The Pathway to Licensing: Selecting and Scaling Safety Analysis in Advanced Reactor Development for Aalo Atomics

Asmaa S. Farag[a,b], Kurt Vedros[a], and Curtis L. Smith[a,c]
[a] Aalo Atomics, Austin, TX, U.S.A., asmaa@aalo.com, kurt@aalo.com, curtis@aalo.com
[b] North Carolina State University, Raleigh, NC, U.S.A, asfarag@ncsu.edu
[c] Massachusetts Institute of Technology, Cambridge, MA, U.S.A, curtis@mit.edu

Abstract: The current U.S. advanced reactor environment increasingly requires developers to navigate multiple regulatory and oversight frameworks while design concepts and facility missions continue to evolve. For emerging reactor programs, the pathway to licensing is rarely linear; it is shaped by project objectives, regulatory authority, design demonstration needs, and the intended function of the facility. This paper presents lessons learned from the licensing strategy development for a sodium-cooled reactor program that evaluated multiple pathways under both the U.S. Department of Energy (DOE) and the U.S. Nuclear Regulatory Commission (NRC). The case study shows that DOE test-authority approaches, NRC research and test reactor approaches, and future NRC commercial licensing each embody different safety philosophies, documentation expectations, and analytical precedents. Accordingly, the depth, scope, and structure of hazard analysis, deterministic safety analysis, and probabilistic risk assessment (PRA) are highly sensitive to three attributes: facility type, mission, and regulatory authority. The paper argues that early separation of near-term demonstration objectives from long-term commercial licensing objectives can reduce rework, that frequent regulator feedback is essential for right-sizing analytical effort, and that PRA is most valuable when facility complexity and licensing decisions require disciplined management of conservatism rather than as a universal early-stage requirement. The experience discussed here contributes to broader efforts toward scalable, risk-informed, and performance-based licensing strategies for advanced reactors.

Keywords: advanced reactors; licensing strategy; DOE authorization; NRC licensing; research reactors; probabilistic risk assessment; sodium reactor

1. INTRODUCTION

Advanced reactor developers in the United States are operating in a policy environment that encourages rapid deployment while still requiring credible, well-documented demonstrations of safety. At the same time, there is no single universal pathway that fits every advanced reactor project. Depending on whether a facility is intended as a research platform, a low-power test, a prototype, a demonstration unit, or a commercial power reactor, the applicable regulatory basis, the expected safety case, and the timing and role of risk information can change materially. The central challenge is therefore not simply how to prepare safety analyses, but how to select a licensing strategy early enough so that those analyses are scaled to the actual mission rather than to an assumed future end state.

For non-light-water reactors, NRC guidance has increasingly emphasized technology-inclusive and risk-informed approaches, including the Licensing Modernization Project (LMP) methodology endorsed in Regulatory Guide (RG) 1.233 [1-3]. In parallel, DOE authorization of nuclear facilities remains grounded in the safety-basis requirements of 10 CFR Part 830 and implementing standards such as DOE-STD-3009-2014 and DOE Order 425.1D [4-6]. These frameworks serve different institutional purposes. The NRC framework is oriented toward the licensing basis and lifecycle operation of civilian reactors under Parts 50, 52, recently 53 and related guidance [7-10], whereas the DOE framework emphasizes documented safety analysis, hazard controls, readiness, and graded application for DOE nuclear facilities [4-6]. Both systems are rigorous, but they do not organize the safety case in the same way.

This paper examines those differences through the experience of Aalo Atomics, a developer of a sodium-cooled advanced small modular reactor. Over roughly two and a half years, the risk assessment program evolved from a commercial multi-module concept, to a DOE-based demonstration pathway, to the implications of NRC-oriented frameworks for future commercialization, and ultimately to a low-power test mission that better aligned with near-term schedule and authorization objectives. The purpose of this paper is not to critique any regulatory body. Instead, it is to document how safety and risk analysis needs adapted as the facility mission changed, how mixed regulatory framing created avoidable inefficiencies, and how clearer separation of near-term and long-term licensing goals improved the coherence of the safety case.

