Why Single-Chat AI Research Fails to Produce Comprehensive Yacht Histories

Evidence from Three Investigations and the Case for a Database-Centered Research Architecture

Technical Editor: Patrick Fitzhorn, PhD
Principal, Yacht Provenance
Fort Collins, Colorado, USA
16 August 2026

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

Abstract

Can a single AI chat reliably produce a comprehensive yacht history and provenance dossier? Two AI agents with that task investigated *Cangarda* and *Moonbeam IV* and did not yield accepted dossiers. Based on those failures, a *Ticonderoga* institutional-learning experiment used cross-conversation persistence of a relational evidence archive, independent detection of database and manuscript defects, and a Conditional Pass for bounded synthesis. It did not complete a full yacht dossier but showed significant promise. The three cases strongly suggest that persistence, organization, evidentiary relationships and retrieval are difficult for single-chat AI agents, while pointing to the need for the necessary, foundational capabilities of an AI-assisted historical research engine. A database-centered architecture with replaceable research agents, a distinct retrieval-and-writing agent and independent review is proposed but has not been validated end to end.

1. The historical commissions

Cangarda (subsequently *Magedoma*) is a large American steel steam yacht built in 1901 by Pusey & Jones at Wilmington, Delaware, to a design attributed to Henry C. Wintringham. Published particulars describe 138 ft (42.06 m) overall, about 126 ft (38.4 m) on deck, a beam of approximately 17 ft 6 in–18 ft (5.33–5.49 m), and draft around 7 ft 8–9 in (2.34–2.36 m). The Naval Marine Archive gives 125 ft 10 in length on deck and 103 ft 3 in waterline length; Sandeman lists 105 ft waterline. These are source-specific values, not a reconciled original survey [NAV; SAN]. A slender steel hull, long wooden deckhouses, funnel and auxiliary spars distinguish her as a gentleman's yacht. Propulsion is by a Sullivan triple-expansion reciprocating steam engine with associated auxiliary machinery; subsequent boiler replacements and extensive steel-hull reconstruction complicate the question of original material survival [P06; P08].

The historical commission extended well beyond identifying the vessel. It required reconstruction of her early Canfield ownership, the distinct reported roles of Willis T. Hanson and George Fulford in a 1904 transaction, later Thousand Islands cruising, the disputed 1927 royal-visit account, Canadian naval use during the Second World War, postwar custody, deterioration and sinking, salvage, major restoration, and later movements. Registered ownership, beneficial interest, physical custody, and use by a naval organization could not be treated as equivalent. Nor could original engines and joinery be conflated with reconstructed hull plating. Conflicting historical particulars, including alternative length definitions, were precisely the kind of evidence requiring critical analysis rather than a smooth unsupported narrative [C18; CT; P06; P08].



Figure 1. *Cangarda* under steam, illustrating the steel hull, long deckhouses, funnel and masts. Reproduced from Robert G. McNeil and Jeffrey Rutherford, "The Restoration of *CANGARDA*," Classic Yacht Symposium (2006), printed p. 1, figure on title page [P06]. Exposure date and reproduction rights not independently verified.

Moonbeam IV is a large Scottish-built sailing yacht designed by William Fife III and built by William Fife & Son at Fairlie for Charles Plumptre Johnson. Sources associate her construction with 1913–1914 and completion or entry into service in 1920. She is a long-overhang, gaff-rigged cutter, with teak planking on steel frames, a long bowsprit, and the distinctive sail number 8. Published overall lengths vary materially: the French maritime heritage register gives 33.16 m overall and 29.09 m hull length, the Comité International de la Méditerranée (CIM) gives 30.16 m overall, and JFA-related accounts give approximately 35 m overall. Reported beam is about 5.05–5.10 m and draft approximately 3.80–3.90 m. The variation is not resolved here; definitions of length, rig projections, date and source must be checked before selecting an authoritative dimension [PMF; CIM; JFA22].



Figure 2. *Moonbeam IV* under sail, showing the cutter rig, long bowsprit and hull overhangs. Image supplied by the Principal; source credit: Boat International, "*Moonbeam IV*," Superyacht Directory, https://www.boatinternational.com/yachts/the-superyacht-directory/moonbeam-iv--66139 [BOAT]. Photographer, exposure date and reproduction rights have not been independently established.

The yacht's provenance raises a different set of challenges from those of *Cangarda*. The investigator needed to establish the precise relationship among the commissioning date, wartime construction interruption, launch and registration; verify reported King's Cup results rather than reproduce racing lore; trace successive owners and flags, including Sutton, Prince Rainier III, disputed Scotti references and the Murray period; and identify the dates and nature of rig changes. The restoration undertaken in Myanmar, and later French work at JFA Yachts and associated yards, required

an item-level distinction between original Fife-era structure and replaced frames, planking, spars, superstructures and interiors. Records about the separate *Moonbeam III* and the yacht *Sunshine* had to be kept distinct from *Moonbeam IV* [M101; M104; JFA22; PMF].

2. Forensic reconstruction of the draft sequence

The numbered files discussed below are not separate commissions or independently authored studies. They are successive research products issued while the same comprehensive yacht-history assignment remained in force. The labels REF (reference and source-register work), CER (continuous-execution research iterations), and MASTER (attempted consolidated dossiers) were names used within the *Cangarda* research archive. A version number identifies a production stage; it does not certify that the version is better or more complete. PDF page counts describe the size of the readable deliverable, not the amount of verified historical knowledge. This distinction is essential to interpreting the observed contractions.

The forensic examination compared extracted PDF text, page counts, recurring subject terms, surviving tables and the structured records bundled with research packages. Some adjacent versions were compared page for page. A separate manifest records hashes and text lengths, permitting a reviewer to identify the actual artifacts rather than rely on the chats' recollections [AUD; MAN]. A shorter report can be an intentionally narrower research memorandum. Its existence is not proof that prior evidence vanished. The relevant question is whether the unchanged commission's comprehensive, cumulative master remained intact and whether a new report incorporated what was already known.

The earliest *Cangarda* reference stages expanded from four pages in REF v1 to fourteen in v2 and eighteen in v3. REF v4 then fell to thirteen pages. The investigation next expanded to a twenty-five-page illustrated dossier (REF v5) and a fifty-page REF v6, which retained the original twenty-five pages and appended twenty-five more. REF v7 was a fourteen-page source-critical treatment centered on the supplied restoration literature. It provided more focused source examination, but no longer reproduced the earlier report's breadth of ownership, movements and related historical material as one readable account.

