YACHT PROVENANCE TECHNICAL MANUSCRIPT SERIES

Number 2 — October 21, 2026

The Governing Principles and Specifications of Yacht Provenance: A Framework for Evidence-Led, AI-Assisted Historical Research

Technical Manuscript No. 2

Technical Editor

Patrick Fitzhorn, PhD

Principal, Yacht Provenance

Yacht Provenance, Fort Collins, Colorado, USA

Fitzhorn@yachtprovenance.com

Publication

Yacht Provenance Technical Manuscript Series

Prepared: October 21, 2026

AI-Assistance Disclosure

This manuscript was prepared with the assistance of an AI Technical Author commissioned by Yacht Provenance, under the technical editorship and scholarly supervision of Patrick Fitzhorn, PhD, Principal of Yacht Provenance.

Abstract

This paper examines how evidence-led historical-research principles can be translated into evolving, testable engineering requirements for AI-assisted yacht provenance investigation. The physical vessel, documentary origins, inferential warrant, independent examination, correction and research continuity impose distinct but interacting obligations. A six-stage derivation model connects scholarly principle, technical obligation, architectural constraint, implementation alternatives, verification evidence and residual limitation. Two constructed examples distinguish implementation conformance from requirements adequacy: protected interview evidence and a source-custody system that passes its tests while overlooking common-source dependence. Requirements-engineering literature supports iterative clarification, inconsistency management and problem–solution co-evolution; the present contribution is a domain-specific integration, not invention of those general methods. The framework specifies inspectable claims but does not establish operational safeguard effectiveness or historical truth.

1. Introduction and the scholarly problem

The six-stage principle-to-verification sequence introduced in this paper is a proposed analytical method for examining requirements in evidence-led, AI-assisted historical research. It draws on established requirements-engineering concepts but is neither a formally derived general framework nor an independently validated method.

The central engineering question is how a scholarly commitment becomes a requirement whose satisfaction can be examined, without mistaking either a passing test or an elegant implementation for historical truth. The analytical sequence used here is: scholarly principle → technical obligation → architectural constraint → alternative implementation mechanisms → verification evidence → residual limitation. These links are not interchangeable. Each represents a different claim requiring its own justification. The paper develops this derivation model and examines the distinct problem of requirements adequacy: whether the specified obligations are sufficient in the first place.

This focus distinguishes the present contribution from the companion studies. Manuscript No. 1 classifies AI failure mechanisms [1]; Manuscript No. 3 develops the eight-node information model [9]; and Manuscript No. 4 analyzes reversible research and independent examination as a failure-informed workflow [13]. The present paper asks why these architectural and procedural choices follow from historical-research obligations, how competing requirements

can be reconciled during design, and what evidence would justify claims of conformance or adequacy. The contribution is a domain-specific requirements derivation and evaluation method, not a new general theory of engineering specifications.

Yacht Provenance investigates the documentary history and provenance of individual recreational yachts. The fundamental subject is a physical vessel, not a name in a registry, a designer attribution, or a preferred account of its past. A yacht may change names, owners, rig, equipment, and configuration while retaining physical continuity. Different vessels may share a name or be conflated in secondary accounts. A persuasive historical narrative may therefore be internally coherent and still attach genuine evidence to the wrong hull. The governing methodology must prevent documentary descriptions from replacing the object of inquiry.

The institutional proposition is: “The vessel is the subject. The evidence directs the investigation. The investigation determines the findings.” Each clause has a technical consequence. The first requires a distinguishable physical referent; the second requires questions and investigative priorities to remain responsive to evidence; the third requires findings to emerge from reasoned assessment rather than retrospective selection of supporting sources. These commitments do not promise that every question can be resolved. An explicitly inconclusive result may be more defensible than an elegant but unsupported attribution.

AI assistance enlarges the researcher's reach but also increases the risk of persuasive error. Language models can organize dispersed sources, propose hypotheses, compare records, and draft accounts. They can also invent content, omit contradictions, confuse similar vessels, or reproduce a false premise across successive participants. Technical Manuscript No. 1 examines thirteen operational families of AI failure [1]. This paper addresses the complementary methodological problem: what conditions make AI-assisted historical findings traceable, contestable, correctable, and suitable for professional examination?

The contribution proposed here is integration, not invention. Source criticism, archival provenance, evidentiary reasoning, systems engineering, and independent verification are established disciplines [2–8]. The YP framework combines them around an individual physical yacht. It is a research method with specified technical safeguards, not a claim that software can automatically discover historical truth. Technical Manuscript No. 3 provides the companion treatment of the eight-node information architecture [9]. The present paper examines the broader scholarly method and the requirements that surround, motivate, and constrain that architecture.

2. Physical-subject identity and historical continuity

Physical-subject identification and historical attribution must be distinguished. The first asks which actual yacht is under investigation and whether a particular item of evidence belongs to it. The second asks what the evidence supports concerning designer, builder, original name, construction, ownership, and chronology. An investigation may have a clearly distinguishable hull while its designer remains unknown. Conversely, a seemingly certain designer attribution has little meaning if the record concerns a different yacht.

Distinguishing evidence can include observed construction characteristics, dimensions, photographs, identifying marks, location, custody, or documented alterations. A reported name or registration entry is useful but not automatically unique. Where two actual yachts plausibly match the initial description, their records must remain separate rather than being silently merged. A multi-vessel comparison must likewise preserve each physical subject and its evidentiary associations. A newly discovered wrong-hull association is different from an ordinary disagreement about the same yacht's historical attributes.