2. REGULATORY AND ANALYTICAL CONTEXT

The U.S. regulatory environment for advanced reactors can be understood as a family of pathways rather than a single track. Commercial reactor licensing under NRC authority has proceeded through established pathways in 10 CFR Parts 50 and 52, while the NRC has also developed a risk-informed, performance-based path through 10 CFR part 53. For non-power reactors, longstanding guidance such as NUREG-1537 and RG 2.6 remains relevant, particularly where the mission is fundamentally experimental or research-oriented rather than commercial electricity generation [10,11].

By contrast, DOE authorization is driven by the safety-basis construct in 10 CFR Part 830. NE-STD-1279-2025 sets requirements for nuclear safety, quality, and radiation control, amongst others. Under this framework, the contractor must establish and maintain a safety basis, identify and analyze hazards associated with the work, derive controls, and demonstrate readiness for startup or restart where applicable [4-6]. DOE-STD-3009-2014 provides the dominant documented safety analysis structure for nonreactor nuclear facilities, with explicit emphasis on hazard analysis, accident analysis, hierarchy of controls, defense-in-depth, and graded application [5]. Even when a facility is reactor-related, the practical authorization workflow can resemble a DOE safety-basis process more closely than an NRC commercial licensing review if the mission is limited, experimental, or demonstrative in character.

After initial hazards analyses, these frameworks ask different initial questions and they also encourage different analytical starting points. In simplified terms, an NRC commercial pathway tends to ask what licensing basis, design criteria, event classes, and programs are required to support the full life of a plant. A DOE authorization pathway tends to ask what work is being performed, what hazards are present, what accidents are credible or bounding for that work, what controls are needed, and whether the organization is ready to operate safely. The distinction is subtle but consequential. An organization that starts from the wrong question may produce technically sophisticated work that nevertheless does not match reviewer expectations.

3. CASE-STUDY OVERVIEW: PROGRAM EVOLUTION AND STRATEGY SELECTION

The original vision considered by Aalo reactor project was a commercial pod concept comprising multiple Aalo-X reactor units. Under that vision, a first-of-a-kind demonstration reactor was seen as a bridge toward eventual commercialization. Early in strategy development, the team assumed that pursuing an initial DOE-based demonstration might provide a more accessible near-term route while still positioning the program for an eventual NRC commercial license. That assumption was not unreasonable. DOE pathways can be attractive to companies because they may appear more compatible with rapid iteration, test-focused missions, and staged technology maturation.

The difficulty arose when near-term DOE objectives and long-term NRC commercial objectives were treated as if they could be satisfied through a largely common analytical package. In practice, the team often tried to frame DOE submittals using concepts and expectations associated with future NRC licensing, including risk-informed constructs inspired by LMP. The analysis intention was constructive: preserve continuity with the ultimate commercial safety case and avoid redoing work later. However, the result was often friction at the expense of efficiency. Documentation became more elaborate than was necessary for the DOE decision at hand, while still not organized completely in the way a future

NRC application would require. The team was therefore exposed to the cost of both approaches without gaining the full benefit of either.

As the program matured, an important strategic reframing occurred. Instead of using the "first facility" as a partial surrogate for the full commercial reactor, the near-term mission was narrowed to a low-power test facility intended primarily to generate institutional learning on fuel and core behavior and to establish experience with the DOE authorization process. This change did not reduce the importance of safety. Rather, it made the safety case more coherent by aligning it with a more specific mission. The simplified mission also sharpened decisions about what analyses were genuinely needed in the short-term and what analyses properly belonged to future commercial licensing phases.

Concurrently, the team still maintained communication with the NRC so that future commercialization discussions would not start from zero. In hindsight, this parallel engagement was valuable when used for familiarization and future planning, but it was less effective at first, when DOE and NRC expectations were attempted as one near-term licensing package. The lesson is that parallel regulatory engagement can be useful, yet the actual authorization basis for a given facility should remain singular and explicit. **Table *1*** illustrates the different pathway characteristics encountered in strategy development.

Table 1: Illustrative comparison of pathway characteristics as encountered in strategy development.

