

Verification and Documentation of Machine- Assisted Work in Architectural Practice

An author-developed voluntary practice framework for licensed architects and architectural firms

DANIEL HUTCHINS, AIA, NCARB

First Public Edition · Version 1.0 · September 4, 2026

Status and Reading Note. *This Guide is an author-developed voluntary professional-practice framework. Part I presents the condition it addresses. Section 0, located after Part VI, states the governing limitations, source taxonomy, and meaning of shall. Publication alone creates no legal obligation and does not define the standard of care.*

Part I The Condition Observed

1.1 Existing Conditions

This document begins the way a forensic report begins: with existing conditions, stated without characterization.

Artificial intelligence is now present in the routine production of architectural instruments of service. Large language models draft specification sections from the firm's master specifications template. They summarize building codes and return compliance conclusions. They screen submittals and draft the architect's review responses. They produce responses to requests for information, code studies, design narratives, and the construction-phase correspondence on which later disputes turn. It is present in firms of every size, and the speed and unanimity of the profession's institutional response, described below, is itself a measure of how widespread it has become. It arrived from the bottom up, practitioner by practitioner, task by task, under ordinary deadline pressure. And it arrived faster than any deliberative body could be expected to respond to it.

The profession's institutions have responded through materials of different kinds and legal effect. AIA has issued professional guidance and a firm toolkit; NCARB has published an architectural regulatory position and model provisions; ICOR has issued cross-profession regulatory principles; and the NSPE Board of Ethical Review has issued an advisory engineering ethics opinion. These materials are not equivalent sources of law, but they converge on a practical theme: AI does not displace professional judgment or responsibility, and machine-assisted output requires qualified human review appropriate to the work. They do not establish one nationally binding verification method for architects.

The cited national materials leave important operational questions open. They do not provide a common definition of verification, a uniform test criterion, or one record specification for architectural practice.

This Guide addresses that gap by proposing a method that firms may evaluate and adapt to their jurisdictions, contracts, services, and risk profiles.

From that condition, the thesis of this Guide:

Artificial intelligence has not altered the architect's standard of care—it has created new ways to breach it, at a speed that has outrun the profession's capacity to verify what it seals. That differential between production and verification, unmeasured and undocumented in much of practice today, is the exposure this Guide exists to close.

1.2 The Inspection Before the Failure

Every forensic engagement begins the same way: with documents that looked complete on the day they were sealed. My practice begins after something has gone wrong. I work in the record: reconstructing how a deficiency entered the instruments of service, who had the duty to catch it, and what a reasonably prudent professional would have done differently. Failures follow patterns, and the patterns are visible in the documents long before the claim arrives.

This document applies that discipline in the opposite direction. It is an inspection conducted before the failure rather than after: it examines today's ordinary work product through the eyes that may examine it in a future deposition. Part I establishes what has not changed (the duty), what has (the production), and where the two now diverge. Part II states the proposed principles; Part III converts them into controls; Part IV specifies the record; Part V addresses the limited significance of adoption and documentation when practice is examined; and Part VI addresses implementation. Section 0 then states the Guide's governing limitations and definitions before the summary and appendices.

1.3 What the Seal Attests

The measure of an architect's conduct is the standard of care: the skill and diligence that a reasonably prudent architect would exercise under similar circumstances. It is a technology-neutral measure. It does not ask what instrument produced a drawing; it asks whether the professional responsible for it exercised the judgment the public is entitled to rely upon. The seal is that reliance made physical: a licensed individual's personal attestation of responsible control over the instruments of service, given under the authority of a license the state can revoke.

Responsible control is defined differently across licensing jurisdictions, but the variations do not erase its practical core. The architect must exercise genuine professional authority over the work and possess the knowledge and ability needed to direct, evaluate, coordinate, and take responsibility for it. This Guide calls the corresponding practice standard reasonable control. Reasonable control does not require the architect to perform every task personally or understand every internal operation of every tool. It requires control proportionate to the tool's opacity and the consequences of error: the architect must be able to interrogate the process, test the output, recognize its limits, and withhold approval when confidence is not justified. Whether after-the-fact review is legally sufficient depends on the jurisdiction and the circumstances. Nominal supervision and uncritical adoption, however, do not supply the professional judgment that responsible control demands.

1.4 The Anatomy of a Failure

Take a familiar example. A code study is bound into a sealed permit set. It cites a provision of the building code for the required fire-resistance rating of the structural frame. The provision as cited does not exist in the adopted edition: the LLM summarized a blend of the base code and another jurisdiction's amendment, and stated it with perfect confidence. The study was accepted under deadline. No protocol required anyone to trace the citation to the adopted edition, so no one did. The professional who sealed the set never saw the provision, only the conclusion. The set was permitted; the error is now load-bearing.

Many months later that citation is found: by a plans examiner on a revision, by a contractor pricing the correction, or by an expert retained after the dispute has already begun. The forensic reconstruction that follows is the one my profession performs every day: how the deficiency entered, what review occurred, and whether the conduct met the standard of care. Here, the project evidence answers the second question: no one traced the citation to the adopted edition, and the architect who sealed the set saw only the conclusion. The absence of a required process did not itself establish the failure; it left a material verification step unassigned, and the facts show that it was not performed.

The same anatomy repeats across the production of practice. A machine-drafted specification section cites a withdrawn standard or a discontinued product, and is incorporated without section-level review. A submittal review is generated by an LLM and issued under a reviewer's stamp that no qualified professional actually applied. A piece of construction-phase correspondence is drafted in seconds, sent in minutes, and becomes the written record on which the dispute turns. None of these is a novel theory of liability. Each is the same failure the profession has always been answerable for: work released without the exercise of reasonable professional care. What is new is the rate. The instrument that produces the work now outruns the process built to check it, and a failure that once occurred at the speed of human error now occurs at the speed of machine-assisted human output.

1.5 The Questions

Now apply the inspection to your own practice. For the last permit set your firm sealed: which portions were machine-assisted? Who verified them, and against what authority: the adopted code edition, the current standard, the actual product data? Could you produce that record if a plans examiner asked? If your carrier asked at renewal? If opposing counsel asked, under oath, with the set on the table?

Many firms could not answer these questions from their documents today. That is not a failure of diligence; the profession's architects are not less careful than they were five years ago. It is a failure of definition, and it is a failure the recent guidance has narrowed without closing. Recent professional and regulatory materials commonly emphasize qualified human review, validation, or oversight. They do not prescribe one nationally uniform architectural verification test or record specification. Without a defined internal method, a firm may have difficulty demonstrating what review actually occurred.

The deposition question is foreseeable, even if the governing definition is not yet settled: What method did the firm use, who applied it, and what records show that the work was actually reviewed? A firm with a written methodology that was adopted, followed, and documented can answer from a contemporaneous record. A firm without one may have to reconstruct its controls after the fact, when missing records and unexplained decisions can themselves become subjects of inquiry.

That is the protective value of this Guide. It is not a safe harbor, a guarantee against liability, or proof that professional obligations were satisfied. It is a disciplined means of reducing an avoidable vulnerability: the gap between a firm's assertion that responsible control was exercised and its ability to demonstrate how.

As professional guidance develops, published methodologies may become relevant through firm adoption, incorporation into agreements or policies, recognition in professional practice, or reliance by a properly qualified witness. Publication alone does not create professional custom or establish the standard of care. But it can sharpen the questions asked about what risks were foreseeable, what controls were available, and why a firm chose the process it did. The legal significance of any methodology will depend on the governing jurisdiction, the circumstances of the project, and evidence that the methodology was actually implemented. Those uncertainties are reasons to exercise reasonable control now, not reasons to postpone it.

1.6 What Reasonably Prudent Now Means

The structural truth beneath all of this is that when the cost of producing work product collapses, professional value shifts toward verification. An LLM can draft a specification section; only a qualified professional can evaluate it, integrate it with the rest of the work, and accept responsibility for its use. An LLM can identify or summarize a code provision; reliance requires verification against the adopted edition, applicable amendments, and project-specific facts. The profession's distinctive value increasingly lies in the disciplined judgment, verification, coordination, and accountability that make its work reliable.

Prudent use lies neither in abstention from these tools nor in reliance on them without verification. For a firm that chooses to conform its practice to this Guide, the operative questions are whether machine assistance is limited to defined tasks; whether material machine contributions are verified by appropriately qualified professionals against authoritative sources before reliance; whether confidential information is protected; and whether contemporaneous records allow material decisions and checks to be reconstructed. These are not proposed as new professional duties. They translate familiar elements of responsible control into an operational method for the tools now in use.

Verification and documentation serve different purposes. Verification supports the reliability of the work; documentation supports the firm's ability to demonstrate how that reliability was assessed. The absence of a record does not establish that no review took place, but it leaves the firm dependent on memory, reconstruction, and assertion precisely when its process is most likely to be questioned. Guidance that calls for review without specifying how that review should be documented therefore stops short of an operational method. This Guide supplies that method: a proportionate, contemporaneous record of what machine assistance contributed, what was checked, against which sources, by whom, and with what result. Such a record is neither a safe harbor nor a guarantee against error or liability. It is evidence of a deliberate professional process—and the difference between claiming that reasonable control was exercised and being prepared to demonstrate it.

Every building site carries one instrument built for proof: the benchmark, the surveyor's datum, set before the first footing is poured, from which every elevation on the project is measured. No one asks the builder to promise that the floors are level; the datum makes the claim checkable, and checkability is what makes it credible. The profession's new means of production needs the same thing: not a new duty, but a fixed reference against which performance of an existing duty can be measured and shown. That is the work this Guide proposes.

The duty has not changed. The proof of it must.

Part II The Seven Principles

Seven principles govern conformance to this Guide whenever machine assistance contributes to an instrument of service that a licensed professional will seal. Each principle is grounded in an existing professional obligation or in the cited model provisions, regulatory guidance, and allied-profession ethics authority. This Part converts those sources into a defined architectural method for machine-assisted work: what qualified professional

judgment must accomplish before the work is relied upon or sealed. Each principle is stated, explained, and linked to the provisions that put it to work. Section 2.8 explains why the method is structured as a defined test rather than left to undifferentiated professional diligence. That section is not an eighth principle; it is the evidentiary basis for the seven.

2.1 Responsible Control

For conformance to this Guide, the architect shall exercise responsible control over every material machine-assisted contribution to an instrument of service. Generated output enters the professional workflow as unverified draft material. It has no authority, reliability, or professional standing of its own.¹

What it means. Responsible control is substantive, not nominal. The adopted and model provisions cited here converge on a clear core: control over the work and professional knowledge sufficient to direct, evaluate, integrate, and take responsibility for it. Illinois requires control over and detailed professional knowledge of technical submissions during their preparation, and expressly rejects mere review or correction without control over the content throughout preparation. California likewise defines responsible control through control over content during preparation. NCARB requires ultimate authority together with the knowledge and ability to oversee, delegate, and integrate design and technical decisions.²

The 2025 NCEES Model Law makes the test especially concrete: responsible charge requires full professional knowledge and control, including authority to review, change, reject, or approve work in progress and final work; awareness of the project's parameters and limitations; the ability to answer questions about the relevant decisions; and acceptance of full responsibility.³ A professional can nominally supervise a process and still fail that test.