This distinction matters after research has begun. A revised builder attribution may require reassessment of one claim. Evidence that a drawing belongs to another hull may affect a whole network of claims, review results, and report passages. The appropriate response is to identify and hold the affected dependencies while preserving original sources and unrelated findings. Wholesale deletion loses history; silent continuation preserves error.

The normal research scope concerns the yacht as a whole: materially relevant origins, construction, designer, builder, ownership, use, alterations, and historical identity. A client's particular concern can determine priority, but not the truth of a finding or the entire scholarly scope. Comprehensive coverage is a duty to consider relevant historical domains and explain material omissions; it is not a promise to recover every fact.

3. Documentary sources, custody, and provenance

Historical evidence is not interchangeable with a digital file. A Source is an identifiable authored or documentary origin; a manifestation is a particular scan, edition, web capture, or rendering; an Evidence Item is a locatable observation within that manifestation. A consequential citation should identify the relevant page, folio, drawing region, timestamp, or equivalent coordinate. A mutable URL is a locator rather than an immutable record of what the researcher examined.

Capture metadata, preserved bytes, and cryptographic digests help establish whether a digital manifestation has changed. They do not establish that the historical document was genuine when created or that its content was accurate. Capture integrity, historical authenticity, reliability, relevance, completeness, and permission to disclose are different questions. They should not be collapsed into a single unsupported confidence number. A source can be authentic but irrelevant to the commissioned yacht, or relevant but restricted from public reproduction.

Source dependence is crucial. Several articles repeating one yard ledger are not several independent confirmations. Copies, citations, reproductions, and derivative summaries should be connected to their

informational origins.

Provenance scholarship supplies formal ways to trace derived data [4], while archival research emphasizes documentary context [5]. Neither replaces case-specific judgment about which physical yacht the evidence concerns.

Original digital captures should remain distinguishable from enhancements, transcriptions, translations, annotations, and summaries. Physical originals require appropriate descriptions of custody, condition, identifying features, and handling constraints. Missing pages, suspect alterations, and uncertain transfers should be recorded as limitations, not concealed.

Interviews require parallel care: verbatim recording, contemporary notes, later recollection, hearsay, interviewer interpretation, and AI summary are different evidentiary artifacts. Agreement among witnesses may reflect a common narrative rather than independent corroboration.

The first practical consequence is that evidentiary provenance must be represented at several levels. A source's asserted origin, the chain by which a particular manifestation was obtained, and the investigator's reasons for associating it with the commissioned yacht are separate claims. A scan obtained from a reputable repository may have excellent capture provenance yet reproduce a misidentified photograph. Conversely, a privately held drawing with uncertain custody may contain consequential information requiring cautious examination rather than categorical rejection. The record must preserve these distinctions so that a reviewer can challenge one without automatically rejecting or affirming the others. This is consistent with archival authenticity scholarship [5] and the source–manifestation separation in Manuscript No. 3 [9].

A second consequence concerns evidentiary independence. Counting documents is not equivalent to counting independent observations. A builder attribution printed in a broker listing, copied into an online registry, and repeated in a magazine article may be one claim distributed through three manifestations. The question is whether each account rests on independently obtained information, not whether the words appear in distinct files. When genealogy is uncertain, the uncertainty must be recorded; the accounts must not automatically be counted as corroborating witnesses. Source genealogy is an analytical requirement, not merely a bibliographic convenience.

A third consequence is preservation of negative and limited evidence. An unsuccessful archival search is meaningful only when its scope is reconstructable: repository, collection boundaries, finding aids, search terms, date, access restrictions, and unexamined holdings. This can support a judgment that a reasonable search occurred, but not the universal proposition that the missing document never existed. Likewise, an interview can establish what a witness recalled on a particular date without independently establishing the event remembered. Explicit limitations prevent absence of evidence from silently becoming evidence of absence.

4. From observation to accountable findings

As examined in Technical Manuscript No. 3 [9], the eight conceptual research objects are Entity, Case, Matter, Source, Evidence Item, Assertion, Assessment, and Conclusion. Entity represents a real-world referent; Case bounds an

investigation; Matter records a research question. Source and Evidence Item identify documentary origins and observations. Assertion is a testable proposition. Assessment records the reasoning applied to evidence and alternatives. Conclusion is a case-specific adopted judgment. These distinctions are conceptual requirements, not a prescription for exactly eight physical database tables.

The central evidentiary chain, developed in Manuscript No. 3 [9, Secs. 3–5], is Source → Evidence Item → Assertion → Assessment → Conclusion. A designer's name in a drawing title block is an observation. “Designer

X designed this yacht” is an assertion. An assessment must ask whether the drawing concerns the right hull, whether the title block is original, whether the drawing was executed in construction, whether revisions exist, and whether competing sources contradict the proposed attribution. A conclusion should not be inferred merely because an assertion has documentary support.

Assessments should identify supporting and contrary evidence, source genealogy, assumptions, alternative hypotheses, temporal applicability, and reasons for confidence or doubt. Unsupported numerical probabilities should not be manufactured. An inconclusive outcome is a legitimate finding about the present evidence. A negative search is bounded by the repository, query, access, date, and coverage; it is not automatic proof that no document exists.

Rejected and superseded interpretations should remain traceable. A finding adopted in one case does not become unquestionable global truth in another. The information architecture preserves the reasoning history and permits correction without overwriting original evidence. A database can enforce referential relationships and versioning; it cannot decide whether a faded hull number truly identifies the vessel. Scholarly judgment remains indispensable.