Pathway framing	**Primary near-term objective**	**Dominant analysis emphasis**	**Observed program implication**
DOE demonstration framed through future NRC concepts	Advance toward early deployment while preserving commercial continuity	Hybridized package mixing hazard analysis, deterministic studies, and NRC-oriented risk-informed constructs	High analytical burden and ambiguous reviewer expectations
DOE authorization aligned to DOE safety-basis precedents	Authorize a test-oriented facility under a clearly defined mission	Hazard categorization, DSA structure, accident analysis, controls, readiness	More coherent documentation and better alignment with regulator feedback
Future NRC commercial licensing planning	Prepare for long-term power-reactor commercialization	Licensing basis development, design criteria, PRA-informed event selection and SSC classification	Best treated as a separate but informed downstream effort
Potential NRC research/test reactor route	Directly license a non-power or experimental facility under NRC authority	Non-power reactor application and emergency planning guidance	May avoid dual-framework translation, but depends on mission and leadership strategy

4. WHY MIXING FRAMEWORKS BECAME INEFFICIENT

One of the clearest lessons from the program was that mixing DOE and NRC constructs too early can be expensive for a startup developer. The cost is not only financial. It appears in schedule expansion, analyst rework, documentation churn, shifting internal expectations, and difficulty prioritizing what must be technically mature at each phase. When a small organization tries to completely satisfy two partially overlapping but structurally different regulatory logics at once, it can create a false sense of progress because the volume of analysis increases while decision clarity decreases.

The underlying issue is that the two frameworks use different organizing principles. 10 CFR 53, LMP, and related NRC guidance are well suited to shaping a commercial licensing basis by linking licensing basis events, safety functions, and classification of structures, systems, and components to a broader risk-informed architecture [1-3]. DOE safety-basis processes, in contrast, expect a disciplined progression from hazard identification to accident analysis to control selection and readiness under a graded approach [4-6]. These are compatible safety philosophies at a high level, but they are not interchangeable document templates. Attempting to use one as a surrogate for the other can lead to analyses that are individually sound yet collectively inefficient.

This experience suggests that developers should distinguish between analytical reuse and analytical translation. Reuse is desirable: inventories, phenomenology, source-term assumptions, thermal-hydraulic insights, and reliability information generated for one pathway can often inform another. Translation is more difficult: a document or analytical structure written for one regulatory decision does not automatically serve the next decision without significant reframing. The most robust strategy is therefore to separate the pathway-specific safety case from the underlying technical knowledge base. In other words, preserve the science and engineering, but tailor the licensing narrative to the regulator and mission being used.

5. MISSION DEFINITION AS THE DRIVER OF ANALYSIS DEPTH

A second major lesson was that facility mission is the single strongest determinant of how much analysis is necessary and what kind of analysis provides decision value. At the beginning of the program, the “first facility” was often treated as a precursor to a commercial reactor and therefore burdened with expectations associated with future power production, long-term operability, and commercial deployment. Once the mission was narrowed to a low-power test and demonstration role, several aspects of the safety case became more stable. Hazard categorization changed, the bounding accident set became easier to define, the operational states of interest became more limited, and the documentation focused on the actual work being authorized rather than on a hypothetical future plant.

This point is especially important for advanced reactor startups, because the same core technology can support multiple legitimate missions across its maturation pathway. A test reactor intended to generate physics data, confirm fuel behavior, or exercise organizational readiness is not analytically identical to a prototype intended to demonstrate integrated system performance, and neither is it analytically identical to a commercial plant expected to deliver power at scale. Conflating these missions tends to import requirements prematurely. Conversely, over-simplifying a facility beyond its real mission can undercut credibility. The correct approach is not minimalism but fit-for-purpose analysis.

Within the Aalo case, mission definition also influenced whether deterministic methods alone were adequate for a given decision. For the simplest, lowest-power mission, deterministic hazard analysis and bounding accident evaluation supplied much of the near-term regulatory value. As the mission moved closer to commercial relevance, however, the importance of PRA increased because the design space contained more interacting systems, more operational modes, and more decisions where a purely bounding deterministic posture would have imposed substantial penalties on design practicality or commercial viability.

