Transactions of the Korean Nuclear Society Autumn Meeting
Gyeongju, Korea, October 29-30, 2026

Assessing Licensing Management Maturity in Nuclear Power Plant Projects: A Preliminary Study

Nguyen Hoang Anh[a*], Jung Woo-yong[a]
[a]*KEPCO International Nuclear Graduate School, 658-91 Haemaji-ro, Seosaeng-myeon, Ulju-gun, Ulsan 45014, Republic of Korea*
[*]*Corresponding author: hoanganh.narime@gmail.com*

***Keywords :** nuclear licensing, maturity model, regulatory coordination, licensing management, design science

1. Introduction

Nuclear power plant projects are enormous megaprojects with long schedules and very high capital costs. In the detailed preparation work before construction, regulatory licensing is a difficult factor that has a foundational impact on the entire project. When the licensing process goes wrong, the consequences are counted in years and billions of dollars. Practice from the most recent first-of-a-kind (FOAK) projects in Europe shows that regulatory licensing can deeply affect project effectiveness. This disparity leads to a situation where academic research points out problems in licensing but lacks tools to help handle those problems. Assessment of organizational capability is the core predictor allowing owners and contractors to map out the optimal preparation path for handling licensing issues.

Licensing and licensing management are two completely different concepts. Licensing is the work of the regulatory body: they review the submission and decide to grant or deny the license [1]. Licensing management, on the other hand, is the work of the project developer: it includes the organizational processes, institutional mechanisms, communication methods, documentation systems, and risk management methods aimed at obtaining and maintaining approvals from the regulatory body. These two concepts are often equated with each other, thereby causing an important confusion. When a project is delayed during the process of working with the regulatory body, people often assume the cause is that the regulatory body is too strict, whereas in many cases it is because the developer has not organized itself methodically enough to meet requirements that had already been published in advance.

What the field lacks is a systematic, validated way to assess licensing management. Two established frameworks come close in spirit but neither occupies the space. The Capability Maturity Model Integration (CMMI) offers a proven structure for grading organizational process capability. It contains no nuclear-specific content [2]. The IAEA Milestones approach gives newcomer countries a staged model of national readiness [3]. But it works at the national and program level. Licensing management assessment needs a project-level, process-specific instrument. General project-management maturity models apply to any large project. None of them addresses the regulatory interface that defines nuclear licensing.

The practical consequences are twofold. A developer entering an unfamiliar regulatory system has no structured way to benchmark its own readiness. An organization seeking to target technical assistance cannot easily separate a program's licensing strengths from its weaknesses.

This paper reports the preliminary stage of a study that addresses this gap. It presents the Licensing Management Maturity Model (LMMM), a five-level framework with six domains and eighteen indicators, constructed as a design-science artifact. The results given here come from the construction and instrument-design stage. The comparative case application is still in progress, and its scoring is not reported. What is reported is the derivation of the model, its structure, the rules governing how its descriptors are written, and the protocol for reviewing and applying it. The paper is therefore best read as a statement of an instrument and its warrant, offered for comment before the case results are fixed.

2. Methods

2.1 Design-science construction

Accordingly, the LMMM is constructed on the design-science procedure for maturity model development proposed by Becker et al. [4]. Each step of the procedure is mapped against the activities actually carried out, and the location where the supporting evidence is stored is recorded. The aim is to ensure that the model's domains and indicators are the traceable output of a documented process, rather than a list compiled at the authors' discretion. Through this procedure, the LMMM is designed to maximize content validity – that is, to measure accurately the concept it is intended to capture.

2.2 Derivation of the indicator set

The six domains follow from a theoretical synthesis of three bodies of work: the nuclear regulatory literature, the maturity-model literature, and the theory of regulatory coordination. Each domain corresponds to a mechanism through which licensing management affects project outcomes – reducing information asymmetry, lowering transaction costs, sustaining institutional legitimacy, managing regulatory uncertainty, providing organizational capacity, or enabling learning.

The indicators are derived differently, and the derivation runs from documented cause to indicator. From the empirical literature on licensing-related delay and cost escalation, nine recurring drivers were isolated, each attested in more than one source: requirements changing during review; insufficient or late pre-application engagement; documentation deficiency and submission rework; unresolved design-regulatory interfaces on novel designs; review duration underestimated in the schedule; contractor non-conformance with licensing requirements; insufficient capacity at the regulatory interface; consent and permitting delays; and failure to learn across projects and reviews. Every indicator is required to address at least one driver, and the rule applied during construction was strict: an indicator that could not be traced to a documented driver was removed.

Applying that rule produced a candidate set larger than the study could assess responsibly. Because every indicator must be scored for every case, and every score requires corroborated public evidence, indicator count translates directly into evidentiary workload. A set large enough to exhaust the available evidence would produce a model that is comprehensive on paper and ineffective in practice. The set was therefore reduced to eighteen indicators, three per domain, under two criteria applied in order. The retained set must collectively cover all nine drivers, so that no documented cause of delay is left unaddressed; subject to that, preference was given to indicators that are distinct rather than sub-cases of another retained indicator, and to indicators more likely to be evidenced from public sources.

