

FISSION FOR ALGORITHMS

The Undermining of Nuclear
Regulation in Service of AI

Dr. Sofia Guerra, FREng and Dr. Heidy Khlaaf

November 2025

AINOW

Table of Contents

Introduction and Overview	3
Key Developments in Civilian Nuclear Regulation	4
Relegating Nuclear Thresholds for the AI Arms Race	
Critiquing the Critiques: The LNT Model and ALARA Principle	6
The Use of Generative AI and Risks to Nuclear Safety, Security, and Nonproliferation	12
Nuclear Licensing: Processes for Permitting Nuclear Facilities	21
Oversimplification and Risks of Using Generative AI for Nuclear Licensing	24
Risks to Nuclear Safeguarding and Proliferation	
New Nuclear Technologies and Overstated Promises to Power AI	30
Concluding Remarks	

Introduction and Overview

As the AI industry's insatiable energy demands collide with power grid infrastructure limits—with an expected 160 percent increase in data center power demand due to generative AI by 2030¹—AI companies have set their sights on nuclear energy as a source from which they can extract a colossal five to fifty gigawatts of additional power by 2028.² AI labs have thus begun mounting pressure to accelerate the deployment of nuclear energy sources, with major nuclear initiatives underway in an attempt to meet this recent surge in demand. These AI demands are currently infeasible, as nuclear development timelines—often ten to twenty years—are out of step with the pace of AI deployment, with large conventional nuclear reactors only capable of producing up to one gigawatt of power (i.e., 1 GW[e]) per unit. This discrepancy between the AI industry's energy demands and the lack of technical feasibility to construct nuclear plants at the pace of these stringent (and often contrived) timescales has created a chasm that is ultimately leading to a slew of efforts to fast-track nuclear timelines that raise serious safety and oversight concerns. Although the nuclear sector has the opportunity to expedite global decarbonization efforts, the monopolization of nuclear energy to explicitly power AI raises serious concerns about whether the risks associated with nuclear facilities and unsubstantiated, fast-tracked initiatives can be justified if they are not to the benefit of civil energy consumption, and if they further entrench power asymmetries that may lead to nuclear destabilization and proliferation.

This report taxonomizes and assesses these nuclear “fast-tracking” initiatives. We examine their feasibility and their impact on nuclear safety, security, and safeguards³—and more largely society's potential exposure to radiation levels—across three primary categories:

1. Policy initiatives seeking to lower regulatory practices including long-established nuclear-safety and risk-analysis approaches, safety culture, acceptable risks, and thresholds in order to reduce timescales for the construction of civil and defense nuclear facilities
2. The use of generative AI to expedite regulatory processes such as nuclear licensing and commissioning for both civil and defense nuclear facilities
3. The promotion of advanced and new nuclear technologies that are contingent on novel or unmaterialized approaches and infeasible timescales

¹ Goldman Sachs, “AI Is Poised to Drive 160% Increase in Data Center Power Demand,” May 14, 2024, <https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand>.

² See Anthropic, “Anthropic's Recommendations to OSTP for the U.S. AI Action Plan,” March 6, 2025, <https://www.anthropic.com/news/anthropic-s-recommendations-ostp-u-s-ai-action-plan>; and David Meyer, “OpenAI Reportedly Wants to Build 5-Gigawatt Data Centers, and Nobody Knows Who Could Supply That Much Power,” *Fortune*, September 27, 2024, <https://fortune.com/2024/09/27/openai-5gw-data-centers-altman-power-requirements-nuclear>.

³ The objective of safeguards is to deter the spread of nuclear weapons by the early detection of the misuse of nuclear material or technology.



Arkansas Nuclear One (Photo from Nuclear Regulatory Commission/Flickr)

First, policy initiatives are being introduced to lower regulatory oversight in order to expedite the construction of civil nuclear facilities. AI labs' assertions regarding the urgency of immediate energy needed for AI has put unprecedented pressure on regulators to reconsider well-established nuclear-safety and risk-analysis approaches, such as the linear no-threshold (LNT) model for radiation exposure and the "as low as reasonably achievable" (ALARA) risk principle with a lack of well-researched and tried-and-tested alternatives to replace these standards.^{4,5} These initiatives are simultaneously accompanied by the reduced independence of nuclear regulatory bodies,⁶ justified by alleged national-security imperatives tied to a purported AI arms race. However, such a politicization of

⁴ White House, "Ordering the Reform of the Nuclear Regulatory Commission," May 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/05/ordering-the-reform-of-the-nuclear-regulatory-commission>.

⁵ In the UK, a regulator review is underway, but the conclusions and recommendations have not been finalized yet. See Department for Energy Security & Net Zero and Ministry of Defence, *Nuclear Regulatory Taskforce: Interim Report*, August 11, 2025, <https://www.gov.uk/government/publications/nuclear-regulatory-taskforce/nuclear-regulatory-taskforce-interim-report>.

⁶ White House, "Ordering the Reform of the Nuclear Regulatory Commission," May 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/05/ordering-the-reform-of-the-nuclear-regulatory-commission>.

nuclear regulation will ultimately lead to the skewing of cost-benefit analysis that may result in increased risk tolerances to society's potential exposure to radiation levels. That is, the use of nuclear energy to power the development of generative AI further increases the risk of the public's exposure to ionizing radiation without a clear or substantiated benefit to justify it. Furthermore, the unprecedented trend of AI labs directly investing in the very nuclear providers they intend to utilize to exclusively power their data centers may lead to conflicts of interest that compromise nuclear licensee readiness and expectations in terms of organizational capabilities and safety culture.

Second, AI-based proposals (and even deployment) have been put forward by AI labs,⁷ nuclear providers,⁸ and licensees⁹ to use large language models (LLMs) to generate nuclear regulatory and licensing documents in hopes of expediting nuclear licensing and commissioning processes. Such efforts purportedly claim that generative-AI will “analys[e] historic nuclear licensing data [that] allows licensing engineers to draft new permitting documents more quickly, ready for review and refinement.”¹⁰ Yet claims that this use of AI “enables a faster and more cost-effective pathway” are not only unsubstantiated, but are impossible to actualize to meet the objective of the licensing process: to reason and understand the safety of the plant, to explore trade-offs between approaches and architecture, and to communicate why the plant is safe. Using AI will not support achieving these objectives—but it may lead to the compromise of nuclear safety and security, given that the most minute mistake introduced within the nuclear licensing process can have catastrophic and cascading consequences, compromising nuclear safety and potentially exposing society to radiation levels. The lack of security observed in commercial LLMs and their vulnerable supply chain may also lead to the compromise of the operation of existing nuclear plants, and undermine the safeguarding of measures intended to avert nuclear proliferation. Giving AI models access to sensitive nuclear data, as necessitated by these proposals, poses a novel risk of nuclear weapons proliferation and jeopardizes the ability of states to honor their international legal obligations to use nuclear material and technology only for peaceful purposes. More generally, the absence or unclear control of access to nuclear-technology information may allow nation-states or agents lacking the know-how to build nuclear weapons to be able to do so.

Third, advanced nuclear technologies such as Small Modular Reactors (SMRs), Advanced Modular Reactors (AMRs), and even nuclear fusion are being touted as alternatives that would immediately alleviate the nuclear timescale bottlenecks presented by conventional nuclear reactors. However, SMRs are a relatively novel technology, with only sixty-two SMRs being in the design prototyping phase, while

⁷ Nelli Babayan, “Microsoft AI for Nuclear Licensing,” Microsoft, September 17, 2024, <https://www.nrc.gov/docs/ML2426/ML24263A264.pdf>.

⁸ Westinghouse, “Redefining the Future of Nuclear Power with AI,” accessed October 26, 2025, <https://westinghousenuclear.com/innovation/westinghouse-ai>.

⁹ Nuclear News, “INL to Use Microsoft’s AI to Streamline Nuclear Licensing,” NuclearNewswire, July 18, 2025, <https://www.ans.org/news/2025-07-18/article-7204/inl-to-use-microsofts-ai-to-streamline-nuclear-licensing>.

¹⁰ Lloyd Register, “Lloyd’s Register to Use Generative AI to Advance the Application of Nuclear Technology in Maritime in Collaboration with Microsoft,” March 6, 2025, <https://www.lr.org/en/knowledge/press-room/press-listing/press-release/2025/lloyds-register-to-use-generative-ai-to-advance-the-application-of-nuclear-technology-in-maritime-in-collaboration-with-microsoft>.

only five are in actual operation following decades of development and construction.¹¹ Estimates of viability for SMRs as a deployable technology may take several years, with a lack of certainty that they will achieve the same economies of scale that conventional nuclear plants provide.¹² There have also been no significant scientific advancements that would prove the feasibility of nuclear fusion, let alone the design and deployment of a functioning plant within the coming years. Ultimately, claims that advanced nuclear technologies can be implemented by 2028¹³ are optimistic, implying a dramatic (and potentially dangerous) acceleration of licensing, regulatory, commissioning, and plant-construction timelines. These claims may put undue pressure on regulators¹⁴ to license novel commercial reactor designs in manufactured or infeasible timescales, which calls into question the safety of said designs.

This work thus examines these three initiatives and their feasibility, impact, and the risks they pose on two levels: increases in civilian exposure to ionizing radiation, and nuclear destabilization and proliferation for nation-states. We first provide a brief and non-exhaustive historical overview of relevant civilian nuclear regulations, followed by an analysis of how a purported AI arms race is being weaponized to dangerously discard, without evidence, the very risk and safety thresholds—that is, LNT and ALARA—established amid the nuclear arms race during the Cold War. Then, we survey the larger process of producing and reviewing nuclear licensing documentation to commission and operate nuclear plants, and the efficacy and risks in utilizing LLMs for said processes and safety argumentation (e.g., safety cases) that may lead to unsafe plants. This includes the exploration of how LLMs can be compromised through a wide array of vulnerabilities that bring not only the safety of their use into question, but their cybersecurity readiness (or lack thereof), which may also compromise national security and nuclear safeguarding. Finally, we explore how, alongside advocacy for subverting well-established nuclear safety norms, tech firms have promoted experimental advanced nuclear technologies that have been predicated on unsubstantiated claims requiring either technological breakthroughs that have yet to come to fruition, or an unsafe acceleration of nuclear timelines.

We conclude that these “fast-tracking” initiatives create dual risks for the public: the use of LLMs in nuclear infrastructure likely leading to the public’s increased risk of exposure to ionizing radiation and nuclear proliferation; and the accelerated efforts to utilize nuclear energy to power the development of these LLMs, further elevating these risks. If these initiatives continue to be pursued, their lack of safety may lead not only to catastrophic nuclear consequences, but also to an irreversible distrust within public perception of nuclear technologies that may inhibit the support of the nuclear sector as part of our global decarbonization efforts in the future.

¹¹ International Atomic Energy Agency, *Small Modular Reactors Technology Catalogue: 2024 Edition*, June 2025, https://aris.iaea.org/Publications/SMR_catalogue_2024.pdf; Nuclear Energy Agency, The NEA Small Modular Reactor Dashboard, April 2023, https://www.oecd-neo.org/icms/pl_78743/the-nea-small-modular-reactor-dashboard.

¹² Nuclear Energy Agency, “The Challenges and Opportunities in Financing Small Modular Nuclear Reactors,” OECD, May 21, 2021, https://www.oecd-neo.org/icms/pl_59235/the-challenges-and-opportunities-in-financing-small-modular-nuclear-reactor.

¹³ Stephen Nellis, “Helion Energy Starts Construction on Nuclear Fusion Plant to Power Microsoft Data Centers,” Reuters, July, 30, 2025, <https://www.reuters.com/business/energy/helion-energy-starts-construction-nuclear-fusion-plant-power-microsoft-data-2025-07-30>.

¹⁴ Nuclear News, “DOE Fast Tracks Test Reactor Projects: What to Know,” NuclearNewswire, August 12, 2025, <https://www.ans.org/news/article-7273/ten-companies-named-for-fasttracked-reactor-pilots-what-to-know>.