6. THE ROLE OF PRA

The experience described here does not support the simplistic conclusion that PRA is either always essential or always unnecessary. Rather, its value depends on where the project sits on the mission-and-complexity ladder. For a narrow test objective under DOE authorization, deterministic analysis dominates because the principal licensing decision concerns whether identified hazards are understood, controlled, and ready for operation. In this setting, a full-detail plant-level PRA may offer less incremental value than disciplined hazard screening, well-justified bounding assumptions, and a strong readiness process.

That observation should not be misread as anti-PRA. In fact, the opposite lesson emerged as the program considered more integrated and commercially relevant reactor concepts moving forward. As complexity grows, the number of potentially interacting initiating events, system dependencies, operator actions, and defense-in-depth questions also grows. A purely deterministic approach risks becoming either excessively conservative or internally inconsistent. PRA is particularly powerful under those conditions because it provides a systematic method to prioritize contributors, evaluate the risk significance of design choices, and avoid spending the same analytical effort on events or systems that do not deserve equal treatment [1-3,12].

Accordingly, the most defensible position is that PRA should be introduced where it improves decision quality, not merely where it signals analytical sophistication. Developers benefit when they identify the threshold at which risk information changes a licensing or design decision. Below that threshold, deterministic approaches may be sufficient. Above it, PRA becomes essential not only for licensing but also for internal capital allocation, design optimization, and the long-term maintainability of the safety case.

7. IMPORTANCE OF FREQUENT REGULATOR INTERACTION AND STAGED DOCUMENTATION

Another strong lesson was the benefit of frequent engagement and staged document submission. In the DOE setting, iterative review of conceptual, preliminary, and more mature documents substantially improved the team's understanding of what reviewers were actually looking for. This mattered because even a technically strong engineering team can overshoot or undershoot regulator expectations when it lacks pathway-specific experience. In several instances, the team had prepared analyses at a level of detail that were either more of what was needed or not yet supported by the maturity of the facility concept. In other instances, the regulator's focus lay in areas that the team had initially underestimated.

The broader lesson is that frequent feedback is not merely an administrative convenience; it is a mechanism for right-sizing the safety case. Especially for emerging programs, regular interaction helps narrow the gap between internal assumptions and external review criteria, thereby improving schedule predictability.

For future applicants, the practical implication is that conceptual, preliminary, and final safety documentation should not be treated as redundant versions of the same report. Each stage has a distinct function. Conceptual documents should align mission, scope, and bounding hazards. Preliminary documents should test the emerging control strategy and identify where additional design maturity is required. Final documents should demonstrate internal consistency, readiness, and implementation. When used well, this staged process is one of the most efficient ways to reduce later rework.

8. DESIGN EVOLUTION AND SUPPLY-CHAIN REALISM

Licensing strategy did not evolve in isolation from engineering realities. For example, one of the most consequential programmatic changes was a shift in fuel selection. The early concept had favored uranium-zirconium-hydride-based fuel because of its attractive safety characteristics and reactor-physics benefits in the original design framing. Over time, however, the project moved toward uranium dioxide fuel to better align with achievable supply-chain conditions, available fabrication pathways, and schedule constraints. This was not a minor substitution; it had implications for phenomenology, available historical data, licensing references, and the maturity of the supporting technical basis.

The key lesson is that licensing strategy must reflect what can be built and supported within the program timeline. A safety case anchored to an elegant but impractical design basis can become unstable if procurement realities force late design changes. Conversely, a design change that improves manufacturability and data availability may strengthen the overall licensing position even if it requires revisiting portions of the safety analysis. The choice of fuel illustrates a broader principle: regulatory strategy, design maturity, and supply-chain credibility are coupled variables. When one changes materially, the others should be reevaluated explicitly rather than assumed to remain valid.