2.3 Maturity level structure

The five-level structure is adapted from CMMI [2], retaining its developmental logic: practices must be consistently performed before they can be standardized, and standardized before measurement and systematic improvement become meaningful. In the licensing context, Level 1 describes activity that is ad hoc and dependent on individuals; Level 2, activity performed consistently but reactively, in response to regulatory prompting rather than in anticipation of it; Level 3, activity governed by defined, documented procedures applied across the organization; Level 4 adds quantitative management, in which the process is measured and the measurements inform decisions; and Level 5 adds systematic improvement driven by those measurements.

Two consequences of adopting this structure are acknowledged rather than assumed away. The levels are ordinal: a score of 4 is higher than a score of 2, but the interval between them is undefined, and domain scores are reported as unweighted means for descriptive purposes only, never entered into parametric analysis. And the developmental logic is inherited from CMMI rather than independently demonstrated for this domain; whether licensing management capability in fact develops in this order is an empirical question the present study does not answer.

2.4 Instrument review and case protocol

The indicator set is reviewed by an informed panel of graduate students and recent graduates from a Korean nuclear institution. Panel members hold engineering or nuclear-related degrees. Few or none have worked in nuclear licensing, and the design does not pretend otherwise.

This limits what the panel can be asked to do. In the measurement literature, content validity depends on judgments from people with substantive expertise in the construct being measured. A panel without licensing experience cannot supply that. Collecting relevance ratings from such a panel and reporting the resulting index as content validity would give the number a validating appearance it has not earned. The design therefore assigns the panel a narrower task. Whether an indicator belongs in the model rests on its documented derivation from established causes of licensing delay — that is, on the literature rather than on the panel. Whether an indicator can be applied consistently is something the panel is well placed to test. Interpretive consistency is accordingly the main quantitative output of this strand. It is measured through six calibration vignettes written to sit at specified levels, four of them placed at the Level 2 / Level 3 boundary, where the descriptors are hardest to separate. Relevance ratings are also collected, but only as a secondary and indicative check. They are reported throughout as informed-panel relevance indices, never as content validity indices, and each is accompanied by a chance-corrected statistic.

Two limits should be stated plainly. A documented derivation makes the selection of indicators traceable; it does not establish that the selection is correct. And agreement among readers shows that the descriptors are clear, not that the levels they describe are the right ones. Review by experienced licensing practitioners remains the necessary next step, and the iterative structure of the design-science procedure adopted here anticipates it.

The model is then applied to four NPP projects, selected purposively to maximize variation in regulator maturity, developer experience, design novelty, and national policy context. Two rules govern the scoring.

First, evidence about a regulator's maturity is not evidence about a developer's maturity. An international peer review of a regulatory body describes that body, not the applicant whose licensing management this model assesses. Regulator-level evidence therefore characterizes context and never contributes to a developer's score.

Second, where the available evidence will not support a higher rating, the lowest level the evidence can bear is assigned. Where the evidence cannot distinguish between levels at all, the indicator is marked Not Determinable rather than assigned a level by inference.

Cross-case analysis proceeds by pattern matching and cross-case synthesis [5]. No inferential test linking maturity scores to project outcomes is attempted; four cases could not support one.

3. Preliminary Results

3.1 Structure of the model

The model comprises six domains, eighteen indicators, and ninety maturity descriptors. Table I summarizes the domains and what each assesses.

Table I: The six domains of the LMMM

Code	Domain and what it assesses
RIM	Regulatory Interface Management – the processes through which the developer manages its relationship with the regulatory body.
DKM	Documentation and Knowledge Management – the systems through which licensing information is produced, controlled, retrieved, and retained.
SPE	Stakeholder and Public Engagement – engagement with parties other than the safety regulator, including host communities and non-nuclear permitting authorities.
SRM	Schedule and Risk Management – the integration of licensing into project planning and risk management.
OCG	Organizational Capacity and Governance – whether the developer has the structure, competence, and governance to sustain licensing management.
ALCI	Adaptive Learning and Continuous Improvement – whether the organization learns from its own licensing experience and that of others.

3.2 How the descriptors are written

Applying the model consistently depends on three conventions, which are held constant across all eighteen indicators. The assessor assigns the highest level for which every element of the descriptor is satisfied; levels are cumulative, and partial satisfaction of a higher descriptor is recorded in the basis-of-assignment field rather than reflected in the score. The discriminator between Levels 2 and 3 is documentation and organizational scope: Level 2 describes a practice performed consistently by those who happen to perform it, Level 3 the same practice governed by a documented procedure applying across the organization. This is the boundary at which assessors most often diverge, and it is the boundary the calibration vignettes chiefly test. Levels 4 and 5 require evidence of measurement and of change respectively: Level 4 requires that defined indicators are collected and reported, not merely that data exists somewhere, and Level 5 requires evidence that the process was changed in response, traceable to the measurement that prompted it. A statement of intention to improve does not satisfy Level 5.

