JEARON WONG / AGENTIC LIFECYCLE GOVERNANCE INDUSTRY SERIES

Agentic AI Auditability & Assurance White Paper 2026

A Lifecycle Evidence Guide for Audit, Assurance, and Enterprise AI Governance

Jearon Wong - Protocol Architect for the Agent Era

Authority	Responsibility	Agent Work	Evidence	Closure

DOCUMENT ID	VERSION	DATE
AIAAWP-2026-v0.1	v0.1 Public Research Edition	May 18, 2026
STATUS	**SERIES**	**VISUAL SOURCE OF TRUTH**
Public research edition; HTML/PDF/manifest/checksum available; not legal advice, not an audit standard, not certification, not an assurance opinion	Agentic Lifecycle Governance Industry Series	HTML/PDF primary; editable derivative

PUBLICATION BOUNDARY

Public Research Edition Status

This artifact is a public research edition available on site with HTML, PDF, manifest, and checksum records. It is not an audit standard, certification, assurance opinion, legal compliance proof, regulator approval, procurement recommendation, vendor ranking, or endorsement.

It applies the GAIC main white paper publication standard to AIAAWP while preserving AIAAWP's public-facing title, subtitle, and document ID.

Publication Contents

HTML anchors are active in the public research edition. PDF rendering uses the shared whitepaper A4 print profile.

Front Matter

Chapters 0-8

Chapters 9-16

Appendices

Sources

Relationship to the Agentic Lifecycle Governance Industry Series

The Agentic Lifecycle Governance Industry Series is fixed as a 3+3 structure: three core industry white papers and three practitioner guides. The Global AI Compliance White Paper 2026 defines the compliance object layer for agentic lifecycle governance. This paper defines the audit evidence object layer. Guide 1 will later translate this paper into audit-ready technical implementation planning. Guide 2 will later translate the Global AI Compliance White Paper 2026 and this paper into compliance operating-model planning. The later insurability white paper will address insurability. This public research edition publishes the auditability white paper as part of the on-site white paper series; it does not replace those later assets.

Executive Thesis

AI agent auditability cannot be built on logs alone. An agentic system becomes auditable only when its lifecycle work can be reconstructed through responsibility-linked evidence chains. **Agentic AI Auditability** is the ability to reconstruct, test, and evidence agentic lifecycle work across authority, responsibility, tools, outcomes, exceptions, and remediation. This definition is author synthesis in this paper grounded in audit evidence language, AI governance guidance, provenance/logging concepts, privacy/evidence-retention guidance, and GAIC's MRO object layer.

Source Note

Big Four sources are used as market context only. Professional audit and assurance sources support terminology, evidence concepts, internal-control language, and assurance boundaries. AI governance/public guidance sources support governance context. Evidence/provenance/observability/logging sources support the logs-versus-evidence distinction. Privacy sources support evidence minimization, selective disclosure, and retention tension. GAIC source truth supports MROs, Validation Lab boundaries, RCCS-M/ALCS context, and companion-paper boundaries. Inline source IDs in this paper map to the package source register and package citation map.

Boundary Note

This public research edition is not an audit standard, certification, legal compliance proof, assurance opinion, regulator-approved method, procurement recommendation, vendor ranking, Big Four endorsement, audit body endorsement, MPLP requirement, or MPLP industry-standard claim. AARM is a proposed readiness model, not a score, benchmark, certification, assurance result, or vendor assessment.

0. Executive Summary

Chapter purpose: State the auditability problem, the object shift, the evidence-chain answer, the AARM bridge, and the non-standard boundary. **Reader question:** Why does agentic AI need a new auditability frame?

Key Claims

- AI agent auditability cannot be built on logs alone.
- Agentic systems become auditable only when lifecycle work can be reconstructed through responsibility-linked evidence chains.
- Auditability shifts the review object from model output alone to lifecycle work across authority, responsibility, tools, outcomes, exceptions, privacy treatment, and remediation.
- GAIC MROs can be translated into audit evidence objects and evidence requests.
- AARM is a proposed readiness model for describing auditability readiness; it does not issue assurance, certification, or legal compliance proof.

Source Grounding Note

Market context is grounded in BF-01, BF-02, and BF-05. Audit evidence language is grounded in AUD-01 and AUD-02. AI governance context is grounded in AI-01 and AI-08. Observability context is grounded in EVID-02. MROs, Validation Lab boundaries, and companion-paper boundaries are GAIC-derived source truth.