Architectural sealing provisions are built around identifiable professional relationships: work prepared by the architect, by individuals under the architect's responsible control, or by another architect whose work is reviewed and coordinated or integrated.⁴ An LLM occupies none of those roles. It is not a professional collaborator, subordinate, or outside architect. It can produce content, but it supplies no professional judgment, knowledge, or accountability. Its output is raw material. Before the firm relies on it, a qualified professional must bring it

¹ NCARB *Model Law* § 103(16) (July 2024), at 8; NCEES *Model Law* § 110.20(E) (Aug. 2025), at 4.

² 225 ILCS 305/14; Cal. Bus. & Prof. Code § 5535.1; NCARB *Model Law*, supra note 1, § 103(16), at 8.

³ NCEES *Model Law*, supra note 1, § 110.20(E)(1–4), at 4; see also NSPE Position Statement No. 10-1778, *Responsible Charge*, as quoted in NSPE Board of Ethical Review, Case 24-2, *Use of Artificial Intelligence in Engineering Practice* (July 18, 2024), at 7.

⁴ NCARB *Model Regulations* R401.1(1)(a)–(d) (July 2024), at 39.

within the firm’s controlled process by understanding, verifying, and deliberately integrating it into the work.

What it requires.

- Treat AI output as unverified draft material. Before a material contribution enters sealed work or a gated material professional communication, a qualified professional shall understand the content, integrate it into the work, and verify it under § 2.2.
- Do not adopt AI output merely because it is fluent, plausible, consistent with expectations, or has received a general review. Apply the content-specific control and verification required by this Guide.
- Create the contemporaneous record required by this Guide: what the machine contributed, what was checked, against which sources, by whom, and with what result.

In practice. The controls at §§ 3.A–3.E; the seal analysis at § 5.2; the pre-sealing checklist at Appendix E.

2.2 Verification

Every material machine-assisted contribution that will be relied upon in sealed work, and every material professional communication outside the sealed set identified in § 3.G, shall be verified before reliance. A verification is complete only when three elements are present: a stated criterion, an authoritative source, and a recorded result.

What it means. The cited materials address review in different ways. AIA professional guidance recommends review and validation by qualified professionals.⁵ ICOR’s regulatory public statement says human oversight should not be automated.⁶ The NCEES policy statement calls for critical assessment, independent checks, and documented methods and data.⁷ NSPE’s advisory engineering ethics opinion concludes, on its stated facts, that the engineer’s sealing of AI-assisted work without responsible charge was unethical.⁸ These

⁵ AIA, *Guidance for the Responsible Use of AI by Architecture and Design Firms* (updated Oct. 22, 2025), item 1, at 2.

⁶ ICOR, *Public Statement on the Role of Artificial Intelligence in the Design Professions* (Jan. 2026), Human Oversight Is Essential.

⁷ NCEES, *Manual of Policy and Position Statements* (Aug. 2025), PS 6.10, *Responsible Use of Artificial Intelligence in Engineering and Surveying*, at 42.

⁸ NSPE Board of Ethical Review, Case 24-2, *Use of Artificial Intelligence in Engineering Practice* (July 18, 2024), conclusion 2, at 8.

sources do not have equivalent legal effect, and none defines the complete test proposed here. For conformance to this Guide, § 2.2 defines that test.

Verification is tripartite. A **criterion**: the specific claim, stated so it can pass or fail (“does the cited section exist in the edition our jurisdiction adopted, and does it say what we rely on it to say”). An **authoritative source**: the adopted code edition, the current standard, the manufacturer’s current data, the governing contract, and never the AI’s restatement of any of them. A **recorded result**: the outcome, captured as the work is done. A check missing any of the three is a review. The ethics record shows the difference: the same engineer, on the same project, kept responsible charge of the document he verified thoroughly and lost it over the one he only skimmed.

What it requires.

- No material machine-assisted contribution enters sealed work, or a gated material professional communication, untested.
- The test is against the primary document; a second AI’s agreement is not a source.
- The result is recorded in the same motion, in the two-layer form Part IV specifies.
- Verification addresses the proposition relied upon, not merely whether a quoted sentence or cited source exists. Where material, it tests applicability to project facts, completeness, relevant exceptions, and coordination with dependent information.
- The depth and granularity of verification shall be proportionate to the consequence of the assertion and the use made of it. Proportionality changes the extent of the test; it does not permit a material assertion entering sealed work to remain unverified.

In practice. The protocols at Appendix C, surface by surface; the record at Part IV; the evidence for why review without a defined test fails at § 2.8.

2.3 Provenance

The architect shall be able to establish the origin of every material machine-assisted assertion relied upon in sealed work or a gated professional communication: which content was machine-assisted, and against which authority each machine-supplied fact was verified.

What it means. Provenance is the chain of custody for the facts an instrument asserts. ICOR’s regulatory public statement calls for the origin and history of data to be tracked and documented.⁹ NSPE’s advisory engineering ethics opinion concludes, on its stated facts, that

⁹ ICOR, *Public Statement*, supra note 6, Data Quality and Integrity.

an AI-drafted document’s omission of citations to the technical authority on which it depended was deficient, independent of whether its conclusions were correct.¹⁰ The practical concern is that AI-generated prose can sever an assertion from its source: a code requirement without the section, a material property without the standard, or a dimension without the calculation. An assertion whose origin cannot be established cannot be shown to have been verified.

What it requires.

- Every load-bearing assertion traces to its source: the code conclusion to the adopted edition and section, the standard to its status, the product to current data, the contract statement to the document and section.

In practice. The project context and verification-entry fields at § 4.3 satisfy this principle by construction; the record specification is Appendix D.

2.4 Competence

The architect shall not employ an AI tool in the production of an instrument of service without competence in that tool: knowledge of its capabilities, its limitations, and its characteristic failure modes, sufficient to evaluate its output.

What it means. The duty of competence arises from applicable licensing rules. NCARB’s model rule requires the knowledge and skill necessary in the specific technical areas involved.¹¹ Allied engineering ethics opinions have long treated production tools as aids to, not substitutes for, professional judgment; ICOR’s public statement describes tool competence in terms of capabilities and limitations.¹² What is new is how this tool fails. An LLM may invent a citation to authority that does not exist. On direct, verifiable legal questions, purpose-built professional research tools hallucinated at rates between 17 and 33 percent, and a general-purpose model at 43 percent, in the cited preregistered evaluation. The specific rates will move with model generations. Two risks remain material: a system may accept a false premise rather than correct it, and it may not reliably report that it has done so. Professional research tools performed better but still returned misleading or false information on roughly one query in five; their subtler error was inapplicable authority—a

¹⁰ NSPE BER Case 24-2, *supra* note 8, conclusion 1, at 8.

¹¹ NCARB *Model Rules of Conduct* (rev. July 2023), Rule 1.3, at 6.

¹² NSPE BER Case 90-6, *Use of CADD System* (Nov. 9, 1990), at 2; accord NSPE BER Case 98-3 (Jan. 22, 1999), at 4; ICOR, *Public Statement*, *supra* note 6, Professional Competence and Continuing Education.

real source from the wrong jurisdiction or a superseded text.¹³ NCARB’s position encourages transparency about how AI tools operate.¹⁴ This Guide applies these non-equivalent sources functionally; it does not claim they impose one uniform rule in every jurisdiction. NCARB’s position also cautions regulators against unnecessarily limiting technological advances.¹⁵ Accordingly, the competence condition for conformance is practical: know what the tool does reliably, where it fails, and how its failures present.

What it requires.

- Treat any claim about a tool’s accuracy, reliability, or fitness for the intended task as a claim to be tested, never a premise to be accepted.
- Maintain the knowledge as a firm asset; § 6.5 treats verification capacity as maintained, not assumed.

In practice. The competence provision at Appendix A, § A.9; the failure modes by surface at Part III.

2.5 Confidentiality

The architect shall not submit confidential, privileged, proprietary, personal, or security-sensitive information to an AI tool unless the submission is authorized and the tool, account, integration, and configuration have been approved for that data class. An open interface shall not be treated as confidential merely because access to it requires an account.

What it means. An NSPE advisory engineering ethics opinion concluded, on its stated facts, that uploading a client’s information into an open AI interface was tantamount to placing it in the public domain without the client’s consent. AIA professional guidance recommends that firms not upload confidential client or firm data into unvetted systems and that they understand a platform’s data-storage and model-training practices. Neither source, by itself, establishes a uniform architectural rule. Applicable law, contract, client restrictions, privilege, data-protection duties, professional obligations, vendor terms, and the actual system configuration must be evaluated. For conformance to this Guide, approval is data-

¹³ Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning and Daniel E. Ho, *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, 22 J. Empirical Legal Stud. 216 (2025), at 216, 223, 225, 230–31.

¹⁴ NCARB, *Position on the Use of Artificial Intelligence in the Architectural Profession* (updated Apr. 2026), AI tools should provide transparency.

¹⁵ NCARB, *Position on AI*, supra note 14, Regulators should not limit technological advances.

class- and configuration-specific; a vendor name alone does not establish that a particular use is vetted.¹⁶¹⁷

What it requires.

- A written register stating, for each approved tool and integration, the owner, approved data classes and uses, provider and material subprocessors, required configuration, storage location where known, retention and training terms, deletion and incident provisions, intellectual-property/output terms, review date, and approval status.
- A defined boundary between each approved instance or integration and any open or unapproved interface, enforced through access controls, configuration, training, and periodic review.

In practice. The controls at § 3.F; the tool register and boundary provisions at Appendix A, § A.8.

2.6 The Record

The architect shall maintain a contemporaneous, retrievable record sufficient to reconstruct, for any gated use, the project and instrument context, which content or verification population was machine-assisted, how it was verified, against what authority, by whom, when, and with what disposition.

What it means. NCARB Model Law § 403 is a model provision; it becomes binding only through jurisdictional adoption. It proposes records adequate to demonstrate responsible control and expressly addresses submissions prepared by integrating the work of others.¹⁸ Part IV proposes the content and timing of a record for machine-integrated work. The significance and admissibility of a contemporaneous record depend on the governing rules and foundation; the cited evidence rules describe possible routes for business records and authentication, not an automatic result.¹⁹

What it requires.

- Open the project-level record when machine assistance is first used for a gated purpose. Record the project and instrument context once, and update it whenever the scope, tool, source set, or intended use materially changes.

¹⁶ NSPE BER Case 24-2, *supra* note 8, at 6.

¹⁷ AIA, *Guidance for the Responsible Use of AI*, *supra* note 5, item 3, at 2.

¹⁸ NCARB *Model Law*, *supra* note 1, § 403, at 18; 225 ILCS 305/14.

