YACHT PROVENANCE TECHNICAL MANUSCRIPT SERIES

Number 1 · August 18, 2026

Operational and Scholarly AI Failure Mechanisms: From Four Initial Classes to Thirteen Operational Families

Technical Manuscript No. 1

Technical Editor

Patrick Fitzhorn, PhD

Principal, Yacht Provenance

Yacht Provenance, Fort Collins, Colorado, USA

Fitzhorn@yachtprovenance.com

Publication

Yacht Provenance Technical Manuscript Series

Manuscript Status: Approved

Version: 1.0

Prepared: August 18, 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

Four broad classes—Mind, Frame, Interaction, and Institutional Error—identified during early Yacht Provenance (YP) design motivated a targeted examination of refereed AI, systems, human-factors, and provenance scholarship. Seventeen operational failure-mechanism groupings are mapped to those lenses, differentiated by causal origin, diagnostic evidence, and intervention, and consolidated into thirteen overlapping operational families. Each family is then examined through a YP-specific relevance scenario and an individual-family removal counterfactual; seven additional cases challenge difficult classification boundaries. Each family makes a credible, distinctive contribution under the stated criteria. The analysis supports an operationally economical classification for the mechanisms examined, not a uniquely optimized number, literature-wide completeness, or demonstrated safeguard effectiveness.

1. Four foundational classes and the YP research problem

Yacht Provenance investigates the documentary history and provenance of individual recreational yachts. A physical hull must be distinguished from its reported names, designers, builders, plans, ownership, and alterations. Original drawings, registry records, photographs, oral testimony, and secondary accounts can be incomplete or mutually dependent. Errors may arise within an AI participant, in the governing premise, through interactions, or in institutional authority and recovery.

Mind Error encompasses individual misrepresentation, unsupported inference, context degradation, planning, or action selection. For example, a successor may retain a proposed builder while omitting that the attribution is disputed [1–3]. Frame Error occurs when a false premise directs otherwise coherent work: a correctly interpreted drawing may concern a different yacht. Interaction Error arises through shared assumptions, handoffs, dependencies, or mistaken agreement among Minds [7]. Institutional Error involves authority, durable accountability, evaluation, and recovery: an objection may have no owner or a corrected conclusion may leave a stale approval active [5,8,10].

These four classes are broad explanatory lenses, not mutually exclusive scientific kinds. They guide discovery but do not identify sufficiently distinct diagnostic questions or interventions. Once the thirteen operational families are developed, the four cease to function as a second governing classification layer.

2. Literature examination and analytical procedure

This study is a targeted purposive synthesis, not a systematic review or blinded derivation. Refereed journal articles and conference proceedings were selected for relevant mechanisms in factuality, inference, context, agents, tools, security, coupled systems, human reliance, evaluation, environmental change, and provenance [1–3,5–17]. Inclusion required a plausible causal route to YP evidence integrity, reasoning, continuity, authority, or recovery. General societal consequences without such a mechanism were excluded; broader risk scholarship is comparative background [11]. Neither exhaustive database screening nor a recorded sampling frame was performed, so no literature-wide completeness claim follows.

The procedure is: (A) map each identified mechanism to one or more of the four initial lenses; (B) distinguish causal origin, observable signature, and candidate response; (C) consolidate into operational families; (D) evaluate each family for a credible YP failure, material consequence, and intelligible detection or containment response; and (E) remove each family counterfactually and examine what the strongest reassignment would obscure. These are author-conducted analytical tests, not executed fault-injection experiments. YP requirements [YP1] are institutional sources, not independent scientific evidence.

3. Mechanisms mapped into the four initial classes

3.1 Classification rules

Mind denotes individual information processing or action selection; Frame denotes a governing premise or objective; Interaction requires transfer, mutual influence, or component coupling; Institutional involves authority, oversight, measurement, maintenance, or durable responsibility. Overlap is assigned only where distinct causal stages are implicated. A hallucination is not automatically a frame failure: it becomes one when an invented statement is adopted as the governing premise.