Author Synthesis Note

Agentic AI Auditability, Agentic Audit Object, Audit Evidence Chain, lifecycle-responsibility-linked agent work, and AARM are author synthesis in this paper. They are not presented as externally issued audit standards or certification schemes.

Main Text

AI agent auditability cannot be built on logs alone. Logs matter. Traces matter. Workflow histories, tool-call records, telemetry, and incident records all help reconstruct technical activity. But agentic auditability asks a different question: not merely what happened inside a system, but whether the work can be reviewed as lifecycle responsibility.

The shift is subtle and material. Traditional AI governance often begins with the model, the dataset, the output, the control environment, and the policy framework. Those layers remain necessary.

An agentic system, however, does not only produce outputs. It delegates work, calls tools, transfers tasks across agents, accepts or disputes outcomes, escalates exceptions, stores evidence, crosses vendor or processor boundaries, and may require remediation. A log can record that a tool call occurred. By itself, it usually does not establish who authorized the action, which human role owned the work, whether the agent acted inside a delegated authority boundary, whether the outcome was accepted, what exception occurred, or whether remediation closed.

This paper defines **Agentic AI Auditability** as the ability to reconstruct, test, and evidence agentic lifecycle work across authority, responsibility, tools, outcomes, exceptions, and remediation. This is an authored

definition. It is grounded in professional audit evidence language, AI governance guidance, provenance concepts, observability/logging documentation, privacy guidance, and GAIC's existing lifecycle responsibility object layer.

It is not a legal definition, audit standard, assurance engagement requirement, certification criterion, or regulator-approved compliance test.

The central contribution of this paper is the Audit Evidence Object layer. The Global AI Compliance White Paper 2026 defined Missing Regulatory Objects for agentic lifecycle governance. This paper translates those objects into auditability terms: evidence requests, walkthrough structures, responsibility maps, evidence partitioning, privacy-aware disclosure, third-party review boundaries, and readiness levels.

The Agentic Audit Object is the unit of review. The Audit Evidence Chain is the responsibility-linked structure that connects lifecycle work to authority, role, tool action, evidence pointer, accepted outcome, exception state, and closure state.

This matters because audit and assurance readers increasingly face AI systems whose behavior cannot be understood through model cards, output samples, policy attestations, or raw telemetry alone. Big Four and professional sources already discuss AI assurance, trusted AI, responsible AI, audit transformation, governance, risk, controls, and agentic AI adoption. Those discussions create the market context for this paper.

This paper does not claim that any firm endorses this framework. Instead, it offers an object-model layer adjacent to those discussions: a disciplined way to ask what must exist before agentic work can be reconstructed.

The paper's recurring distinction is simple:

- Logs record activity.
- Observability explains system behavior.
- Audit evidence supports responsibility review.
- Responsibility-linked evidence chains make agentic lifecycle work reconstructable.

The distinction does not make logs unimportant. Logs and traces are often evidence ingredients. They may identify time, sequence, tool, service, user, or execution path. But audit evidence requires more than event capture. It requires relevance, reliability, sufficiency in context, relationship to review objective, and linkage to responsibility. For agentic systems, that means linking activity to lifecycle work units, human roles, agent roles, delegated authority, accepted outcomes, exception handling, privacy treatment, and remediation closure.

This paper also introduces the **Agentic Auditability Readiness Model (AARM)**. AARM describes whether an agentic system is unobservable, log-visible, trace-linked, evidence-structured, auditability-ready, or assurance-ready for planning purposes. AARM is deliberately bounded. It is not a score, benchmark, certification, audit opinion, assurance conclusion, legal compliance proof, regulator approval, procurement recommendation, or vendor ranking. Its purpose is to help organizations and reviewers discuss readiness before they overstate assurance.

For enterprises, this paper prepares the path to two practitioner guides. Guide 1 will translate auditability into technical architecture for CIO, CTO, platform, security, and engineering teams. Guide 2 will translate auditability into compliance operating-model design for CCO, governance, legal, risk, internal audit, and policy-to-evidence teams. This paper defines what audit-ready means; the later guides will address how to build and govern toward it.

For audit and assurance firms, this framework can be used as a discussion and readiness structure. It can support evidence request framing, walkthrough scoping, object-model review, and internal method conversations. It does not replace professional methodology, professional judgment, engagement acceptance, independence rules, attestation criteria, assurance standards, or audit procedures. The paper is useful only if its boundary is kept visible.