The pattern recurred under the CER designation. CER v8 contained sixteen pages; v9 reached twenty; v10 retained those twenty pages and added eight. CER v11 then contracted to eleven pages. Its extracted text lacked the names 'Hanson' and the date '1927', even though the 1904 transaction and the royal-cruise question had appeared in other research products. Those topics later reappeared. CER v12 expanded to twenty-six pages, v13 contracted to seventeen, and v14 reached twenty-one. The MASTER sequence then grew through twenty-seven, thirty-five and forty-one pages, followed by the forty-six-page v18 and a sixteen-page v19 selected critical edition [AUD; CT].

These repeated contractions are not equivalent events. REF v7, CER v11 and v19 were presented as source-focused or selected editions. They may contain legitimate corrections and improved qualifications. But they were not replacements for the complete history that had been continuously requested. The significant failure is the lack of a stable controlling manuscript that preserved and reconciled the historical coverage already developed while incorporating the improved source criticism. A local improvement in fidelity was repeatedly accompanied by a loss of breadth in the next deliverable.

CSV means comma-separated values, a plain-text format for structured tables. During the investigation, the chat created research registers: chronology files listed dated events; ownership files listed reported owners or custodians and periods; voyage files recorded passages, movements or operations; evidence and claim files linked assertions to reported sources. A 'row' is one table entry, not necessarily one verified fact. The number of rows indicates the amount of recorded material, not its accuracy or independence.

Early REF packages preserved a thirteen-row chronology and seven-row ownership register. REF v5 and v6 included a fifty-five-row chronology. Later CER and MASTER packages repeatedly retained twenty-three chronology entries, thirteen ownership entries and seventeen voyage entries. In several packages these records were byte-identical even while the prose PDF changed substantially [AUD; CT]. Their survival is critical evidence against the simplistic explanation that context compression erased every earlier discovery. At least some prior information remained available in external files.

Yet preserved data are useful only if the chat systematically consults them, updates them and checks each new narrative against them. The repeated omission of developed subjects from successor PDFs despite surviving registers shows a failure to make the registers authoritative. The registers themselves were not independently verified against original archives, and their counts cannot be used as a completeness score. The later thirty-three-page reconciled timeline, assembled from multiple version lineages, illustrates the burden of reconstructing a cumulative record after the main drafting process failed to maintain one [CT].

The *Moonbeam IV* sequence followed a different trajectory. Early three- and four-page investigations offered summaries and archival leads. More developed fifteen- and seventeen-page reports eventually separated chronology, owners, voyages, rig and technical changes, visual evidence, bibliography and unresolved questions. But narrower five-, six- and nine-page supplements and a three-page evidence memorandum also appeared under the continuing comprehensive commission. Some numbered exports were textual duplicates. The exact order of a few generically named files cannot be proved without the complete execution log; their differing scopes can be compared directly [AUD].

A twenty-five-page integrated report eventually became the unchanged front section of a 101-page volume. The seventy-six additional pages consisted predominantly of individual event worksheets, source dossiers and open-issue worksheets. The subsequent 104-page version retained all 101 earlier pages unchanged and appended three pages; the 109-page version retained all 104 pages unchanged and appended five more. Page-for-page text comparison establishes this inheritance. It does not establish that the underlying historical research improved by eighty-four pages.

The most consequential late addition concerned a possible 2023 collision and mast damage, together with registration information and more recent racing evidence. The casualty was explicitly unverified, so it should not have been presented as a proven mast replacement. Nevertheless, it should have entered the controlling chronology and rig discussion as a qualified lead with explicit implications and required sources. Instead the original sections remained unchanged while the new information sat in an appendix [M101; M104; M109]. The report grew, but its principal historical account became stale relative to its own evidence.

Table 1 provides an auditable transition-level summary. “Omitted” denotes absence from the successor readable PDF, not proven deletion from ZIPs or registers. “Appended” denotes addition without modification of predecessor page text. Exact timestamps and complete conversational causation are not recoverable from the PDF artifacts alone.

Pair	Pages	Observable change	Interpretation
C REF v5→v6	25→50	First 25 pages retained; 25 appended	Append-only; no replacement of base
C REF v6→v7	50→14	Broad report replaced by source-critical extract	Readable coverage reduced; archive loss unproved
C CER v9→v10	20→28	First 20 pages retained; eight appended	Append-only
C CER v10→v11	28→11	Hanson and 1927 absent from successor text	Readable-topic omission; later recovery
C CER v12→v13	26→17	Separate domains recast as compact report	Reduced coverage; some source qualifications improved
C v17→v18	41→46	First 41 pages retained; five appended	Append-only
C v18→v19	46→16	Selected critical edition	Non-equivalent document type; not storage deletion
M 25→101	25→101	Main 25 pages unchanged; 76 worksheets added	Length expansion without main-text revision
M 101→104	101→104	All 101 pages unchanged; three added	Unintegrated source addendum
M 104→109	104→109	All 104 pages unchanged; five added	New casualty/race leads not integrated

Together, the sequences show non-monotonic progress: a new draft did not reliably preserve the coverage, qualifications and newly discovered evidence of its predecessors. **The commission never changed, and repeated reminders did not reliably alter the next product.** The artifacts demonstrate failure of cumulative document production and research governance. They do not expose the model's active context, the exact timing of context compression, or the causal contribution of each prompt. Context drift, selective retrieval, recent-task dominance and ineffective use of external registers remain plausible explanations rather than independently observed internal events.

3. Thirteen operational failure families: detailed assessment

The following thirteen categories are not new classifications invented for this paper. They were defined in Yacht Provenance Technical Manuscript 'Operational and Scholarly AI Failure Mechanisms: From Four Initial Classes to Thirteen Operational Families,' 18 August 2026 [1]. That manuscript derived the operational families from four broad lenses addressing individual-chat behavior, governing frames, interactions and institutional error—and a targeted examination of AI, systems, human-factors and provenance literature. The families overlap: one faulty assertion may involve misrepresentation, unwarranted inference, failed retention and inadequate verification at different stages. Here they are applied retrospectively to observed yacht-research artifacts. A category is not counted as demonstrated merely because it appears in the taxonomy.

For this retrospective classification, “demonstrated” means a defect directly established in an examined artifact or paired comparison; “strongly indicated” means supported by multiple artifacts but without full causal traces; “plausible” means consistent with the documents but not uniquely established; and “not demonstrated” means the available record cannot support the proposed mechanism. A defective output must not be equated automatically with an internal model mechanism. The matrix below separates observations, alternatives and limits.