The eight-node scheme should not be mistaken for a complete ontology of maritime history or a proof that any particular database is correct. Manuscript No. 3 compares it with database provenance, CIDOC CRM, temporal data models, engineering product-data traceability, and W3C PROV-DM [9, Secs. 3.9–5]. Its distinctive role is to preserve separations among physical referents, bounded inquiries, durable questions, documentary origins, located observations, propositions, evaluations, and case-specific judgments. The governing methodology additionally requires source acquisition discipline, independent examination, reasoned stopping, publication checks, and correction procedures. An information model can represent these activities and their results; it cannot substitute for their competent execution.

This division of labor is essential. A perfectly linked chain can preserve an erroneous identification, and a database may faithfully retain an assessment that overlooked contrary evidence. Conversely, a scholarly objection may reveal a dependency never encoded in the data model. Structural conformance, source fidelity, interpretive warrant, and effective independent challenge are therefore separate claims requiring separate forms of examination [9, Secs. 5, 8–10].

5. Reversible research and professional sufficiency

Discovery, Resolution, and Completion Assessment form a reversible cycle. Discovery gathers sources, pursues leads, and records negative searches. Resolution tests assertions and competing explanations. Completion Assessment evaluates coverage and the consequences of stopping. New evidence, contextual challenge, or a material omission may return work to an earlier stage. The process is not a one-way sequence that becomes correct merely by reaching its final box.

Each Matter should retain its question, relation to the whole-yacht coverage plan, hypotheses, search history, source leads, unresolved issues, and next action. A question about a rudder alteration may lead to original design plans, photographs, previous-owner testimony, and later repair records. The emerging evidence can change the question itself. A research process organized solely around the original query would be unable to

follow that development.

Contextual challenge is continuous. A material interpretation should be tested against credible alternatives, missing source classes, apparent corroboration of common origin, and unsupported inferential steps. Challenges require reasoned disposition through inquiry, correction, qualification, or documented rejection. This investigative self-challenge is distinct from the later isolated independent examination.

Stopping decisions present an unavoidable tension: the possible archive is unbounded, while a professional investigation has finite time and resources. No universal source count establishes completeness. A defensible completion assessment considers relevant domains, significant negative searches, unresolved leads, plausible further opportunities, evidence limitations, and the consequences of remaining uncertainty. An inconclusive historical account may represent sufficient inquiry; answering one preferred question does not necessarily do so.

The stages are disciplined modes of work rather than rigid chronological gates. Discovery and Resolution interact: evaluating a drawing may reveal a previously unknown yard, creating a new archival lead; an ownership chronology may expose a period of uncertain hull identity, requiring renewed subject examination. Completion Assessment may return the inquiry to Discovery when a material source class has not been considered. Reversibility is not indecision. It prevents procedural progress from substituting for evidentiary sufficiency.

Professional sufficiency depends on the consequences of uncertainty. An unresolved decorative fitting may have little bearing on a vessel's origins, whereas an unresolved hull-number conflict may affect nearly every attribution. Equal research effort is not warranted in both cases. Completion assessment must consider the importance of the Matter, credibility and accessibility of remaining leads, cost of pursuit, and risk that the present account would mislead a reader. These considerations support proportional stopping without pretending that history can be exhaustively searched.

A useful challenge is a record that answers a client's initial designer question while leaving material ownership, refit, and builder leads unexamined. Declaring completion solely because the initial question was answered confuses task satisfaction with whole-yacht coverage. The converse challenge is an investigation that never stops because another source is logically possible. Both outcomes fail the balance between disciplined scope and professionally defensible sufficiency.

6. Independent source-facing examination

Independent review is not synonymous with a second agreeable narrative. Two AI participants can share training biases, contaminated packets, prior answers, or a false governing premise. Their agreement does not establish independent evidence. The relevant safeguards concern what the reviewer could see, which sources were examined, and whether an independent judgment was recorded before exposure to the investigator's conclusions [1,7].

The method separates Phase A and Phase B. Phase A begins with a neutral source-facing packet that excludes the investigator's adopted account and previous review conclusions. The reviewer investigates, records alternative findings and omissions, and seals that work. Phase B then compares the independently recorded result with the investigator's narrative. The order matters: comparison before independent examination can encourage anchoring and suppress genuine disagreement.

Isolation is an access-control and provenance condition, not merely an instruction to ignore earlier answers. Packet versions, permitted sources, access records, and the Phase A seal must be identifiable. A reviewer who can retrieve prohibited narrative material through a shared search index may have a contaminated examination. Material changes to the source universe can require renewed source-facing work; trivial formatting changes should not automatically cause a complete scholarly restart.

Objections should name the challenged proposition or omission, relevant evidence, significance, and disposition. Responses may involve renewed investigation, correction, explicit qualification, or reasoned

rejection. Critical unresolved objections should remain visible rather than disappearing into consensus prose. Even correctly isolated review can miss errors or share common inferential weaknesses. Its effectiveness

requires challenging test cases, observed detection and miss rates, and explicit residual limitations, not merely a completed checklist.

The value of isolated examination can be stated in terms of information flow. Before a reviewer formulates an independent account, exposure to the investigator's conclusions creates an anchoring channel. A nominally separate reviewer may reproduce the investigator's structure even while disagreeing on details. Isolation removes that particular influence; it does not remove shared model training, common source errors, or shared assumptions.

Reviewer independence is therefore layered: packet independence, access independence, independent evidentiary examination, and independence of resulting judgment require separate attention.

Detection must also be distinguished from containment. A reviewer may identify a wrong-vessel association yet fail to prevent the flawed attribution appearing in a prepared report. An objection must connect to affected assertions, assessments, and publication artifacts so that discovery triggers an observable correction path. Independent examination thus depends on versioned evidence and dependency management rather than functioning as a detached editorial opinion. Manuscript No. 1 distinguishes representation, inference, interaction, and verification failures because one safeguard cannot cover every pathway [1].