Executive Orientation Table

Executive Orientation Table

Paper element	What it means	Source / synthesis status	Boundary
Agentic AI Auditability	Ability to reconstruct, test, and evidence agentic lifecycle work across authority, responsibility, tools, outcomes, exceptions, and remediation	author synthesis in this paper grounded by AUD-01, AUD-02, AI-01, AI-08, EVID-01, EVID-02, and GAIC-SOURCE	Not audit standard, legal proof, or assurance opinion
Agentic Audit Object	Reviewable lifecycle work object for agentic systems	author synthesis in this paper derived from GAIC MROs and provenance/audit evidence concepts	Not mandatory schema
Audit Evidence Chain	Responsibility-linked chain connecting work unit, authority, role, tool, evidence, outcome, exception, privacy treatment, and closure	author synthesis in this paper	Not a certification criterion
MRO-to-evidence mapping	Translation of GAIC MROs into audit evidence objects and evidence requests	GAIC-derived plus audit/control language	MROs are not legal mandates
AARM	Proposed readiness model for auditability	author synthesis in this paper grounded by R4B/R1 and audit/governance sources	Not score, certification, or assurance result

Cross-Links

- See Appendix A for evidence request categories.
- See Appendix C for the full MRO-to-audit-evidence mapping.
- See Appendix D for the AARM readiness matrix.
- See Appendix F for boundary language.

Boundary Note

This chapter does not claim this paper is published, final, sealed, regulator-approved, audit-body-endorsed, Big-Four-endorsed, legally sufficient, certifying, or assurance-producing.

1. Scope, Audience, and Non-Audit-Standard Boundary

Chapter purpose: Define who this paper is for, what it does, what it does not do, how sources are used, and how GAIC/MPLP/Validation Lab relate to the draft. **Reader question:** Who should use this paper, and what authority must it not claim?

Key Claims

- This paper is a lifecycle evidence guide and auditability readiness framework.
- This paper is not an audit standard, assurance opinion, certification, legal compliance proof, regulator approval, procurement recommendation, vendor ranking, or Big Four/audit body endorsement.
- Big Four sources are market context only.
- GAIC is the source for MROs, AARM planning, Validation Lab boundary, and companion-paper scope.
- MPLP may be discussed as one optional lifecycle protocol path, never as required or industry-standard.

Source Grounding Note

Assurance and attestation boundaries use AUD-03 and BOUND-03. Certification and conformity-assessment boundaries use BOUND-01 and BOUND-02. Governance role context uses AUD-07. Big Four AI assurance context uses BF-02 and BF-03 only as market context.

Author Synthesis Note

This paper's scope, Audit Evidence Object layer, Agentic Audit Object, Audit Evidence Chain, and AARM are author synthesis. They are draft framework constructs, not externally adopted professional requirements.

Main Text

This paper is written for readers who need to understand whether agentic AI work can be reconstructed. Its primary audience includes audit and assurance professionals, internal audit teams, technology risk teams, AI governance committees, enterprise control owners, CIOs, CTOs, CCOs, CROs, legal/privacy stakeholders, and platform leaders responsible for agentic systems. Its secondary audience includes advisory teams, AI assurance practitioners, risk engineering teams, agent runtime builders, procurement and vendor-risk teams, and standards or policy readers exploring agentic AI governance.

The paper's scope is deliberately narrow. It does not try to restate all model governance, responsible AI, cybersecurity, privacy, or enterprise risk management. Those fields remain necessary. This paper focuses on the lifecycle evidence layer that becomes necessary when AI systems act through agents, tools, memory, delegation, multi-agent handoffs, and remediation workflows. The question is not simply whether a model behaved acceptably. The question is whether the lifecycle work can be reconstructed across authority, responsibility, tools, outcomes, exceptions, and closure.

This paper is best understood as a lifecycle evidence guide. It defines the proposed object layer needed for audit-ready agentic work, maps GAIC MROs into audit evidence objects, proposes evidence request

categories, provides walkthrough and partitioning patterns, introduces AARM readiness levels, and preserves boundaries around privacy, third-party validation, and assurance. It is a draft architecture for auditability, not a professional standard.

The source hierarchy matters. Professional audit and assurance sources provide language for evidence, controls, assurance boundaries, internal audit, attestation, and professional limits. AI governance sources provide context for risk management, accountability, monitoring, documentation, and human oversight. Evidence, provenance, observability, and log-management sources help distinguish raw event data from responsibility-linked evidence. Privacy sources ground the tension between reviewability, minimization, selective disclosure, and retention. Big Four sources provide market context: they show that AI assurance, trusted AI, audit transformation, and agentic AI are active enterprise topics. They do not endorse this paper.