ID	Operational failure mechanism	Four-class mapping	Source
M01	Invented or source-unfaithful content	Mind; Frame if adopted as premise	[1,3]
M02	Unsupported inference / unwarranted certainty	Mind; Frame if it governs inquiry	[1,14]
M03	Relevant long-context material omitted or ignored	Mind; Interaction if lost at handoff	[2]
M04	Task or role instruction not followed	Mind; Frame if goal shifts; Institutional if authority fails	[7]
M05	Repeated steps, derailment, premature stop	Mind; Frame if stopping objective faulty	[7]
M06	Information withheld or lost between agents	Interaction; Mind if supplied content ignored	[7]
M07	Reasoning–action mismatch	Mind; Interaction only if interface contributes	[7,9]
M08	Hostile tool-output instruction redirects agent	Interaction; Institutional; Frame if goal captured	[6]
M09	Hidden dependency or feedback propagates defect	Interaction; Institutional if maintenance fails	[5]
M10	Required verification absent or incorrect	Interaction if multi-party; Institutional if oversight fails	[7]
M11	Human overreliance or inappropriate rejection	Mind/human judgment; Institutional if oversight design contributes	[10]
M12	Evaluation misses material behavior or measures wrong target	Institutional; Interaction if coupled scoring	[12,13]
M13	Reliability/uncertainty degrades under distribution change	Mind; Institutional if requalification absent	[14]
M14	Lifecycle coordination or maintenance failure	Institutional; Interaction for coupled systems	[8]
M15	Mistaken derivation lineage creates false corroboration	Interaction; Institutional	[15]
M16	Digital evidence origin or provenance misattributed	Frame; Institutional	[16]
M17	Authenticity asserted without adequate archival context	Frame; Institutional	[17]

3.2 Mapping decisions and findings

M04 and M07 illustrate conditional mapping: role noncompliance need not involve a false frame, and an individual agent’s reasoning–action mismatch need not involve other Minds. M08 is first a trust-boundary breach; it becomes a frame error only if the hostile instruction becomes the operative goal. Thus the four classes are not applied merely because a failure has remote consequences in another domain.

M15–M17 are deliberately expressed as failures, not as names of provenance disciplines. Database derivation provenance [15] informs the failure of treating dependent accounts as independent; digital evidence provenance [16] informs erroneous origin attribution; archival metadata scholarship [17] informs unsupported authenticity judgments. These references do not themselves establish yacht-document authenticity or physical-hull identity. The latter are YP-specific applications requiring separate evidentiary examination.

All seventeen selected mechanisms fit at least one broad lens, but the classes remain too coarse: Mind conflates transcription, inference, and context loss; Institutional conflates measurement, maintenance, and authority. Coverage of these seventeen groupings does not establish that all important published mechanisms were encountered.

4. Differentiation into thirteen operational families

4.1 Criteria and mechanism-to-family synthesis

A distinct family is retained where causal origin, diagnostic evidence, or prevention/detection/recovery response differs materially. Closely related mechanisms are consolidated when finer division would not change an institutional decision. These are YP analytical classifications rather than classifications asserted by the source authors. Compound incidents may carry several labels, provided initiating defect, propagation, and failed barrier remain separable.

Family	Operational name	Mechanisms	Reason for separation / candidate intervention
F01	Knowledge and representation	M01; M16–17 qualified	Misread source content calls for direct comparison with an authoritative original, before evaluating inferences.
F02	Inference and uncertainty	M02; M13	Accurately transcribed evidence may still fail to warrant a confident builder or designer attribution.
F03	Context selection and retention	M03; M06	A missing contradiction or lost qualification requires a packet-to-authoritative-register completeness check.
F04	Objective and specification	M04; M08 conditional	A false subject or designer premise can govern competent downstream work; premise validity requires direct challenge.
F05	Planning and sequential decisions	M05	An appropriate objective can still fail through repeated steps, poor lead priority, or premature STOP.
F06	Tool, interface, and execution	M07; M08 effect	Intended action and actual tool effect may diverge; reconcile commands, outputs, and uncertain side effects.
F07	Security and trust boundaries	M08	Untrusted documentary content may attempt to become an instruction; enforce source/authority separation.
F08	Component interaction	M06; M09	Corrected assertions may leave stale dependent artifacts; trace propagation and invalidate descendants.
F09	Verification and error containment	M10; M15–17 qualified	Dependent sources or inadequate review can permit false corroboration; test source lineage and independent checking.
F10	Human–AI interaction	M11	Fluent but weak claims can distort human reliance; display contrary primary evidence.
F11	Evaluation and measurement	M12	A PASS score can measure procedural completion instead of factual validity; use independently keyed outcomes.
F12	Adaptation and environmental change	M09; M13–14	Changed model/repository behavior may invalidate previous qualification; trigger revalidation.
F13	Organizational and institutional control	M04; M10; M14	Unowned objections, stale holds, and invalid release state require accountable ownership and restoration checks.