¹⁹ Fed. R. Evid. 803(6), 901(a), 902(11) (rules current through Dec. 1, 2025).

- Create verification entries as the tests are performed—not at project close and not after the work is questioned. Each entry shall identify the machine-assisted content or verification population, the method of verification, the authority consulted, the qualified professional performing the review, the date, the findings, and the disposition.
- Do not mistake a prompt history, model output, or vendor-generated activity log for a verification record. Those materials may support the record, but they do not establish that a qualified professional tested the content against appropriate authority.
- Preserve the record with the project files, protect it against unauthorized alteration or loss, and retain it for the period established by governing requirements and the firm’s applicable retention policy.

In practice. Part IV specifies the record; Appendix D is the working specification; § 5.3 shows the record answering questions under oath.

2.7 Boundaries

For conformance to this Guide, AI may assist professional tasks but shall not be the sole or primary basis of a compliance determination, shall not originate a professional review action issued under a seal, and shall not supply sealed content without qualified professional review, integration, and verification.

What it means. These boundaries are conservative conditions of conformance proposed by this Guide. They are based on the premise that compliance determinations and sealed review actions require accountable professional judgment. The cited NCARB model rule addresses reliance on qualified persons concerning building laws.²⁰ It does not expressly classify AI output or establish a nationally applicable prohibition. Whether governing law independently restricts a particular use remains jurisdiction-specific. Under this Guide, AI may assist those tasks but may not supply the sole or primary basis for the professional determination or originate a sealed review action.

What it requires.

- No compliance conclusion rests solely or primarily on AI output.
- No review action is machine-generated and issued under a seal.
- Unclassified uses are treated as gated until classified.

²⁰ NCARB *Model Rules of Conduct*, supra note 11, Rule 1.2, at 6.

In practice. The three tiers at § 3.G; the full classification at Appendix B; the policy provision at Appendix A, § A.4.

2.8 The Evidence Behind the Principles

The principles above define verification as a test rather than leaving its content to a general instruction to exercise professional diligence, and that choice is not stylistic. It rests on evidence from medicine and human-factors research, fields that have measured what happens when professionals review automated output. As of the date of issuance of this Guide, no published study of architects verifying AI-assisted instruments of service has been identified. The adjacent literature shares features relevant to reliance—a licensed professional, an imperfectly reliable tool, and consequences borne by others—but differs in work, evidence base, and institutional setting. It supports the design of verification controls; it does not establish that those controls have been validated in architectural practice.

Researchers use automation bias to describe the tendency to accept an automated aid’s output in place of adequate independent evaluation. It has been observed under both single-task and multitask conditions. Lyell and Coiera found no bias in the low-complexity subset of the single-task group, but bias appeared in multitask subtasks rated low, medium, and high in verification complexity.²¹ Romeo and Conti characterize the broader relationship directionally: as verification becomes more difficult, reviewers are more likely to accept automated output without adequate independent evaluation.²² The evidence makes verification complexity a relevant risk factor, not a threshold below which bias is absent. Applying that adjacent evidence to an architectural code study is an inference, not an empirically tested architectural result.

The broader effect on combined human–AI performance has also been measured. A 2024 meta-analysis pooling 106 experiments found that human–AI combinations performed, on average, worse than the better-performing of the human or the system alone.²³ The losses concentrated in decision-making tasks, and the difference between decision-task losses and creation-task gains was statistically significant.²⁴

²¹ David Lyell and Enrico Coiera, *Automation bias and verification complexity: a systematic review*, 24 J. Am. Med. Informatics Ass’n 423 (2017), at 423, 426, 429.

²² Giuseppe Romeo and Daniela Conti, *Exploring automation bias in human–AI collaboration: a review and implications for explainable AI*, 41 AI & Soc’y 259 (2026), at 259, 262.

²³ Michelle Vaccaro, Abdullah Almaatouq and Thomas Malone, *When combinations of humans and AI are useful: a systematic review and meta-analysis*, 8 Nature Human Behaviour 2293 (2024), at 2293, 2295.

²⁴ Vaccaro et al., *supra* note 23, at 2295.

Two operational implications follow. First, general review is not the same operation as the verification defined at § 2.2, and human–AI combinations do not reliably outperform the better performer alone. Second, task type matters: the meta-analysis found average gains in creation tasks and average losses in decision tasks. Those research categories do not map exactly onto architectural instruments or the Guide’s tiers. This Guide therefore uses destination and reserved professional judgment as operational gates, and uses task type and characteristic failure to shape the verification method. That taxonomy is the Guide’s proposal, not a finding established by the meta-analysis.

Reviewer education alone is not a sufficient control. In the studies reviewed by Lyell and Coiera, the tested interventions did not significantly reduce automation bias. Those interventions included additional practice performing the underlying task manually, exposure to examples of automation failure, and explicit instruction about automation bias and ways to avoid it.²⁵

A later randomized clinical trial showed that substantial automation bias remained possible after physician AI-literacy education. All 44 participants were licensed physicians who had completed a twenty-hour physician AI-literacy course covering LLM capabilities, prompt engineering, and critical evaluation of LLM-generated output. Each physician diagnosed six clinical vignettes with optional access to LLM recommendations. Physicians in the control group received error-free LLM recommendations. Physicians in the experimental group received deliberately erroneous LLM recommendations in three of the six cases. Compared with the control group, the experimental group’s adjusted diagnostic-reasoning score was 14.0 percentage points lower, and its adjusted first-choice diagnostic accuracy was 18.3 percentage points lower.²⁶

The trial did not measure the effect of the physician AI-literacy course because every participant had completed it. Its narrower finding is still important: completing the course did not eliminate susceptibility to erroneous LLM advice. Reviewer education remains valuable, but it cannot substitute for a verification process designed to detect plausible automated error.

²⁵ Lyell and Coiera, *supra* note 21, at 430.

²⁶ Ihsan Ayyub Qazi, Ayesha Ali, Asad Ullah Khawaja, Muhammad Junaid Akhtar, Ali Zafar Sheikh and Muhammad Hamad Alizai, *Automation Bias in Large Language Model–Assisted Diagnostic Reasoning among Physicians Trained in AI Literacy—A Randomized Clinical Trial*, 3 NEJM AI No. 5 (2026), at 6–7, doi:10.1056/Ala2501001.


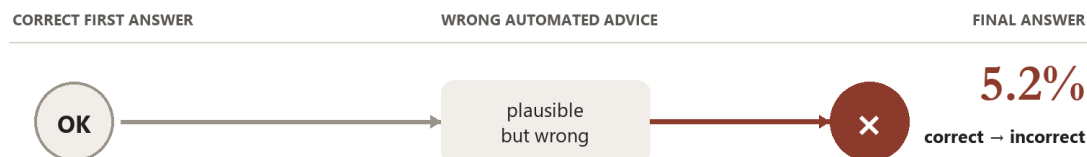


Figure 1. All 44 physicians had completed the same twenty-hour AI-literacy course. The comparison is erroneous versus error-free LLM recommendations, not trained versus untrained physicians. Data: Qazji et al., medRxiv preprint.

An earlier study measured the failure directly, in the smallest unit it has: the professional who had the right answer and gave it up.



Measured across 520 prescribing cases; switching was more frequent among less-experienced reviewers.

Figure 2. The switching rate. Data: Goddard, Roudsari and Wyatt (2014).

Professionals reviewing automated advice changed a correct answer to an incorrect one in 5.2 percent of all cases. The less experienced the reviewer, the more often it happened.²⁷

Training alone is therefore not a complete control. The prior review found no significant reduction from the specified training interventions.²⁸ A separate randomized trial tested a dual behavioral nudge at the moment of reliance: 72 AI-trained physicians completed 432 case evaluations while receiving a benchmark-accuracy cue and a case-specific confidence signal. Compared with an LLM recommendation alone, the dual cue improved diagnostic-

²⁷ Kate Goddard, Abdul Roudsari and Jeremy C. Wyatt, *Automation bias: empirical results assessing influencing factors*, 83 Int'l J. Med. Informatics 368, 370 (2014).

²⁸ Lyell and Coiera, *supra* note 21, at 430.

reasoning accuracy by 7.6 percentage points and first-choice accuracy by 10.9 percentage points.²⁹

Separate studies support controls at the moment of reliance

A literature review and a separate randomized trial answer different questions




Figure 3. The panels report separate evidence: a prior review of training interventions and a separate randomized trial of a dual behavioral nudge. The trial did not test authoritative-source verification. Data: Lyell and Coiera (2017); Qazi et al., *medRxiv preprint* (2026).

That trial tested behavioral cues, not verification against an authoritative source. Its reported benefit was larger among less experienced physicians, a subgroup result.³⁰ It supports the narrower proposition that an in-workflow intervention can change reliance behavior in that setting. The design implication for this Guide is limited: a control placed at the moment of reliance is more likely to be performed within the workflow than one delivered only as prior instruction. Whether the protocols in § 2.2 improve architectural outcomes or remain effective under deadline pressure has not yet been empirically tested.

Part III Reasonable Control by Failure Surface

Part II states the principles. This Part converts them into operational controls organized around the surfaces of practice on which machine assistance can fail. Together, these controls constitute the Guide's method of Reasonable Control: a documented, task-specific means of maintaining responsible control when machine assistance enters professional work.

²⁹ Ihsan Ayyub Qazi, Ayesha Ali, Asad Ullah Khawaja, Muhammad Junaid Akhtar, Ali Zafar Sheikh and Muhammad Hamad Alizai, *Mitigating Automation Bias in Physician-LLM Diagnostic Reasoning Using Behavioral Nudges: A Randomized Controlled Trial*, medRxiv 10.64898/2026.06.01.26354596 (June 2, 2026), at 2, 5–6.

³⁰ Qazi et al., *Behavioral Nudges*, supra note 29, at 15.

Each surface begins with an illustrative failure scenario and then states the applicable requirements in the normative language of § 0.5. The scenarios illustrate characteristic failures; they are not reports of events established in decided cases. Each requirement is a condition of conformance to this Guide. A cited source imposes an external requirement only where the text expressly identifies an applicable enacted law, adopted rule, or contract. Wherever the work product has a defined document structure, the control is addressed to that structure.

3.A Code, Zoning, and Regulatory Analysis

The characteristic failure of this surface is a complete, internally consistent analysis built on the wrong governing premise. A machine-assisted feasibility study for a mixed-use tower calculates floor-area ratio, height, parking, gross area, and yield from the base zoning district while missing an overlay, a parcel-aggregation condition, or the local definition that determines what counts as floor area. Every downstream number reconciles because every number descends from the same error. The code variant is equally precise: the machine cites a real provision from an edition the jurisdiction did not adopt, or applies model text without the local amendment that controls the project. Published evaluation of AI-assisted legal research has documented both unsupported authority and the subtler use of real but inapplicable authority on directly verifiable legal questions.³¹ The regulatory structure makes that failure consequential because the governing requirement is ordinarily the edition adopted by the authority having jurisdiction, as locally amended.³² The failure is not merely a fabricated citation. It is inapplicable authority converted into apparently reliable design information.