The first five Atomic Energy Commissioners at Los Alamos in 1947 (Photo from U.S. Department of Energy/Flickr)

Key Developments in Civilian Nuclear Regulation

At the height of the Cold War and on the heels of the Manhattan Project, the United States passed the Atomic Energy Act of 1946, wherein the Atomic Energy Commission (AEC) was formed to govern the exploitation of a new discovery: uranium fission. Although weapons development was the AEC's first priority, the US Congress embraced the heat of fission as useful for civil purposes. In 1954, the Atomic Energy Act was revised to bring private companies like Atomics International into the development of nuclear power. It is through the Manhattan Project and subsequent efforts with Atomics International to pursue the commercialization of the generation of electricity by nuclear power that in 1969, Chauncey Starr, a nuclear and electrical engineer and one of the pioneers of probabilistic risk analysis,

formalized the first risk-analysis frameworks to evaluate the risk and safety mitigations of developing and deploying powerful civilian nuclear systems.¹⁵

Starr defined *risk analysis* as the study of the relevant cause-and-effect relationships that give rise to safety risks, their magnitudes and distributions, and the identification of mitigations for risk reduction, all which concern the character of risk and the social significance of risk identified.¹⁶ He posited that risk analysis would require the consideration of the societal evaluation of risk and, as such, the interpretation of public attitudes and values. This characterization was intended to emphasize that risk analysis concerns questions that can be addressed scientifically and questions that involve public-policy concerns—where *risk tolerance* characterizes the degree, amount, or volume of risk that a society or individual will withstand; while *risk threshold* measures the level of risk exposure above which action must be taken to address risks proactively, and below which risks may be accepted. Starr noted that risk acceptance, and thus thresholds, are inseparable from risk perception and evaluation.

During the same period, radioactive fallout due to nuclear-weapons testing,¹⁷ such as the 1945 Trinity test nuclear fallout that reached forty-six states over the course of ten days,¹⁸ led to public unease over nuclear energy. As a result of the radioactive fallout from tests—leading to cancer and other illnesses in areas affected by said fallout—the public became increasingly aware of the hazards of radiation, and therefore opposed exposure to radiation independently of its source, whether military or civil. Nuclear weapons testing exposed the hidden dangers of radiation, and the line between military and civilian uses of nuclear technology was unclear: Both relied on splitting atoms, produced radiation, and carried risks of accidents and contamination.

The AEC chairman, Glenn T. Seaborg, thus prioritized restoring the credibility of civilian nuclear power and the public's risk perception of it, and commissioned the production of two safety reports: 1973's *The Safety of Nuclear Power Reactors (Light Water-Cooled) and Related Facilities*¹⁹ and *The Reactor Safety Study, from 1975*.²⁰ The 1973 report, often referred to as WASH-1250, introduced key concepts that remain crucial to the design and safety philosophy of nuclear plants to this day, such as *defense in depth*²¹ and *design-basis accidents, which concern the role of redundant and diverse safety systems in*

¹⁵ Chauncey Starr, "Social Benefit Versus Technological Risk," *Science* 165 (1969): 1232–1238, <https://doi.org/10.1126/science.165.3899.1232>.

¹⁶ Chauncey Starr, Richard Rudman, and Chris Whipple, "Philosophical Basis for Risk Analysis," *Annual Review of Energy* 1 (November 1976): 629–662, <https://doi.org/10.1146/annurev.eg.01.110176.003213>.

¹⁷ J. Samuel Walker and Thomas R. Wellock, *A Short History of Nuclear Regulation, 1946–2024*, U.S. Regulatory Commission, July 2024, <https://www2.nrc.gov/docs/ML2421/ML24211A051.pdf>.

¹⁸ Science & Global Security, "SGS Maps Radioactive Fallout from U.S. Nuclear Weapon Tests, Beginning with July 1945 Trinity Test," July 21, 2023, <https://sgs.princeton.edu/news-announcements/news-2023-07-21>.

¹⁹ U.S. Atomic Energy Commission, *The Safety of Nuclear Power Reactions and Related Facilities*, July 1973, <https://www.nrc.gov/docs/ML1214/ML12143A280.pdf>.

²⁰ U.S. Nuclear Regulatory Environment, *The Reactor Safety Study: The Introduction of Risk Assessment to the Regulation of Nuclear Reactors*, August 2016, <https://www.nrc.gov/docs/ML1622/ML16225A002.pdf>.

²¹ Defense in depth is a key concept in the strategy adopted for nuclear safety worldwide, based on different barriers of protection and additional protective means of ensuring their integrity.

*preventing accidents.*²² WASH-1250 did not, however, attempt to quantify the occurrences of effects of radiation. It was Starr's work that provided the conceptual framework for moving from safety regulations in WASH-1250 to quantify and rank risk numerically in *The Reactor Safety Study*, or WASH-1400; this allowed for the evaluation of the risk and safety mitigations involved in developing and deploying civilian nuclear plants.



The AEC's Emergency Radiation Team (Photo from U.S. Department of Energy, Flickr)

WASH-1400 thus introduced the seminal concept of *probabilistic risk assessment* (PRA), and was one of the first reports published by the newly established US Nuclear Regulatory Commission (NRC) after the US Congress transferred authority to regulate nuclear power plants from the AEC to the NRC in 1975. WASH-1400 laid the foundation for a new approach to risk analysis by quantifying the likelihood

²² Design-basis accidents are the types of accidents that the plant must be able to withstand without harm to people or the environment.

and consequences of severe reactor accidents. The PRA quantitative risk assessment considered the probabilities of accidents at three levels:

1. Core damage and core meltdown
2. Release of radiation from the containment structures of a reactor
3. Illness or death from radiation among people exposed to the radiation and environmental consequences

Furthermore, WASH-1400 shifted the focus of nuclear plant safety: Rather than only considering the worst-case accidents that engineers expected to occur (i.e., design-basis accidents), identifying lower-probability accident sequences—such as cliff-edge cases—allowed plants to be designed more safely through devising systematic system responses (i.e., beyond-design-basis events).

Despite its achievements, WASH-1400 suffered from several technical gaps, such as the level of uncertainty in the data intended to model operational metrics (e.g., the reliability of equipment), while overlooking several causes of accidents (e.g., earthquakes, equipment aging, etc.).²³ These technical gaps ultimately led the NRC to move away from WASH-1400 in 1979. Though controversial and criticized, WASH-1400 marked a fundamental shift toward quantitative risk-based safety that includes estimates of the likelihood of nuclear reactor accidents and their consequences (i.e., PRA), and the identification of corresponding accident sequences and associated preparedness for accident scenarios.²⁴

The adoption of the PRA necessitated determining the probabilities of illness or death from radiation exposure, requiring the establishment of a dose-response model for radiation risk. Furthermore, President Harry S. Truman's decision to use atomic weapons against Japan and the 1945 Trinity test nuclear fallout significantly impacted concerns over and tolerance toward nuclear exposure. Indeed, not only did the types of radioactive isotopes developed for nuclear use not previously exist in nature, but the potential for a much larger portion of the population to be exposed to the effects of nuclear radiation led to worldwide concern regarding nuclear safety. A number of organizations were thus formed to study the impact of atomic radiation. For example, the United Nations Scientific Committee on the Effects of Atomic Radiation (UNSCEAR) was established by the United Nations General Assembly on December 3, 1955 to address concerns related to the impact of radiation on the environment and human health.²⁵

²³ John H. Perkins, "Development of Risk Assessment for Nuclear Power: Insights from History," *Journal of Environmental Studies and Sciences* 4 (2014): 273–287, <https://doi.org/10.1007/s13412-014-0186-8>.

²⁴ This approach, currently used by regulators internationally, considers a broader risk-informed resilience approach, including *beyond-design basis events*.

²⁵ Werner Burkart, "50 Years UNSCEAR," statement, International Atomic Energy Agency, May 26, 2006, <https://www.iaea.org/newscenter/statements/50-years-unscear>.

These initiatives ultimately led to the establishment of radiation dose limits and targets based on the concept of the *linear no-threshold model (LNT)*, which determined that stochastic risk increases linearly with dose and that no level of exposure is entirely risk-free. There is a wide-ranging international consensus by UNSCEAR,²⁶ the International Commission on Radiological Protection (ICRP),²⁷ the NRC,²⁸ and the nuclear community more generally that LNT is based on long-term and detailed studies, and is the most appropriate model compared to other more complex and uncertain models. LNT thus continues to be used by most radiation guidelines, despite limited criticisms that this model overestimates the risks of low-dose radiation exposure.

Despite the establishment of seminal safety frameworks such as risk analysis, PRA, and LNT, on March 28, 1979, one of the most serious nuclear power plant accidents occurred at the Three Mile Island (TMI) Unit 2 nuclear power plant near Harrisburg, Pennsylvania. A combination of equipment failures, design flaws, and operator errors resulted in a partial core meltdown. The accident was caused by a number of relatively minor equipment failures that ultimately cascaded into an accident sequence given that the nuclear operators did not have enough information or training to ascertain the appropriate response, which led them to deliberately suspend their most essential safety system. The review and aftermath of the accident revealed that risk associated with a power plant cannot be estimated by probabilistic analysis of the nuclear design features alone, and that consideration of the possibility of human error in the design and operation of the plant is key for a systematic PSA approach. The realization that human factors and procedures had not been adequately included in safety considerations resulted in the recognition of *safety culture* as an essential part of nuclear safety. That is, the technical reliability of equipment was not enough to prevent accidents: Although the failure of a valve was in part the cause of the accident, several other factors, including inadequate information in the control room leading to the misdiagnosis of the faults, deficiencies in operator training, and a lack of communication and emergency preparedness all exacerbated the accident sequence. In fact, the most catastrophic nuclear accidents such as TMI, Chernobyl, and Fukushima often “bring into sharp focus how a poor safety culture can lead to the most catastrophic consequences.”²⁹

The TMI accident also further enforced the importance and acceptability of PRA as a risk-analysis technique. Before the accident at TMI, recall that the NRC had moved away from the WASH-1400 report and PRA as a strategic tool. Yet the accident sequence that led to the partial core meltdown at TMI demonstrated the value of PRA in identifying potential causes of accidents, and also confirmed that small incidents could cascade into more catastrophic accidents. In fact, the type of accident that occurred had been estimated by WASH-1400 as one of the most probable, and demonstrated the

²⁶ UNSCEAR, *UNSCEAR 2024 Report Volume I*, June 2025, https://www.unscear.org/unscear/en/publications/2024_1.html.

²⁷ International Commission on Radiological Protection (ICRP), *The 2007 Recommendations of the International Commission on Radiological Protection*, ICRP Publication 103, *Annals of the ICRP* 37, no. 2–4 (2007): 1–332, https://journals.sagepub.com/doi/pdf/10.1177/ANIB_37_2-4.

²⁸ Greta Joy Dicus, “Why the NRC Bases Its Regulations on the Linear Non-Threshold Theory,” (speech, 2001 Spring Joint Meetings of the Virginia Chapter of the Health Physics Society and the Virginia Section of the American Nuclear Society, March 24, 2001), U.S. Nuclear Regulatory Commission, <https://www.nrc.gov/docs/ML0110/ML011000432.pdf>.

²⁹ Geoff Gill, “The Development of Safety Culture in the Nuclear Industry,” *Loss Prevention Bulletin* 288 (December 2022): 2–6, https://www.icheme.org/media/19234/lpb288_pg02.pdf.



*President Jimmy Carter's motorcade leaves Three Mile Island Nuclear Power Station after the accident on April 1, 1979
(Photo from Nuclear Regulatory Commission/Flickr)*

benefits of having a tight containment building. Ultimately, the review of the TMI accident led to a number of improvements on how nuclear plants are designed, operated, and regulated. This includes the establishment of nuclear-safety culture as a core characteristic for safety performance of nuclear plants, and a fundamental criterion for nuclear operational readiness.³⁰ Detailed guidance on key characteristics or behaviors required to achieve an effective safety culture have now been developed, for example, by the World Association of Nuclear Operators (WANO),³¹ the International Atomic Energy Agency (IAEA),³² the NRC,³³ and the UK Office for Nuclear Regulation (ONR).³⁴ These guidelines have defined a common set of characteristics for safety culture, including consistent prioritization of safety issues, personal ownership and accountability, learning from incidents and accidents, and improving safety practices to reflect experience.

³⁰ Robert Fisher, *Roadmap to Operational Readiness*, New Unit Assistance Working Group, January 2023, <https://www.wano.info/resource/roadmap-to-operational-readiness>.