4.2 How the distinctions are obtained

Representation, inference, and retention (F01–F03). A shipyard number can be copied incorrectly (F01), copied correctly but used to infer the wrong physical hull (F02), or omitted from the continuation packet (F03). These demand original-source comparison, competing-hypothesis assessment, and context completeness respectively [1–3]. Treating all three simply as Mind Error would fail to indicate which evidence should be examined first.

Objective, plan, execution, and trust (F04–F07). An invalid premise can be pursued with an otherwise sound plan (F04); a valid premise can be undermined by poor lead sequencing (F05); a sound plan can fail at the tool boundary (F06); and a hostile instruction in an acquired PDF can cross the source/authority boundary (F07). An injection may also corrupt the objective or cause an external effect, but those are separate causal stages, not reasons to merge the families [6,7,9].

Propagation, verification, and evaluation (F08, F09, F11). A corrected assertion may remain in a stale report (F08), escape independent examination because two accounts share one source (F09), and receive a passing score because the evaluation measures checklist completion rather than historical truth (F11). Dependency traversal, source-genealogy examination, and outcome-valid measurement are

distinct interventions [5,7,12,13].

Human reliance, environmental change, and institutional control (F10, F12, F13). Overtrust of a persuasive AI narrative calls for human-facing contradictory evidence (F10). A changed model or archive requires requalification (F12). An unowned objection or stale release approval requires durable authority and restoration checks (F13). These concerns may interact but do not reduce to the same diagnostic question [5,8,10,14].

The synthesis yields thirteen intelligible operational questions for the examined mechanisms. The four initial lenses are no longer needed as an operational classification layer. This does not establish that the literature uniquely requires thirteen top-level families.

5. Operational relevance test: YP applicability

5.1 Analytical test and results

Each family is placed in an intended YP activity and assessed for P (credible initiating mechanism), C (material negative consequence), and I (intelligible observable manifestation and candidate response). These are constructed analytical scenarios, not documented YP incidents, independent scores, or executed fault injections. A family passes when all three can be articulated without speculative dependence on a technology YP does not use.

Family	YP activity	Initiating failure	Observable symptom	Material consequence	Candidate response
F01	Source examination	Hull number transcribed incorrectly	Source image disagrees with transcription	Wrong-vessel attribution	Compare transcription to authoritative image/record
F02	Historical resolution	Builder inferred from resemblance alone	Contrary candidate builders not assessed	Unsupported builder attribution	Rival-hypothesis and uncertainty review
F03	Mind succession	Ownership contradiction lost in compression	Unresolved item missing from packet	Disputed transfer presented as settled	Compare packet to authoritative open-matters register
F04	Inquiry framing	Client designer premise accepted as fact	No independent premise or hull-identity challenge	Investigation answers wrong question	Challenge subject identity and attribution premise
F05	Discovery / STOP	Material archive lead prematurely closed	Lead lacks justified disposition	Premature conclusion	Review lead register and stopping rationale
F06	Tool extraction	Extraction tool corrupts official-number digit	Extracted value disagrees with source	Evidence record contaminated	Exact extraction/source comparison and effect reconciliation
F07	Source intake	PDF contains hostile operational instruction	Source text treated as higher-authority command	Unauthorized action or disclosure	Enforce trust boundary and source/instruction separation
F08	Reporting	Corrected attribution leaves old draft current	Dependent report still cites superseded assertion	False client deliverable	Traverse dependencies; invalidate stale artifact
F09	Independent review	Copied accounts counted as independent	Shared origin omitted from source lineage	False corroboration passes	Check source genealogy and independent review basis
F10	Principal review	Fluent account overtrusted	Contrary primary source discounted	Human approves weak conclusion	Present conflicting primary evidence explicitly
F11	QC / evaluation	Checklist PASS despite false attribution	Score ignores factual outcome	False assurance of correctness	Independently keyed historical-outcome test
F12	Extended research	Repository or model behavior changes	Previously qualified extraction differs	Undetected reliability degradation	Change-triggered requalification
F13	Objection and release	Critical objection lacks accountable owner	Hold unresolved while release proceeds	Unauthorized release or invalid continuation	Owner/hold/release audit and recovery check