Controls. The architect shall not rely on machine output as the sole or primary basis for a compliance determination (§ 2.7). Before a material machine-assisted regulatory proposition is incorporated into an instrument of service or relied upon in a gated professional communication, verification shall establish the governing-source stack: the jurisdiction and adopting authority; the adopted edition and effective date; applicable local amendments and administrative criteria; incorporated standards; and the project facts, definitions, exceptions, and dependencies that determine applicability. The verifier shall read the cited provision and each material cross-reference on which the conclusion depends. Where the conclusion depends on interpretation rather than text, the interpretation is professional judgment and shall be made by a qualified professional. The record shall identify the proposition tested, governing sources, material applicability facts, method, verifier, date, disposition, and

³¹ Magesh et al., *Hallucination-Free?*, supra note 13, at 223, 225, 230–31, doi:10.1111/jels.12413.

³² International Code Council, *Code Adoption Resources* (accessed Sept. 2, 2026).

affected downstream calculations, drawings, or decisions. The controls of this surface descend from §§ 2.2 and 2.7. The working protocol is Appendix C.2; the policy provision is Appendix A, § A.5.2.

3.B Specifications

The characteristic failure of this surface is the substitution of product plausibility for assembly compliance. A machine drafts a technically fluent penetration-firestopping section, names a current sealant, and cites ASTM E814 or UL 1479. The product exists and the test method is real, but neither fact establishes that the proposed condition is protected. The controlling evidence is a listed system—or, where appropriate and accepted, a project-specific judgment with a justifiable technical basis—whose conditions match the rated wall or floor, penetrating item, material and size, opening and annular space, sleeve condition, packing, sealant depth, and required rating.³³ A product-only check passes while the assembly fails. Withdrawn standards, discontinued products, and obsolete manufacturer data are simpler variants of the same defect: the section appears complete because its words are plausible while the performance chain has not been established.

Controls. A machine-assisted specification section shall receive section-level professional review before incorporation; it shall not be incorporated on the strength of a whole-document impression. Review shall follow the requirement through the full three-part section. Part 1 references shall be verified against the issuing body’s current record. Part 2 products and manufacturers shall be verified against current manufacturer data. Material system performance shall be verified against the applicable listed, tested, evaluated, or otherwise accepted assembly and the installation conditions stated in Part 3. Where performance belongs to an assembly rather than a product, product data alone is not verification. The review shall also address material compatibility with adjacent systems, drawing coordination, related sections, and the conditions under which the stated performance is achieved. Each material verification shall be recorded under Part IV. A designation, product, assembly, or performance claim that cannot be verified against an appropriate current source shall not remain in sealed work. The working protocol is Appendix C.3; the policy provision is Appendix A, § A.5.3.

³³ ASTM International, ASTM E814-26, *Standard Test Method for Fire Tests of Penetration Firestop Systems* § 5.1 (2026); ASTM International, ASTM E3157-25, *Standard Guide for Understanding and Using Information Related to Installation of Firestop Systems* §§ 1.4, 1.9 (2025); UL Solutions, *Firestop and Joint Application Guide*, at 2, 8–12 (2024).

3.C Submittal and Shop Drawing Review

The characteristic failure of this surface is an approval that reaches farther than the review actually performed. A machine compares a resubmitted electrified-door-hardware package with the specification and the architect's prior comments. It reports every comment resolved and proposes an "approved as noted" disposition. Elsewhere in the package, however, an unmarked revision changes an electrified lock from the specified fail-safe function to fail-secure. The machine has answered the comparison it was given; it has not determined whether the resubmittal disclosed every change, whether the change is a deviation from the Contract Documents, or whether the architect intends to approve that specific deviation. The package may be internally consistent. The failure lies in the scope and meaning of the architect's review action.

Submittals are not Contract Documents. They show how the contractor proposes to conform to the information and design concept expressed in those documents, and the architect's review has a correspondingly limited purpose. Under commonly used general conditions, approval of a submittal does not approve an undisclosed deviation; the contractor must specifically identify deviations and, on resubmittal, revisions other than those requested by the architect.³⁴ Machine comparison can accelerate screening. It cannot determine what the architect has knowingly accepted.

Controls. Machine assistance may inventory specified criteria, compare submitted data, track responses to prior comments, and flag apparent changes. Before taking action, the architect shall determine that the submittal has received the contractor's review required by the governing Contract Documents; that the package is sufficient for the review action sought; that stated deviations have been isolated for express decision; and that revisions outside the architect's prior comments have been identified and reviewed. The architect shall understand the scope of any machine comparison, including pages, fields, criteria, prior versions, and related submittals omitted from it, and shall obtain review by affected consultants where required. A machine may draft disposition language only after the architect has selected the disposition through independent professional judgment; it shall not select or issue the disposition. The record shall identify the version reviewed, comparison scope, prior comments tested, disclosed deviations, material additional revisions, reviewer, and final disposition. References to a machine-generated disposition issued "under the architect's seal" should not be used; the relevant act is the architect's review action or review stamp. The controls of this surface descend from § 2.7 within the review scope established by the

³⁴ AIA Document A201-2017, *General Conditions of the Contract for Construction* §§ 3.12.4–.9, 4.2.7. The project-specific contract and modifications govern.

governing Contract Documents. The working protocol is Appendix C.5; the policy provision is Appendix A, § A.5.5.

3.D Drawing Production, Coordination, and QA/QC

The characteristic failure of this surface is coherent wrongness. A parametric facade revision changes the module, floor-to-floor zone, or panel geometry and automatically updates elevations, schedules, quantities, and fabrication geometry. Every generated view agrees. The revision nevertheless fails at an interface the automation does not govern: slab-edge embeds, perimeter fire containment, corner and transition conditions, movement joints, maintenance-equipment clearance, or construction tolerances. A conventional cross-sheet check may find no discrepancy because the documents inherit the same premise. Consistency among outputs is not verification when the same untested input generated them all.

Controls. Machine-assisted drawing production shall receive coordination verification appropriate to the rate and reach of generated change, not solely milestone review. For a machine-generated or machine-propagated change, verification shall identify the governing input, the intended dependent outputs, and the interfaces not governed by the automation. The test shall confirm both lineage within the dependency chain and coordination at its boundaries. Quantities, dimensions, classifications, keyed references, and schedules generated or propagated by a machine shall be verified against the governing project source before issuance. Agreement among outputs generated from the same source is evidence of propagation; it is not, by itself, evidence that the source or resulting construction is correct. The record shall identify the governing input, affected outputs, excluded or unmanaged interfaces, test performed, verifier, date, and disposition. The controls of this surface descend from §§ 2.1 and 2.2 and govern propagated content under the definition at § 0.3. The working protocol is Appendix C.4; the policy provision is Appendix A, § A.5.4.

3.E Construction-Phase Correspondence

The characteristic failure of this surface is an answer issued in the wrong professional and contractual form. An RFI asks whether a smoke-control duct may shift through a rated shaft wall to avoid a structural beam. A machine retrieves a similar response and drafts an intended practical instruction. It does not recognize that the answer must reconcile the mechanical smoke-control sequence, structural constraint, required shaft protection, damper and access conditions, and an applicable firestop system. Nor does it determine whether the response merely interprets the existing Contract Documents or changes the Work, cost, or time. A technically plausible sentence can therefore become both an incomplete life-safety decision and an unauthorized project change.

The distinction is contractual, not semantic. Common general conditions separate interpretations and RFI responses from Supplemental Instructions, minor changes, Change Orders, and Construction Change Directives; the standard AIA RFI form states that neither the request nor the response authorizes work that increases cost or time.³⁵ The machine may draft words. The architect must determine what professional decision is being made and which project instrument can carry it.

Controls. A machine-assisted RFI response, field directive, or other project correspondence is gated outside the sealed set when it could materially affect life safety, code compliance, building performance, construction cost or time, contractual rights, or a professional determination. Before issuance, the architect shall verify the material factual premises, implicated Contract Documents, contractual authority, affected disciplines, and known effects on related work. The architect shall determine whether the matter can be answered as an interpretation or requires a Supplemental Instruction, minor change, Change Order, Construction Change Directive, proposal request, or other instrument authorized by the governing Contract Documents. Review by affected consultants shall be obtained where the response crosses their scope. Routine minor confirming correspondence remains permitted under § 3.G and does not require the Part IV record. A machine shall not issue correspondence under the architect’s authority without the required professional review. The record shall identify the question, sources reviewed, affected work and disciplines, contractual classification, reviewer, date, and disposition. The working protocol is Appendix C.6; the policy provision is Appendix A, § A.5.6.

3.F Confidentiality and Data Handling

The characteristic failure of this surface is disclosure through scope rather than words. A project team connects a cloud-based assistant to the federated model for a courthouse, data center, health facility, or other security-sensitive project so that a user can ask one coordination question. The prompt contains no confidential detail. The connector nevertheless indexes geometry, room names, security zones, equipment, issue history, embedded metadata, and linked consultant models—potentially through a different account, tenancy, retention rule, processing location, or subprocessor chain than the team assumed. Nothing in the resulting instrument of service reveals the disclosure. The failure occurred at the system boundary before the answer was generated.

³⁵ AIA A201-2017, *supra* note 34, §§ 4.2.8, 4.2.11–14, 7.4; AIA Contract Documents, *Instructions: G716-2004, Request for Information (RFI)*. The project-specific contract governs.

Current professional guidance calls for qualified validation of AI-assisted work, approved systems and uses, protection of confidential information, and understanding how an external AI platform stores, reuses, or trains on submitted information. It also recognizes that AI capabilities may be embedded in modeling, specification, code-checking, and document-management tools rather than presented through a stand-alone prompt interface.³⁶ The prompt is therefore not the data boundary, and the product name is not the unit of approval.

Controls. Before protected information becomes accessible to a machine-assisted system, the firm shall classify the data and confirm that the specific provider, tool, account, tenancy, interface, integration, and configuration are approved for that class and consistent with client restrictions, privilege obligations, contract, and applicable law. Approval shall identify the full data path and access scope, including automatically indexed files, linked models, metadata, logs, extensions, and downstream services. It shall address, as material, provider and subprocessors; identity and access controls; encryption; storage and processing location; retention and deletion; use of submitted information for provider model development; incident notice and response; intellectual-property and output-use terms; and approved integrations. Review of the text a user intends to submit is not a substitute for review of the information the connected system can access.

The licensed professional remains accountable for the professional use of the output. Legal, security, privacy, procurement, and provider-review questions may be performed or advised upon by qualified persons; this Guide does not presume that the licensed professional is the firm's sole expert in those disciplines. The designated professional is responsible for ensuring that the required reviews and approvals occur and are documented.