Family	Artifact signature	Classification	Alternative / limit
F01	Source identity conflicts and erroneous Laurier attribution	Demonstrated source-error exposure	AI-origin fabrication not established
F02	Hanson/Fulford; naval transfer; Moonbeam ownership alternatives	Strongly indicated warrant limits	Some qualifications were preserved
F03	Subjects absent in later readable reports; stale main text	Demonstrated product discontinuity	Context compression not logged
F04	Substitute extracts/worksheets under unchanged commission	Demonstrated unmet objective	Local subtask may be intentional interim work
F05	Repeated rebuilding and worksheet expansion	Strongly indicated ineffective sequencing	Execution trace incomplete
F06	Unretrieved Myanmar paper and catalogue-only images	Demonstrated access-status gaps	Tool malfunction not shown
F07	No hostile instructions identified	Not demonstrated	No security experiment
F08	CSV/PDF divergence; late Moonbeam addenda	Demonstrated integration failure	No complete dependency graph
F09	Derivative source independence unverified	Strongly indicated verification deficit	Claim genealogy incomplete
F10	Repeated Principal reminders and rejection	Demonstrated oversight burden	Human overreliance contradicted by rejection
F11	101-page volume includes 76 pages of apparatus	Demonstrated invalid length proxy	Apparatus may still be useful
F12	No recorded model/tool changes	Not demonstrated	Environment logs absent
F13	No enforced intermediate regression/release gate	Strongly indicated institutional weakness	Principal final hold worked

F01 — Knowledge and representation

The source object itself was unstable in several areas. *Cangarda* was also named *Magedoma*, but a distinct *Magedona* appears in search material. The exclusion of that separate vessel appears in early REF drafts, is absent from the narrow REF v7 text, and returns later. *Moonbeam IV* must be distinguished from *Moonbeam III*; the Myanmar paper also discusses *Sunshine*. A record attached to the wrong hull would contaminate technical and ownership history. The 1927 *Cangarda* passenger attribution to Wilfrid Laurier is impossible because he died in 1919. Later versions recognized this contradiction. However, the surviving report attributes the passenger name to a derivative historical account. The demonstrated AI failure is adopting or transmitting an impossible attribution without sufficient chronological checking; a stronger claim of AI-origin fabrication requires locating its first appearance and examining the purported source. Later versions recognized the contradiction, providing a positive example of correction that must also be tracked through dependent text.

F02 — Inference and uncertainty

The 1904 newspaper report names Hanson as buyer while other histories attribute the yacht to Fulford. Registered title, beneficial interest, agency and delivery are distinct propositions. The wartime naval transfer is variously described as loan or donation. *Moonbeam*'s Sutton transfer and Scotti identity remain disputed, and original race results were not obtained. The chats sometimes preserved these qualifications properly. Their weakness was repeated production of qualified prose without securing or systematically evaluating decisive original records.

F03 — Context selection and retention

The *Cangarda* Hanson and 1927 topics disappear from the readable CER v11 report and later reappear. Similar scope oscillation occurs across REF v6/v7 and v18/v19. *Moonbeam*'s new evidence is present in appended pages but absent from the controlling report. These are observed continuity failures. They are consistent with context compression, selective retrieval, or recent-instruction dominance, but the model's actual context state is not in the artifacts. Surviving CSVs demonstrate that information could exist outside the latest prose while still failing to influence it.

F04 — Objective and specification

The comprehensive yacht history commission was never changed, with a request for only a source extract, an evidence memorandum or worksheets. Nevertheless, those forms repeatedly became the next product. REF v7 followed a 50-page dossier with 14 pages of parsed evidence; *Moonbeam* issued a three-page source memorandum despite having previously generated multi-domain reports. Each may be useful as supporting work, but not as the final deliverable. The observable error is substitution of a local corrective task for the persistent global objective.

F05 — Planning and sequential decisions

Cangarda repeatedly alternated between rebuilds, expansions and contractions; *Moonbeam* generated a large quantity of standardized event and source worksheets while major original-source gaps remained. The sequence did not converge on an accepted comprehensive dossier. The failure is not the use of worksheets, which can support scholarship, but the absence of a controlled plan that makes every new source acquisition or correction improve the same authoritative master.

F06 — Tool, interface and execution

The *Moonbeam* Myanmar paper was indexed but its complete PDF bytes were not preserved, despite its discussion of two different yachts. Original Fife drawings and numerous image references remained catalogue-only. *Cangarda*'s two Classic Yacht Symposium papers were examined and supplied source-specific engineering figures, proving that some direct source examination was feasible. These examples show incomplete acquisition and source-status handling, not necessarily broken tools. Tool failures require logs; the PDFs alone cannot establish them.

F07 — Security and trust boundaries

No examined document demonstrates a malicious prompt injection, an adversarial instruction embedded in a source, or a successful override of the research commission by untrusted content. This error family is not established and should not be counted merely to complete the taxonomy. A future test could use deliberately adversarial mock archival pages, but that experiment was not performed.

F08 — Component interaction

Cangarda's chronology, ownership and voyage registers survived in ZIPs while successive PDFs did not consistently reflect their contents. This is divergence between structured evidence and narrative components. *Moonbeam*'s final addendum raised a possible 2023 collision and mast damage but left earlier rig and chronology sections unchanged. The correct response would be a qualified, cross-referenced contested event and its implications, not a fabricated definitive casualty. The observable defect is lack of dependency-aware propagation.

F09 — Verification and error containment

A derivative account can be copied into a chronology, assigned a citation identifier and repeated in worksheets without any independent witness being added. That is evidentiary authority inflation. The *Cangarda* Laurier contradiction tests whether correction propagates: it is present and recognized in later v17–v19, so some containment worked. *Moonbeam*'s modern Firth painting represents a 1923 race retrospectively; the main report correctly says it is not a race log, but the distinction must persist wherever cited. Independent source genealogy and a claim-level audit remain incomplete.

F10 — Human–AI interaction

The Principal repeatedly reminded both chats of the same comprehensive commission and rejected inadequate products. The chats' assurances that they would correct deficiencies did not reliably predict what the next document contained. This establishes a costly supervisory burden and failure to convert feedback into durable production controls. It does not establish that the Principal relied uncritically on the AI; repeated rejection is evidence of the opposite.