GAIC source truth plays a different role. The Global AI Compliance White Paper 2026 created the MRO object layer and the broader Agentic Lifecycle Governance framing. This paper uses those GAIC-derived objects as its internal foundation. When this paper discusses MRO-01 through MRO-16, RCCS-M/ALCS context, the Evidence-Based Validation Pattern, Validation Lab boundary language, or companion-paper sequencing, it is relying on GAIC source truth, not external audit standards.

MPLP should be handled with care. MPLP may be described as one protocol path whose lifecycle responsibility semantics can help express evidence objects. This paper does not require MPLP, does not rank MPLP above other approaches, does not claim MPLP is an industry standard, and does not claim MPLP proves compliance, auditability, assurance readiness, or enterprise readiness. Guide 1 may later discuss implementation paths, but this paper must remain implementation-neutral.

Validation Lab is also boundary-sensitive. GAIC describes Validation Lab as a non-certifying evidence adjudication example. This paper may use that boundary as an example of third-party evidence review without certification. It must not describe Validation Lab as a certification body, conformity assessment body, regulator, audit firm, assurance provider, legal authority, or proof of compliance.

The safest reading of this paper is this: it provides a structured way to ask what evidence must exist before agentic AI work can be meaningfully reviewed. It does not determine whether that evidence is sufficient for a particular audit, assurance engagement, legal analysis, regulatory process, insurance underwriting decision, procurement decision, or board conclusion. Those determinations belong to qualified professionals operating under their own methods, scopes, independence requirements, criteria, and legal obligations.

Scope Table

Scope Table

Area	This paper does	This paper does not do	Source / boundary grounding
Auditability	Defines proposed lifecycle evidence objects and readiness questions	Issue audit opinions or audit procedures	AUD-01, AUD-02, AUD-03
Assurance	Frames readiness for possible assurance planning	Provide assurance conclusion or engagement standard	AUD-03, BOUND-03
Certification	Distinguishes validation from certification	Certify systems or define certification criteria	BOUND-01, BOUND-02
Legal/privacy	Frames evidence/privacy tension	Provide legal advice or legal compliance proof	PRIV-01 to PRIV-05, AI-09
Big Four context	Uses public materials as market context	Claim endorsement, adoption, or need	BF-01 to BF-05
GAIC/MRO	Uses MROs as GAIC-derived object layer	Claim MROs are law or external standards	GAIC-SOURCE
MPLP	Treats MPLP as optional protocol path	Require MPLP or claim industry-standard status	GAIC-SOURCE

Area	This paper does	This paper does not do	Source / boundary grounding
Validation Lab	Gives non-certifying evidence adjudication example	Claim certification, assurance opinion, or regulator approval	GAIC-SOURCE, BOUND-01, BOUND-02

Cross-Links

- Appendix F contains reusable boundary language.
- Chapter 12 expands the validation/certification boundary.
- Chapter 15 addresses professional use by audit and assurance firms.

Boundary Note

This chapter intentionally repeats non-claim language because this paper's usefulness depends on not overstating its authority.

2. Why AI Agent Auditability Is Different from Model Governance

Chapter purpose: Explain why model governance is necessary but insufficient for agentic auditability. **Reader question:** What changes when AI systems act through agents, tools, memory, delegation, and multi-agent workflows?

Key Claims

- Model governance remains necessary for agentic systems, but it does not fully reconstruct lifecycle work.
- Agentic AI introduces delegated authority, tool action, role separation, handoffs, accepted outcomes, exception states, and remediation closure.
- Auditability must therefore move beyond model/output review into lifecycle responsibility review.

Source Grounding Note

Big Four market context appears in BF-01, BF-04, and BF-05. AI governance context appears in AI-01, AI-02, AI-04, AI-05, and AI-08. Internal audit context appears in AUD-06. GAIC source truth provides MRO and RCCS-M/ALCS lifecycle-object context.

Author Synthesis Note

The claim that "model governance is necessary but insufficient" for agentic auditability is author synthesis in this paper grounded in the shift from model/output governance to lifecycle responsibility objects.

Main Text

Model governance is not obsolete. It remains part of responsible AI practice. Organizations still need to understand model purpose, data, evaluation, limitations, monitoring, robustness, bias, security, human oversight, and change management. Public AI governance frameworks and professional guidance continue to provide important language for risk management, accountability, documentation, monitoring, controls, and internal audit review.

The limitation is not that model governance is wrong. The limitation is that agentic systems create work that is not reducible to model behavior. An agent may plan a sequence, invoke a tool, pass work to another agent, store state, retrieve memory, request human confirmation, act under delegated authority, trigger an external workflow, or remediate an exception. A model output may be only one event in a broader lifecycle.