Isolation can be tested adversarially. A test may plant an investigator's distinctive false attribution in an unintended shared search index; the Phase A environment should not retrieve it. Another may supply two purportedly independent accounts copied from one transcription; their dependence should be discovered or qualified. A third may alter a source packet after a sealed examination; the process should determine whether renewed examination is required. These are proposed test designs, not observed performance results.

7. Reporting, quality control, and correction

Research sufficiency, independent examination, report preparation, quality control, delivery, and client acceptance are different events. A sound historical assessment does not guarantee that a PDF accurately represents it. A technically intact PDF does not prove that the historical conclusion is true. Delivery and acceptance cannot repair an evidentiary defect. These distinctions prevent operational milestones from being mistaken for scholarly validation.

A professional report should identify the physical yacht, present the evidence-supported history, and distinguish observation, assertion, assessment, and conclusion. It should address materially relevant origins, construction, ownership, use, alterations, and historical identity where the evidence permits. Important contradictory evidence, inaccessible sources, negative searches, and unresolved questions should be disclosed. Narrative coherence must not be achieved by silently erasing uncertainty.

The report should be constructed from stable research records rather than reconstructed from a model conversation. Consequential claims must trace to specific assessments and source manifestations. Quotations, captions, names, dates, measurements, and reproduction rights require editorial checks. Quality control should examine the exact artifact version. A changed caption produces a new file, even if the underlying historical interpretation is unchanged; substantive changes may also require renewed scholarly examination.

Corrections after delivery must preserve the issued version and identify what changed. A factual challenge triggers examination of the evidence, not automatic alteration to match a client's preferred account. A changed finding may make dependent assessments, review comparisons, and report passages stale. Previously issued artifacts remain part of the documentary record. Correctability is credible only when the path from earlier to revised findings can be reconstructed.

8. Preservation and research continuity

Research knowledge should survive changes in investigators, AI participants, software, and repositories. It cannot depend on an undocumented predecessor conversation. A durable case record preserves source locations and manifestations, unresolved Matters, negative searches, alternative hypotheses, current and superseded assessments, review objections, and next actions. A successor should reconstruct the research state from inspectable records, not from confident narrative assurance.

Material correction requires dependency analysis. Manuscript No. 3 [9, Sec. 7] develops the associated information-model relationships; the present concern is their role in institutional correction. A revised drawing attribution may affect an Assertion, Assessment, Conclusion, review packet, dossier passage, and released artifact. Known dependencies can be traversed mechanically; unmodeled intellectual dependencies require deliberate examination. Complete automatic discovery of every causal relationship is not guaranteed. The practical obligation is to preserve known relationships, inspect possible hidden effects, and prevent stale downstream material from silently regaining currency.

Originals, derivatives, earlier versions, and current records should remain distinguishable. A superseded artifact is not erased; it is marked as no longer current for its affected purpose. This is especially important for historical work, where the history of interpretation may itself become valuable. Negative searches and abandoned leads should also survive, so successors do not repeatedly pursue failed avenues or mistake silence for absence.

Backup success is not equivalent to recoverability. Restoration checks must recover evidence bytes, provenance relationships, access restrictions, unresolved states, and version history. External operations introduce a different hazard: a timeout after sending a document or creating an accounting effect does not prove that the action failed. Reconciliation before retry avoids accidental duplication. These engineering controls protect research continuity without pretending that

all institutional records belong in one database.

9. Engineering specifications and design evolution

A requirement is not adequately explained by a label such as “traceability” or “independence.” It must identify the protected scholarly property, the situations in which it matters, the observable failure that must be prevented, and the evidence needed to examine conformance. The G/C/M distinction is therefore a relationship of abstraction, not a sequence in which a broad aspiration automatically guarantees a particular design. One outcome can require several constraints; one constraint can be realized through different mechanisms; one mechanism may serve several outcomes. Conversely, a technically successful mechanism may fail to meet the intended scholarly purpose.

An engineering specification can establish a dependable baseline without being exhaustive or immune to ambiguity. Requirements engineering recognizes that clarification, conflict detection, validation, and controlled change continue throughout development [2,3]. Design scholarship describes the co-evolution of problem understanding and candidate solutions [6]. Discovering an unanticipated interaction during implementation does not automatically invalidate the entire specification; it creates an engineering question that must be analyzed and resolved.

Three distinctions help. Governing outcomes (G) state what must be achieved or prevented, such as evidence traceability or uncontaminated review. Architectural constraints (C) establish defined system boundaries. Implementation mechanisms (M) are replaceable ways of satisfying the outcomes and constraints: particular tables, access arrangements, synchronization methods, or backup products. The detail of a mechanism should not be confused with the intellectual necessity of the requirement it serves.

Consider a constructed example. A review workspace appears to isolate Phase A, but a test reveals that a shared search index exposes the investigator's conclusions. The required outcome has failed even if the reviewer promises not to read those results. The engineering response is to examine the exposure, compare

alternative isolation designs, test whether prohibited access is actually prevented, and evaluate cost and recovery consequences. Replacing the search mechanism may satisfy the same requirement; simply abandoning isolation would not.

An inconsistency should be recorded with its affected provisions, consequences, alternatives, and testable resolution. Work unaffected by a bounded problem need not automatically stop. Equally, unresolved defects that undermine a necessary research safeguard cannot be concealed by proceeding to a later workflow state. This is how an evolving engineering design remains compatible with a stable scholarly purpose. A technical control is credible only when its operation can be observed and challenged.