F11 — Evaluation and measurement

Moonbeam's apparent leap to 101 pages largely consisted of 76 pages of standardized apparatus, while its first 25 pages were unchanged. Subsequent 104- and 109-page volumes appended evidence without updating that narrative. *Cangarda*'s length repeatedly fell after longer drafts, but preserved CSVs show that page contraction was not always information destruction. Neither increasing nor decreasing page count measures scholarly progress. Better measures include independently examined sources, source-to-claim locators, unresolved high-impact issues and regression-free narrative integration.

F12 — Adaptation and environmental change

The documents do not identify a specific model upgrade, retrieval change, tool revision or environmental shift that caused a regression. Long-horizon execution may be sensitive to such changes, but the present record does not establish them. A future commission should preserve model/tool configuration and context-summary events with each revision.

The following analysis uses the actual original research instruction supplied by the Principal, not a reconstructed or idealized commission. The vessel particulars differed between the two chats, but the research objective, presentation requirements and treatment of uncertainty were substantively the same.

F13 — Organizational and institutional control

The two chats were assigned the combined responsibilities of evidence acquisition, source classification, historical inference, chronological integration, document production and self-review. The observed release-readiness deficiency is not that the Principal approved unreliable histories: both commissions remained unaccepted, demonstrating an effective final approval hold. The weakness was earlier in the workflow. Cangarda preserved structured chronology, ownership and voyage records in several ZIP packages while producing successor PDFs that omitted developed subjects. Moonbeam IV carried an unchanged controlling report into progressively longer compilations, leaving consequential new evidence in appendices. Neither sequence demonstrates an enforced pre-release check requiring reconciliation against a canonical research record. This is direct evidence of missing or ineffective intermediate production controls, but not experimental proof that a particular alternative institutional design would have succeeded. A better-specified commission could have made these controls explicit; wording alone would not have enforced them [AUD; MAN; C18; M101–M109].

4. Could an institutional error have caused the failures?

This is a substantive question, not simply a question of how forcefully the Principal instructed the chats. A commission can be defective because it leaves scope, evidence standards, authority or completion criteria ambiguous (F04/F13). It

can also be well stated but inadequately implemented. These alternatives lead to different diagnoses. The available record supports the second explanation more strongly than the claim that the Principal failed to describe the work.

Original Cangarda commission:

You are to research the steam vessel Cangara, 126', built 1901 at the yard of Pusey and Jones in Wilmington, DE. You will create a PDF of its comprehensive history and provenance. You do not need to resolve uncertainty, gaps, conflicting information, although you need to keep a record of such. Organize the information by date from earliest to latest. Do not write in a narrative style. Leave placeholders for information desired but not obtainable without outside help (shipyard records, USCG and Canadian registration / documentation information, ...). Proceed with the comprehensive archival investigation. Retain photographic evidence, drawing evidence, other graphical evidence as necessary. Ensure that sources are identified in a bibliography.

4.1 The commission was already substantial and remained unchanged

The two commissions called for comprehensive chronological histories of identified physical yachts, including ownership and custody, operations, design and alterations, photographs and drawings, explicit documentary gaps, and source-traceable findings. Subsequent directions repeatedly required greater depth, chronological integration, preservation of prior research and continued execution. The reports themselves acknowledge these obligations. Thus the observed regressions cannot simply be explained by an absence of any instruction to conduct serious historical research. The chats repeatedly produced narrower source notes or larger worksheet packages while the substantive objective remained the same [AUD; CS; MS].

The instruction establishes an identifiable physical subject; a comprehensive documentary objective; a chronological, non-narrative report; explicit preservation of conflicting information and unresolved gaps; placeholders for records unavailable without external assistance; inclusion of photographic, drawing and other graphical evidence; a bibliography; and a PDF deliverable. It does not ask the researcher to manufacture a continuous ownership chain, resolve every contradiction or obtain private archives. Indeed, the permission to preserve uncertainty is an explicit safeguard against unsupported historical certainty. The spelling “Cangara” in the quoted instruction does not change the identified subject, which is additionally specified by type, approximate length, build date and yard. A competent researcher should recognize the naming issue, verify the hull identity and continue with qualified findings.

These instructions are a credible professional research brief. “Comprehensive” is not a fixed page count, and the commission does not supply a source-access threshold, required research-register schema, claim-level acceptance tests, version-control procedure or independent reviewer. Those omissions create operational latitude, but they do not make the central objective obscure. In ordinary scholarly practice, the investigator is expected to develop a suitable method, explain what was and was not examined, and maintain a cumulative record. The repeated appearance of competent chronologies, source tables, technical observations and visual material in the surviving drafts confirms that the chats understood much of the assignment.

4.2 Would a more tightly specified commission have guaranteed success?

No such conclusion is supported. A more explicit initial commission could plausibly have reduced ambiguity about the expected dossier, its coverage, citation precision, minimum source examinations, research milestones, figure treatment and acceptance tests. But the repeated failure to incorporate preserved CSV registers into later *Cangarda* reports, and the preservation of an unchanged *Moonbeam IV* narrative through successive append-only expansions, are not explained by ambiguous prose instructions alone. They show a breakdown between research state and manuscript production. The record contains no controlled comparison in which the same yacht was investigated under two alternative commission texts, so the effect of better wording cannot be measured.

Could a reasonable commissioner expect a tool-enabled AI chat to accomplish this work? A reasonable person could expect it to understand the brief, investigate accessible public material, extract and compare documents, organize an event chronology, identify archival gaps and produce a cited illustrated PDF. That expectation is supported by tasks the two chats demonstrably performed. It would be a stronger and unverified expectation to assume that one continuing conversation can autonomously coordinate all these activities through reliable final acceptance without persistent

controls. The two cases show a separation between constituent-task competence and end-to-end commission reliability.

The evidence also constrains the counterfactual. A better-written commission might have mandated a master evidence database, explicit completion rubric, regression checks and a review gate. Such requirements might have made failure more detectable and some regressions less likely. But the chats were repeatedly reminded of the unchanged objective and still produced narrower substitutes or unintegrated appendices. It is therefore unsupported to conclude that stronger wording alone would have guaranteed a successful dossier. Compliance must be observable in the artifacts, not inferred from a promise to comply.

4.3 The missing institutional mechanisms

A stronger commission should specify an authoritative external research record, mandatory source-manifest preservation, claim identifiers and genealogy, explicit unresolved alternatives, version-by-version regression checks, staged deliverables, and a prohibition against describing an intermediate extract as the comprehensive successor master. Each new draft should be tested against a coverage matrix: dated chronology, ownership, voyages or races, technical and physical continuity, graphical evidence, source bibliography and gaps. Material omissions should block release unless explicitly justified; new evidence should trigger review of every dependent section. Independent source-facing examination must precede acceptance. These controls are not merely better adjectives in a prompt; they are executable institutional procedures with identifiable responsibility.