For auditability, that broader lifecycle matters. If an agent drafts a message, the model output may be reviewed. If an agent sends the message through a CRM tool, updates a customer record, triggers a refund, routes a contract, or changes a production configuration, the relevant review object is no longer only the output. It is the work unit: who initiated it, what authority permitted it, what agent and tool executed it, what evidence was retained, who accepted the result, what exception occurred, and how any remediation closed.

This creates an object gap. Traditional governance may document the model, evaluate outputs, approve use cases, and monitor performance. Agentic auditability must also document authority boundaries, role mappings, tool-action evidence, responsibility transfer, evidence partitioning, privacy treatment, and closure. These are lifecycle objects, not merely model attributes.

The distinction becomes clearer when human oversight is considered. A human-in-the-loop control may be meaningful only if the loop is tied to an identifiable responsibility state. Who was the human? What role did they occupy? What were they asked to confirm? Was the confirmation advisory, blocking, or final? Was it tied to the delegated authority boundary? Did it create an accepted outcome state? Did it leave evidence that another reviewer can reconstruct? Without these objects, "human oversight" may be real operationally but weak evidentially.

Agentic systems also complicate responsibility because agent roles are not human roles. Human-like names, personas, and task labels may help interfaces, but they can blur governance. An "analyst agent" is not the same thing as an analyst. It is a bounded execution role with tool permissions, instructions, constraints, evidence obligations, and escalation paths. A human or organizational role still owns intent, authority, acceptance, and remediation.

The point of this paper is to place model governance beside lifecycle evidence architecture. Model governance asks whether the model and its use are governed. Agentic auditability asks whether the agentic work can be reconstructed. Both are needed. The second cannot be inferred from the first.

Table 1: Traditional AI Audit vs Agentic AI Auditability

Table 1: Traditional AI Audit vs Agentic AI Auditability

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Model documentation

AGENTIC AUDITABILITY FOCUS	Lifecycle work reconstruction
MISSING LIFECYCLE OBJECT	Agentic Audit Object
EVIDENCE NEEDED	Work unit ID, scope, evidence chain, accepted outcome
SOURCE / SYNTHESIS NOTE	Mixed: AI-01, AI-08, GAIC-SOURCE, author synthesis in this paper
BOUNDARY	Not a new audit standard

Output review

AGENTIC AUDITABILITY FOCUS	Accepted outcome review
MISSING LIFECYCLE OBJECT	Accepted outcome state
EVIDENCE NEEDED	Reviewer role, acceptance/rejection/dispute state, evidence pointer
SOURCE / SYNTHESIS NOTE	GAIC-derived + AUD-01/AUD-06
BOUNDARY	Acceptance is not legal proof

Monitoring

AGENTIC AUDITABILITY FOCUS	Authority and exception monitoring
MISSING LIFECYCLE OBJECT	Authority drift record
EVIDENCE NEEDED	Baseline authority, observed action, escalation, closure evidence
SOURCE / SYNTHESIS NOTE	AI-02, EVID-04, GAIC-SOURCE
BOUNDARY	Drift is not automatically a legal breach

Access control

AGENTIC AUDITABILITY FOCUS	Delegated authority review
MISSING LIFECYCLE OBJECT	Delegated authority boundary
EVIDENCE NEEDED	Scope, condition, expiry, tool action, escalation
SOURCE / SYNTHESIS NOTE	AI-08, AUD-06, GAIC-SOURCE
BOUNDARY	IAM permission is not full business authority

Control testing

AGENTIC AUDITABILITY FOCUS	Lifecycle responsibility walkthrough
MISSING LIFECYCLE OBJECT	Responsibility map and evidence chain
EVIDENCE NEEDED	Human role, agent role, tool record, partitioned evidence
SOURCE / SYNTHESIS NOTE	AUD-04, AUD-05, EVID-01
BOUNDARY	Walkthrough is not formal audit procedure

Table note: Mixed source-grounded, GAIC-derived, and author-synthesis table. It is not a certification table, audit procedure, procurement tool, or vendor comparison.

Cross-Links

- Chapter 3 defines the audit object shift.
- Chapter 5 defines the Agentic Audit Object Model.
- Appendix C maps MROs to audit evidence objects.

Boundary Note

This chapter does not attack model governance or claim model governance is useless. It argues that agentic auditability requires an additional lifecycle responsibility layer.