The firm shall maintain the tool register specified in Appendix A, § A.8. The register shall identify the accountable owner, approved uses and data classes, provider and material subprocessors, required configuration, data scope and known location, retention and provider-development terms, deletion and incident provisions, intellectual-property and output terms, approval date, reviewer roles, next review date, and current status. Approval of one account, tenancy, or integration does not approve every interface bearing the same product name. The firm shall suspend or narrow use when a material change invalidates the approval.

³⁶ AIA, *Guidance for the Responsible Use of AI*, supra note 5, at 2–3; AIA AI Task Force, *AI Firm Toolkit* (Aug. 2026).

3.G The Task Taxonomy

The controls above are organized by failure surface. This section classifies the tasks themselves so that a firm can sort a proposed use of machine assistance before it begins rather than diagnose it after it fails.

Every use of machine assistance in practice falls into one of three tiers. A permitted task produces no material machine-assisted contribution to sealed work and no material professional communication outside the sealed set. Routine internal drafting, formatting, indexing, administrative work, and minor confirming correspondence are ordinarily permitted; the Guide's verification record is not required, although ordinary professional, contractual, confidentiality, and quality-control duties still apply. A gated task produces output that may materially enter sealed work, or a professional communication outside the sealed set that could materially affect life safety, code compliance, building performance, construction cost or time, contractual rights, or a professional determination. It is permitted only behind the applicable § 2.2 verification and Part IV record. A prohibited use falls within § 2.7. An unclassified use shall be treated as gated until classified.

The same software and project information can produce all three tiers. If a machine indexes photographs and room labels for an internal existing-conditions survey log, the task is ordinarily permitted. If it extracts dimensions or observed conditions into an existing-conditions plan that may enter sealed work, the task is gated. If it is allowed to resolve ambiguous geometry, infer a concealed condition, or represent the accuracy of the survey without the professional judgment and verification required by this Guide, the use is prohibited. The tier follows the output's destination, material consequence, and role in professional judgment—not the product name, prompt, or apparent simplicity of the task.

Classification shall answer four questions: What exact output will the machine produce? Where can that output go or be relied upon? What material consequence can it have? Does reliance on it supply or displace professional judgment reserved by § 2.7? If material machine-assisted content will enter, or is a candidate to enter, an instrument of service that a licensed professional will seal, the task is at least gated. If the output remains outside the sealed set, it is gated only when it is a material professional communication within the defined categories. Routine minor correspondence does not become gated merely because an architect sends it. If the use would assign to a machine a professional determination reserved by § 2.7, the use is prohibited. For a gated task, the characteristic failure then determines the applicable protocol.

This taxonomy is the Guide's operational proposal, not an empirical conclusion or a nationally uniform legal classification. Its allocation reflects current professional guidance

that AI remains a tool, that professional judgment and legal responsibility remain with the licensed practitioner, and that unclear or consequential uses require heightened review.³⁷ Sealed destination supplies the primary gate. Materiality performs the narrower function of preventing consequential professional communications outside the sealed set from escaping the gate; it does not convert every architectural decision or email into a verification event.

Delegated design offers a limited process analogy. Where design is delegated under the governing Contract Documents, the architect first states the performance and design criteria the delegated work must satisfy, and reliance is tied to the services, certifications, and approvals of an appropriately licensed design professional.³⁸ The analogy does not determine the legal status of machine output. Because machine output does not arrive under another design professional’s seal or responsibility, this Guide requires qualified verification rather than professional reliance on the machine.

The tiers state where this Guide’s controls attach. Sealed work remains the primary trigger. The limited outside-the-seal rule reaches only material professional communications. Within the gated tier, verification depth and granularity are proportionate to consequence and use, but a material assertion entering sealed work is not left unverified. The full classification is Appendix B; a firm adopts its own through Appendix A, § A.4.

Classify the task before machine assistance begins

Sealed work is the primary gate; Section 3.G supplies the limited communication gate




Figure 4. Sealed work is the primary trigger. Outside the sealed set, the limited gate applies only to a professional communication that could materially affect a § 3.G consequence category. Reserved professional judgment closes the use to AI. Derived from § 3.G.

³⁷ NCARB, *Position on AI*, supra note 14; AIA AI Task Force, *AI Firm Toolkit*, supra note 36.

³⁸ AIA A201-2017, supra note 34, § 3.12.10.1. The project-specific contract governs.

Part IV The Record

Part II states that a verification yields a recorded result, and that the architect shall maintain a record sufficient to reconstruct what was machine-assisted and how it was verified. This Part specifies that record: what it must establish, when it must be made, what it must contain, how long it must be kept, and to whom its contents are owed. The specification is short by design. A record that a practice cannot maintain under deadline is not a control; it is an aspiration that produces nothing on the day it is needed. These are author-proposed operational controls unless a jurisdiction, rule, or contract independently imposes one.

The distinguishing feature of this Part is that it is written from the position of the person who will later read the record (a plans examiner, an underwriter, a licensing board, or opposing counsel) rather than from the position of the person keeping it. Each field below is included because, as a practice-management matter, omitting it limits what the record can later show.

4.1 What the Record Must Establish

The record shall be sufficient to establish, for any gated use, which content or defined population was machine-assisted; the project, instrument, version, and tool context; what test was performed before reliance; who performed and approved it; and the result, status, exceptions, corrections, and reverification where applicable.

The cited sources do not have identical status. NCARB Model Law § 403 proposes records adequate to demonstrate responsible control over submissions prepared by the architect or by integrating the work of others; NCARB's Model Regulations and Illinois law provide related provisions within their respective scope.³⁹ Section 4.1 uses those sources to support the record's proposed purposes. They do not prescribe this Guide's complete field set, aggregation rules, sampling limits, or reverification triggers. A record establishes only the information it actually captures.

4.2 Contemporaneity

The record shall be made at or near the time of the verification it documents, in the ordinary course of practice. A record assembled after a claim has been asserted does not satisfy this Part.

³⁹ NCARB *Model Law*, supra note 1, § 403, at 18; NCARB *Model Regulations*, supra note 4, at 39; 225 ILCS 305/14.

Contemporaneous recordkeeping may assist later authentication and hearsay analysis, but admissibility is not automatic. Federal Rule of Evidence 803(6) requires each condition of the business-record exception. The timing, knowledge, regularly conducted activity, and regular-practice elements must be supported by qualifying testimony or certification, and the opponent may show that the source, method, or circumstances indicate a lack of trustworthiness. Rule 902(11) can provide a method of self-authentication when its requirements and notice provisions are met; it does not decide relevance, hearsay, trustworthiness, or evidentiary weight.⁴⁰ State proceedings may apply different rules. Preservation, spoliation, and adverse-inference questions arise under separate jurisdiction-specific law and depend on the facts; this Guide states no general adverse-inference rule.

For conformance to this Guide, the verification record is created as part of the work rather than assembled after a dispute. That design may help support the timing and regular-practice elements of a later foundation, but it does not establish them automatically. A later reconstruction may be more difficult to substantiate and may carry different weight, but timing alone does not determine admissibility.

The Guide’s operational test is whether the firm can retrieve the stated record without reconstructing it from memory or metadata; that is a record-system assessment criterion, not a legal test for contemporaneity.

4.3 Minimum Record Content

The record shall use two linked layers. Project and instrument context may be recorded once and reused. Each verification, aggregated population, or approved sample shall have its own entry. Together the layers shall identify what was tested, against what, how, by whom, when, and with what disposition and status.

A specific, retrievable entry is more useful than a generic notation such as “checked,” but its legal significance depends on the applicable foundation, context, and governing rules. The fields below are a practice specification, not a statement that every jurisdiction requires each field.

Existing regulation demonstrates the use of named project-record fields in at least one jurisdiction. New York requires, for evaluated third-party documents, a record identifying the project and documents evaluated, the source, the party for whom the evaluation was

⁴⁰ Fed. R. Evid. 803(6), *supra* note 19; Fed. R. Evid. 901(a), 902(11).

conducted, the date, and the evaluator’s seal and signature.⁴¹ The broader two-layer specification below is this Guide’s proposal and is not attributed to the New York rule.

Provenance is supported by the context and entry fields taken together. NSPE’s advisory engineering ethics opinion concluded, on its stated facts, that machine-drafted content omitting citations to pertinent technical authority was deficient, independent of whether its conclusions were correct.⁴² That source supports the need to reconnect an assertion to authority; it does not prescribe the complete record below.

4.3.1 Project and Instrument Context

The following context shall be recorded once for the project, instrument, issuance, or other bounded body of work to which verification entries link:

- project identifier; deliverable or instrument title; revision or version; and issue date or status;
- the task classification and governing verification protocol;
- the tool and provider; product, model, and version to the extent known; and any configuration material to the output or review; and
- a durable reference to the relevant input and output where appropriate, such as a controlled-repository identifier, file reference, or hash.

Data minimization governs this layer. The record does not require indiscriminate retention of prompts, privileged communications, confidential project data, personal information, or proprietary material. When full input or output retention is inappropriate, the firm shall retain a controlled reference or description sufficient to identify the content and verification without unnecessarily duplicating protected material.

4.3.2 Verification Entry

Each verification entry shall state:

- the content, assertion, item, or defined population verified, with its location in the instrument or communication;
- the pass criterion—the proposition the content must satisfy;
- the authoritative source or other governing evidence, with jurisdiction where applicable, edition or version, date, and pinpoint;

⁴¹ 8 NYCRR § 29.3(a)(3).

⁴² NSPE BER Case 24-2, *supra* note 8, conclusion 1, at 8.

- the verification method and evidence sufficient to show how the criterion was tested;
- the verifier’s identity, role, qualification relevant to the subject, and verification date or timestamp;
- the disposition: confirmed, corrected, removed, escalated, or unresolved; and
- any correction and reverification, approval or status, and linked exception or unresolved item.

4.3.3 Aggregation, Sampling, and Reverification

Entries may be aggregated only when the defined population shares the same task type, pass criterion, authoritative source and version, verification method, required verifier qualification, and disposition. The record shall identify the population and state whether review was complete or sampled.

Sampling may be used for repeated content under a documented, risk-based plan. It shall not be used to bypass verification of a unique material assertion affecting life safety, code compliance, building performance, contractual rights, or another professional determination. A failed sample shall trigger an expanded review appropriate to the failure and a record of the resulting correction or escalation.

Reverification is required when a change to the machine-assisted content, authoritative source or edition, material project facts, or a dependent instrument could affect the prior result. The original entry is preserved; the new entry is linked to it and states the superseding status. Exceptions and unresolved items remain visible until closed by an identified person or carried expressly into the release decision.

4.4 Retention

The record shall be retained as a distinct project-record category under a written schedule approved for the firm’s work and jurisdictions. The schedule shall consider applicable licensing requirements, contract duties, limitation and repose periods and their triggering events, insurer guidance, privacy and data-minimization duties, and legal holds. This Guide establishes no universal retention period, trigger, or safe harbor.