³¹ Geoff Gill, "The Development of Safety Culture in the Nuclear Industry," *Loss Prevention Bulletin* 288 (December 2022): 2-6, https://www.icheme.org/media/19234/lpb288_pg02.pdf.

³² International Atomic Energy Agency, "Safety Culture," leaflet, <https://www-ns.iaea.org/downloads/ni/safety-culture/safety-culture-leaflet.pdf>.

³³ U.S. Nuclear Regulatory Commission, "Safety Culture Policy Statement," accessed October 26, 2025, <https://www.nrc.gov/about-nrc/safety-culture/sc-policy-statement.html>.

³⁴ Office for Nuclear Regulation, *Licensing Nuclear Installations*, November 2021, <https://www.onr.org.uk/media/30nh5c0f/licensing-nuclear-installations.pdf>.

Safety and probabilistic risk analysis, radiation dose limits, and safety culture have thus been established as some of the most defining pillars of civilian nuclear safety and regulation. These key developments together ensure that the day-to-day operations are safe for workers and the public (e.g., “as low as reasonably achievable,” or ALARA) and that catastrophic risks are kept within tolerable societal limits. These approaches in congruence limit the likelihood and amount of dose exposure by considering dose limits, technological feasibility, economic cost, and social and environmental factors. Yet despite their importance for numerous principles central to the safety of civilian nuclear regulation, these are the very key pillars that are likely to be compromised due to the recent slew of efforts to fast-track nuclear licensing and construction timelines. In the following sections, we further examine these nuclear-safety principles in detail against proposed nuclear “fast-tracking” initiatives, and the potential impact on nuclear safety and security.

Relegating Nuclear Thresholds for the AI Arms Race

An emerging theme evident in these initiatives has been the pullback of independence traditionally possessed by nuclear regulators, allowing the US government to enforce positions that may have otherwise been independently deemed perilous for nuclear safety and security.

As the AI industry’s growth model hinges on the assertion that infinitely increasing scale purportedly leads to more powerful AI, AI companies such as OpenAI have announced investments that allocate a \$100 billion investment into data center infrastructures for model training.³⁵ These data center investments are ultimately what have led AI executives to demand, often from the US government, that a colossal five to fifty gigawatts of additional power be provided by 2028 to support the data

centers that uphold the paradigm of “scale is all you need.” In seeking to fast-track nuclear licensing and construction timelines to bridge the discrepancy between the AI industry’s energy demands and the lack of both precedent and feasibility to construct nuclear plants in the desired timelines, the US government and other nation-states have begun introducing proposals to revise, and often undermine, nuclear regulatory thresholds to expedite the development of civil and defense nuclear facilities. An emerging theme evident in these initiatives has been the pullback of independence traditionally possessed by nuclear regulators, allowing the US government to enforce positions that may have otherwise been independently deemed perilous for nuclear safety and security. Indeed, since the US

³⁵ Kate Brennan, Amba Kak, and Sarah Myers West, “Artificial Power: AI Now 2025 Landscape,” AI Now Institute, June 3, 2025, <https://ainowinstitute.org/2025-landscape>.

Congress transferred authority to regulate nuclear power plants from the AEC to the newly established NRC in 1974, the NRC has operated as an independent regulator shielded from political or industry influence. As of February 2025, however, the White House has enabled the Office of Management and Budget to oversee the regulatory process of previously independent agencies via the Ensuring Accountability for All Agencies executive order³⁶ that ultimately disrupted the autonomy many independent agencies, including the NRC, previously maintained.

The White House, emboldened by its new ability to politically steer the NRC's regulatory actions and decision-making, issued another executive order on May 23, 2025. Ordering The Reform of the Nuclear Regulatory Commission contained several provisions requiring the NRC to establish arbitrarily shortened deadlines for decisions on nuclear licensing and construction permits, including "no more than 18 months for final decision on an application to construct and operate a new reactor of any type," regardless of whether a safety record has been established for prospective designs or previously unevaluated safety mechanisms.³⁷ The executive order further demanded agency review of all the extensive NRC regulations, while recommending the dismantling of one of the aforementioned key safety pillars of nuclear safety, the linear no-threshold model and its consequential standard ALARA. Yet through a set of unsubstantiated claims, the executive order characterized both LNT and ALARA as "unscientific," disregarding international consensus based on long-term scientific studies, while providing no scientifically backed alternative.

The pretense of an AI arms race is therefore being used to discard the very risk and safety thresholds established by the nuclear-arms-race era amid the threats of the Cold War

The executive order also asked NRC to consult with the Department of Defense (DOD) and the Department of Energy (DOE) about the determination of radiation exposure limits; the two agencies may lack the requisite expertise and are incentivized to speed the pace of AI adoption, thereby compromising NRC's independence. The

NRC's deference to the DOE's goals is particularly evident in the DOE's announcement of the Reactor Pilot Program that paves a new pathway allowing for reactor authorization outside national labs using a DOE authorization process. The DOE-issued Request for Application noted that "reactors built and operated pursuant to the DOE pilot program will not require Nuclear Regulatory Commission licensing," and that "DOE-approved reactor designs can and will be fast tracked for future NRC licensing."³⁸ In fact, an addendum to the recent memorandum of understanding between the NRC and

³⁶ White House, "Ensuring Accountability for All Agencies," February 18, 2025, <https://www.whitehouse.gov/presidential-actions/2025/02/ensuring-accountability-for-all-agencies>.

³⁷ White House, "Ordering the Reform of the Nuclear Regulatory Commission."

³⁸ U.S. Department of Energy, "U.S. Department of Energy Reactor Pilot Program," accessed October 27, 2025, <https://www.energy.gov/ne/us-department-energy-reactor-pilot-program>.

the DOE states that any NRC reviews of DOE-approved reactors will focus on new risks or safety issues, such as design changes in new applications.³⁹

The deference of radiation exposure limits and licensing to the DOD and the DOE is not only political, but also reveals the underlying motives for undermining existing nuclear standards when considering two other executive orders released in congruence: the Deploying Advanced Nuclear Reactor Technologies For National Security and the Reinvigorating the Nuclear Industrial Base executive orders. Citing national security imperatives, the former order promotes the acceleration of advanced nuclear technologies through provisions that endorse categorical exclusions under the National Environmental Policy Act (NEPA) for construction of nuclear reactors on federal sites, barring any review by the NRC. Moreover, the order requires the DOE to designate AI data centers operating within DOE facilities as critical defense facilities, while designating DOE-owned sites for deployment of advanced nuclear reactor technologies to power AI infrastructure. The latter order emphasizes “a global race to dominate in artificial intelligence” and outlines a target of five gigawatts of power uprates to existing nuclear reactors and ten new large reactors with complete designs under construction by 2030—echoing the call for five to fifty gigawatts of additional power by AI companies—and a coordination with the DOD to assess the feasibility of restarting or repurposing closed nuclear power plants as energy hubs for military microgrid support.

The pretense of an AI arms race is therefore being used to discard the very risk and safety thresholds established by the nuclear-arms-race era amid the threats of the Cold War, bringing into question how contradicting these fundamental thresholds is “ushering in nuclear renaissance.”⁴⁰ The purported cost-benefit justification for this discrepancy is that the accelerated and ubiquitous adoption of AI above all else provides a marker of the US’s technological advantage and defense prowess over China and other adversaries. Not only are these presumptions dubious given the historical inaccuracies and questionable efficacy of AI-based systems,^{41,42,43,44,45,46} but they now concede the determinations of nuclear risk tolerances and thresholds to partial actors such as AI companies and the DOD, reflecting

³⁹ Nuclear Regulatory Commission, Addendum No. 9 to the Memorandum of Understanding Between U.S. Department of Energy and U.S. Nuclear Regulatory Commission on Nuclear Energy Innovation for Coordinating DOE and NRC Technical Expertise and Knowledge on Advanced Nuclear Reactor and Advanced Reactor Fuel Technologies, October 24, 2025, <https://www.nrc.gov/docs/ML2530/ML25303A288.pdf>.

⁴⁰ Department for Energy Security and Net Zero, “Golden Age of Nuclear Delivers UK-US Deal on Energy Security,” press release, GOV.UK, September 15, 2025, <https://www.gov.uk/government/news/golden-age-of-nuclear-delivers-uk-us-deal-on-energy-security>.

⁴¹ Heidy Khlaaf and Sarah Myers West, “Safety Co-Option and Compromised National Security: The Self-Fulfilling Prophecy of Weakened AI Risk Thresholds,” *arXiv preprint*, April 21, 2025, <https://arxiv.org/pdf/2504.15088>.

⁴² Heidy Khlaaf, Sarah Myers West, and Meredith Whittaker, “Mind the Gap: Foundation Models and the Covert Proliferation of Military Intelligence, Surveillance, and Targeting,” *arXiv preprint*, October 18, 2024, <https://arxiv.org/abs/2410.14831>.


⁴³ Haydn Belfield and Christian Ruhl, “Why Policy Makers Should Beware Claims of New ‘Arms Races,’” *Bulletin of the Atomic Scientists*, July 2022, <https://thebulletin.org/2022/07/why-policy-makers-should-beware-claims-of-new-arms-races>.

⁴⁴ Stephen Cave and Seán S. ÓhÉigeartaigh, “An AI Race for Strategic Advantage: Rhetoric and Risks,” in *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, December 27, 2018, <https://dl.acm.org/doi/10.1145/3278721.3278780>.

⁴⁵ Heather M. Roff, “The Frame Problem: The AI ‘Arms Race’ Isn’t One,” *Bulletin of the Atomic Scientists* 75, no. 3 (May 2019): 95–98, <https://doi.org/10.1080/00963402.2019.1604836>.

⁴⁶ Helen Toner, Jenny Xiao, and Jeffrey Ding, “The Illusion of China’s AI Prowess,” *Foreign Affairs*, June 2023, <https://www.foreignaffairs.com/china/illusion-chinas-ai-prowess-regulation-helen-toner>.

an adoption of risk scales that is fundamentally at odds with democratic deliberative norms and in contradiction with long-term scientific studies and consensus.



A construction photo of the Fast Flux Test Facility, a nuclear test reactor owned by the U.S. DoE (Photo from U.S. Department of Energy, Flickr)

Critiquing the Critiques: The LNT Model and ALARA Principle

The claims that the Ordering the Reform of the Nuclear Regulatory Commission executive order puts forward regarding LNT and ALARA as being “unscientific” likely derive from long-standing biological, epidemiological, and ethical critiques that have yet to offer a materialized alternative. Indeed, estimations of cancer risk at low doses of ionizing radiation are complex, as there is well publicized research that claims that risks of cancer after low-dose radiation are much higher than those provided by international scientific consensus, while others claim much lower risks at low radiation doses, and

even no risk at all, because of an assumed low-dose threshold for the process of cancer induction.⁴⁷ However, recent criticism against the use of LNT and ALARA have often cherry-picked selective data points from the latter studies, which regard the following:

- *Epidemiological uncertainty*: Evidence for harm at very low doses (i.e., <100 millisieverts [mSv]) is weak or statistically inconclusive. Some studies suggest a threshold such as the Hormesis model.⁴⁸
- *Historical bias*: Some scholars argue that the early adoption of LNT in the 1956 report on Biological Effect of Atomic Radiation (BEAR) may have been ideologically driven, given that the report did not consider all available scientific evidence.⁴⁹
- *Biological repair mechanisms*: Critics argue that DNA repair, apoptosis, and adaptive responses may be able to mitigate low-dose radiation effects—which LNT does not consider—and that repair mechanisms can activate if radiation exposure takes place during a longer time interval.⁵⁰

Conversely, the international scientific consensus on radiation risk has been built over many decades by a large cohort of international scientists and organizations who have in fact considered the strengths and weaknesses of all available evidence, including how alternatives like threshold or hormesis models lack robust low-dose human data and thus regulatory consensus. Organizations such as UNSCEAR, ICRP, and the IAEA have repeatedly concluded that the LNT model should remain the default framework for radiation protection due to its practical and reasonable nature,⁵¹ and that the associated radiation risk coefficients are “coherent with radiobiological knowledge, adhere to epidemiological information, and incorporate ethical judgements on the relative harm associated with different health effects.”⁵² In addition, the US National Science And Technology Council have acknowledged the lack of conclusive evidence on alternatives to LNT and ALARA, and have put forward a yet to be executed strategy to reach conclusive evidence on low dosage models.⁵³

The White House’s disregard for international consensus around the LNT model has consequently led the Ordering the Reform of the Nuclear Regulatory Commission Executive Order to recommend disbanding the ALARA standard given that it is predicated on LNT. ALARA is a principle of optimization

⁴⁷ Public Health England, “Ionising radiation: estimation of cancer risk at low doses,” GOV.UK, September 4, 2008, <https://www.gov.uk/government/publications/ionising-radiation-estimation-of-cancer-risk-at-low-doses/ionising-radiation-estimation-of-cancer-risk-at-low-doses>.