3. The Audit Object Shift: From Model Output to Lifecycle Work

Chapter purpose: Define the proposed review object for agentic auditability. **Reader question:** What exactly should audit stakeholders reconstruct?

Key Claims

- Agentic auditability requires a shift from model/output review to lifecycle work review.
- The Agentic Audit Object is a proposed review object for lifecycle-responsibility-linked agent work.
- Lifecycle-responsibility-linked agent work connects intent, authority, role, agent action, tool action, evidence, outcome, exception, privacy treatment, and closure.

Source Grounding Note

Audit evidence language is grounded in AUD-01 and AUD-02. Governance context is grounded in AI-01 and AI-08. Provenance concepts are grounded in EVID-01. MRO source truth is GAIC-derived.

Author Synthesis Note

Agentic Audit Object and lifecycle-responsibility-linked agent work are author-synthesis constructs. They are not claimed as existing professional standards, legal categories, or regulator-defined objects.

Main Text

An audit object is the thing being reviewed. In many AI governance conversations, the object is a model, dataset, output, use case, control, system, or policy. Those objects remain important. But agentic systems create another review object: lifecycle work.

Lifecycle work is not a single output. It is the chain of actions and decisions by which an agentic system receives intent, acts under authority, uses tools, transfers work, records evidence, produces a result, reaches acceptance or dispute, handles exceptions, and closes remediation. If that chain cannot be reconstructed, the system may be observable but not audit-ready.

This paper calls the proposed review unit the **Agentic Audit Object**. It is a structured representation of lifecycle-responsibility-linked agent work. It is not a mandatory data schema. It is a conceptual object model that helps audit, assurance, governance, and technology teams ask whether the relevant lifecycle work can be reviewed.

The Agentic Audit Object exists because agentic systems blur the boundary between process and output. A model-generated draft may be an output. A tool action that sends the draft, records the transaction, updates a system, triggers a workflow, or changes a downstream state is lifecycle work. The auditability question changes from "Was the output acceptable?" to "Can the lifecycle work that produced and accepted that outcome be reconstructed?"

The object shift also changes the meaning of responsibility. Responsibility cannot be inferred from a username, service account, model name, trace ID, or tool log alone. A responsibility-linked review object should identify the human role responsible for intent, authority, review, acceptance, exception handling, and remediation. It should separately identify the agent role and tool surface that executed the work. It should preserve the difference between technical execution and business responsibility.

Provenance concepts help explain why this structure matters. Provenance language distinguishes entities, activities, agents, and relationships. Audit evidence language asks whether evidence is relevant and reliable for the review objective. this paper combines those ideas with GAIC's MRO layer: lifecycle work must be addressable, linked, partitioned, and reviewable.

The Agentic Audit Object therefore contains more than a log. It contains the work unit, the authority boundary, the human role, the agent role, the tool-action record, the evidence pointer, the accepted outcome state, the exception state, the remediation closure state, and the privacy/selective disclosure profile. Those fields may live in different systems. The model is not saying every enterprise must store them in one database. It says that the relationships must be reconstructable.

This object shift is also where this paper becomes distinct from the Global AI Compliance White Paper 2026. The Global AI Compliance White Paper 2026 identifies the missing regulatory objects for agentic lifecycle governance. This paper asks how those objects become evidence. The same MRO can be read as a governance gap, an engineering object, or an audit evidence object. This paper focuses on the third reading.

Agentic Audit Object Overview

Object component	Review question	Evidence examples	Source / synthesis status	Boundary
Lifecycle work unit	What work is under review?	Work unit ID, scope, task intent, lifecycle stage	author synthesis in this paper grounded by EVID-01 and GAIC-SOURCE	Not mandatory schema
Authority boundary	Was the action authorized within scope?	Delegation record, condition, expiry, escalation path	GAIC-derived + AI-01/AI-08	Not proof of legal delegation
Human role responsibility	Who owned intent, review, acceptance, or closure?	Role map, review record, acceptance owner	GAIC-derived + AUD-06/AUD-07	Not legal liability assignment
Agent role responsibility surface	What agent role executed or transformed the work?	Agent role ID, constraints, capability boundary	GAIC-derived author synthesis + EVID-01	Agent is not a legal person
Tool-action evidence	What external or consequential action occurred?	Tool call record, affected system, reversibility, rollback evidence	GAIC-derived + EVID-02/EVID-03	Not legal liability conclusion
Accepted outcome	Was the result accepted, disputed, rejected, or remediated?	Acceptance state, reviewer role, dispute reason	GAIC-derived + AUD-01/AUD-04	Not compliance proof
Exception / remediation closure	How was deviation handled and closed?	Exception record, corrective action, recheck, closure owner	GAIC-derived + EVID-04/AUD-04	Not legal settlement

Table note: Mixed GAIC-derived and author-synthesis table grounded in audit evidence, provenance, and governance sources. It is not an audit standard or legal template.