9. ALTERNATIVE PATHWAY NOT FULLY EXPLORED: NRC RESEARCH/TEST REACTOR ROUTE

In retrospect, one pathway that may have warranted deeper early evaluation was a direct NRC research or test reactor route for an appropriately defined facility. The purpose of raising this point is not to argue that such a path would necessarily have been superior; rather, it is to note that it might have reduced the burden of translating between DOE and NRC expectations if the initial mission truly fit a non-power reactor posture. NRC guidance for non-power reactors is mature and well established through NUREG-1537 and RG 2.6 [10,11]. For some advanced reactor developers, these pathways may provide a clearer alignment between regulator and mission than a DOE pathway that is simultaneously being used as a stepping stone toward later commercial licensing.

Whether that option is attractive in practice depends on leadership experience, project timeline, site conditions, investor expectations, and the intended use of the facility. In the Aalo case, leadership perspectives and prior institutional experience shaped which paths were emphasized. That is common across startups. The lesson is not that one leadership view was right or wrong, but that teams should explicitly test pathway assumptions against the actual mission and against in-house regulatory experience. A path that appears familiar because senior personnel have used part of it before may still become difficult if the organization lacks experience translating between multiple systems at once.

10. POLICY ACCELERATION AND THE VALUE OF A NARROWER FIRST MISSION

Recent federal policy developments increased the perceived urgency of reaching first criticality for advanced reactor demonstrations. In 2025, the White House issued executive actions aimed at accelerating reactor testing and DOE subsequently established a Reactor Pilot Program intended to expedite advanced reactor testing outside the national laboratories [13,14]. These developments created strong incentives for developers to identify pathways capable of supporting rapid execution. Ambitious policy targets can compress timelines, attract private capital, and expand the workforce base for nuclear innovation, but they also place a premium on mission discipline.

Within that environment, narrowing the first facility's purpose was strategically beneficial. A low-power test mission focused on discrete learning objectives was more plausibly aligned with an accelerated schedule than a near-term attempt to demonstrate the full envelope of a commercial reactor. The resulting facility shifted from the original commercial product. Its value lay in creating an executable first step: establishing the safety-basis process, demonstrating organizational safety and operational readiness, manufacturing readiness, advancing fuel and core understanding, and building institutional experience that would later support commercial deployment.

This distinction matters because rapid criticality and commercial readiness are not identical objectives. A demonstration optimized for speed can still provide substantial long-term value if the organization is disciplined about what conclusions may legitimately be carried forward into the commercial safety case. The risk arises only when the early test mission is oversold as a full proxy for a later plant that differs materially in power, operating envelope, or deployment basis.

11. RECOMMENDATIONS FOR STARTUP GOVERNANCE OF LICENSING STRATEGY

From a governance perspective, the case suggests that licensing strategy should be revisited at formal decision points rather than only when external pressure forces a change. At minimum, those decision points should include major mission changes, power-level changes, fuel or core configuration changes, site changes, and transitions from test objectives to demonstration or commercial objectives. Each decision point should ask whether the current hazard basis, accident set, regulator interface plan, and use of PRA remain appropriately scaled.

A practical governance mechanism is a pathway review board or equivalent function that includes design, safety analysis, licensing, operations, and program management. Its purpose is not bureaucratic oversight for its own sake, but preservation of alignment. Such a function can make explicit which analyses are being performed for immediate regulatory value, which are being performed to retire technical uncertainty, and which are being deferred because they do not yet improve a near-term decision. For advanced reactor startups, that discipline may be one of the few reliable defenses against strategic over analysis.

12. DISCUSSION

The broader significance of this case is that advanced reactor licensing strategy should be approached as a scaling problem. The question is not simply which regulator is involved, but how much safety-analysis architecture is warranted for the mission, maturity, and hazard profile of the facility being proposed. A mature, risk-informed commercial licensing basis built around risk-informed principles is entirely appropriate for a commercial plant. A DOE-style documented safety analysis is entirely appropriate for a DOE-authorized nuclear facility. Problems arise when a project implicitly assumes that the same near-term documentation package can serve both purposes without substantial translation.

At a technical level, the findings support a graded philosophy. Hazard analysis, deterministic accident evaluation, control selection, and readiness are foundational at all stages. PRA enters most strongly where system interactions, lifecycle operability, and the need to calibrate conservatism appropriately make explicit risk ranking valuable. This is consistent with the general direction of advanced-reactor regulation, but it also cautions against using risk-informed language as a substitute for pathway-specific discipline.