9.1 Specification quality and the limits of completeness

Engineering literature distinguishes an intended outcome from its requirement wording and from the solution selected to satisfy it. The historical IEEE 830 recommended practice identifies correctness, unambiguity, completeness, consistency, verifiability, modifiability, and traceability as desirable characteristics of a software requirements specification [11]. ISO/IEC/IEEE 29148:2018 places requirements processes and information products throughout the systems and software lifecycle and explicitly discusses iterative and recursive application [10]. These are general engineering principles; neither establishes that a particular Yacht Provenance safeguard works. The full licensed standard was not examined for this paper; statements here are confined to its published scope.

Completeness has several meanings. Syntactic completeness asks whether a requirement includes its necessary conditions and terms. Coverage completeness asks whether the set addresses identified needs and hazards. Design completeness would imply that every implementation choice and interaction was already settled. The first two can be examined and improved; the third is not a realistic premise for a complex evolving system. A specification can be a dependable baseline while leaving technical alternatives open. Nor does a complete identifier inventory demonstrate collective adequacy for the risks of historical research.

A requirement should permit an observable challenge. 'Preserve documentary provenance' is an important objective but an inadequate stand-alone test. A testable elaboration requires each consequential citation to identify its source, preserved manifestation or access limitation, and reproducible locator. A reviewer can then challenge an untraceable quotation. Likewise, 'ensure independent review' requires elaboration into packet isolation, source-facing examination, objection records, and tests of actual access. INCOSE guidance emphasizes characteristics including necessity, clarity, feasibility, consistency, and verifiability [12].

9.2 Inconsistency as an engineering condition

Finkelstein and colleagues examine inconsistency handling in multi-perspective specifications and explain why demanding instantaneous absolute consistency may be impractical and may discard valuable information [3]. Their work does not suggest that contradictions are harmless; rather, inconsistencies should be identified, localized, reasoned about, and managed. For Yacht Provenance, two requirements may appear to compete: retaining interview material for reproducibility while restricting sensitive access; preserving publication history while preventing stale conclusions from appearing current; or maintaining research continuity while isolating a reviewer. The engineering task is to find a design satisfying the relevant outcomes or expose a genuine incompatibility for controlled resolution.

The type of inconsistency matters. Two requirements may be logically incompatible; two mechanisms may merely compete for resources; a requirement may lack a temporal condition; or a test may misinterpret the obligation it examines. These demand different remedies. Ambiguity can be clarified by defining circumstances. Mechanism conflict can be resolved through design substitution. Genuine conflict among required outcomes demands explicit trade-off analysis, not silent deletion. An unresolved issue should remain visible in the engineering record and affect work whose validity depends on it, rather than automatically stopping unrelated activity.

9.3 Problem–solution co-evolution and controlled change

Dorst and Cross describe a co-evolutionary relation between problem formulation and candidate solutions, based on observations of experienced designers [6]. Designing source custody may reveal that a single concept of 'document' fails to distinguish an authored original, revised edition, cropped reproduction, and downloaded file. Designing reviewer isolation may reveal a shared-index dependency invisible when the original access boundary was described. Solution work improves understanding of the problem. These discoveries can refine implementation and clarify requirements without negating the scholarly objective.

A controlled change record should identify the triggering observation, affected obligation, alternatives, consequences, verification method, and disposition. It should distinguish changing the intended outcome from replacing the mechanism used to satisfy it. The crucial test is whether the revised system meets the substantive obligation under challenging conditions. Changing a checklist from FAIL to PASS without correcting the exposed defect is not engineering resolution.

9.4 Worked interaction of requirements

Consider a constructed case in which a private interview recording supports a disputed ownership transfer. Evidence preservation favors retaining original audio and a precise transcript locator. Privacy and use restrictions constrain access and quotation. Independent review requires source-facing examination, and continuity requires a later researcher to recover the material's evidentiary status. Unrestricted duplication sacrifices confidentiality; deletion sacrifices reproducibility. A more adequate design distinguishes custody from publication, preserves the original under controlled access, provides permitted review access with recorded limitations, and publishes only defensible quotations or paraphrases. Tests should examine unauthorized access, locator reproducibility, packet integrity, and recovery after interruption.

This example illustrates why specification quality cannot be judged solely by reading individual sentences. System behavior depends on requirements interacting at their boundaries. Traceability should connect each substantive outcome to design mechanisms, verification challenges, and evidence of observed behavior. A passing document review establishes neither a passing implementation test nor historical truth.

9.5 Complete derivation: confidential interview evidence

The six-stage derivation model is proposed here as a domain-specific analytical method, not as a formally derived or independently validated engineering framework. Its stages synthesize established requirements-engineering ideas to expose the relationship between scholarly purposes, technical constraints, implementation choices, observable tests, and remaining uncertainty. The worked derivation illustrates applicability; it does not establish completeness, predictive performance, or general effectiveness.

Consider the constructed interview example as a complete requirements derivation. The scholarly principle is that consequential historical findings remain re-examinable while evidence is handled in accordance with its access and reproduction limitations. The derived technical obligation is twofold: preserve the original recording and a stable, time-coded locator for every consequential quoted or paraphrased observation; and prevent unauthorized access or public disclosure. Neither obligation can simply override the other. The historical question remains examinable even when the public report cannot reproduce the recording.

The architectural constraints are separation of source custody from publication copies; explicit manifestation identity; a durable relationship from the cited Evidence Item to the protected recording; enforceable access boundaries for source-facing examination; and recorded rights and access limitations. Two possible mechanisms illustrate substitution. Mechanism A uses a restricted repository with role-scoped access and a separately generated redacted review packet. Mechanism B uses encrypted evidence storage and an isolated review workspace that grants time-limited, logged access to the protected recording. Either may be suitable if it preserves the same obligations; neither is made valid by its product name or architecture diagram alone.