4.4 Institutional failure versus successful oversight

The Principal repeatedly rejected incomplete products. That was an effective final institutional safeguard, not evidence of uncritical reliance. The deficiency lay earlier, where the chats themselves were permitted to issue purportedly improved drafts without proving that the new product preserved or superseded the previous research state. F13 therefore describes inadequate internal process assurance, while F04 describes the observed failure to satisfy the standing objective. F03, F05, F08 and F11 describe interacting mechanisms through which the defect manifested. The two cases support an institutional-design recommendation; they do not prove that a revised commission alone would have produced success.

Both chats combined researcher, source registrar, author, editor and self-verifier roles without a demonstrated mandatory source-facing release gate. Their products repeatedly fell short of commission requirements. The Principal's refusal to accept them shows that the ultimate institutional approval safeguard worked. The missing control was earlier: automatic comparison with the prior authoritative research state, mandatory reconciliation of new evidence, and independent examination before a draft was presented as progress.

The F13 finding should therefore be divided into two questions. Commission formulation was largely adequate as a statement of purpose and evidentiary boundaries, although it did not prescribe a full production-control architecture. Execution assurance was inadequate within the single-chat workflow: no consistently enforced master-record comparison, source-to-claim audit, correction dependency check or independent pre-release examination prevented the next deficient draft. The Principal’s repeated refusal to accept the results was an effective final safeguard, not evidence of institutional acquiescence. This distinction assigns responsibility without presuming that every implementation detail must be specified by the commissioning authority.

4.5 Literature and methodological context

The long-context literature provides a reason to investigate retention, not a finding that compression caused any particular omission here. Liu et al. [L1] demonstrated position-dependent performance when relevant evidence occurs within long inputs. FActScore [L2] evaluates long-form generation through atomic factual propositions and is relevant to the distinction between fluent prose and verifiable claims. RARR [L3] studies retrieval-augmented attribution and revision, suggesting an auditable alternative to unsupported self-correction. AgentBench [L4] examines agent performance across interactive tasks, but its benchmark results cannot be transferred directly to yacht historiography. PROV-DM [L5] formalizes derivation and activity lineage and is relevant to tracking a historical claim from source to register to final dossier. These sources motivate safeguards; none establishes a causal diagnosis of the two yacht chats.

A fully reproducible causal study would require the complete instruction-execution chronology, the exact versioned source packets, source-to-claim links, the context presented at each execution, and an independent historical-source audit. The present study instead supports reproducible document-level findings through the artifact manifest and transition matrix. The distinction is important: a researcher can demonstrate that a topic disappeared from a successor PDF without demonstrating that it disappeared from the chat's accessible state.

The failures interact rather than occur as thirteen independent events. A source identity or warrant problem (F01/F02) enters a chronology; later drafting may retrieve only part of the prior record (F03); a recent corrective task may displace the standing commission (F04); output grows through worksheets or contracts into source extracts (F05/F11); evidence and narrative diverge (F08/F09); repeated Principal intervention is required (F10); and internal release checks do not prevent the next inadequate product (F13). This is a plausible explanatory sequence, not a reconstruction of hidden model states. Two selected unsuccessful yacht commissions are enough to demonstrate those commissions failed and to identify concrete recurring artifact defects. They cannot estimate the failure rate of all AI yacht histories or establish context compression as the unique cause. A definitive causal study requires conversation traces, full claim-level version diffs, and matched comparisons with an externally maintained evidence register.

5. The Ticonderoga institutional-learning experiment

Tioga II, subsequently Ticonderoga, is a 1936 L. Francis Herreshoff sailing yacht built by Quincy Adams Yacht Yard. Its documented research leads include Herreshoff Design No. 66, a distinct earlier Tioga (Design No. 50), conflicting launch dates, the Noyes–Johnson ownership question, reported Coast Guard service, the official 1965 Transpac elapsed-time record, and extensive later restorations. The third investigation tested institutional controls for accumulating and using such evidence. Completion of the comprehensive yacht dossier was not its experimental success criterion.



Figure 3. Tioga II / Ticonderoga under sail, showing her two-masted rig and hull profile. Photograph supplied by the Principal; source: Classic Sailboats, "L. Francis Herreshoff — Ticonderoga," https://classicsailboats.org/l-francis-herreshoff-ticonderoga/ [CS-TICO]. Photographer, exposure date, and reproduction rights not independently established.

Research foundation and independent verification

The Principal and Commission Agent established a continuing research charter, a mandatory persistence demonstration, and an independent technical review function. The Historical Researcher created an SQLite archive with sources, claims, evidence links, events, ownership, operations, technical configurations, graphics, issues, activities, and revisions. The accepted Phase I archive was TICO-P1-20261025-RESTORE-002 (database SHA-256 f35ebf4c764c99850ac019dab72eca1c77515657ece5684935e146146f76b333). A separate conversation retrieved it from ChatGPT Library and verified the database checksum, integrity and relationships without relying on the original research session. This established present cross-conversation recovery, not guaranteed future availability or historical correctness [T1].

The first Phase II successor, TICO-P2-20261025-RESEARCH-001 (database SHA-256 cd666adf72518116e10c86ef5bdd7aef604c8f7d0c7a192c7cc205cad70bb0e5), added six sources, six claims and six evidence links. The Independent Technical Reviewer compared all twelve SQLite tables field by field and found no deleted or modified Phase I records. The reviewer also examined the Mystic Seaport catalogue, collection finding aid, WoodenBoat register and Classic Sailboats account directly, and confirmed that catalogue-level information was not misrepresented as inspection of original plans. The Final Gate A Review, pp. 1–3, recommended conditional acceptance, with documentary-capture, synthetic-segregation and authority requirements [T1].

The researcher then developed historical-only views and evidence-domain triggers to prevent synthetic assertions entering historical outputs. The reviewer independently retrieved candidate TICO-P2-20261025-CORRECTIONS-002 (database SHA-256 2cb7ad470662d3762d7a15a99edbc56a32bbb464dac684c62a24e5db5ab1f8c6) and discovered that missing source or claim classification records could bypass the intended trigger rejection through SQLite NULL behavior. This was a demonstrated enforcement defect, not a conjecture about AI reasoning. The reviewer withheld promotion of the candidate; a bounded correction was subsequently prepared as ENFORCEMENT-003. The Principal restricted further software development because assurance work was beginning to displace the yacht investigation [T4].