5.2 Interpretation

All thirteen constructed scenarios satisfy P, C, and I as author-conducted analytical judgments. Their value is that each makes a specific adverse pathway and corresponding diagnostic response understandable to a YP investigator. They span source acquisition and interpretation, research planning, context continuity, external tools, multi-Mind reporting, independent review, Principal judgment, quality assessment, environmental change, and institutional release. They do not demonstrate frequency, real-world incidence, safeguard performance, or independent classification reliability.

The relevance test establishes why each family matters to YP, but not why it must remain separate from the other twelve. That is the distinct question examined in Section 6.

6. Counterfactual removal and separation

6.1 Counterfactual method

For each family Fx, temporarily remove its label, identify the strongest available substitute among the other twelve, and reassign the mechanisms. The test asks what remains describable, what diagnosis or control trigger becomes less explicit, and what consequential YP action might thereby be omitted or misdirected. A family has operational value if its removal materially weakens discrimination, even where broad descriptive coverage survives. This is not an exhaustive optimization over alternative hierarchical or merged taxonomies.

Removed	Best substitutes	Coverage remaining	Operational consequence of removal
F01	F02/F03	Wrong attribution still describable	Direct source-fidelity check loses a dedicated trigger; investigator may debate inference before checking
F02	F01/F04	False conclusion still describable	Competing hypotheses and calibrated warrant become implicit; accurate evidence may be mistaken for s proof.
F03	F01/F08	Omission or handoff error still describable	Compression and packet completeness lack a specific audit trigger; unresolved contradictions may vanis succession.
F04	F02/F05	Wrong answer or bad plan still describable	Independent premise challenge is obscured; a well-executed investigation can pursue the wrong hull.
F05	F04/F06	Ineffective work still describable	Lead prioritization and STOP rationale lack a dedicated review; promising archival leads may close pre
F06	F05/F08	Tool error still describable	Actual external effects may not be reconciled with intent; duplicate or uncertain actions risk being repea
F07	F06/F13	Unauthorized effect still describable	Source/authority boundary loses a distinct control; hostile PDF instructions may be treated as ordinary t
F08	F03/F09/F13	Handoff or review failure still describable	Downstream dependency traversal may not occur; corrected claims can remain in reports or approvals.
F09	F08/F11	Error escape still describable	Source-lineage and reviewer-independence challenge loses a specific trigger; copied sources may pass a corroboration.
F10	F04/F13	Human acceptance still describable	Overreliance/disuse loses explicit human-facing evidence design; persuasive weak findings may be appr
F11	F09/F13	Bad review still describable	Validity of the metric itself is obscured; checklist completion may substitute for historical truth.
F12	F03/F08/F13	Changed output still describable	Requalification after environmental change lacks an explicit trigger; old tests may remain incorrectly au
F13	F08/F09/F10	Objection or release error still describable	Accountable ownership and release/restore authority disperse; unresolved holds may lack enforceable d

6.2 Difficult separation pairs

F01/F02. A misread original record is tested against the source; an accurately read record used for an unsupported attribution is tested against competing hypotheses. Combining them could preserve both as submechanisms, but would no longer direct the initial diagnostic question as explicitly. F04/F05. Premise validity and lead-sequence adequacy are logically independent: a plan audit cannot correct a false subject, and a premise check cannot detect a premature STOP.