⁴⁸ Edward J. Calabrese, “Hormesis: a revolution in toxicology, risk assessment and medicine,” *EMBO Reports* 5 (October 2004): S37–S40, <https://pmc.ncbi.nlm.nih.gov/articles/PMC1299203>.

⁴⁹ Edward J. Calabrese, “On the Origins of the Linear No-Threshold (LNT) Dogma by Means of Untruth, Artful Dodges and Blind Faith,” *Environmental Research* 142 (2015) 432–442, <https://atominights.com/wp-content/uploads/LNT-and-NAS-Environ.-Res.-1.pdf>.

⁵⁰ Edward J. Calabrese and Paul B. Selby, “Muller Mistakes: The Linear No-Threshold (LNT) Dose Response and US EPA’s Cancer Risk Assessment Policies and Practices,” *Chemico-Biological Interactions* 383 (2023): 110653, <https://doi.org/10.1016/j.cbi.2023.110653>.

⁵¹ United Nations Scientific Committee on the Effects of Atomic Radiation, *Effects of Ionizing Radiation*, 2006, https://www.unscear.org/unscear/uploads/documents/unscear-reports/UNSCEAR_2006_Report_Vol.I.pdf.

⁵² United Nations Scientific Committee on the Effects of Atomic Radiation, *Sources, Effects and Risks of Ionizing Radiation*, 2012, <https://www.unscear.org/unscear/en/publications/2012.html>.

⁵³ Executive Office of the President, “Radiation Biology: A Response to the American Innovation and Competitiveness Act,” National Science and Technology Council, January 2022, <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/01/LDR-Report-2022.pdf>.

that—with consideration of dose limits, technological feasibility, economic cost, and social and environmental factors—necessitates that all exposures to ionizing radiation of any member of the public are kept as low as reasonably achievable through the examination of the benefits from risk acceptance against the cost and utility of any further risk reduction. Some critics, arguing that the LNT model is overly cautious, have claimed that ALARA is unproportional in its risk reduction and is thus responsible for prolonged nuclear construction and licensing timelines. Such arguments ultimately misconstrue the principles of risk proportionality in nuclear safety, and overlook numerous other complex factors that contribute to these timeline delays.⁵⁴

The intended objective of ALARA is to reach an “acceptable” level of risk—that is, that a risk’s corresponding losses are considered acceptable by a society given the low radiological significance of the releases against high costs that may be associated with reducing them further. This “acceptable” level of risk needs to be below the dose limit; some interpretations of ALARA (e.g., *as low as reasonably practicable*, or ALARP) regard this dose limit as the upper bound of the “tolerable” level of risk.⁵⁵ Note



Workers from the Division of Radiological Health monitor the air for radiation off-site from a Nevada nuclear test location, 1950 (Photo from The U.S. Food and Drug Administration/Flickr)

⁵⁴ Josh T. Smith, “Brief: Streamline U.S. Nuclear Regulation for Energy Abundance,” Abundance Institute, April 2025, https://api.slash.am/storage/v1/object/public/page_uploads/ed9aed01-09a1-4d95-be93-405cceb50dd/articles/nuclear-regulation-energy-abundance/1744930150357/AbundanceInstitute_2025_Energy_Publications_APR_StreamlineUSNuclearRegulationForEnergyAbundance.pdf.

⁵⁵ European ALARA Network, “Developing ALARA culture,” *Nuclear Engineering International* (November 2012): 20–21, <https://www.eu-alara.net/index.php/activities/eand-documents-and-publications/docman-menu/sg-dec2012/160-developing-alara-culture-nuclear-engineering/file.html>.

that for ALARP, risk “tolerability” is distinct from risk “acceptability.”⁵⁶ Above a certain level—the *dose limit*—risks need to be controlled so that the likelihood of exposure above those limits cannot be justified under any ordinary circumstances.⁵⁷ Below such levels, operations are allowed to take place provided that the associated risks have been made as low as reasonably achievable. ALARA is thus an obligation of means in that it depends on processes and judgments, and is not a given value of exposure,⁵⁸ as is often misconstrued by critics. The acceptable level of exposure is ultimately circumstantial given considerations of technological feasibility, economic cost, and social and environmental factors, including the recognition that some marginal risk reduction may be unjustifiably costly, while operating under the assumption that stochastic radiation risk increases linearly with dose.

As an example, the normal operation of any plant imposes doses of radiation on workers and some members of the public that are greater than what the general public is allowed exposure to. Many nation-states legally follow the ICRP’s specified limits of 20 mSv a year on average over a five-year period, with no more than 50 mSv in any one year for nuclear personnel⁵⁹ and, in the case of members of the public, 5 mSv in any one year. This corresponds to levels above which the risks have been deemed intolerable. The difference between these exposures and those resulting from an accident (i.e., an uncontrolled release) is that the former are actual and continuously measurable releases during

normal operation, whereas the latter are a question of probabilities.

Although ALARA is an integral part of the overall design of the plant, concerns that it could result in a practice of dose minimization rather than risk-informed optimization, while inaccurately attributing all nuclear costs and timelines to the principle of optimization, misrepresent the complexity of nuclear safety and licensing processes and the ALARA principle itself.

Critics often fail to consider that the control of radiation exposure to nuclear personnel and the public goes *beyond* the principles of optimization of protection and safety within the bounds of normal facility operations.⁶⁰ Adequate protection is also required for fault and accident prevention and conditions, where the risk of uncontrolled radiation releases resulting from accidents must be considered during the design and licensing of nuclear plants through techniques such as *defence in*

⁵⁶ For ALARP, risk tolerability regards the willingness to tolerate a risk so as to secure benefits under the assumption that said risks are sufficiently controlled. On the other hand, an “acceptable” risk means that a risk’s corresponding losses are considered acceptable by a society given existing social, economic, and environmental conditions.

⁵⁷ International Atomic Energy Agency, *Regulatory Control of Radioactive Discharges to the Environment*, October 2018, https://www-pub.iaea.org/MTCD/Publications/PDF/PUB1818_web.pdf.

⁵⁸ European ALARA Network, “Developing ALARA Culture.”

⁵⁹ International Commission on Radiological Protection, *Application of the Commission’s Recommendations for the Protection of People in Emergency Exposure Situations*, ICRP Publication 109 (2009), <https://www.icrp.org/publication.asp?id=icrp%20publication%20109>.

⁶⁰ International Atomic Energy Agency, *Regulatory Control of Radioactive Discharges*.

depth,⁶¹ where multiple barriers are in place to prevent the likelihood and consequences of possible fault conditions.

Although ALARA is an integral part of the overall design of the plant,^{62,63} concerns that it could result in a practice of dose minimization rather than risk-informed optimization,⁶⁴ while inaccurately attributing all nuclear costs and timelines to the principle of optimization,⁶⁵ misrepresent the complexity of nuclear safety and licensing processes and the ALARA principle itself.

Operationalization concerns that dose minimization rather than a risk-informed optimization is being pursued, or where overly cautious approaches are adopted to secure regulatory approval,⁶⁶ should be addressed practically where needed, rather than through the erosion of well-established standards that would lead to adverse downstream repercussions. Otherwise, the lack of scientific consensus on unmaterialized alternatives to LNT and ALARA (e.g., controlled dosage and timing) and the deference of the determinations of nuclear tolerances and thresholds to the interests of partial actors such as AI companies and the DOD,⁶⁷ are likely to lead to skewed risk tolerances that adversely affect nuclear safety and security.

Government claims that “harnessing the power of commercial nuclear to meet rising energy demand and fuel the AI revolution” and corresponding policy initiatives ultimately seek to position nuclear infrastructure as an extension of AI infrastructure in service of alleged national security imperatives.

Government claims that “harnessing the power of commercial nuclear to meet rising energy demand and fuel the AI revolution”⁶⁸ and corresponding policy initiatives ultimately seek to position nuclear infrastructure as an extension of AI infrastructure in service of alleged national security imperatives. Yet incorporating the interests of political and partial actors invested in perpetuating an AI Arms race into the cost-benefit justification and determinations of nuclear-risk proportionality will ultimately

⁶¹ International Nuclear Safety Advisory Group, *Defense in Depth in Nuclear Safety: INSAG-10*, report, 1996, https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1013e_web.pdf.

⁶² Industry Radiological Protection Co-ordination Group (IRPCG), *The Application of ALARP to Radiological Risk: A Nuclear Industry Good Practice Guide*, 2012, [https://nuclearinst.com/write/MediaUploads/SDF%20documents/IRPCG/Application of ALARP to Radiological Risk.pdf](https://nuclearinst.com/write/MediaUploads/SDF%20documents/IRPCG/Application%20of%20ALARP%20to%20Radiological%20Risk.pdf).

⁶³ International Commission on Radiological Protection, *Application of the Commission's Recommendations for the Protection of People in Emergency Exposure Situations*, ICRP Publication 109 (2009), <https://www.icrp.org/publication.asp?id=icrp%20publication%20109>.

⁶⁴ American Nuclear Society, “Risks of Exposure to Low-Level Ionizing Radiation: Position Statement #41,” November 2020, <https://cdn.ans.org/policy/statements/docs/ps41.pdf>.

⁶⁵ Neil Chilson and Josh T. Smith, “Comment on Request for Information on the Development of an Artificial Intelligence (AI) Action Plan,” *The Abundance Institute*, March 14, 2025, <https://files.nitrd.gov/90-fr-9088/Abundance-Institute-AI-RFI-2025.pdf>.

⁶⁶ Department for Energy Security & Net Zero and Ministry of Defence, *Nuclear Regulatory Taskforce: Interim Report*, August 11, 2025, <https://assets.publishing.service.gov.uk/media/6899da57e7be62b4f064320e/nuclear-regulatory-taskforce-interim-report-2025.pdf>.

⁶⁷ Victor Gilnsky, “Congress Wants to Turn the Nuclear Regulator into the US Industry's Cheerleader—Again,” *Bulletin of the Atomic Scientists*, November 21, 2024, <https://thebulletin.org/2024/11/congress-wants-to-turn-the-nuclear-regulator-into-the-us-industrys-cheerleader-again>.

⁶⁸ Department for Energy Security and Net Zero, “Golden Age of Nuclear Delivers UK-US Deal.”

lead to an overindexing on often unsubstantiated benefits of AI while undermining the risks to the general population. Previously, the determinations of safety thresholds for safety-critical systems within the US have largely been established through broad democratic deliberations involving elected representatives, regulators, and expert consensus-building that draws on both societal and individual value risks.⁶⁹ As a result, nuclear risks have been deemed justifiable—presuming that the public’s risk to exposure to ionizing radiation is reduced in line with ALARA—due to the capacity nuclear power provides in flexibly meeting much-needed civilian grid demands (e.g., millions of homes) at low CO₂ emissions in line with climate targets.

In shifting the risk calculus and tolerances of nuclear infrastructure to instead be in service of AI through top-down executive orders at odds with democratic deliberative norms, the public’s risk to exposure to ionizing radiation is relegated to the purported and unmaterialized benefits of AI, without any of the benefits that civil nuclear capacity was initially intended to provide.