Verification must be partitioned. A locator test requests the cited segment and compares it to the retained recording. An access test attempts retrieval from an unauthorized workspace, including through shared indexes, caches and derivative files. A continuity test replaces the research participant and reconstructs the evidentiary relationship from durable records. A disclosure test examines the exact public artifact for prohibited material. Failure conditions include an unresolvable citation, exposure of restricted audio, loss of the interview context after succession, or a published quotation unsupported by the retained segment. The tests must include negative controls so that ordinary permitted access is not incorrectly blocked.

Even successful execution would establish only specified conformance under the tested conditions. It would not establish that the interviewee remembered accurately, that the witness was firsthand, or that the recording was historically authentic. Those remain questions of source criticism and historical inference. This separation between operational verification and evidentiary warrant is central to the proposed derivation model.

9.6 A conformance pass that exposes inadequate requirements

A second constructed case shows why verifying an implementation is not the same as validating its requirements. Suppose the written source-custody requirements demand preserved captures, timestamps, digests and precise citations, but say nothing about relationships among copies and derivative publications. An implementation satisfies every specified custody check. Three newspaper articles, each perfectly captured and cited, nevertheless reproduce one erroneous yard ledger. The system records three distinct files and the investigator treats them as three independent confirmations of the builder. The implementation conforms to the stated requirements while the requirements fail to address a consequential form of documentary dependence.

The inadequacy is discovered by challenging the scholarly principle rather than retesting file preservation. The principle is evidentiary independence: corroborative weight must reflect informational origins, not document count. A revised technical obligation therefore requires known and suspected derivation relationships to be representable and inspectable, including explicit uncertainty when genealogy is unresolved. Candidate mechanisms include typed source-to-source links, an evidence-genealogy graph, or a source-dependence register with queryable relationships. A

discriminating validation fixture supplies several publications with a common origin and tests whether the research record exposes dependence and prevents unqualified independent-corroboration claims.

The change does not abandon the principle or demonstrate that every unknown copying relationship can be detected. It improves the requirements set after an adequacy failure was observed. This is problem–solution co-evolution in a specific historical-research setting [2,3,6]: implementation and evaluation reveal a missing requirement, which is then formulated, traced to its scholarly purpose, and challenged again. Requirements adequacy remains open to further evidence even when each current requirement is individually clear and verifiable.

10. Integrated classes of technical requirements

10.0 The principle-to-verification mapping

The following mappings illustrate the proposed six-stage analytical method. They are reasoned examples of requirements derivation and possible tests, not empirical validation of the method or evidence that the corresponding safeguards have passed implementation testing.

The model is most useful as a many-to-many trace rather than a linear compliance checklist. Subject identity yields constraints on distinguishing physical referents and testing wrong-hull associations. Documentary provenance yields custody, manifestation, locator and genealogy obligations. Independent examination yields packet and access separation requirements. Correction yields version and dependency-invalidation obligations. Continuity yields recoverable records and restoration challenges. A requirement may support more than one principle, and a single scholarly principle may demand several independent checks. The matrix below presents representative derivations, not an exhaustive implementation inventory.

Scholarly principle	**Technical obligation / constraint**	**Discriminating challenge**	**Residual limit**
Correct physical subject	Separate candidate hulls; preserve contested source links	Two same-name vessels; reject automatic merge	Unknown identifying evidence
Evidentiary independence	Record known/suspected source derivation	Three copies of one ledger; reject triple corroboration	Undetected common origins
Independent examination	Isolated neutral packet and sealed source-facing phase	Inject prior narrative via shared index	Shared model assumptions
Correctable findings	Versioned dependencies and scoped invalidation	Change wrong-hull association after report	Unmodeled intellectual links
Continuity	Recoverable source, matter and revision state	Replace participant; reconstruct unresolved dispute	Unrecorded context

This mapping also exposes gaps. If a scholarly principle has no discriminating verification challenge, it is not yet operationally examinable. If a test has no traced scholarly purpose, it may reward a convenient metric rather than protect the research. If a mechanism has no identified constraint, it may be an accidental technical commitment. These are diagnostic questions for requirements validation, separate from whether the selected implementation passes its tests.

The substantive requirements form interdependent classes rather than a series of independent virtues. Subject identity keeps evidence associated with the correct yacht. Scope and coverage prevent inquiry from being reduced to a selected answer. Source custody and genealogy preserve what was examined and whether purported corroboration is independent. Evidentiary object distinctions prevent observations from being promoted directly to findings. Reversible research and contextual challenge permit questions to change as evidence emerges.

Independent source-facing examination challenges the investigation without prior-answer contamination. Objection records preserve disagreement. Reporting requirements connect the public account to stable evidence and disclose unresolved limitations. Exact-artifact checks prevent a verified research state from being confused with an unexamined document revision. Correction and dependency analysis identify affected findings and retain superseded records. Continuity and restoration requirements permit work to survive interrupted systems or changed participants.

These classes can be connected to identifiable requirement families: YP-SUB and YP-RES concern physical subjects and inquiry; YP-EVD, YP-CUS, and YP-INT concern evidence and custody; YP-ASM and YP-DATA concern accountable inference and conceptual distinctions; YP-INV concerns research execution; YP-TR and YP-VAL concern transitions and correction; YP-RPT concerns reporting; YP-SEC and YP-BCP concern isolation, integrity, and continuity. The family identifiers provide internal traceability; the public argument rests on the

obligations themselves rather than a reproduced clause catalog.