F08/F09. Propagation can occur before independent examination; verification can fail because of common-source dependence even without a handoff defect. Dependency invalidation and source-genealogy challenge are therefore different interventions. F09/F11. A

reviewer can miss an error under a valid evaluation metric, while an invalid metric can misreport success even when one reviewer happens to catch the error. The measured process and the measure's validity must remain distinct [12,13].

6.3 Result and limits

Each individual removal weakens at least one operationally consequential diagnostic or control trigger in the present classification. No family is an obvious candidate for deletion under these criteria. This supports retaining the thirteen distinctions as an operationally economical working taxonomy; it does not show that all thirteen must be separate top-level families. Another taxonomy could preserve some distinctions as subcategories. No exhaustive merger search, blinded external scoring, or empirical validation was conducted.

7. Adversarial boundaries and published multi-agent check

Author-designed boundary case	Candidate families	Distinctive concern
Authentic drawing attributed to wrong hull	F02/F04/F09	Authenticity of paper does not prove identity of subject.
Forged drawing attributed to correct hull	F01/F09	Correct vessel association does not establish document authenticity.
Two reviewers rely on one erroneous source	F08/F09/F10	Procedural separation does not ensure evidentiary independence.
External send succeeds; acknowledgment lost	F06/F08/F13	Uncertain side effect requires reconciliation before retry.
Corrected claim leaves approval current	F08/F09/F13	Dependency and release state require invalidation.
Contradiction supplied but ignored	F03/F09/F11	Availability of context does not prove effective use.
Client imposes false designer premise	F02/F04/F10	Premise challenge must precede downstream inference.

These seven cases are author-designed challenges, not independently sampled published failures or executed experiments. They test classification boundaries and do not identify a demonstrated need for F14. Separately, Cemri et al. [7] report fourteen multi-agent failure modes across system design, inter-agent misalignment, and task verification. Examples include history loss (F03), repeated steps or premature termination (F05), inter-agent information failure (F08), reasoning–action mismatch (F06), and incorrect verification (F09/F11). Those YP assignments are the present author's analytical crosswalk, not classifications endorsed by the source authors.

8. Architectural implications and evidentiary limits

The thirteen families motivate distinct candidate interventions: source comparison (F01); rival-hypothesis review (F02); context and succession audit (F03); premise challenge (F04); lead and stopping review (F05); external-effect reconciliation (F06); instruction trust boundaries (F07); dependency invalidation (F08); source-lineage examination (F09); human-facing contrary evidence (F10); outcome-valid evaluation (F11); change-triggered requalification (F12); and accountable objections, holds, and recovery (F13). The YP governing specification describes related requirements [YP1]. Their presence in a specification is not evidence of effective implementation.

Demonstrating safeguard performance would require controlled evidence fixtures, independently established expected outcomes, repeated trials, miss and false-alarm rates, restoration checks, and residual-risk disposition. Direct scholarly grounding is strongest for factuality, context, tool/security, human reliance, and evaluation. Predictive-model uncertainty and distribution shift are partly analogical for historical-research LLMs; archival and database provenance address different evidentiary problems and cannot be conflated with physical-yacht identity. The paper makes classification and design claims, not empirical reliability claims.

9. Conclusions

Classification developed. The analysis develops a thirteen-family, literature-informed operational taxonomy from a targeted inventory of seventeen mechanism groupings. The four initial classes are useful as broad discovery lenses but are too coarse to distinguish source fidelity, inferential warrant, context integrity, objective validity, planning, execution, trust boundaries, propagation, verification, human reliance, evaluation, adaptation, and institutional authority. Once those distinctions are formulated, the four classes are not needed as a second governing layer.

What the analysis demonstrates. The seventeen examined groupings can be accommodated within the thirteen families under explicit causal and operational criteria. Each family has a credible, consequential YP-specific failure scenario and an intelligible candidate diagnostic response. Each individual-family removal weakens at least one distinct control-design question under the adopted classification. In particular, F04 and F05 distinguish an invalid objective from a defective plan; F08 and F09 distinguish propagation from failed containment; and F09 and F11 distinguish verification performance from validity of its evaluation. These distinctions make the taxonomy useful for selecting the right investigative challenge or safeguard rather than merely labeling an adverse outcome.