Historical accumulation, manuscript regression, and bounded synthesis

The historical branch progressed through HISTORY-004, POST1936-005 and INTEGRATED-006, adding original-design leads, conflicting 1936 launching dates, ownership uncertainties, reported Coast Guard service, official 1965 Transpac results and later restoration evidence. The independent focused review of Versions 004–006 found that sources increased 11→16→18, claims 15→26→28, evidence links 15→26→28 and open issues 2→6→6. It found no deleted or modified preexisting historical rows. Yet predecessor database-hash metadata were inconsistent in Versions 004 and 005, and this historical branch did not incorporate the separate enforcement branch. The reviewed four-page 1932–2022 manuscript also compressed earlier documentary analysis from a five-page 1932–1936 chapter. The reviewer recommended Major Revision: preserve the four-page synthesis as a working product but restore historical substance from the retained evidence [T2].

The Principal commissioned a diagnostic task using a fixed corpus: no new historical searches, no new source acquisition, and no database development. The existing researcher produced an eight-page document titled Tioga II / Ticonderoga: Cumulative Documentary History and Provenance, 1932–2022, together with a 23-subject coverage matrix. The independent reviewer found all 23 subjects represented and issued a Conditional Pass. Conditions were correction of numbered citations inconsistent with the bibliography, integration of later history into the main chronological treatment rather than supplementary sections, and more exact regression-matrix locators and evidentiary-qualification checks. The test demonstrated bounded synthesis capability, not an accepted comprehensive dossier; a separate retrieval-and-writing agent was not tested [T3]. Representation of all 23 matrix subjects establishes only subject-level coverage; it does not demonstrate adequate depth for each subject or completeness of the comprehensive historical commission.

The Principal and Commission Agent were active throughout: they established the persistence gate, interpreted independent reviews, prioritized archival work, limited database engineering and designed the bounded test. The reviewer supplied more than approval or rejection: it independently inspected records, identified a concrete SQL control defect, distinguished erroneous metadata from actual record loss, and diagnosed manuscript regression. This

added institutional control and supervisory cost simultaneously. Ticonderoga succeeded as an institutional-learning experiment because it exposed separable capabilities and their limits; completion of its comprehensive yacht dossier was not the experimental success criterion.

6. Persistence, organization, relationships and retrieval: the core of the research engine

The three cases suggest that the central architectural requirement is not continuity of an individual AI conversation but persistence, organization, relational integrity and retrievability of the accumulated evidentiary record. This is a stronger proposition than merely placing a database beside a research chat. The repository becomes the enduring institutional research object from which research can be resumed, audited, corrected and written, while chats become replaceable participants. The Ticonderoga archive showed that recorded claims and relationships could survive even when the manuscript compressed earlier scholarship. That divergence is the direct empirical motivation for making evidence management central rather than auxiliary.

The evidentiary foundation: persistence, organization, and relationships

Persistence requires more than generating downloadable CSV or SQLite files. The institution must know which archived state is current, be able to retrieve its bytes independently of the producing conversation, validate their integrity, preserve predecessors and identify successor changes. Ticonderoga demonstrated retrieval from ChatGPT Library in another conversation and preservation of examined records across research Versions 004–006. It also demonstrated that checksums and ZIP integrity are necessary but insufficient: inconsistent predecessor metadata and divergent research/enforcement branches can leave multiple internally valid but institutionally competing states.

A historical research repository must distinguish the vessel, its names, persons, roles, dated events, sources, individual documentary manifestations, claims, observations, technical configurations, photographs, original drawings and unresolved questions. These are not interchangeable text fragments. For example, Harry Noyes as reported commissioning patron, Robert F. Johnson as listed shipyard customer, and a registered owner are different propositions even if later evidence establishes that some roles belong to the same individual. Likewise an archival finding aid identifying Herreshoff Design No. 66 is not examination of the drawing itself. The organization of evidence protects such distinctions before any prose is generated.

Relational integrity is the scholarly core of the model. Each consequential claim should point to its actual source manifestation and locator; an observation should identify what was read, photographed or transcribed; a historical event should identify which claims support or contradict it; and a proposed conclusion should retain its derivation and contrary evidence. A later researcher may add an original registry without deleting earlier secondary testimony. An accepted interpretation should have a documented rationale, reviewer status and dependency links, so that correction of a source can identify affected chronology entries, figure captions, ownership statements and manuscript passages. Database foreign keys protect structural relationships; they cannot by themselves establish historical truth or source independence.

Retrieval, authority, and independent scholarly production

Retrieval must recover not merely plausible passages but the complete relevant set of accepted, disputed and unresolved evidence for a requested historical period or topic. A dossier author should receive an identified evidence snapshot, its coverage inventory, source locators, conflicting assertions and explicit gaps. The author should not silently select only convenient facts or infer certainty from missing rows. The bounded Ticonderoga test showed that retrieving a fixed corpus could improve cumulative synthesis, while its remaining citation and integration deficiencies show why retrieval must be paired with verification. Reliable comprehensive retrieval from a fully populated institutional repository remains to be demonstrated.

Before manuscript generation, the retrieval-and-writing agent shall produce a versioned evidence-coverage packet tied to one approved database snapshot. Its coverage inventory shall enumerate each applicable historical domain—identity and design, construction, ownership and custody, voyages and racing, technical evolution, repairs and

restoration, graphical evidence, and present status—and identify the retrieved source observations and claims supporting each domain. Where a domain lacks evidence, the packet shall explicitly identify the gap rather than silently omit the domain.

Retrieval completeness shall be assessed against the database's applicable records, not by the number of passages retrieved. The packet shall include all material contradictory claims and unresolved issues linked to the selected vessel and period; list intentionally excluded records with reasons (for example, wrong vessel, synthetic test data, duplicate source manifestation, or out-of-scope period); and provide counts and stable identifiers for included and excluded claims. Any limitation caused by inaccessible source manifestations shall remain visible.

Before acceptance of the manuscript, the Independent Reviewer shall compare the evidence-coverage packet with the approved snapshot and then compare the manuscript with the packet. The reviewer shall verify that consequential included evidence, contrary assertions, unresolved gaps, and material exclusions are represented faithfully, with source locators and qualifications intact. A manuscript failing this coverage audit shall return for correction rather than being accepted on the basis of apparent narrative completeness.