In shifting the risk calculus and tolerances of nuclear infrastructure to instead be in service of AI through top-down executive orders at odds with democratic deliberative norms, the public’s risk to exposure to ionizing radiation is relegated to the purported and unmaterialized benefits of AI, without any of the benefits that civil nuclear capacity was initially intended to provide.⁷⁰ This political roll-back of LNT and ALARA without substantiated alternatives could thus result in cost-benefit justifications that skew the estimates of the impact of radiation

releases associated with nuclear reactor accidents (i.e., for PRA), leading to less rigorous plant and system designs that would diminish the safety protections currently in place. This would result not only in increasing exposure to radiation during normal operation, but also in the relaxation of safety barriers and associated mitigations for accident scenarios that may result in accidents with catastrophic consequences. Indeed, the nuclear industry’s remarkable safety record⁷¹ is not due to the inherently low risk of ionizing radiation, nor to the simplicity of harvesting energy through fission, but to the very safety practices and regulations now being relegated. It has been well established that the erosion of safety processes and complacency can only result in further accidents, as was the case with the Space Shuttle *Challenger* in 1986.⁷² As such, regulatory approaches have built upon the lessons learned from major nuclear accidents that have been consistently linked to overconfidence in design and erosion of

⁶⁹ Khlaaf and West, “Safety Co-Option and Compromised National Security.”

⁷⁰ Ivan Penn and Karen Weise, “Big Tech’s A.I. Data Centers Are Driving Up Electricity Bills for Everyone,” *New York Times*, August 14, 2025, <https://www.nytimes.com/2025/08/14/business/energy-environment/ai-data-centers-electricity-costs.html>.

⁷¹ Nuclear Energy Institute, “Safety: Is Nuclear Energy Safe? The Answer Is Unequivocally Yes,” <https://www.nei.org/fundamentals/safety>.

⁷² Diane Vaughan, *The Challenger Launch Decision Risky Technology, Culture, and Deviance at NASA*, Enlarged Edition, (Chicago: University of Chicago Press, 2016): <https://press.uchicago.edu/ucp/books/book/chicago/C/bo22781921.html>.

safety discipline over time. Unfortunately, the introduced executive orders and similar executive measures elsewhere risk repeating these mistakes again.

Beyond undermining nuclear regulatory thresholds, the US executive orders boast of employing emerging yet unsubstantiated technologies to accelerate the approval of new reactor designs, while also endorsing manufactured approval timelines that may put an unprecedented pressure on regulators to deviate from safety processes to meet contrived timeboxing. These efforts have resulted in unprecedented initiatives that will impact nuclear safety, security, and safeguards, which we explore further in the following sections.

The Use of Generative AI and Risks to Nuclear Safety, Security, and Nonproliferation

As the AI industry's insatiable energy demands collide with infrastructure limits, the mounting pressure to accelerate the deployment of nuclear energy sources⁷³ has led to major nuclear development efforts to meet this recent surge in demand. Notably, Google and Microsoft have entered power-purchase agreements with Kairos Power,⁷⁴ and Constellation Energy⁷⁵ and Helion Energy,⁷⁶ respectively; while Amazon has agreed to a memorandum of understanding (MOU) with X-Energy and Korea Hydro & Nuclear Power (KHNP).⁷⁷ Despite such lucrative deals, and the concerning executive orders that undermine nuclear regulatory norms while endorsing manufactured approval timelines to expedite the development of civil and defense nuclear facilities, the discrepancy between the AI industry's energy demands and the lack of technical feasibility to construct nuclear plants persists. AI executives now maintain that the timelines for *licensing* nuclear plants remain out of step with their immediate need to extract additional power to support AI data centers, with claims that "licensing is the single biggest bottleneck for getting new [nuclear] projects online."⁷⁸

Nuclear licensing refers to the well-established process that requires nuclear operators to demonstrate, over the lifetime of a nuclear plant, that the risks arising from their activities are

⁷³ Lee Harris, "Microsoft Bets on Artificial Intelligence to Power a Nuclear Resurgence — and More AI," *Financial Times*, November 27, 2024, <https://www.ft.com/content/ec111b47-faf7-48dc-a16e-f1247dfe65ba>.

⁷⁴ World Nuclear News, "Google and Kairos Power Team Up for SMR Deployments," October 15, 2024, <https://www.world-nuclear-news.org/articles/google-and-kairos-power-team-up-for-smr-deployments-in-us-first>.

⁷⁵ Aaron McMillan, "Inside Microsoft & Constellation's Power Purchase Agreement," *Procurement Magazine*, July 2, 2025, <https://procurementmag.com/news/microsoft-constellation-restarting-a-nuclear-reactor>.

⁷⁶ Sissi Cao, "Microsoft Inks Fusion Power Deal With Helion, Another Sam Altman Company," *Observer*, April 10, 2023, <https://observer.com/2023/05/microsoft-fusion-power-deal-helion-sam-altman>.

⁷⁷ Power Technology, "X-Energy, Amazon, KHNP and Doosan Enter MoU for SMRs Deployment," August 26, 2025, <https://www.power-technology.com/news/xenergy-amazon-khnp-doosan-mou>.

⁷⁸ Nelli Babayan, "Microsoft AI for Nuclear Licensing."

adequately controlled, while providing a clear account of the measures in place for managing those risks. In hopes of expediting the efforts required for licensing approvals, AI companies have now put forward a slew of AI-based licensing efforts, namely the use of generative AI models to generate nuclear regulatory and licensing documents. These efforts will purportedly “analys[e] historic nuclear licensing data [that] allows licensing engineers to draft new permitting documents more quickly, ready for review and refinement,” as noted by the Lloyd’s Register⁷⁹ collaboration with Microsoft to utilize OpenAI’s models to advance the deployment of nuclear in maritime applications. Other efforts, such as Atomic Canyon’s collaboration with both the Idaho National Laboratory (INL)⁸⁰ and the Oak Ridge National Laboratory (ORNL),⁸¹ claim that generative AI and LLMs will “streamline licensing for new

Yet claims that this use of generative AI ‘enables a faster and more cost-effective pathway’ to nuclear licensing are not only unsubstantiated; they also misconstrue the purpose of the licensing process, and raise serious safety and oversight concerns.

nuclear plants.”⁸² Westinghouse, a leading nuclear company, has developed the generative AI model *bertha*TM for “AI-powered licensing support” that has been trained on their own proprietary nuclear data.⁸³ Yet claims that this use of generative AI ‘enables a faster and more cost-effective pathway’ to nuclear licensing are not only unsubstantiated; they also misconstrue the purpose of the licensing process, and raise serious safety and oversight concerns.

Contrary to what these proposals tout, nuclear licensing is not a bureaucratic process with the objective of merely developing laborious documentation—it’s a process that spans the entire life cycle of the nuclear plant. While licensing documentation is an important part of the regulatory process, the production of such documentation is not its primary aim. Instead, the production of the documentation is a means of achieving a fundamental principle: that those responsible for the risks should understand and control them. As such, the most minute mistakes introduced within the nuclear licensing process—mistakes that would likely occur with the use of generative AI—can thus have catastrophic and cascading consequences, and compromise nuclear safety thresholds and society’s potential exposure to radiation levels. In the next section, we provide a brief overview of the licensing process, followed by an examination of the intrinsic contradiction the use of generative AI poses within it.

⁷⁹ World Nuclear News, “Lloyd’s Register to Use AI to Aid Maritime Nuclear Licensing,” March 6, 2025, <https://www.world-nuclear-news.org/articles/lloyds-register-to-use-ai-to-aid-maritime-nuclear-licensing>.

⁸⁰ Nuclear News, “Atomic Canyon Partners with INL on AI Benchmarks,” NuclearNewswire, September 11, 2025, <https://www.ans.org/news/2025-09-11/article-7362/atomic-canyon-partners-with-inl-on-ai-benchmarks>.

⁸¹ Oak Ridge National Laboratory, “Oak Ridge National Laboratory, Atomic Canyon to Accelerate Nuclear Licensing with AI,” July 22, 2025, <https://www.ornl.gov/news/oak-ridge-national-laboratory-atomic-canyon-accelerate-nuclear-licensing-ai>.

⁸² Ben Geman, “New AI-Nuclear Partnership Aims to Streamline Plants’ Licensing,” MSN, July 22, 2025, <https://www.msn.com/en-us/news/other/new-ai-nuclear-partnership-aims-to-streamline-plants-licensing/ar-AA1J3Yus>.

⁸³ Westinghouse, “Redefining the Future of Nuclear Power.”

Nuclear Licensing: Processes for Permitting Nuclear Facilities

Although the details of the licensing process differ from country to country,^{84,85,86} a nuclear licensee and its nuclear operations must meet safety principles set out in international standards and national regulatory guidance, such as those from the IAEA and those from the relevant national regulatory bodies (e.g., ONR, NRC). These fundamental safety principles often derive from the objective to “protect people and the environment from harmful effects of ionizing radiation,”⁸⁷ provide clear direction for the design of safety systems, and are well established across the nuclear industry. The licensing process is typically divided into several steps, with an authorization required at the end of each step: site and design review, construction, commission and operation, ongoing oversight, and decommissioning. This allows the regulator to maintain oversight throughout the process. In practice, licensees must show regulators not only that technical systems meet the expectations of national and international standards, but also that safety has been considered at every stage—from design and construction through operation to decommission—in order to receive and maintain regulatory approval.



*Containment dome for a test reactor at Brookhaven National Laboratory in New York
(Photo from U.S. Department of Energy/Flickr)*

⁸⁴ U.S. Nuclear Regulatory Environment, “How We Regulate,” accessed October 27, 2025, <https://www.nrc.gov/about-nrc/regulatory.html>.

⁸⁵ Office for Nuclear Regulation, “How We Regulate,” accessed October 27, 2025, <https://www.onr.org.uk/our-work/how-we-regulate>.

⁸⁶ Government of Canada, “Regulatory Framework Overview,” accessed October 27, 2025, <https://www.cnscc-cscn.gc.ca/eng/acts-and-regulations/regulatory-framework>.

⁸⁷ International Atomic Energy Agency, *Fundamental Safety Principles*, November 2006, https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1273_web.pdf.

For example, the UK nuclear Safety Assessment Principles (SAPs) include the following fundamental principle for understanding and controlling risks: “FP.4 - Dutyholders must demonstrate effective understanding and control of the hazards posed by a site or facility through a comprehensive and systematic process of safety assessment.”⁸⁸ As the safety assessment is performed, the licensee will identify possible hazards associated with the plant, along with measures to control them. The assessment of hazards and associated risks considers those arising both from normal operation (e.g., using ALARA) and fault and accident conditions,⁸⁹ typically performed through a combination of different kinds of analysis, such as design-basis analysis, PRA,⁹⁰ and beyond-design-basis accidents, including severe accident analysis (SAA). Architectural and operational solutions are explored to ensure that hazards have been reduced and controlled to an acceptable level using techniques such as *defence-in-depth*, which seek to prevent the escalation of deviations from normal operation and maintain effectiveness of barriers between radiation and people and the environment. Similarly, the NRC assessment of public safety and its regulations uses a risk-informed performance-based approach that considers questions such as: *What can go wrong? What are the consequences? How likely is it that something will go wrong? What performance is needed?*⁹¹ The regulatory approach requires that the licensee provide “reasonable assurance of adequate protection to public health and safety” and that it does not endanger defence or security.⁹² Indeed, one of the most important objectives of the licensing process is to reason and understand the risks of the nuclear plant design and the safety of the plant, to explore trade-offs between approaches and architecture, and to communicate why the resulting plant is safe through the aforementioned safety and probabilistic risk analyses.

Finally, a key part of the licensing process is the formal written demonstration that the nuclear facility abides by fundamental safety principles and can thus be operated safely—that is, that the proper operating conditions have been achieved that prevent accidents and mitigate their consequences, resulting in protection of people, society, and the environment from undue radiation risks.⁹³ This body of documentation records the understanding and reasoning achieved during the licensing process described earlier aimed at demonstrating the claim that *the plant is safe*. Documentation requirements vary, depending on local and national regulatory requirements and expectations. For example, in the US, the expectations for the safety analysis report (SAR) for licence applications is described in 10 CFR Part 50.34.⁹⁴ In the UK, regulations set high-level safety goals and principles through the construction

⁸⁸ Office for Nuclear Regulation, *Safety Assessment Principles for Nuclear Facilities: 2014 Edition, Revision 1*, January 2020, <https://www.onr.org.uk/media/pobf24xm/saps2014.pdf>.

⁸⁹ Office for Nuclear Regulation, “Safety Assessment Principles (SAPs),” accessed October 27, 2025, <https://www.onr.org.uk/publications/regulatory-guidance/regulatory-assessment-and-permissioning/safety-assessment-principles-saps>.