The cited retention sources differ in status and jurisdiction. NCARB Model Law § 403 proposes retaining records adequate to demonstrate responsible control for at least five years; Texas and New York impose their own adopted requirements within their stated

scope.⁴³ A firm must identify the rule governing its work. Keeping a record for a stated number of years does not, by itself, establish that the record was adequate to demonstrate responsible control.

Limitation and repose rules vary by jurisdiction, claim, duration, and triggering event. Substantial completion is relevant in some regimes but is not a universal trigger.⁴⁴ Some jurisdictions may lack a construction-specific statute of repose, and risk-management recommendations must be read within their stated assumptions.⁴⁵ A firm may use a documented project- or jurisdiction-specific schedule, or a uniform schedule shown to be adequate for all work it covers. The selection should be reviewed with qualified counsel and the firm's broker or insurer as appropriate; neither a board minimum nor a single risk-management rule can be assumed sufficient nationwide.

The verification record shall not be treated as an ordinary disposable draft. The written schedule shall state how it is classified, where it is stored, and how secure disposition occurs when the schedule expires. A legal hold suspends routine destruction for the affected records regardless of the ordinary schedule. The policy shall identify who may issue and release a hold, the systems, backups, and provider-held copies within scope, and the documented process for disposition after the hold is released. Retention of full prompts or protected source material is required only when appropriate under the schedule, data-minimization rule, and legal hold.

4.5 Disclosure and the Three-Duty Reconciliation

For content within this Guide's gated tier, verification and recordkeeping are conditions of conformance. Client disclosure is a distinct question governed by contract, applicable law or rule, and the circumstances. These duties are not substitutes for one another.

Three non-equivalent sources speak to disclosure, and at first they seem to point in different directions. An NSPE advisory engineering ethics opinion concludes, on its stated facts, that absent a contractual requirement there was no professional or ethical obligation to disclose use of an AI tool to the client, while noting that ethical principles favor transparency where the tool played a substantial role.⁴⁶ ICOR's regulatory public statement calls for data

⁴³ NCARB *Model Law*, supra note 1, § 403, at 18; 22 Tex. Admin. Code § 1.103(c); 8 NYCRR § 29.3(a)(3).

⁴⁴ 8 NYCRR § 29.3(a)(3); Tex. Civ. Prac. & Rem. Code § 16.008(a); accord Tex. Civ. Prac. & Rem. Code § 16.009; 22 Tex. Admin. Code § 1.103(c).

⁴⁵ ACEC Risk Management Committee and NSPE Professional Liability Committee, *Document Retention Guidelines—A Risk Management White Paper* (2016), at 3, 5, 6.

⁴⁶ NSPE BER Case 24-2, supra note 8, conclusion 3, at 8.

provenance to be tracked and documented.⁴⁷ AIA professional guidance recommends disclosure to clients and stakeholders of when and how the tools are used.⁴⁸ Their legal effect and professional scope differ. The practical tension can be analyzed by identifying the party, source, and independently applicable duty.

To the client, disclosure is first a question of the governing agreement and independently applicable law or rule. Where the agreement requires disclosure, that requirement governs; where it does not, the architect may disclose and, where machine assistance played a substantial role, generally should consider doing so. The NSPE advisory engineering ethics opinion discussed above concluded, on its stated facts, that no disclosure duty arose absent a contractual requirement while separately requiring verification. It is an allied-profession advisory opinion, not architectural law.

To a licensing board, any applicable record duty is independent of what the client contract requires or what the client knows. Whether such a duty applies, its scope, and its relationship to the seal depend on the governing jurisdiction. An agreement that no client disclosure is required does not resolve independently applicable board requirements.

To the work itself, the need for verification may arise from the applicable standard of care, a licensing rule, a contract, or a firm's own conformance commitment. Client consent does not by itself make unverified content reliable or relieve independently applicable duties. NCARB's position identifies whether, when, and how AI use should be declared as an open regulatory question, so firms should monitor jurisdiction-specific developments.⁴⁹

Part V When the Record Is Examined

The preceding Parts state what this Guide asks of a conforming firm. This Part addresses the limited ways a written methodology and contemporaneous records may assist when the firm's conduct is examined. Whether the Guide, its requirements, or evidence of conformance is relevant, admissible, or persuasive depends on governing law and the facts.

5.1 What Conformance Can and Cannot Show

Adoption of this Guide may provide evidence that a firm selected a stated methodology. Project records and other evidence may help show whether that methodology was followed. Neither adoption

⁴⁷ ICOR, *Public Statement*, supra note 6, Data Quality and Integrity.

⁴⁸ AIA, *Guidance for the Responsible Use of AI*, supra note 5, item 4, at 2.

⁴⁹ NCARB, *Position on AI*, supra note 14, areas for further exploration.

nor conformance establishes that the standard of care was met, and non-adoption does not establish that it was breached.

Begin with the question as it actually arrives. A claim has been filed. The architect's conduct will be measured against the standard of care, meaning what a reasonably prudent architect would have done in the same circumstances. The firm's lawyer will ask a version of one question early: can we show that the way you worked is the way careful firms work? That is where a written, followed methodology earns its keep, and the law has settled rules for what the showing is worth.

Counsel may ask what methods the firm used and whether those methods were reasonable under the circumstances. Evidence of professional custom can be relevant but is not conclusive.⁵⁰ A firm-developed methodology may document the firm's chosen process; it is not, by itself, proof of professional custom. The baseline obligation remains the skill and knowledge that members of the profession in good standing ordinarily bring to the work.⁵¹ A customary practice may also be unreasonable under the circumstances of a particular case.⁵²

The American Law Institute's approved Restatement position on medical malpractice illustrates a narrower point: compliance with a guideline issued by an authoritative body may support, but does not compel, a finding of due care, while a guideline alone does not establish breach without qualified expert testimony.⁵³ That rule concerns authoritative medical guidelines and does not establish the status of this Guide. The analogy supports caution in describing what any private methodology can prove.

Possible relevance of voluntary guidance. Voluntary guidance may be considered with other evidence when it has been adopted by a firm, incorporated into an agreement, recognized in professional practice, or used by a qualified expert with an adequate foundation. Publication alone does not establish recognition, custom, admissibility, or the standard of care. This Guide therefore makes no claim that a firm's decision not to adopt it is evidence of negligence.

⁵⁰ *Restatement (Third) of Torts: Liability for Physical and Emotional Harm* § 13 (Am. L. Inst. 2010).

⁵¹ *Restatement (Second) of Torts* § 299A & cmt. b (Am. L. Inst. 1965).

⁵² *Restatement (Second) of Torts* § 295A & cmt. c (Am. L. Inst. 1965).

⁵³ *Restatement of the Law Third, Torts: Medical Malpractice* § 6(b) & cmt. f (Am. L. Inst., Tentative Draft No. 2, 2024), at 39, 44.

5.2 The Seal

The legal effect and wording of a seal are jurisdiction-specific. Within this Guide, sealing marks the licensed architect's acceptance of the professional responsibility required by the governing jurisdiction; machine assistance does not displace that responsibility.

The specific representations made by a seal depend on the governing jurisdiction. NCARB's model law proposes one formulation.⁵⁴ NCARB's regulatory position and ICOR's public statement emphasize that sealing remains a licensed act and human oversight remains essential, whatever the tools.⁵⁵ Those sources do not make their formulation binding nationwide.

Selected model and state provisions describe circumstances in which an architect may seal work prepared by others⁵⁶ or rely on information supplied by outside parties, subject to conditions such as review, coordination, integration, professional knowledge, or reasonable trust.⁵⁷ Delegated design provides a related contractual example: the architect first specifies governing criteria, and the delegated design arrives under another design professional's seal.⁵⁸ An allied-profession model similarly frames responsible charge in terms of professional knowledge and control.⁵⁹ These sources were not drafted to determine how generative AI output should be categorized, and they do not establish that every outside source supplies a warranty. This Guide therefore adopts a conservative practice rule: AI output is treated as unverified draft content and may enter sealed work only through qualified professional review, integration, and verification under §§ 2.1 and 2.2. That rule is a proposed control, not a statement that existing law has classified AI output nationwide.

5.3 The Deposition

The firm's methodology shall enable it to answer, from the contemporaneous record and without reconstruction, the questions a competent examiner will ask about machine assistance in a sealed instrument.

The questions are not speculative; they follow from the seal representation, and any competent examiner reaches them within minutes. Whether AI was used in preparing the

⁵⁴ NCARB *Model Law*, supra note 1, § 401(2), at 17.

⁵⁵ NCARB, *Position on AI*, supra note 14, AI is a tool, not a replacement for professional judgment; ICOR, *Public Statement*, supra note 6, Human Oversight Is Essential.

⁵⁶ NCARB *Model Regulations*, supra note 4, R401.1(1)(a)–(d), at 39; 225 ILCS 305/14.

⁵⁷ NCARB *Model Rules of Conduct*, supra note 11, Rule 5.2, at 7.

⁵⁸ AIA A201-2017, supra note 34, § 3.12.10.1, at 16.

⁵⁹ NCEES *Model Law*, supra note 1, § 110.20(E), at 4.

instrument. Which portions. What was done to verify the machine-supplied content before sealing. Against what source. When. By whom, and with what qualifications. Whether the firm had a written methodology at the time. Whether it was followed on this project. And whether the architect can produce the record.

A firm operating to this Guide should be able to answer these questions from contemporaneous project records. Without a defined methodology or record, a witness may have more difficulty reconstructing and substantiating what review occurred. That absence does not establish that the review was inadequate. The practical value of a specific record is narrower: it connects the firm’s account to the work, source, verifier, and disposition.

Expert testimony, if offered, is governed independently by Rule 702 and controlling jurisdictional law.⁶⁰ Those authorities do not validate this Guide or its methodology and do not shift any party’s burden of proof.

5.4 Firm Conduct and Personal Exposure

An architect in charge of a firm’s architectural practice is answerable for the firm’s conduct.

The NCARB Model Rules propose that an architect in charge of a firm’s architectural practice is accountable under those rules when the firm violates them.⁶¹ Whether and how that proposition applies depends on the jurisdiction’s adopted rules and the facts. The practical point is organizational: the controls of Part III are addressed to individuals, while the record system, tool register, and methodology are firm-level instruments. Conformance therefore requires clear principal-level responsibility rather than assuming the NCARB model language applies automatically.

5.5 The Renewal

The firm shall be able to describe its methodology, its verification protocol, and its record system in the terms professional liability underwriters ask about them, and to substantiate the description.

A 2026 broker survey reports that professional liability insurers have begun addressing design-firm AI adoption and that governance, controls, transparency, and accountability may affect underwriting discussions.⁶² Sections 2.2 and Part IV address controls and records, § 2.4 addresses competence, and § 4.5 addresses disclosure. Those provisions identify subjects

⁶⁰ Fed. R. Evid. 702 (amended Dec. 1, 2023).

⁶¹ NCARB *Model Rules of Conduct*, supra note 11, Guiding Principle F, at 5.