Cross-Links

- Chapter 5 expands the object model fields.
- Chapter 6 maps each MRO into audit evidence objects.
- Appendix B provides a walkthrough template.

Boundary Note

The Agentic Audit Object is a proposed object model in this paper for reviewability. It does not create legal categories, professional requirements, certification criteria, or mandatory implementation schema.

4. Why Logs Are Not Audit Evidence Chains

Chapter purpose: Ground the core this paper claim that logs and traces are useful but insufficient for agentic auditability. **Reader question:** What can logs show, and what do they fail to prove?

Key Claims

- Logs, traces, metrics, and workflow histories can support auditability.
- They do not automatically prove authority, responsibility, accepted outcome, exception handling, privacy treatment, or remediation closure.
- Audit Evidence Chains increase evidentiary value by linking activity to lifecycle responsibility metadata.

Source Grounding Note

Audit evidence language is grounded in AUD-01 and AUD-02. Observability and logging context is grounded in EVID-02 and EVID-03. Provenance concepts are grounded in EVID-01. Incident/governance context is grounded in EVID-04. AI Act logging context may be referenced as legal context only through AI-06.

Author Synthesis Note

"Logs are not audit evidence chains" is author synthesis in this paper. It is grounded in the distinction between observability/logging and professional audit evidence concepts.

Main Text

Logs are necessary but not sufficient. That sentence is the safest starting point for agentic auditability. The point is not to diminish logs, traces, metrics, monitoring systems, workflow histories, or observability platforms. They are often the first evidence ingredients available. Without them, reconstruction may collapse into narrative memory.

The problem is that logs usually describe activity, not responsibility. A log may show that a service called a tool at a time. A trace may show that a workflow crossed several components. A metric may show latency, error rate, or throughput. An incident record may show that something failed. Those are useful signals. They do not automatically tell a reviewer whether the action was authorized, which human role owned it, which agent role was acting, what business scope applied, whether the tool action was reversible, whether the outcome was accepted, what privacy treatment applied, or whether remediation closed.

Audit evidence is tied to a review objective. In professional audit language, evidence is evaluated for relevance and reliability within context. For agentic systems, the review objective is not merely to explain system behavior. It is to reconstruct lifecycle responsibility. That means the evidence must connect the technical event to authority, role, tool action, outcome state, exception state, privacy treatment, and closure.

Observability explains how systems behave. Auditability asks whether lifecycle work can be reviewed. The same trace can be valuable for both, but it plays different roles. In observability, the trace helps engineers understand execution path and performance. In auditability, the trace is an input to a larger evidence chain. It must be joined with responsibility records, authority records, evidence pointers, acceptance records, and exception/remediation records.

Consider a tool call. A log may show the tool name, timestamp, endpoint, response code, and service identity. The auditability question is broader: Who authorized the tool action? Was the agent operating inside delegated scope? Which data or system was affected? Was the action reversible? Did a human review occur before or after the action? Was the result accepted, disputed, or remediated? Was sensitive data minimized or redacted in review evidence? Was there a closure state?

The same distinction applies to incident records. An incident log may show that an error occurred. It may not show whether the affected outcome was withdrawn, corrected, re-reviewed, accepted, escalated, or closed. It may not show whether the issue was due to authority drift, responsibility transfer, cross-project reuse, tool-action side effects, privacy leakage, or vendor/runtime substitution. Without linkage to lifecycle objects, the incident is visible but not fully reconstructable.

This paper calls the needed structure an **Audit Evidence Chain**. An Audit Evidence Chain is not just a sequence of logs. It is a responsibility-linked reconstruction path. It ties a lifecycle work unit to the relevant evidence objects: authority boundary, human role, agent role, tool action, evidence pointer, accepted outcome, exception, remediation, privacy treatment, and third-party review boundary where relevant.

This structure also protects against over-retention. A naive response to auditability gaps is to collect everything. That can create privacy and confidentiality risks. A better approach is to preserve evidence pointers, partition keys, disclosure profiles, redaction rules, retention logic, and integrity records so reviewers can reconstruct what they are authorized to inspect without unnecessary exposure.