⁹⁰ Here we use probabilistic risk analysis as equivalent to probabilistic safety analysis.

⁹¹ U.S. Nuclear Regulatory Environment, “Risk Assessment in Regulation,” accessed October 27, 2025, <https://www.nrc.gov/about-nrc/regulatory/risk-informed.html>.

⁹² U.S. Nuclear Regulatory Environment, *10 CFR Part 50—Domestic Licensing of Production and Utilization Facilities*, Accessed on October 27, 2025, <https://www.nrc.gov/reading-rm/doc-collections/cfr/part050/full-text.html>.

⁹³ International Atomic Energy Agency, *IAEA Nuclear Safety and Security Glossary: 2022 (Interim) Edition*, October 2022, <https://www.iaea.org/resources/publications/iaea-nuclear-safety-and-security-glossary>.

⁹⁴ U.S. Nuclear Regulatory Environment, *10 CFR Part 50—Domestic Licensing of Production and Utilization Facilities*, § 50.54, accessed October 27, 2025, <https://www.nrc.gov/reading-rm/doc-collections/cfr/part050/full-text#part050-0054>.

of a safety case—a structured argument supported by evidence intended to justify that a system is acceptably safe for a specific application in a defined operating environment—that is driven by the licensee’s safety management systems and approaches. A licensee is responsible for providing a comprehensive, evidence-based safety case that demonstrates that the fundamental safety principles have been met, and that the risks have been reduced *as reasonably practicable* (given that ALARP rather than ALARA is utilized in the UK).

Oversimplification and Risks of Using Generative AI for Nuclear Licensing

Given that nuclear licensing documentation requires precise and rigorous legal and engineering outputs—ranging from the production and assessment of safety cases, the engineering substantiation of the design, to the construction and decommissioning of the plant itself—producing highly structured licensing documents is not a box-ticking exercise, as often implied by generative AI proposals put forward. Despite no substantiated evidence that generative AI is remotely capable of achieving the precision and engineering capabilities required for licensing activities, proposals boasting of its capabilities to draft new permitting documents persist. Propositions put forward thus present an overly simplified and imprecise interpretation of the nuclear licensing process in an attempt to shoehorn AI model use and present generative AI as a plausible solution where it is in fact not fit for purpose.

Consider the example used within Microsoft’s proposal documents, where an OpenAI model is prompted with “Can you please describe this image? It’s in a renewal application for a nuclear plant” and “Given this table and understanding of it, can you create a similar table but for the project in Pennsylvania?”⁹⁵ In both of these prompts, the user seemingly does not have capacity or understanding of fundamental nuclear information to be able to interpret a diagram of an “intake well” or a simple table-mapping permit and authorization steps required, respectively. Yet no such unqualified individual

Propositions put forward thus present an overly simplified and imprecise interpretation of the nuclear licensing process in an attempt to shoehorn AI model use and present generative AI as a plausible solution where it is in fact not fit for purpose.

would be permitted to work with or on nuclear licensing activities given the stringent qualifications required for nuclear personnel to perform their duties, also known as a *suitably qualified and experienced person* (SQEP). These prompts are more akin to a layman using generative AI to understand rudimentary nuclear matters, rather than expediting an expert’s workflow. The same Microsoft proposal additionally claims that AI models “can be

⁹⁵ Nelli Babayan, “Microsoft AI for Nuclear Licensing.”

extended to multiple regulatory regimes, so long as training data exists for it, and can generate content in those formats, and in the language required for those application documents.” This demonstrates not only a shallow understanding of nuclear licensing documentation that merely reduces it to “formatting” and “language,” but also that claims that the availability of data would guarantee reliable or accurate outcomes are unfounded.

Research has consistently demonstrated the lack of accuracy of generative AI and LLMs. In an OpenAI study that measured factuality under the constrained setting of fact-seeking queries with a single, verifiable answer, OpenAI’s GPT-4o and Anthropic’s Claude both scored less than 50 percent.⁹⁶ Furthermore, recent AI evaluations have confirmed that newer generative AI models are producing less accurate results and higher *hallucination rates*,^{97,98,99} a term used to describe responses generated by AI that contain false or misleading information presented as fact. OpenAI itself has recently conceded to the widely accepted notion that hallucinations cannot be eliminated from LLMs due to the inherent nature of deep neural networks (DNNs), which underpin their architectural foundation.^{100,101} Even for lower-risk tasks, such as providing summarization for scientific research, studies have indicated “a strong bias in many widely used LLMs towards overgeneralizing scientific conclusions, posing a significant risk of large-scale misinterpretations of research findings.”¹⁰² In safety-critical settings, such as in healthcare, recent research has demonstrated that the high LLM accuracy demonstrated on medical benchmarks (e.g., MedQA) are likely due to models exploiting statistical patterns in their training data, given that small wording tweaks in prompts cut accuracy by up to 38 percent on validated questions.¹⁰³

⁹⁶ Jason Wei et al., “Measuring Short-Form Factuality in Large Language Models,” *arXiv*, November 7, 2024, <https://doi.org/10.48550/arXiv.2411.04368>.

⁹⁷ Jeremy Hsu, “AI Hallucinations Are Getting Worse – and They’re Here to Stay,” *New Scientist*, May 9, 2025, <https://www.newscientist.com/article/2479545-ai-hallucinations-are-getting-worse-and-theyre-here-to-stay>.

⁹⁸ Maxwell Zeff, “OpenAI’s New Reasoning AI Models Hallucinate More,” *TechCrunch*, April 18, 2025, <https://techcrunch.com/2025/04/18/openais-new-reasoning-ai-models-hallucinate-more>.


⁹⁹ Megan Morrone, “Exclusive: Popular Chatbots Amplify Misinformation,” *Axios*, September 4, 2025, <https://www.axios.com/2025/09/04/popular-chatbots-amplify-misinformation>.

¹⁰⁰ Gyana Swain, “OpenAI Admits AI Hallucinations Are Mathematically Inevitable, Not Just Engineering Flaws,” *ComputerWorld*, September 18, 2025, <https://www.computerworld.com/article/4059383/openai-admits-ai-hallucinations-are-mathematically-inevitable-not-just-engineering-flaws.html>.

¹⁰¹ El-Mahdi El-Mhamdi et al., “On the Impossible Safety of Large AI Models,” *arXiv*, September 30, 2022, <https://doi.org/10.48550/ARXIV.2209.15259>.

¹⁰² Uwe Peters and Benjamin Chin-Yee, “Generalization Bias in Large Language Model Summarization of Scientific Research,” *Royal Society Open Science* 12 (March 2025): 241776, <https://royalsocietypublishing.org/doi/epdf/10.1098/rsos.241776>.

¹⁰³ Suhana Bedi et al., “Fidelity of Medical Reasoning in Large Language Models,” *JAMA Network Open* 8, 8 (2025): e2526021, <https://doi.org/10.1001/jamanetworkopen.2025.26021>.



A control room at Brookhaven National Laboratory (Photo from Brookhaven National Laboratory/Flickr)

These low accuracy rates are a far cry from the 95–99 percent safety reliability rates expected for nuclear systems, which should make nuclear regulators and licensees hesitate before utilizing inaccurate LLMs to generate safety-case documentation for the very systems requiring precise proof that they are acceptably safe with high reliability rates. Research from the National Aeronautics and Space Administration (NASA) on adjacent safety-critical fields has concluded that there currently exists no substantiated evidence that LLMs can appropriately generate arguments or parts of arguments within safety cases. Through a thorough literature review, the NASA researchers demonstrate that no studies assess either the relative costs of various methods of generating or assessing assurance arguments, or the impact of automation on human performance of supervisory tasks. The researchers additionally refute the argument that fine-tuning LLMs on appropriate data is a sufficient mitigation for the aforementioned flaws, that one “cannot expect the generation of similar-sounding text to produce potential defeaters that account for new knowledge about how safety arguments might be wrong.” They ultimately conclude that safety-case automation through the use of LLMs “does not *eliminate* human involvement but rather *changes the nature* of that involvement, with the resulting potential that the change might set the human up to fail.”¹⁰⁴

¹⁰⁴ Dr. Mallory S. Graydon and Dr. Sarah M. Lehman, *Examining Proposed Uses of LLMs to Produce or Assess Assurance Arguments*, Report, National Aeronautics and Space Administration, March 2025, <https://ntrs.nasa.gov/api/citations/20250001849/downloads/NASA-TM-20250001849.pdf>.

Indeed, an LLM could generate text that is convincing and appears regulatory-compliant, but includes subtle errors, omissions, and misinterpretations of nuclear regulations, safety standards, and dose limits resulting in flawed conclusions that lead to major accidents with catastrophic consequences. Even if such LLM outputs were to be reviewed by humans, humans are known to exhibit *automation bias*, meaning a propensity to trust AI outputs over other sources of information.¹⁰⁵ In the context of nuclear licensing, reviewing convincing text with subtle errors and omissions would thus remain challenging, even given an advanced level of nuclear expertise. The effort required to review LLM or AI-based outputs in a manner that would counteract automation bias, while meticulously ensuring the precision and correctness of error-prone outputs, would likely lead to more labor and longer timelines than when composing licensing documentation without AI-based assistance. Similar conclusions have been reached regarding productivity gains of experienced programmers utilizing LLM-based tools: AI made them 19 percent slower.¹⁰⁶ Furthermore, overreliance on LLMs and AI-generated outputs are likely to undermine human accountability while reducing critical oversight in decision-making. The delegation of organizational knowledge, personal ownership, accountability, and learning from incidents are antithetical to safety culture itself, and are likely to result in major accidents. As previously noted, safety culture is a core characteristic of safety performance, whose weakness has been significantly linked to nuclear accidents.

Safety concerns aside, the use of generative AI within safety-critical applications inherently expands the attack vectors of the infrastructure they interface with in a manner that extends beyond the traditional cybersecurity threats considered within nuclear regulatory assessments. Such expanded vectors of attacks include theoretical and practical demonstrations of “jailbreaks” and adversarial attacks that aim to craft inputs that manipulate a model to produce intentionally erroneous outputs or subvert the model’s safety filters and restrictions.^{107,108} Other new and undetectable attack vectors include poisoning web-scale training datasets and “sleepers agents” within commercial generative AI models, which may intentionally or inadvertently assist the subversion of models used within military applications and ultimately compromise their behavior.^{109,110} Through such attacks, hostile actors could gain insights into sensitive national nuclear infrastructure, such as proprietary-design, nuclear-fuel-cycle, or enrichment-pathways data that could aid nuclear proliferation. Adversarial actors could also compromise the safety of nuclear facilities through intentionally poisoning data or models to produce licensing documentation that appears to support the demonstration that a nuclear facility is safe, when such is not the case, thus increasing a society’s potential exposure to radiation levels.

¹⁰⁵ Kate Goddard et al., “Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators,” *Journal of the American Medical Association* 16, 19 (June 2011): 121–127, <https://pmc.ncbi.nlm.nih.gov/articles/PMC3240751>.

¹⁰⁶ Joel Becker et al., “Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity,” *Model Evaluation & Threat Research*, July 10, 2025, <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study>.

¹⁰⁷ El-Mahdi El-Mhamdi et al., “On the Impossible Safety of Large AI Models.”

¹⁰⁸ Baoyuan Wu et al., “Attacks in Adversarial Machine Learning: A Systematic Survey from the Life-Cycle Perspective,” *arXiv*, February 19, 2023, <https://doi.org/10.48550/arXiv.2302.09457>.

¹⁰⁹ Nicholas Carlini et al., “Poisoning Web-Scale Training Datasets is Practical,” in *2024 IEEE Symposium on Security and Privacy (SP)* (San Francisco, CA: IEEE, 2024), 407–25, <https://doi.org/10.1109/SP54263.2024.00179>.

¹¹⁰ Evan Hubinger et al., “Sleepers Agents: Training Deceptive LLMs that Persist through Safety Training,” *arXiv*, January 10, 2024, <https://doi.org/10.48550/arXiv.2401.05566>.