An operational model would maintain one Principal-approved authoritative evidence snapshot, identified by immutable version ID, creation time, content checksum, predecessor and approval record. Research agents would submit additions or proposed corrections as candidate transactions or branches, never silently overwrite the accepted state. An independent review process would check source locators, classification, record-level differences and conflicts; a designated authority would accept or reject the candidate. Accepted snapshots would be immutable and reproducible, and manuscript versions would record exactly which snapshot they used. Divergent branches would be reconciled explicitly, preserving their provenance and rejected or superseded assertions. This protocol responds directly to the branch divergence observed in Ticonderoga; it is a proposed design, not a control already shown to operate end-to-end.

Under the proposed model the Principal and Commission Agent define scope and acceptance criteria; one or more replaceable research agents examine and deposit evidence; an evidence-governance function maintains the authoritative state; a separate retrieval-and-writing agent constructs the chronological dossier from a frozen snapshot; and an independent reviewer compares the manuscript with the same evidence snapshot. The review loop identifies omissions, misleading certainty, source misattribution, failed chronology integration and loss of qualifications, then requires targeted revision and revalidation. Agent succession need not imply loss of the institution's accumulated research because the repository, rather than any chat, retains it.

The experiments support treating the governed evidentiary repository as foundational infrastructure. They do not prove that a sophisticated database alone creates scholarship, that multiple agents will outperform one in every case, or that the proposed retrieval-and-writing agent can produce an accepted comprehensive dossier. The Ticonderoga experiment demonstrated structural preservation and bounded synthesis by the same research chat, not a complete multi-agent production chain. Moreover, its SQL NULL defect and branch divergence show that centralizing evidence introduces technical obligations that must be bounded by scholarly purpose. Database engineering is central when it establishes persistence, organization, relationships and retrieval; open-ended feature development that does not improve those functions remains a distraction.

Criterion	Cangarda	Moonbeam IV	Ticonderoga
Persistent evidence record	CSV registers survived but inconsistently governed prose	Worksheets and sources accumulated; incomplete integration	Versioned SQLite with cross-conversation recovery
Independent review during research	Not demonstrated	Not demonstrated	Demonstrated; found database and manuscript defects
Cross-conversation recovery	Not established	Not established	Demonstrated for accepted Phase I archive
Manuscript regression	Demonstrated	Demonstrated	Demonstrated despite database preservation

Comprehensive dossier accepted	No	No	No; not institutional test criterion
Experimental contribution	Retrospective failure diagnosis	Retrospective failure diagnosis	Bounded institutional-learning success

7. Conclusions

Limitations and future evaluation

These are three purposively selected cases, not a randomized or controlled sample. The Ticonderoga interventions changed several variables simultaneously. No prevalence rate for AI yacht research failure can be inferred. Context compression and instruction drift are plausible contributors but cannot be uniquely reconstructed from the surviving PDFs and database snapshots. The additional governance also carries measurable costs in supervision, verification, version reconciliation and interrupted execution.

A future evaluation should hold the historical task and acceptance rubric constant while comparing a single-chat workflow with an evidence-centered workflow. It should measure independently examined sources, source-to-claim accuracy, correction propagation, retained historical coverage, reviewer findings, human supervisory effort and final dossier acceptance. A separate retrieval-and-writing chat remains a proposed component to be tested.

The three Yacht Provenance investigations suggest that comprehensive AI-assisted historical research may be more reliably organized as a sequence of specialized, institutionally governed functions than as a single continuing conversational commission. Historical evidence acquisition and preservation, manuscript synthesis, and independent scholarly verification are distinct activities with different operational requirements.

The central architectural requirement emerging from these investigations is persistence, organization, relational integrity and retrievability of historical evidence. The *Cangarda* and *Moonbeam IV* investigations showed that continuing conversational research could yield fragmented or non-cumulative manuscripts even when supporting research records survived. *Ticonderoga* demonstrated that an external relational archive could preserve sources, claims, evidence links and unresolved issues across conversations while the generated manuscript still regressed.

The *Ticonderoga* experiment demonstrates that cumulative research records can survive across conversations and that bounded synthesis from a preserved corpus can recover substantial historical content under independent review. These findings support, but do not yet validate, a database-centered architecture in which replaceable research agents contribute to an authoritative evidence repository and separate retrieval-and-writing agents construct manuscripts through iterative independent review.

This is an institutional model, not an established end-to-end solution. The research database must be governed as a scholarly evidentiary system, not merely treated as reliable storage: source genealogy, contested assertions, version authority, reviewed snapshots and retrieval completeness remain essential. A future yacht investigation must test the separate retrieval-and-writing function directly, with independent source-facing review, measurable supervisory cost and final dossier acceptance criteria.

Bibliography and documentary references

[COM] Principal. Original Cangarda research commission, reproduced in Section 4.1; Moonbeam IV commission equivalent in scope with vessel particulars changed. Supplied by the Principal.

[TM1] Yacht Provenance. Operational and Scholarly AI Failure Mechanisms, Technical Manuscript No. 1, approved Version 1.0 (2026).

[AUD] Yacht Provenance. Successive-Draft Evolution of the Cangarda and Moonbeam IV Commissions, forensic evidence memorandum (2026).

[MAN] Yacht Provenance. YP_TM5_Artifact_Evolution_Manifest.csv; 46 PDF instances, including duplicate exports, with document paths, page counts, extracted text lengths, and SHA-256 hashes.

[C18] Cangarda Master Working Investigation, Version 18, 46 pages. [CT] Cangarda Reconciled Timeline, 33 pages.

[M101] Moonbeam IV 101-page Research Draft. [M104] Updated Research Volume, 104 pages. [M109] Cumulative Research Volume, 109 pages.

[T1] Yacht Provenance. Gate A — Final Independent Technical Review, Tioga II / Ticonderoga (1936), 24 October 2026, pp. 1–3. Reviewed TICO-P1-20261025-RESTORE-002 (SQLite SHA-256 f35ebf4c764c99850ac019dab72eca1c77515657ece5684935e146146f76b333) and TICO-P2-20261025-RESEARCH-001 (SHA-256 cd666adf72518116e10c86ef5bdd7aef604c8f7d0c7a192c7cc205cad70bb0e5).

[T2] Yacht Provenance. Focused Independent Examination — Manuscript Regression and Archive Continuity, Versions HISTORY-004, POST1936-005, and INTEGRATED-006, two-page independent review supplied by the Principal. Compared 11/16/18 sources, 15/26/28 claims, and manuscript regression.