⁶² Ames & Gough, *2026 A/E Professional Liability Survey Results*, amesgough.com (survey of 15 insurers) (as of July 25, 2026).

a firm may be asked to explain; they do not make this Guide an insurer requirement or underwriting benchmark.

A documented methodology may help a firm provide an underwriter with a more specific account of its controls. An underwriter may evaluate that account together with the firm's actual practices, services, claims experience, tools, and requested coverage. This Guide is one possible framework for organizing the description; it is not necessarily an underwriting benchmark. Separately, the coverage itself is moving: standardized generative-AI exclusion endorsements took effect in general liability in January 2026, and the professional lines have begun to follow, with at least one carrier introducing an absolute AI exclusion across its errors-and-omissions, directors-and-officers, and fiduciary products while others have signaled they will underwrite and price the exposure instead.⁶³ A firm should read its current policy for exclusions addressed to AI or automated systems, put the embedded-assistance question to its broker in writing, and substantiate its methodology rather than merely describe it. The Conformance Statement at § 6.4 provides the instrument for substantiating the methodology.

Part VI Implementation and Conformance

The preceding Parts define this Guide's method and record for demonstrating responsible control. This Part addresses implementation at three practice scales and defines what conformance permits a firm to say about its practice. The implementation premise is direct: a methodology cannot protect a firm merely by existing on paper. It must be designed to survive deadline pressure, followed in practice, and documented as the work occurs. When those conditions are met, the methodology and contemporaneous record provide concrete evidence of how the firm controlled machine-assisted work—evidence the firm can use to manage risk, improve practice, and account for its decisions in regulatory review, underwriting, a claim, or litigation. The legal effect of that evidence depends on actual implementation, the project facts, and governing law.

6.1 The Sole Practitioner

The sole practitioner is the simplest organizational case because responsibility is not divided among offices or departments. Conformance still requires a written methodology; a classification of recurring uses; the applicable verification protocols; the two-layer record of

⁶³ Matthew Lerner, *Insurers, Brokers Adjust as AI Exclusions Emerge*, Business Insurance (Apr. 7, 2026); Don Jergler, *Insurer Interest in AI Exclusions Growing as Risk Becomes Omnipresent*, Insurance Journal (July 22, 2026).

§ 4.3; a confidentiality and tool-vetting process; competence maintenance; and a written retention and legal-hold schedule. These controls may be implemented in a compact project log and a small set of written procedures, but the completeness of the required fields does not change with firm size.

The same confidentiality, competence, version-control, legal-hold, assessment, and conformance-accuracy requirements apply to a sole practitioner. They may be satisfied through the practitioner's own documented procedures and qualified outside legal, insurance, security, privacy, or technical advice where needed. Nothing in this Part exempts the sole practitioner from provisions that apply to the work performed.

6.2 The Small and Midsize Firm

At the scale of the small and midsize firm, the record becomes a workflow question: the verification is performed where the professional already works, and the record is made in the same motion. A designated professional is responsible for maintaining the methodology and ensuring that qualified legal, security, privacy, procurement, vendor, and professional reviews are assigned; the designee need not personally possess every specialty. The tool register shall contain the fields in § 3.F and be reviewed on a stated cycle and upon material change. Verification records shall reside in one named, retrievable system linked to the project and instrument context of § 4.3. The platform is not prescribed; completeness, contemporaneity, access control, and retrieval are.

6.3 The Multi-Office Firm

The multi-office firm must address two problems that the smaller practice does not face: consistency and jurisdiction. Consistency first. The methodology should be one document firm-wide, with office-level supplements only where a jurisdiction requires more; divergent office practices are the deposition problem in § 5.3 multiplied: the examiner asks why the firm's own offices could not agree on what verification required. A designated professional should be named for each office; one firm-level designee shall be responsible for the whole. Jurisdiction second. Retention under § 4.4 shall be keyed to the governing requirements and exposure periods applicable to the firm's work. A multi-office firm may use a documented project- or jurisdiction-specific schedule, or a uniform schedule shown to be adequate for all work it covers. The firm's written policy shall state the selected rule and its basis, and the firm shall apply it consistently.

6.4 The Conformance Statement

A firm may state conformance to this Guide only by version, in writing, and after a documented assessment supports the statement for the stated scope, date, and exceptions. The statement is the firm's self-attestation of methodology and implementation; it is not an author-issued certification or a warranty of work.

The Conformance Statement is a short executed instrument identifying the Guide version and effective date, the organizational and service scope assessed, the assessment date, and stated exceptions or unresolved items. It states that the firm maintains the methodology, verification protocols, task classification, tool register, competence process, record system, and retention/legal-hold controls the identified version specifies. Under § 0.8, review against the current version is recurring. The statement neither warrants that any instrument is error-free nor covers work outside its stated scope. It should be executed only when the firm can substantiate it with evidence.

6.5 Assessment Against the Guide

Before fieldwork begins, an assessment shall define in writing the Guide version and assessment date; offices, services, project types, systems, and time period in scope; assessor identity and qualifications; independence and conflicts; privilege and confidentiality arrangements; intended recipients and distribution; and known limitations. Where the assessor is internal, the memorandum shall state the reporting line and any responsibility for the controls being assessed. Because adjacent evidence identifies inexperience with the task as a driver of automation bias, competence maintenance remains within the assessment scope.⁶⁴

The assessment examines the written methodology; the inventory and classification of actual recurring and embedded uses; confidentiality governance and the tool register; verification protocols as performed; the completeness, contemporaneity, status, and retrievability of records; retention and legal-hold controls; and competence maintenance. Interviews may orient the work but do not substitute for documentary and project evidence.

The assessor shall define the population, sampling method, sample size, risk stratification, exclusions, and evidence-sufficiency threshold. Sampling shall include material and higher-consequence uses and shall not be represented as proof concerning items outside the

⁶⁴ Goddard, Roudsari and Wyatt, *supra* note 27, at 370; Moustafa Abdelwanis, Hamdan Khalaf Alarafati, Maram Muhanad Saleh Tammam and Mecit Can Emre Simsekler, *Exploring the risks of automation bias in healthcare artificial intelligence applications: A Bowtie analysis*, 5 J. Safety Science and Resilience 460, 463 (2024).

sample. A failed sample, unexplained exception, or missing record shall trigger an expansion or limitation appropriate to the condition, documented in the findings.

Each finding shall state the condition observed, evidence reviewed, Guide section, consequence or severity, required or recommended correction, responsible owner, target date, and any limitation or unresolved issue. Corrective action is closed only after evidence of implementation and, where appropriate, reperformance or follow-up testing.

The assessment's product is a findings memorandum with an executive statement of scope and limitations, detailed findings, exceptions, and a corrective-action record. It is not a maturity score, peer ranking, assurance opinion, or certification. Any Conformance Statement is the firm's own self-attestation under § 6.4 and Appendix F; an assessor's work does not create author approval or third-party status under this Guide.

§ 0 Governing Provisions

This section states the Guide's status, scope, defined terms, relationship to cited authority, normative language, version controls, and author disclosure. It governs the interpretation of every provision of the Guide and supplies limitations applicable throughout.

Reading Rule. *This Guide is an author-developed voluntary framework. It does not create legal obligations, define the standard of care, or represent consensus or endorsement by the cited organizations. A shall statement defines a condition of conformance to this Guide only. External duties arise, if at all, from the applicable jurisdiction, contract, engagement, or firm policy. The source-status distinctions in § 0.4 apply throughout.*

0.1 Purpose and Scope

This Guide proposes a framework for machine assistance whose output may enter instruments of service and, in the limited circumstances defined at § 3.G, material professional communications outside the sealed set. It defines a method for verifying material machine-assisted contributions, a corresponding record, and implementation controls. Its requirements govern only a claim of conformance to this Guide; they do not state external legal requirements unless an identified law, rule, or contract independently does so.

The Guide takes no position on which tools a firm should use, does not assess a firm's readiness to adopt them, and is not a technology guide or curriculum. Parts II through IV draw principles from the cited authorities and guidance and then add the author's proposed

definitions, controls, and records. The Guide identifies these additions as proposed practice controls rather than existing legal requirements.

0.2 Audience

The Guide is addressed to the licensed architect in responsible control of instruments of service and to the principals of architectural firms. Practitioners in engineering, landscape architecture, and interior design may find parts of the framework relevant, but this Guide does not claim that its requirements transfer to those professions without modification. Their governing laws, rules, contracts, and professional terminology must be evaluated independently.

0.3 Defined Terms

- **Responsible control; responsible charge.** *Related professional duties with a common operational core: the professional must possess the knowledge and authority required to direct, evaluate, approve or reject, and take responsibility for the work. Architectural provisions commonly use **responsible control**; engineering provisions commonly use **responsible charge**. The governing definition remains profession- and jurisdiction-specific. This Guide draws from their shared operational core only within each source’s stated scope.*⁶⁵
- **Machine assistance.** *The contribution of content, analysis, or disposition to professional work product by an artificial intelligence (“AI”) system, principally today a large language model (“LLM”), whose output is generated rather than retrieved or deterministically computed, whether the system stands alone or is embedded in ordinary production software. Deterministic computation is not machine assistance: building information modeling, parametric propagation, and analysis run to a stated method produce computed content, not generated content. Content a system propagates (a dimension, quantity, or reference carried deterministically across a document set from a governing change) is governed by § 3.D alone.*
- **Verification.** *A test with a stated pass criterion, performed against an authoritative source, yielding a recorded result. Defined at § 2.2.*
- **Authoritative source.** *The primary document that governs an assertion: the code edition adopted by the authority having jurisdiction as locally amended; the current published standard; the manufacturer’s current data; the governing contract document. An AI’s restatement of a source is not a source.*

⁶⁵ NCARB *Model Law*, supra note 1, § 103(16), at 8; NCEES *Model Law*, supra note 1, § 110.20(E), at 4.

- **The record.** *The contemporaneous, two-layer documentation specified in Part IV: project and instrument context recorded once, with a verification entry for each verification or documented verification population. Where federal vocabulary is useful, this Guide adopts terms from NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (July 2024), including confabulation, data provenance, and human–AI configuration.*
- **Instrument of service.** *As defined in the governing agreement; presumptively, the drawings, specifications, and other documents prepared by the architect for the project.*
- **Material; materiality.** *Capable, in context, of materially affecting life safety, code compliance, building performance, construction cost or time, contractual rights, or a professional determination. Materiality depends on consequence and intended use; it is not a generalized label for every professional communication.*
- **Professional representation.** *A communication issued under the architect’s authority that states or interprets a professional or contractual determination. Outside the sealed set, only a material professional representation is gated under § 3.G; routine internal drafting, administrative work, and minor confirming correspondence remain permitted under this Guide.*

0.4 Relationship to Existing Authority

Existing authorities state essential principles. This Guide makes them executable. The sources collected below establish responsible control, qualified review, independent verification, provenance, confidentiality, recordkeeping, and professional accountability at sealing. NCEES Position Statement 6.10 calls for independent checks and documented methodologies, data sources, and assumptions. NCARB Model Law § 403 supplies the architectural foundation for records adequate to demonstrate responsible control over original and integrated work. This Guide joins those propositions with the contributions of the other cited sources and carries them into a single architecture-specific operating method. It defines when controls attach, what verification requires, which authoritative source the test must use, what the contemporaneous record must contain, and how a firm implements the method through policy, task classification, failure-surface protocols, and pre-sealing review. The table identifies each source’s contribution, the operational question it leaves open, and the precise provisions through which this Guide answers it.

