include developing NRC guidance and participating with standards development organizations to inform the drafting and revision of AI standards and guidance documents. No changes were made to the AI Strategic Plan in response to this comment. |
| 104 | X-Energy | Clevin Canales | ML22175A206 section 4, “Strategic Goals” (pages 4-1 to 4-5) | ML22339A038 | (Page 4-5) Could you provide more information on the pilot projects and proofs of concept that the NRC plans to undertake? | Specific details regarding potential pilot projects and proofs of concept are not available at this time. The NRC staff plans to develop these details during the implementation of the plan. The NRC staff plans to engage with stakeholders in public forums about potential pilot studies or proofs of concept. No changes were made to the AI Strategic Plan in response to this comment. |

Draft Artificial Intelligence Strategic Plan (ML22175A206) Public Comment Table

| No. | Affiliation | Commenter Name | Comment Reference | Comment Source | Comment/Basis | NRC Response |
|-----|-------------|----------------|-----------------------------|-----------------------------|--|---|
| 105 | X-Energy | Clevin Canales | ML22175A206 | ML22339A038 | <i>(General) As the AI and Autonomy Levels approach Level 3 and 4 (well into the future), what activities in the Strategic Plan will ensure the NRC is prepared to properly evaluate and assess the trustworthiness of the AI system vs. the human, given that neither can be perfect.</i> | The NRC will address such level of detail during implementation of the AI Strategic Plan. In particular, the NRC will consider AI trustworthiness in the human-machine interaction as part of the development of the framework considered in Strategic Goal #1, "Ensure NRC Readiness for Regulatory Decision-Making." No changes were made to the AI Strategic Plan in response to this comment. |