Representative verification challenges make the requirements falsifiable. Two copies of one ledger should not count as independent corroboration. A later wrong-hull discovery should hold affected claims without destroying unrelated evidence. A reviewer exposed to an investigator's account before sealing Phase A should not be treated as isolated. A changed release file should not inherit an earlier exact-byte quality result. A replacement researcher should recover open questions and contrary evidence without relying on predecessor memory. These are specified challenges, not claims that every control has already passed operational testing.

The technical point is their interaction. A perfectly preserved source can still be attributed to the wrong hull; a sound assessment can still be omitted from a successor packet; an independent review can still be invalidated by changed evidence; a corrected conclusion can still appear in a stale report. Effective research integrity requires the links among controls to remain inspectable.

10.1 The requirement classes form an assurance chain

The research requirements can be understood as a sequence of separately supportable claims. Subject identity establishes which physical yacht is studied. Evidence custody establishes what material was obtained and where an observation appears. Inference establishes how the observation bears on a proposition. Independent examination establishes whether alternatives and omissions were challenged. Reporting establishes whether the public account faithfully represents the research. Correction and continuity establish whether new evidence can be traced through dependent artifacts. A defect in an earlier class may undermine later classes without being repaired by flawless execution of later procedures.

For example, a correctly transcribed drawing attributed to the wrong hull passes a transcription check but fails subject association. An elegant assessment can therefore be wrong despite impeccable citation format. A report faithfully reproducing that assessment may be editorially accurate yet historically misleading. One overall quality score can conceal these distinct failures. Separate evidence is needed for separate claims, with explicit disclosure of unresolved limitations.

10.2 Representative verification challenges

The following examples specify observable expectations, not completed tests. Challenge A supplies two yachts sharing a historical name and a photograph belonging to only one: expected behavior is separate subject identities and a qualified documentary link; automatic merging is failure. Challenge B supplies three articles copied from one ledger: expected behavior is visible dependence; counting three independent confirmations is failure. Challenge C supplies a design drawing and later photograph showing different rudders: expected behavior is time-qualified interpretation without an invented alteration date. Challenge D exposes the Phase A reviewer to a prohibited earlier narrative: expected behavior is detection of isolation failure and renewed examination where consequential. Challenge E changes an attribution after publication: expected behavior is preservation of the old release, identification of stale dependencies, and examination of the revised account. Challenge F interrupts a system after an uncertain external send: expected behavior is effect reconciliation before retry, not an assumption that timeout proves failure.

These challenges make the integrated framework falsifiable at particular interfaces. Results cannot be inferred from a requirement's existence or from conceptual adequacy of the eight-node model. Tests must run against actual systems with recorded inputs, expected outcomes, observations, deviations, and limits. Where no independent historical answer key exists, evaluation should distinguish preservation of uncertainty and provenance from arrival at a historically correct but currently unknowable answer.

10.3 Proportionality and the burden of rigor

Every additional distinction has a cost. Preserving manifestations, competing assertions, review packets, dependency records, and superseded releases increases storage and editorial work. Excessive process can crowd out investigation, especially in a small institution. The engineering response is to distinguish indispensable outcomes from mechanisms that can be simplified, and to ask which controls have meaningful diagnostic value. A compact safeguard that reliably exposes a wrong-hull premise can be more useful than an elaborate checklist rewarding procedural completion. The framework should remain open to evidence about its own operating costs and effectiveness.

11. Limits and evaluation of effectiveness

The distinction between requirements validation and implementation verification can be stated operationally. Verification asks whether the implemented system meets specified obligations under identified conditions. Validation asks whether those obligations, taken together, are fit for the scholarly risks the institution actually faces. A source-custody system can pass its preservation tests while missing source genealogy; an isolation system can block direct access while leaking through a shared index; a complete workflow can preserve a mistaken physical-subject premise. A credible evaluation therefore tests both the adequacy of the requirement set and the conformance of its realization.

Procedural rigor cannot manufacture missing historical evidence. A ledger may have been destroyed, a witness may be mistaken, and a designer's drawing may never have been executed. Source criticism can qualify these limitations but cannot eliminate them. Independent examination may fail to notice the same contradiction as the investigator. A cryptographic digest can establish file integrity without establishing authenticity or truth. An orderly database can preserve a false claim as faithfully as a true one.

Three different claims must not be conflated. Catalog coverage asks whether the relevant obligations have been described. Requirements adequacy asks whether those obligations address the institution's actual research risks. Operational effectiveness asks whether implemented procedures reliably detect, prevent, contain, or

correct failures. The first requires traceability; the second requires substantive critical examination; the third requires evidence from tests and actual use. A manuscript explaining a framework establishes neither a universally complete specification nor a proven accuracy improvement.

Evaluation should distinguish compliance with workflow from correctness of historical outcomes. Challenging cases need independent reference evidence where available, known injected defects, tests of source dependence and physical identity, measurements of missed errors and false alarms, and checks of restoration after interruption. Some historical questions have no known ground truth; evaluation must disclose that limit rather than invent an answer key. The thirteen-family analysis in Manuscript No. 1 provides a useful failure-oriented companion to such testing [1].

There are also proportionality limits. A one-person research institution cannot adopt every safeguard at unlimited complexity or cost. A defensible design should distinguish indispensable scholarly outcomes from incidental technical mechanisms, preserve observable controls, and evaluate the burden of recovery and continuing operation. More process is not automatically more accuracy. The test is whether the process makes consequential errors more detectable and correction more reliable without overwhelming the inquiry it is intended to support.

An evaluation program should distinguish three levels of evidence. Conceptual analysis can show that a framework has a place to represent contradiction, an unexamined lead, or a superseded conclusion. Controlled fixtures can show whether an implementation preserves those distinctions under specified challenges. Experience in actual investigations can expose failures not anticipated by fixtures. None alone proves universal historical accuracy. A scholarly report should identify which level supports each performance claim.