Source	Foundation established by the source	Operational question left open	This Guide's answer
<i>AIA, Guidance for the Responsible Use of AI</i> (2025)	AI output should be reviewed and validated by qualified professionals before reliance.	What qualifies as review and validation, and what demonstrates that they occurred?	Section 2.2 defines verification as a test with a stated criterion, an authoritative source, and a recorded result. Part IV and Appendix D specify the contemporaneous record. Appendix C applies the test to the principal failure surfaces of architectural practice.
<i>AIA, AI Firm Toolkit</i> (2026)	Responsible implementation requires governance, verification, firm policy, standardized workflows, and increasing organizational maturity.	Which uses require controls, where is the gate, and what firm process follows once the gate is crossed?	Section 3.G and Appendix B classify uses as permitted, gated, or prohibited. Section 2.2 establishes the verification required at the gate. Part IV and Appendix D establish the record, and Appendix A converts the complete system into an adoptable firm policy.
<i>NCARB, Position on AI</i> (2024, revised 2026)	Professional judgment, sealing responsibility, and accountability remain with the licensed architect regardless of the tools used.	How does an architect exercise responsible control over machine-assisted content, and what must be established before sealing?	Section 2.1 defines responsible control over every material machine-assisted contribution. Section 4.5 reconciles verification, recordkeeping, and disclosure. Section 5.2 carries those requirements to the seal, and Appendix E converts them into a pre-sealing decision.
<i>NCARB Model Law</i> § 403 (2024)	An architect should maintain records adequate to demonstrate responsible control over original work and work produced by integrating the contributions of others.	What must a machine-assistance record contain, when must it be created, and how must it address aggregation, sampling, reverification, and retention?	Section 2.6 and §§ 4.1–4.4 define the record's purpose, timing, content, aggregation rules, reverification triggers, and retention framework. Appendix D supplies the working record specification, and Appendix A, §§ A.6–A.7 , installs it within firm policy.
<i>NCEES Model Law</i> § 110.20(E) (2025)	Responsible charge rests on professional knowledge and control, supported by identifiable professional responsibilities rather than nominal supervision.	How do professional knowledge and control operate when material content originates through a machine rather than another professional?	Section 2.1 translates knowledge and control into an architecture-specific rule for material machine-assisted contributions. Section 2.2 requires verification before reliance. Section 5.2 and Appendix E apply both requirements at sealing.
<i>NCEES Position Statement 6.10</i> (2025)	Professionals should critically assess AI output, perform independent checks, and produce verifiable results supported by documented methodologies, data sources, and assumptions.	What constitutes an independent check, what makes its result verifiable, and what documentation must the project record preserve?	Section 2.2 defines the criterion-source-result test. Sections 4.2–4.3 and Appendix D establish the documentation required to demonstrate it. Appendix C supplies the architecture-specific protocols by which the test is performed.
<i>ICOR Joint Statement</i> (2026)	Human responsibility and oversight cannot be automated, and the	How is provenance traced from a material assertion to its	Section 2.3 defines provenance as a traceable chain from each material machine-assisted assertion to its

Source	Foundation established by the source	Operational question left open	This Guide's answer
	provenance of information should be tracked and documented.	controlling authority, when is that chain required, and how is it preserved?	authoritative source. Section 3.G and Appendix B determine when the control attaches. Part IV and Appendix D preserve the chain in the project record.
NSPE BER Case 24-2 (2024)	Responsible professional use of AI implicates verification before sealing, responsible charge, confidentiality, technical authority, and disclosure.	What makes verification comprehensive in the production and sealing of architectural instruments of service?	Section 2.2 and Appendix C establish a repeatable architecture-specific verification method. Sections 2.3–2.5 integrate provenance, competence, and confidentiality. Appendix E carries the method through the final pre-sealing review.

Note: The cited materials differ in legal status and professional scope. Model provisions become binding only through applicable jurisdictional adoption; professional guidance, regulatory positions, implementation resources, public statements, and advisory ethics opinions have the effect assigned by governing law and context. Engineering and surveying materials are used as allied-profession sources, not as statements of architectural law. This table identifies the Guide's synthesis and operational contribution; it does not treat the cited sources as legally equivalent or claim that the Guide alters governing law or independently determines the standard of care.

0.5 Normative Language

Within this Guide, *shall* identifies a condition of conformance, *should* identifies a recommendation, and *may* identifies permission. These verbal conventions are adapted from the *ISO/IEC Directives, Part 2* (9th ed. 2021), clauses 7.2–7.4. They do not convert a provision of this Guide into law or an external professional requirement. Within the published version history, section identifiers will not be reused for unrelated provisions. Retired provisions will remain visible in the public change log as reserved or superseded.

0.6 Status, Reliance, and Limitations

This Guide is an author-developed voluntary professional-practice framework. It is not a law, licensing-board rule, consensus standard, insurer requirement, or determination of the legal standard of care, and it has not been issued or endorsed by the organizations whose materials it cites. Publication alone imposes no obligation to follow it. A firm may voluntarily adopt some or all of its provisions, and obligations may arise independently through applicable law, a licensing rule, a contract, an engagement, or the firm's own policy. Conformance means only that the firm has assessed its stated practice against the identified version and scope of this Guide. Conformance does not establish that the standard of care was met, and non-adoption does not establish that it was breached. The relevance or weight

of this Guide in any dispute depends on governing law, the facts, professional acceptance, and the basis on which it is offered. This Guide is not legal advice and does not warrant error-free work.

0.7 About the Author and Disclosure

The author is a licensed architect practicing in forensic architecture, retained in construction disputes to examine how deficiencies enter instruments of service. He may provide continuing-education instruction and implementation consulting to architecture firms on the subject matter of this Guide. This Guide is published without charge, and conformance requires retention of no consultant: Appendix A supplies a model firm policy; Appendix B the task classification; Appendix C the verification protocols; Appendix D the record specification; Appendix E the pre-sealing checklist; Appendix F the conformance and proposal language; and Appendix G the glossary. A separate Source Verification Supplement reproduces the cited source pages held in the author's corpus. No fee payable to the author is contingent on any finding as to a firm's conformance, and the author issues no certification, registration, mark, or statement of conformance.

This is a single-author work. Unless separately identified, the cited organizations have not reviewed or endorsed the Guide's interpretations or proposed controls. Sources cited in this edition were reviewed through August 31, 2026. Corrections and later versions are governed by § 0.8.

0.8 Version Control and Maintenance

Each edition of this Guide is identified by a version number and effective date. The canonical publication page is <https://reasonablecontrol.org/guide>. That page identifies the current edition and provides access to each dated, version-numbered publication. Errata, the change log, and superseded editions are maintained at <https://reasonablecontrol.org/versions>. An erratum identifies the affected section, the correction, and its publication date. A later version does not retroactively alter a firm's conformance statement or project record. Any statement of conformance shall identify the exact version used and the date through which the statement was assessed.

The Guide at a Glance

Summary. The seven principles, the verification test, the two-layer record, and the task tiers appear below. Each is stated in full in the cited section.

The seven principles (Part II):

- **Responsible control**—This Guide treats generated output as unverified draft material, not as an authoritative or independently reliable source. A qualified professional must understand, integrate, and verify each material machine-assisted contribution before relying on it or sealing the instrument in which it appears (§ 2.1).
- **Verification**—Every material machine-assisted contribution entering sealed work and every material professional communication gated under § 3.G must be verified before reliance. The test addresses the proposition relied upon and, where material, its applicability to project facts, completeness, relevant exceptions, and coordination with dependent information (§ 2.2).
- **Provenance**—Every material machine-assisted assertion relied upon in sealed work or a gated professional communication must trace to an authoritative source. An assertion whose origin cannot be established cannot be shown to have been verified (§ 2.3).
- **Competence**—A professional must know what the tool does reliably, where it fails, and how its failures present (§ 2.4).
- **Confidentiality**—Protected information may be submitted only when the submission is authorized and the specific tool, account, integration, and configuration are approved for the applicable data class (§ 2.5).
- **The record**—The record is created contemporaneously and must be sufficient to reconstruct which content or verification population was machine-assisted, how it was verified, against what authority, by whom, when, and with what disposition (§ 2.6).
- **Boundaries**—AI may assist professional tasks, but it may not serve as the sole or primary basis for a compliance determination, originate a professional review action issued under a seal, or supply sealed content without qualified professional review, integration, and verification (§ 2.7).

Verification is a test with three elements (§ 2.2). The **criterion** states a proposition that can pass or fail. The **authoritative source** is the primary document governing the assertion: the adopted code edition as locally amended, the current standard, the manufacturer’s current data, or the governing contract document. An AI restatement is not a source. The **recorded result** captures the disposition as the test is performed.

The record has two linked layers (§ 4.3; Appendix D). The project and instrument context is recorded once. Each verification entry identifies what was tested and where; the pass criterion; the authoritative source and pinpoint; the method and evidence; and the verifier’s identity, role, relevant qualification, and timestamp. It also records the disposition, approval or status, exceptions, corrections, and reverification.

Every use of machine assistance falls into one of three tiers (§ 3.G; Appendix B). A permitted use produces no material contribution to sealed work and no material professional communication outside the sealed set. A gated use either produces output that may materially enter sealed work or produces a professional communication outside the sealed set that could materially affect life safety, code compliance, building performance, construction cost or time, contractual rights, or a professional determination. It requires verification and the Part IV record. A prohibited use falls within § 2.7. Routine minor confirming correspondence remains permitted. An unclassified use is treated as gated until classified.

Retention (§ 4.4). The firm keeps the verification record as a distinct project-record category under a written schedule that addresses applicable licensing and contractual duties; statutes of limitation and repose, including their triggering events; insurer guidance; privacy and data-minimization duties; and legal holds. This Guide establishes no universal retention period or trigger.

System test (§ 4.2; § 6.5; Appendix D, § D.4). For a named instrument or material professional communication, the firm must be able to retrieve the linked context and verification entries, including status, exceptions, corrections, and reverification, without reconstructing the process from staff memory or file metadata.

Implementation and conformance (Part VI; Appendices A–G). The appendices provide the model firm policy (Appendix A), task classification (Appendix B), verification protocols (Appendix C), record specification (Appendix D), pre-sealing checklist (Appendix E), illustrative conformance, proposal, and renewal language (Appendix F), and glossary (Appendix G). The same substantive requirements apply at every firm size. Conformance is self-declared; the author does not issue certifications (§ 0.7; § 6.4).

Daniel Hutchins, AIA, NCARB