Attempts to address these challenges have been unsuccessful,¹¹¹ as research has persistently shown that it is always possible to construct attacks that are transferable across all generative AI models.¹¹² Any fine-tuning or guardrails introduced as a way to enable accurate performance or security protections can therefore be bypassed. Potential existential limitations in combating these novel attack vectors also arise due to the lack of traceability of human labor and unknown data sources across the supply chain of commercial generative AI models repurposed for nuclear applications. Indeed, traceability, a core requirement of nuclear systems necessitating the tracking and documenting of all artifacts throughout the development life cycle, is required to guarantee that no aspect of the development pipeline is compromised to ensure a system's security and fitness for use. Current generative AI models as they stand may thus introduce greater risks to the critical infrastructures in which they are embedded that are disproportionate to the benefits that any use of AI may produce.

Risks to Nuclear Safeguarding and Proliferation

The uses of generative AI within nuclear licensing additionally have far-reaching consequences that extend beyond the safety and security risks of nuclear infrastructure, and propagate to critical nuclear safeguarding and nonproliferation efforts. Nuclear safeguards are technical and legal measures aimed at preventing the proliferation of nuclear weapons at the state level by ensuring that countries comply

Along with the potential for adversarial actors to compromise nuclear secrets due to generative AI models' lack of security, the control of sensitive nuclear data by AI companies through models residing within their cloud infrastructure raises serious concerns over whether their possession of such data may lead to nuclear destabilization and proliferation, and further entrench power asymmetries.


with international obligations not to use nuclear materials from civil nuclear programs for non-peaceful purposes. Safeguards are primarily enforced through the Treaty on the Non-Proliferation of Nuclear Weapons (Non-Proliferation Treaty, or NPT) which is "regarded as the cornerstone of the global nuclear non-proliferation regime and an essential foundation for the pursuit of nuclear disarmament."¹¹³ The NPT aims to prevent the spread of nuclear weapons and technology and promotes cooperation in the peaceful uses of nuclear energy, with existing commitments made by China,

¹¹¹ Deep Ganguli et al., "Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned," *arXiv*, August 23, 2022; last revised November 22, 2022, <https://doi.org/10.48550/arXiv.2209.07858>.

¹¹² Andy Zou et al., "Universal and Transferable Adversarial Attacks on Aligned Language Models," *arXiv*, July 27, 2023; last revised December 20, 2023, <https://doi.org/10.48550/arXiv.2307.15043>.

¹¹³ United Nations Office for Disarmament Affairs, "Treaty on the Non-Proliferation of Nuclear Weapons," accessed October 27, 2025, <https://disarmament.unoda.org/en/our-work/weapons-mass-destruction/nuclear-weapons/treaty-non-proliferation-nuclear-weapons>.

France, Russia, the UK, and the US to ensure that nation-states that do not have nuclear weapons do not acquire them. Yet nuclear licensees, regulators, and AI companies have failed either to recognize or to publicly acknowledge how the use of generative AI may in fact lower the barrier to nuclear secrets know-how, and render some safeguarding and non-proliferation efforts moot.



*Researchers working on a computer supporting nuclear design, 1969
(Photo from Lawrence Livermore National Laboratory/Flickr)*

Along with the potential for adversarial actors to compromise nuclear secrets due to generative AI models' lack of security, the control of sensitive nuclear data by AI companies through models residing within their cloud infrastructure raises serious concerns over whether their possession of such data may lead to nuclear destabilization and proliferation, and further entrench power asymmetries. Consider that Microsoft's proposals have noted that generative AI "requires more than just historic licensing data, it also needs real-time and project specific data,"¹¹⁴ signaling that AI providers are necessitating that regulators and licensees hand over nuclear secrets for purported

¹¹⁴ Nelli Babayan, "Microsoft AI for Nuclear Licensing."

model improvements that have yet to come to fruition. This is further reflected in the Lloyd's Register¹¹⁵ collaboration with Microsoft to utilize OpenAI's models to advance the deployment of nuclear in maritime applications, and OpenAI's partnership with US National Laboratories for nuclear weapons security,¹¹⁶ both of which would require that sensitive nuclear secrets be provided to OpenAI either through model training or as inputs. Yet the knowledge, technology, and materials to build nuclear power plants and nuclear weapons are analogous, which is why a number of *export controls* and guidelines related to the transfer of nuclear equipment, materials, software, and data exist to prevent them from contributing to the development of nuclear weapons.¹¹⁷ As such, using "project-specific data" in generative AI to support nuclear licensing or defense and weapons applications implies that nuclear know-how is made available to AI providers and their models, despite AI companies' lack of experience and nuclear expertise to classify nuclear secrets or manage its misuse. These concerns are further exacerbated by tech-firm-linked energy start-ups like Oklo, which is backed by OpenAI's Sam Altman, having access to the US government's weapons-grade plutonium stockpile.¹¹⁸

Furthermore, if AI models are trained on sensitive nuclear data, generative AI models will embed said representations of this training data and, subsequently, potential nuclear capabilities within their functionality, further compromising export control and safeguarding of nuclear secrets and technology,

As such, using "project-specific data" in generative AI to support nuclear licensing or defense and weapons applications implies that nuclear know-how is made available to AI providers and their models, despite AI companies' lack of experience and nuclear expertise to classify nuclear secrets or manage its misuse.

while exacerbating vulnerabilities that allow for the extraction of model data through observed model predictions alone.¹¹⁹

Despite such risks, legislatures and nuclear regulators have not addressed the lack of adequacy of current safeguards and export controls that would accompany the use of generative AI models for nuclear applications. Ultimately, the management of safeguards, export controls, and the NPT need to be implemented by complex legislative arrangements of the states that are party to the NPT and their nuclear regulators, and that are subject to IAEA safeguards.

¹¹⁵ World Nuclear News, "Lloyd's Register to Use Generative AI," March 6, 2025, <https://www.world-nuclear-news.org/articles/lloyds-register-to-use-ai-to-aid-maritime-nuclear-licensing>.

¹¹⁶ Hayden Field, "OpenAI Partners with U.S. National Laboratories on Scientific Research, Nuclear Weapons Security," NBC News, January 30, 2025, <https://www.nbcnews.com/tech/tech-news/openai-partners-us-national-laboratories-research-nuclear-weapons-secu-rcna190008>.

¹¹⁷ Nuclear Suppliers Group, *Guidelines for Transfers of Nuclear-Related Dual-Use Equipment, Materials, Software, and Related Technology (INFCIRC/254, Part 2)*, accessed October 27, 2025, <https://nuclearsuppliersgroup.org/index.php/en/guidelines/nsg-guidelines/guidelines-part-2>.

¹¹⁸ Martha Muir, "US Offers Nuclear Energy Companies Access to Weapons-Grade Plutonium," *Financial Times*, October 21, 2025, <https://www.ft.com/content/2fbbc621-405e-4a29-850c-f0079b116216>.

¹¹⁹ Nicholas Carlini et al., "Extracting Training Data from Diffusion Models," in *Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23)* (Anaheim, CA: USENIX, 2023), 5253–5270, <https://www.usenix.org/system/files/usenixsecurity23-carlini.pdf>.

Overall, AI companies' access to nuclear secrets along with their lucrative nuclear power-purchase agreements surface a concerning trend toward tech firms accessing and monopolizing scarce public energy,¹²⁰ critical infrastructure, and nuclear material that is in opposition to both the public's and nation-states' benefit. In particular, the acquisition of power outputs from nuclear civil infrastructure by AI companies raises serious concerns over whether risks associated with nuclear facilities and unsubstantiated fast-tracked initiatives can be justified if it's not to the benefit of civil energy consumption. The unfettered and potentially uncontrolled access to nuclear know-how perpetuates this power asymmetry further and raises nuclear proliferation risks given AI companies' incessant proposals to gain access to nuclear secrets under the guise of fine-tuning generative AI models, while not having addressed AI's crucial safety and security vulnerabilities that may perpetuate proliferation. This lack of concern is indicative of AI companies' lack of expertise in classifying nuclear secrets and managing its misuse. These risks are further exacerbated by AI companies' accelerated investments in SMRs and their pressure on several nation-states (e.g., the US and UK)¹²¹ to reduce and accelerate regulation for the approval of novel yet unproven reactor designs, without either they or their subsidiaries having experience in operating reactors.¹²²

New Nuclear Technologies and Overstated Promises to Power AI

Alongside advocacy for subverting well-established nuclear safety norms, tech firms have looked to experimental advanced nuclear technologies such as SMRs, AMRs, or even nuclear fusion as alternatives to purportedly alleviate the timescale bottlenecks presented by conventional nuclear reactors. The majority of power-purchase agreements made by tech firms have focused on the development and deployment of SMRs to satiate AI datacenters' energy consumption via investments in energy start-ups such as Kairos Power,¹²³ Helion,¹²⁴ X-Energy,¹²⁵ and Oklo.¹²⁶ A total of \$3 billion has been raised in private investments, with an additional \$6 billion committed through US government

¹²⁰ Brennan et al., "Artificial Power: AI Now 2025 Landscape."

¹²¹ Department for Energy Security and Net Zero, "Golden Age of Nuclear Delivers UK-US Deal."

¹²² Allison Macfarlane, "Trump Just Assaulted the Independence of the Nuclear Regulator. What Could Go Wrong?" *Bulletin of the Atomic Scientists*, February 21, 2025, <https://thebulletin.org/2025/02/trump-just-assaulted-the-independence-of-the-nuclear-regulator-what-could-go-wrong>.

¹²³ World Nuclear News, "Google and Kairos Power team up for SMR deployments," October 15, 2024, <https://www.world-nuclear-news.org/articles/google-and-kairos-power-team-up-for-smr-deployments-in-us-first>.

¹²⁴ Sissi Cao, "Microsoft Inks Fusion Power Deal."

¹²⁵ X-Energy, "X-Energy, Amazon, Korea Hydro & Nuclear Power, and Doosan Enerbility Announce Partnership to Scale Advanced Nuclear Energy for AI Infrastructure," August 25, 2025, <https://x-energy.com/media/news-releases/x-energy-amazon-korea-hydro-amp-nuclear-power-and-doosan-enerbility-announce-partnership-to-scale-advanced-nuclear-energy-for-ai-infrastructure>.

¹²⁶ Jamie Smyth, "Sam Altman-Led Nuclear Start-Up Signs Major AI Power Supply Deal," *Financial Times*, December 18, 2024, <https://www.ft.com/content/137ff7ad-bc84-421d-bd8c-f7264adfe6e4>.

agencies, including the DOD and DOE.¹²⁷ These investments have been predicated on claims that new nuclear technologies will be in operation by 2028,¹²⁸ implying either technological breakthroughs that have yet to come to fruition, or an unsafe acceleration of nuclear timelines.¹²⁹

Although a portion of companies involved in the design and construction of SMRs—such as Rolls Royce, Électricité de France, and Korea Hydro & Nuclear Power (KHNP)—are building on their experience in traditional nuclear technologies, the energy start-ups garnering the most investment and touting as-yet-to-be-materialized claims have no nuclear operational experience. These start-ups’ lack of operational experience and pledged infeasible timelines impact their ability to meet organizational capabilities¹³⁰ and safety-culture¹³¹ expectations, as required by the IAEA. Consider as an example that the NRC, which seldom rejects formal nuclear license applications, denied Oklo’s original application because Oklo failed to provide enough information to proceed with review activities in spite of “multiple information requests, audits and public meetings.”¹³² This not only demonstrates Oklo’s lack of understanding of what is required to support a license application, but their lack of organizational competencies need to be built according to the demands of nuclear technology.¹³³ Oklo has since indicated their intention to submit a license application and has had a number of pre-submittal discussions with the NRC, in spite of changing the concept design substantially since the original submission. However, Oklo has subsequently joined the DOE Reactor Pilot Program, indicating that they may no longer require NRC licensing. Nevertheless, it is these very start-ups that have pressured governments and their regulators to license their unproven commercial reactor designs in manufactured or infeasible timescales that are ultimately at odds with well-established nuclear and safety processes.