At a programmatic level, the case highlights the importance of strategy decisions that are sometimes treated as secondary to engineering. Mission definition, fuel/design choice, site and regulator interface, and documentation staging are not peripheral management concerns; they shape the content and usefulness of the safety case itself. For startups in particular, the opportunity cost of getting these strategic choices wrong can be very high because teams rarely have the staffing margin to absorb repeated analytical restarts.

13. CONCLUSION

The Aalo experience indicates that the pathway to licensing for advanced reactors should be selected and scoped with the same discipline applied to core neutronics, thermal-hydraulics, or materials selection. The principal conclusions are as follows. First, the required depth and structure of safety analysis are highly sensitive to facility mission, regulatory authority, and intended end use. Second, attempting to satisfy DOE authorization needs and complete future NRC commercial licensing expectations through a single early-stage analytical framework can create inefficiency. Third, frequent regulator interaction and staged documentation materially improve the ability to right-size analysis and reduce rework. Fourth, PRA provides its greatest value as system complexity and commercial relevance increase. Fifth, major design choices such as fuel selection should be evaluated not only for engineering performance but also for data pedigree, and licensing stability.

More broadly, this case supports a scalable and performance-based view of advanced reactor development. A strong licensing strategy does not attempt to prove everything at once. It aligns each facility with a clear mission, applies the appropriate safety framework for that mission, preserves technical knowledge for later reuse, and avoids burdening early demonstrations with analytical commitments that properly belong to later commercial phases. Such an approach can improve both technical rigor and execution speed while maintaining respect for the distinct roles of DOE and the NRC in the evolving U.S. advanced reactor landscape.

REFERENCES

[1] U.S. Nuclear Regulatory Commission, Regulatory Guide 1.233, Guidance for a Technology-Inclusive, Risk-Informed, and Performance-Based Methodology to Inform the Licensing Basis and Content of Applications for Licenses, Certifications, and Approvals for Non-Light-Water Reactors, 2020.

[2] Nuclear Energy Institute, NEI 18-04, Rev. 1, Risk-Informed Performance-Based Guidance for Non-Light Water Reactor Licensing Basis Development, 2019.

[3] U.S. Nuclear Regulatory Commission, Advanced Reactor Application Guidance / Licensing Modernization Project guidance webpage, accessed 2026.

[4] 10 CFR Part 830, Nuclear Safety Management, U.S. Department of Energy.

[5] U.S. Department of Energy, DOE-STD-3009-2014, Preparation of Nonreactor Nuclear Facility Documented Safety Analysis, 2014.

[6] U.S. Department of Energy, DOE Order 425.1D, Verification of Readiness to Start Up or Restart Nuclear Facilities, 2010.

[7] 10 CFR Part 50, Domestic Licensing of Production and Utilization Facilities, U.S. Nuclear Regulatory Commission.

[8] 10 CFR Part 52, Licenses, Certifications, and Approvals for Nuclear Power Plants, U.S. Nuclear Regulatory Commission.

[9] U.S. Nuclear Regulatory Commission, Part 53 – Risk-Informed, Technology-Inclusive Regulatory Framework for Commercial Nuclear Plants, accessed 2026.

[10] U.S. Nuclear Regulatory Commission, NUREG-1537, Guidelines for Preparing and Reviewing Applications for the Licensing of Non-Power Reactors, Parts 1 and 2, 1996 and later updates.

[11] U.S. Nuclear Regulatory Commission, Regulatory Guide 2.6, Emergency Planning for Research and Test Reactors and Other Non-Power Production and Utilization Facilities, Revision 2, 2017.

[12] ASME/ANS RA-S-1.4-2021, Probabilistic Risk Assessment Standard for Advanced Non-Light Water Reactor Nuclear Power Plants, 2021.

[13] The White House, Reforming Nuclear Reactor Testing at the Department of Energy, Executive Order, May 23, 2025,Webpage.

[14] U.S. Department of Energy, U.S. Department of Energy Reactor Pilot Program, program description webpage, accessed 2026.