Transfer from general engineering literature has limits. Requirements scholarship provides concepts of consistency, traceability, verification, and change control; it does not determine whether a maritime attribution is true. Historical evidence may be incomplete, temporally ambiguous, or inaccessible to repeat testing. General AI research identifies hazards but does not supply institution-specific error rates. The framework is an argued response to those hazards. Claims of improved detection, reduced false attribution, or reliable recovery require direct evaluation rather than analogy.

The framework also faces an epistemic limit: a meticulously preserved reasoning chain can make an error auditable without making it obvious. A reviewer may share an investigator's mistaken belief that a famous yacht name denotes one hull, or never encounter the contrary record. Engineering creates opportunities for correction; it does not remove the need for competent historical judgment. Mature evaluation should record misses and false alarms, test difficult identity and dependence cases, and disclose uncertainty rather than converting procedural conformity into truth.

12. Conclusions

The distinctive engineering result is a traceable method of translating historical-research principles into testable obligations while keeping the requirements themselves open to critical examination. The chain from principle through obligation, constraint, mechanism and observed test result makes implementation claims inspectable; the residual-limitation step prevents those claims from becoming assertions of historical truth. The conformance-with-inadequate-requirements counterexample shows why a technically correct implementation can still fail the research purpose. Accordingly, requirements development and evaluation must be iterative, with source criticism and domain-specific failure challenges informing each other.

Yacht Provenance's governing methodology begins with a particular physical vessel, not a predetermined historical attribution. Evidence directs the questions and constrains the conclusions. Documentary sources, their manifestations, observed content, propositions, assessments, and adopted findings must remain distinguishable. Whole-yacht coverage, reversible inquiry, source genealogy, and recorded uncertainty provide the basis for a defensible historical account.

Independent source-facing examination, durable objection records, exact-artifact quality control, correction lineage, and recoverable research state extend that discipline across participants and time. Engineering

specifications translate the scholarly obligations into observable technical requirements while permitting implementation methods to evolve. Their acceptance as design baselines does not imply that ambiguity or conflict can never arise; it requires such discoveries to be addressed through disciplined analysis rather than concealed or improvised resolution.

The resulting framework is an integration of established scholarly and engineering obligations for AI-assisted yacht history. It does not claim to invent source criticism or to prove that every safeguard works. Its enduring contribution will depend on the quality of the historical investigations it supports, the openness of its findings to challenge, and the institution's capacity to preserve and correct its own record.

References

[1] Yacht Provenance, 2026, “Operational and Scholarly AI Failure Mechanisms: From Four Initial Classes to Thirteen Operational Families,” Yacht Provenance Technical Manuscript Series, No. 1.

[2] Nuseibeh, B., and Easterbrook, S., 2000, “Requirements Engineering: A Roadmap,” Proceedings of the Conference on the Future of Software Engineering, pp. 35–46.

[3] Finkelstein, A., Gabbay, D., Hunter, A., Kramer, J., and Nuseibeh, B., 1994, “Inconsistency Handling in Multi-Perspective Specifications,” IEEE Transactions on Software Engineering, 20(8), pp. 569–578.

[4] Buneman, P., Khanna, S., and Tan, W.-C., 2001, “Why and Where: A Characterization of Data Provenance,” International Conference on Database Theory, pp. 316–330. doi:10.1007/3-540-44503-X_20.

[5] Pacheco, A., Da Silva, C. G., and De Freitas, M. C. V., 2023, “A Metadata Model for Authenticity in Digital Archival Descriptions,” Archival Science, 23, pp. 629–673.

[6] Dorst, K., and Cross, N., 2001, “Creativity in the Design Process: Co-evolution of Problem–Solution,” Design Studies, 22(5), pp. 425–437.

[7] Parasuraman, R., and Riley, V., 1997, “Humans and Automation: Use, Misuse, Disuse, Abuse,” Human Factors, 39(2), pp. 230–253.

[8] Liu, N. F., et al., 2024, “Lost in the Middle: How Language Models Use Long Contexts,” Transactions of the Association for Computational Linguistics, 12, pp. 157–173.

[9] Yacht Provenance, 2026, “An Eight-Node Information Architecture for Yacht History and Provenance Research,” Yacht Provenance Technical Manuscript Series, No. 3, October 21. Technical Editor: Patrick Fitzhorn, PhD.

[10] ISO/IEC/IEEE, 2018, ISO/IEC/IEEE 29148:2018, Systems and Software Engineering—Life Cycle Processes—Requirements Engineering. IEEE/ISO, doi:10.1109/IEEESTD.2018.8559686. (Technical standard; published scope consulted.)

[11] IEEE, 1998, IEEE Std 830-1998, IEEE Recommended Practice for Software Requirements Specifications. IEEE. (Historical recommended practice, cited for specification-quality characteristics, not as a current governing standard.)

[12] INCOSE Requirements Working Group, 2023, Guide to Writing Requirements, Version 4, overview. International Council on Systems Engineering. (Professional guidance, not refereed research.)

[13] Yacht Provenance, 2026, “Failure-Informed Architecture for Reversible AI-Assisted Research: An Integrated Research and Review Framework,” Yacht Provenance Technical Manuscript Series, No. 4, October 21. Technical Editor: Patrick Fitzhorn, PhD.

Copyright © 2026 Yacht Provenance LLC. All rights reserved.

No part of this technical manuscript may be reproduced, distributed, transmitted, translated, or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without the prior written permission of the copyright owner, except as permitted by applicable copyright law.

For permission requests, write to the copyright owner at the address below:

Fitzhorn@yachtprovenance.com