[T3] Yacht Provenance. Bounded Scholarly Synthesis — Final Independent Review, two-page report supplied by the Principal. Evaluated the eight-page Tioga II / Ticonderoga: Cumulative Documentary History and Provenance, 1932–2022, and its 23-subject coverage matrix. Disposition: Conditional Pass, with citation, integration, and locator corrections required.

[T4] Yacht Provenance. Independent Corrective Archive Verification, candidate TICO-P2-20261025-CORRECTIONS-002 (SQLite SHA-256 2cb7ad470662d3762d7a15a99edbc56a32bbb464dac684c62a24e5db5ab1f8c6); independent negative test identified missing-classification/NULL enforcement defect. Reviewer findings supplied by the Principal.

[P06] McNeil, R. G., and Rutherford, J. (2006). “The Restoration of CANGARDA.” Classic Yacht Symposium, printed p. 1 (Figure 1).

[P08] Rutherford, J., and McNeil, R. G. (2008). “A Return to the Edwardian Era—Completing CANGARDA, Initiating CORONET.” Classic Yacht Symposium.

[BOAT] Boat International. “Moonbeam IV,” The Superyacht Directory. https://www.boatinternational.com/yachts/the-superyacht-directory/moonbeam-iv--66139 (Figure 2, supplied by Principal; image rights unverified).

[CS-TICO] Classic Sailboats. “L. Francis Herreshoff — Ticonderoga.” https://classicsailboats.org/l-francis-herreshoff-ticonderoga/ (Figure 3, supplied by Principal; image rights unverified).

[KPBS] Robinson, J. (2015). “ULTIMATE RESTORATIONS: Cangarda: The Last American Steam Yacht.” KPBS.

[NAV] Naval Marine Archive. “The Steam Yacht Cangarda: Ready for Restoration.” https://navalmarinearchive.com/research/pdf/magedoma_cangarda.pdf.

[SAN] Sandeman Yacht Company. “Pusey & Jones 126 ft Steam Yacht 1901/2009.” https://www.sandemanyachtcompany.co.uk/yacht/618/CANGARDA.

[PMF] Patrimoine Maritime et Fluvial. “Moonbeam IV.” https://www.patrimoine-maritime-fluvial.org/navires/moonbeam-iv/.

[CIM] Comité International de la Méditerranée. “Moonbeam IV.” https://www.cim-classicyachts.org/en/yachts/moonbeam-iv.

[JFA22] “Moonbeam IV at JFA Yachts” (2022), shipyard-supplied particulars. https://lulop.com/en_EN/post/show/242077/moonbeam-iv-at-jfa-yachts.html.

[CS; MS] Cangarda Commission Self-Assessment and Moonbeam IV Critical Self-Assessment (24 October 2026), retrospective testimony, not independent verification.

[PR] Independent Peer Review Report, TM5 Draft 4 (2026), Major Revision.

[PR10] Independent Technical Peer Review, TM5 Draft 10 (2026), Major Revision, supplied by Principal.

[L1] Liu, N. F., et al. (2024). “Lost in the Middle: How Language Models Use Long Contexts.” Transactions of the ACL 12:157–173. doi:10.1162/tacl_a_00638.

[L2] Min, S., et al. (2023). “FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation.” EMNLP 2023. arXiv:2305.14251.

[L3] Gao, L., et al. (2023). “RARR: Researching and Revising What Language Models Say, Using Language Models.” ACL 2023. arXiv:2210.08726.

[L4] Liu, X., et al. (2024). “AgentBench: Evaluating LLMs as Agents.” ICLR 2024. arXiv:2308.03688.

[L5] Moreau, L., and Missier, P., eds. (2013). PROV-DM: The PROV Data Model. W3C Recommendation.

Appendix A — Reproducibility and artifact identification

The accompanying file YP_TM5_Artifact_Evolution_Manifest.csv records 46 examined PDF instances (including duplicate exports). Its fields are container/archive name, internal document path, PDF page count, extracted text character count and SHA-256 of the PDF bytes. ZIP package names identify the version. Duplicate exports do not represent independent research. The extraction method was PDF text extraction, supplemented by selected page inspection and comparison of preserved CSV register contents. Exact historical source verification and full conversation reconstruction remain outside this audit.

CANGARDA_CER_v10.zip / CANGARDA_CER_v10/report.pdf: 28 pp.; SHA-256 557e14e0f4305a18915b77f5ea04b27d6c691af078b16448ba5bc9521929a8d1.

CANGARDA_CER_v11.zip / CANGARDA_CER_v11/report.pdf: 11 pp.; SHA-256 464d2f75e6b552f150e70a48fc27532101912396648a358dce869a7b14fe6e1e.

CANGARDA_CER_v12.zip / CANGARDA_CER_v12/report.pdf: 26 pp.; SHA-256 d29f001cb774753d435fb98813363eae4f962386ef87b84b45e11722ec8d14a9.

CANGARDA_CER_v13.zip / CANGARDA_CER_v13/report.pdf: 17 pp.; SHA-256 602ad4cf57279d557fc7dcea69882c61a0c8cc6a3c8d88fc7091b3f8109e1f4b.

CANGARDA_MASTER_v17.zip / CANGARDA_MASTER_v17/report.pdf: 41 pp.; SHA-256 4f008a892dac3b7ebbf0b4aa7f249de7a2e2ed2d88960e25991301622e6930d2.

CANGARDA_REF_v6.zip / CANGARDA_REF_v6/investigation/report.pdf: 50 pp.; SHA-256 e965fd7b8c809e82c079ced7cf7e5b129366db962fb9858eafac478c716f3a20.

CANGARDA_REF_v7_parsed_sources.zip / CANGARDA_REF_v7/investigation/parsed_evidence_report.pdf: 14 pp.; SHA-256 6cf2eb08e0f33e4fb3556cc8faada831fd8b182ec550c6860fcaf1732357be4a.

The full manifest is a companion dataset; this appendix reproduces representative identifiers, not all 46 rows.

[PR10] Yacht Provenance, Independent Technical Peer Review of Technical Manuscript No. 5, Expanded Draft 10, Major Revision, 2026 (review supplied by the Principal).

YACHT PROVENANCE — TECHNICAL MANUSCRIPTS

Copyright and Permissions

Copyright © 2026 Yacht Provenance LLC. All rights reserved.

No part of this manuscript may be reproduced, distributed, transmitted, or otherwise used without the prior written permission of the copyright owner, except as permitted by applicable copyright law.

For permission requests, contact: Fitzhorn@yachtprovenance.com