Operational economy and limits. Thirteen families offer a manageable, sufficiently differentiated classification for the mechanisms examined, with no demonstrated need for F14 in this inventory or the stated boundary challenges. This is a bounded analytical judgment, not proof that thirteen is uniquely optimal or that the selected literature is exhaustive. The relevance and removal tests are author-conducted conceptual analyses, not independent experiments. Neither comprehensive coverage of all future or published AI failures nor effectiveness of YP's proposed controls is established. The appropriate institutional next step is to use the taxonomy to design and independently test fault-detection, containment, and recovery mechanisms, while revisiting its boundaries if a material unclassified failure is discovered.

References — refereed scholarship

[1] Ji, Z., et al., 2023, "Survey of Hallucination in Natural Language Generation," ACM Computing Surveys, 55(12), Article 248. doi:10.1145/3571730.

[2] 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. doi:10.1162/tacl_a_00638.

[3] Manakul, P., Liusie, A., and Gales, M., 2023, "SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models," Proceedings of EMNLP 2023, pp. 9004–9017. doi:10.18653/v1/2023.emnlp-main.557.

[4] Bommasani, R., Liang, P., and Lee, T., 2023, "Holistic Evaluation of Language Models," Annals of the New York Academy of Sciences, 1525(1), pp. 140–146. doi:10.1111/nyas.15007. Comparative background.

[5] Sculley, D., et al., 2015, "Hidden Technical Debt in Machine Learning Systems," Advances in Neural Information Processing Systems, 28.

[6] Debenedetti, E., et al., 2024, "AgentDojo: A Dynamic Environment to Evaluate Prompt Injection Attacks and Defenses for LLM Agents," Advances in Neural Information Processing Systems, 37. doi:10.52202/079017-2636.

[7] Cemri, M., et al., 2025, "Why Do Multi-Agent LLM Systems Fail?," Advances in Neural Information Processing Systems, 38. doi:10.52202/085713-4082.

[8] Amershi, S., et al., 2019, "Software Engineering for Machine Learning: A Case Study," ICSE-SEIP, pp. 291–300. doi:10.1109/ICSE-SEIP.2019.00042.

[9] Qin, Y., et al., 2024, "ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs," International Conference on Learning Representations.

[10] Parasuraman, R., and Riley, V., 1997, "Humans and Automation: Use, Misuse, Disuse, Abuse," Human Factors, 39(2), pp. 230–253. doi:10.1518/001872097778543886.

[11] Weidinger, L., et al., 2022, "Taxonomy of Risks Posed by Language Models," ACM FAccT, pp. 214–229. doi:10.1145/3531146.3533088.

[12] Ribeiro, M. T., Wu, T., Guestrin, C., and Singh, S., 2020, "Beyond Accuracy: Behavioral Testing of NLP Models with CheckList," ACL, pp. 4902–4912. doi:10.18653/v1/2020.acl-main.442.

[13] Liang, P., et al., 2023, "Holistic Evaluation of Language Models," Transactions on Machine Learning Research.

[14] Ovadia, Y., et al., 2019, "Can You Trust Your Model's Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift," Advances in Neural Information Processing Systems, 32.

[15] Buneman, P., Khanna, S., and Tan, W.-C., 2001, “Why and Where: A Characterization of Data Provenance,” International Conference on Database Theory, LNCS 1973, pp. 316–330. doi:10.1007/3-540-44503-X_20.

[16] Turner, P., 2005, “Digital Provenance—Interpretation, Verification and Corroboration,” Digital Investigation, 2(1), pp. 45–49. doi:10.1016/j.diin.2005.01.002. Analogical evidence.

[17] 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. doi:10.1007/s10502-023-09422-w.

Institutional reference — not scholarly evidence

[YP1] Yacht Provenance, Unified Governing Specification with R1 Appendix for Acceptance, especially YP-AI-101–102, YP-TEST-043–044, YP-RES-001–005, YP-TR-021, YP-EVD-001–005, and YP-VAL-001–002. Requirements do not establish control effectiveness.

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