3.3 An illustrative indicator

Table II reproduces the descriptors for RIM-1, the pre-application engagement process, which addresses the delay driver of technical positions left unsettled until formal submission. It is offered as an illustration of the descriptor style: each level states an observable condition, and the documentary evidence an assessor would typically examine – regulator meeting minutes and correspondence logs, published engagement plans, regulator statements on engagement quality – is recorded with the indicator.

Table II: RIM-1 pre-application engagement, maturity descriptors

Level	Observable condition
1	Contact with the regulator occurs only when a submission is made; technical positions are first tested during review.
2	Pre-application meetings take place but are arranged case by case on individual initiative, with no standing agenda and informal or absent records.
3	A documented engagement plan identifies the topics requiring early alignment, the participants, and the meeting cadence. Outcomes are minuted and positions recorded in referenceable form.
4	Engagement is planned against the submission schedule and its effect is measured, for example the proportion of pre-agreed positions subsequently challenged in formal review.
5	Those measurements change the engagement plan itself: topics that repeatedly generate late challenge are moved earlier or handled differently, with the revision documented.

3.4 An anticipated finding about evidence

One consequence of the descriptor design can be anticipated before the case scoring is complete. Across all six domains, the Level 4 and Level 5 descriptors turn on measurement and improvement records that developer organizations rarely publish. A substantial proportion of the seventy-two indicator-by-case cells is therefore expected to be recorded as Not Determinable at those levels. This is reported as a finding about the limits of public-source assessment rather than concealed by inference.

4. Discussion

The principal contribution of this stage is the development, refinement, and operationalization of a maturity assessment model in a field where such models have not previously seen wide application. The broader claim is that maturity models — well established in software engineering and in project management generally — can be applied effectively to the management of the nuclear regulatory interface. The condition for doing so is that the assessment indicators must rest on observable documentary evidence rather than on organizational self-reporting. It is this discipline in the use of evidence that allows the model to support comparison across cases in a field where access to primary data is tightly restricted.

Three limitations are stated plainly, because no amount of procedural care can remove them. First, content validity remains untested: the review panel, though technically knowledgeable, does not include specialists working directly in licensing. Full validation with a panel of practicing licensing specialists is identified as a task for a later stage. Second, construct validity and predictive validity remain untested: this stage of the research does not address whether LMMM scores correspond to any independently measurable organizational property, or whether they predict project outcomes. Third, the developmental logic of the level structure is inherited from other models and has not been demonstrated specifically for the domain of licensing management.

For practitioners, the immediate application of the model is diagnostic rather than comparative. A project developer can apply the eighteen indicators to its own current processes, identify the domains sitting at Level 1 or Level 2, and prioritize the documented procedures needed to move to Level 3 — the level at which, on the evidence assembled here, most of the common causes of delay begin to be addressed systematically. A capacity-building program can use the structure to organize an organizational gap analysis, and an international organization can use it to target technical assistance more precisely than general statements of readiness allow.

The research is also timely. The expansion of nuclear power into newcomer countries is generating demand for structured tools to assess regulatory and licensing readiness. At the same time, the emergence of licensing processes for small modular reactors (SMRs) raises new questions about how licensing management maturity must develop when both the design and the review process are unfamiliar. The present study focuses on conventional large reactors, but the model has been deliberately designed at a level of generality broad enough to support later application in the SMR context.

5. Conclusions

A preliminary maturity model for nuclear licensing management has been constructed as a design-science artifact. The model comprises six domains, eighteen indicators traceable to nine documented causes of licensing delay, and five cumulative levels adapted from CMMI. The descriptor rules, panel review protocol, and evidence rules for case scoring were all fixed before implementation began.

Work in progress comprises the panel's assessment of interpretive consistency and the comparative scoring of four NPP projects against publicly available evidence. Results will be reported with agreement statistics for each domain.

Two extensions are anticipated beyond the present scope. The first is a validation round with practitioners working directly in licensing. The second is application of the model to a newcomer program, where diagnostic rather than comparative use of the instrument would deliver the greatest practical value, and where Not Determinable entries — a common outcome of assessment based on public sources — could be replaced by direct evidence from the organization itself.

ACKNOWLEDGMENT

This research was supported by the 2026 Research Fund of the KEPCO International Nuclear Graduate School (KINGS), the Republic of Korea.

REFERENCES

[1] International Atomic Energy Agency, Governmental, Legal and Regulatory Framework for Safety, IAEA Safety Standards Series No. GSR Part 1 (Rev. 1), IAEA, Vienna, 2016.

[2] CMMI Product Team, CMMI for Development, Version 1.3, CMU/SEI-2010-TR-033, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2010.

[3] International Atomic Energy Agency, Milestones in the Development of a National Infrastructure for Nuclear Power, IAEA Nuclear Energy Series No. NG-G-3.1 (Rev. 1), IAEA, Vienna, 2015.

[4] J. Becker, R. Knackstedt, J. Pöppelbuß, Developing maturity models for IT management, Bus. Inf. Syst. Eng. 1 (2009) 213–222.

[5] R.K. Yin, Case Study Research and Applications: Design and Methods, 6th ed., SAGE Publications, Thousand Oaks, CA, 2018.