AI labs and their nuclear energy start-up subsidiaries have consistently bolstered US government initiatives to reduce regulation and hasten approval of their often unproven reactor designs. Through the previously aforementioned executive orders and other legislative initiatives such as the Fusion Energy Act of 2023 and the Accelerating Deployment of Versatile, Advanced Nuclear for Clean Energy (ADVANCE) Act of 2024, the NRC has been persistently directed to identify and report on licensing commercial “fusion machines” and SMRs, including streamlining considerations. Concerns have previously been raised that these start-ups wish “their reactors could be exempted from the requirements that all other designs before them have had to meet” while lobbying the NRC to “trust

¹²⁷ Martha Muir, “US and Investors Gambling on Unproven Nuclear Technology, Warn Experts,” *Financial Times*, October 5, 2025, <https://www.ft.com/content/8a18e722-3efa-404e-9f2a-709eed877f18>.

¹²⁸ Stephen Nellis, “Helion Energy Starts Construction on Nuclear Fusion Plant.”

¹²⁹ Nuclear News, “DOE Fast Tracks Test Reactor Projects: What to Know,” *NuclearNewswire*, August 12, 2025, <https://www.ans.org/news/article-7273/ten-companies-named-for-fasttracked-reactor-pilots-what-to-know>.

¹³⁰ International Atomic Energy Agency, *Mapping Organizational Competencies in Nuclear Organizations: Nuclear Energy Series No. NG-T-6.14*, December 2020, https://www-pub.iaea.org/MTCD/Publications/PDF/PUB1844_web.pdf.

¹³¹ IAEA, “Culture for Safety: Nuclear Safety and Security Programme,” leaflet, n.d., accessed October 27, 2025, https://www.iaea.org/sites/default/files/culture_for_safety_leaflet.pdf.

¹³² World Nuclear News, “NRC denies Oklo licence application,” January 7, 2022, <https://www.world-nuclear-news.org/Articles/NRC-denies-Oklo-licence-application>.

¹³³ International Atomic Energy Agency, *Mapping Organizational Competencies in Nuclear Organizations*.

their simplistic computer models of reactor performance and [...] deploy their untested technology across the country.”¹³⁴ As previously discussed, these executive and legislative measures have ultimately sought to pull back the independence traditionally possessed by nuclear regulators, allowing the US government to enforce positions that may otherwise have been independently deemed perilous for nuclear safety and security. Moreover, these initiatives have failed to consider that SMRs are still an experimental technology, and that there have also been no significant scientific advancements that would prove the feasibility of nuclear fusion, let alone the design and construction of a functioning plant by 2028.

Indeed, SMRs are a relatively novel technology whose designs are often still experimental and in development. As per the IAEA, sixty-two SMRs are in the design phase, while five are in construction and only four are in actual operation, following years of development and construction.¹³⁵ Similarly, the NEA Small Modular Reactor Dashboard: Third Edition¹³⁶ identified 127 SMRs; from the seventy-four SMRs assessed in this dashboard, seven are designs under construction or operating as a first of a kind

Overall, the unsubstantiated claims and infeasible timescales put forward by many proponents of advanced nuclear technologies have consequences beyond missed timelines. They ultimately skew policymakers’ perceptions on the actualization of safe plant construction and operation, leading to regulatory changes that subvert well-established nuclear safety norms.

(FOAK). As is well known in the nuclear industry, FOAK are slow projects with significant delays and overspent, as design, construction, and licensing issues are continuously identified and addressed. Construction is typically only streamlined when it moves from FOAK to NOAK (nth of a kind), where cost and duration decreases as the number of the same SMRs are constructed. Recent studies indicate that the difference in cost and duration for the first plant to later developments may be reduced to one third.¹³⁷ Although SMRs may have a number of benefits in the future, their viability as a deployable technology may take several years, with a

lack of certainty that they will achieve the same economies of scale that conventional nuclear plants provide, especially those designs that generate less power.¹³⁸ This has already led some SMR vendors, such as NuScale and Holtec, to double their module sizes from the original designs. The Oklo Aurora microreactor has increased from 1.5 MW to 15 MW and may even go to 50 MW, while the General

¹³⁴ Allison Macfarlane, “Trump Just Assaulted the Independence of the Nuclear Regulator.”

¹³⁵ International Atomic Energy Agency, *Small Modular Reactors Technology Catalogue*.

¹³⁶ Nuclear Energy Agency, “The Challenges and Opportunities in Financing Small Modular Nuclear Reactors,” OECD, May 21, 2021, https://www.oecd-neo.org/icms/pl_108326/the-nea-small-modular-reactor-dashboard-third-edition.

¹³⁷ Chandrakanth Bolisetti et al., *Quantifying Capital Cost Reduction Pathways for Advanced Nuclear Reactors*, Report, Prepared for U.S. Department of Energy Systems Analysis & Integration Campaign, June 6, 2024, https://indigitalibrary.inl.gov/sites/sti/sti/Sort_109810.pdf.

¹³⁸ Nuclear Energy Agency, “The Challenges and Opportunities in Financing Small Modular Nuclear Reactors.”

Electric-Hitachi BWRX-300 and Westinghouse AP300 are both starting out at the upper limit of what is considered an SMR.¹³⁹

Proponents of SMRs have additionally argued that the risks associated with SMRs are less than those presented by traditional nuclear plants,¹⁴⁰ justifying “streamlining” initiatives to reduce regulation and hasten their approval. These reduced risks are purportedly due to SMRs’ relatively smaller size, which would entail lower radiation release in the case of an incident; and the incorporation of passive safety measures such as the reliance on gravity and natural air circulation to cool the reactor.¹⁴¹ Yet these measures do not dissipate either the scale or the variation of risk that would accompany the constructions and operation of SMRs. For one, as with the use of any enriched uranium in a nuclear plant, SMRs will produce nuclear waste in need of management and disposal. Some studies have demonstrated that SMRs may in fact create greater and more complex nuclear waste per unit of energy produced than large power plants.¹⁴² Furthermore, SMRs’ modularity and their associated economies of scale imply that several reactors are likely to be deployed within the same site, countering the decreased radiation argument, since the same initiating event could cause several of the modules to fail simultaneously. As with the Fukushima accident where a loss of power led to catastrophic failures in several power plant units, it is even more likely that such common-cause failures will cascade to multiple SMR reactors given their highly coupled modular design. As numerous

Moreover, seeking to accelerate the deployment of underdeveloped nuclear technologies may be to the detriment of their development, as the lack of certainty with their safety is likely to sow distrust within public perception, bringing into question the safety of these nuclear technologies and undermining the benefits they may bring once matured.

SMR designs have yet to be trialed and improved as a result of operating experience, limited data additionally increases the uncertainty and the potential for subtle unexpected cascading events that might have catastrophic consequences.

Uncertainties of whether SMRs will achieve the economy of scale of larger plants are also likely to lead to safety cost-cutting measures. SMR proponents such as ORNL have already proposed the reduction or elimination of many well-established safety features required for operating reactors like

¹³⁹ Ed Lyman, “Five Things the ‘Nuclear Bros’ Don’t Want You to Know about Small Modular Reactors,” *The Equation*, April 30, 2024, <https://blog.ucs.org/edwin-lyman/five-things-the-nuclear-bros-dont-want-you-to-know-about-small-modular-reactors>.

¹⁴⁰ IX Power Limited, “Written Evidence Submitted by IX Power™ Limited d/b/a IX Power Machines™ (SNP0013),” April 14, <https://committees.parliament.uk/writtenevidence/50216/html>; Manuel Herrera, “Are Small Modular Reactors


Europe’s Energy Salvation?” *Istituto Affari Internazionali* 22, 57 (November 2022), <https://www.iai.it/sites/default/files/iaicom2257.pdf>.

¹⁴¹ International Atomic Energy Agency, *Advances in Small Modular Reactor*

Technology Developments: 2020 Edition, 2020, https://aris.iaea.org/Publications/SMR_Book_2020.pdf.

¹⁴² Stanford University, “Stanford-Led Research Finds Small Modular Reactors Will Exacerbate Challenges of Highly Radioactive Nuclear Waste,” *Stanford Report*, May 30, 2022, <https://news.stanford.edu/stories/2022/05/small-modular-reactors-produce-high-levels-nuclear-waste>.

reinforced concrete containment structures, motor-driven emergency pumps, and rigorous quality assurance standards for backup safety equipment.^{143, 144} Yet it remains unclear if such changes would be certain to lower the overall cost of the licensing and construction of SMRs. Furthermore, SMRs' passive safety features are unlikely to be an adequate replacement for scaling back some well-established safety mechanisms. As demonstrated by the NRC's review of NuScale's SMR design, the passive emergency systems could deplete cooling water of boron, which is needed to keep the reactor safely shut down after an accident.¹⁴⁵ Technology readiness aside, progress toward deployment of SMRs and AMRs depends on several factors, such as licensing, siting, supply chain, and fuel allocation, all which contribute to the complexity and timescales of their construction.



Researchers at Brookhaven National Laboratory (Photo from Brookhaven National Laboratory/Flickr)

These uncertainties and novel risks that accompany the development and operation of SMRs indicate that significantly longer timescales are needed to appropriately assess and develop the viability, safety, and efficacy of SMRs as a deployable technology, putting them out of step with the pace of AI deployment. **Overall, the unsubstantiated claims and infeasible timescales put forward by many proponents of advanced nuclear technologies have consequences beyond missed timelines. They ultimately skew policymakers' perceptions on the actualization of safe plant construction and operation, leading to regulatory changes that subvert well-established nuclear safety norms.** As

¹⁴³ Lyman, "Five Things the 'Nuclear Bros' Don't Want You to Know."

¹⁴⁴ Alexander J. Huning et al., "An Introduction to Microreactor Licensing Basis Events," 2023, <https://www.osti.gov/servlets/purl/1993692>.

¹⁴⁵ Shanlai Lu, "Evaluation of NuScale Post ECCS Actuation Boron Dilution Events," Division of Safety Systems, Office of Nuclear Reactor Regulation, U.S. Nuclear Regulatory Commission, June 6, 2020, <https://www.nrc.gov/docs/ML2023/ML20232D086.pdf>.

demonstrated by the impacts of the discussed executive orders, these initiatives place nuclear regulators under undue pressure to meet contrived timeboxing to approve novel reactor designs, not only undermining the role of the regulator, but perpetuating the unsafe acceleration of nuclear construction. The subversion of well-established nuclear safety norms to hasten the deployment of advanced nuclear technologies in service of AI raises serious concerns over whether risks associated with experimental nuclear technologies can be justified given the risks they would pose to civil society. **Moreover, seeking to accelerate the deployment of underdeveloped nuclear technologies may be to the detriment of their development, as the lack of certainty with their safety is likely to sow distrust within public perception, bringing into question the safety of these nuclear technologies and undermining the benefits they may bring once matured.**

Concluding Remarks

Despite the slew of “fast-tracking” efforts underway to accelerate the construction and use of nuclear energy to power AI datacenters, these initiatives are ultimately infeasible in the face of the substantiated requirements needed to safely construct and operate nuclear plants. This discrepancy between the AI industry’s energy demands and infeasibility to rapidly construct nuclear plants has instead led to ideological justifications—whether it be through the weaponization of a purported AI arms race, or unsubstantiated claims regarding the efficacy of AI models and advanced nuclear technologies—that seek to discard well-established nuclear safety principles and wide-ranging international consensus to hasten the deployment of nuclear reactors.

In shifting the risk calculus and tolerances of nuclear infrastructure to be in service of AI, the public’s risk of exposure to ionizing radiation is relegated to unmaterialized benefits of AI, without any of the benefits that civil nuclear capacity is intended to provide. This shift has largely been perpetuated by the pullback of independence of nuclear regulators, allowing the US government to enforce positions that may otherwise have been independently deemed perilous for nuclear safety and security. Furthermore, potential control of sensitive nuclear data by AI companies and the monopolization of nuclear energy to explicitly power AI raises serious concerns over whether their possession of such data and operations may lead to nuclear destabilization and proliferation, thereby further entrenching power asymmetries.

Overall, these “fast-tracking” initiatives raise serious concerns over whether risks associated with the hastened design, construction, and licensing of nuclear facilities can be justified if they exacerbate risks that would lead to an increase in civilian exposure to ionizing radiation, or nuclear destabilization and proliferation for nation-states. If these initiatives continue to be pursued, their lack of safety may lead not only to catastrophic nuclear consequences, but also to an irreversible distrust within public perception of nuclear technologies that may inhibit the support of the nuclear sector as part of our global decarbonization efforts in the future.