Table 2: Logs vs Audit Evidence Chains

Table 2: Logs vs Audit Evidence Chains

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Tool-call log

WHAT LOGS/TRACES CAN SHOW	Tool, endpoint, timestamp, service, response
WHAT THEY DO NOT AUTOMATICALLY PROVE	Business authority, human owner, reversibility, accepted outcome
REQUIRED LINKAGE	Delegated authority record, human role map, tool-action evidence, outcome state
RELATED MRO	MRO-02, MRO-05, MRO-04
SOURCE / SYNTHESIS NOTE	EVID-02/EVID-03 + GAIC-SOURCE

Trace span

WHAT LOGS/TRACES CAN SHOW	Execution path, service dependencies, timing
WHAT THEY DO NOT AUTOMATICALLY PROVE	Lifecycle responsibility or review sufficiency
REQUIRED LINKAGE	Work unit ID, agent role, evidence pointer, role map
RELATED MRO	MRO-01, MRO-03, MRO-08
SOURCE / SYNTHESIS NOTE	EVID-01/EVID-02 + author synthesis in this paper

Workflow state

WHAT LOGS/TRACES CAN SHOW	Step completion, state transition
WHAT THEY DO NOT AUTOMATICALLY PROVE	Whether state is a governance acceptance state
REQUIRED LINKAGE	Acceptance criteria, reviewer role, dispute/remediation state
RELATED MRO	MRO-04, MRO-16
SOURCE / SYNTHESIS NOTE	AUD-01/AUD-04 + GAIC-SOURCE

Incident log

WHAT LOGS/TRACES CAN SHOW	Error event, detection, response action
WHAT THEY DO NOT AUTOMATICALLY PROVE	Remediation closure or legal fault
REQUIRED LINKAGE	Exception owner, corrective action, recheck, closure acceptance
RELATED MRO	MRO-07, MRO-16
SOURCE / SYNTHESIS NOTE	EVID-04/AUD-04

Access log

WHAT LOGS/TRACES CAN SHOW	User/service access event
WHAT THEY DO NOT AUTOMATICALLY PROVE	Delegated business authority or scope
REQUIRED LINKAGE	Authority boundary, confirmation rule, risk class, expiry
RELATED MRO	MRO-02, MRO-07
SOURCE / SYNTHESIS NOTE	AI-01/AI-08 + author synthesis in this paper

Retention log

WHAT LOGS/TRACES CAN SHOW	Data/evidence storage or deletion event
WHAT THEY DO NOT AUTOMATICALLY PROVE	Whether retention/disclosure is legally sufficient
REQUIRED LINKAGE	Privacy treatment, minimization rule, retention basis, disclosure profile
RELATED MRO	MRO-10, MRO-12, MRO-13
SOURCE / SYNTHESIS NOTE	PRIV-01 to PRIV-05

Table note: Mixed source-grounded and author-synthesis table. Logs are treated as useful inputs, not as useless artifacts. The table is not a vendor critique or audit procedure.

Cross-Links

- Appendix A lists evidence request categories.
- Chapter 9 addresses evidence partitioning.
- Chapter 10 addresses minimization and selective disclosure.

Boundary Note

This chapter does not claim observability vendors fail, logs are useless, or telemetry can never support audit evidence. It claims only that raw logs do not automatically become responsibility-linked audit evidence chains.

5. The Agentic Audit Object Model

Chapter purpose: Define the proposed object model for audit-ready agentic lifecycle work. **Reader question:** What makes lifecycle work audit-ready?

Key Claims

- Audit-ready agentic lifecycle work should be addressable, reconstructable, partitioned, privacy-aware, and reviewable.
- The Agentic Audit Object Model binds lifecycle work unit, authority boundary, human role responsibility, agent role responsibility, tool-action evidence, evidence pointer, accepted outcome, exception state, remediation closure, and privacy treatment.
- The model is a proposed review architecture, not a mandatory schema or standard.

Source Grounding Note

The model uses audit evidence language from AUD-01, control vocabulary from AUD-04, AI governance context from AI-01, and provenance structure from EVID-01. MRO dependencies are GAIC-derived.

Author Synthesis Note

The Agentic Audit Object Model is author synthesis in this paper. It is a conceptual object model for auditability, not a required implementation format.

Main Text

The Agentic Audit Object Model translates lifecycle work into a reviewable structure. It does not require that every enterprise use the same database table, evidence pack format, or protocol. It asks whether the relationships needed for review can be reconstructed.

The first requirement is an addressable lifecycle work unit. A reviewer needs to know what work is under review: the task, workflow, agent run, delegated action, tool call, handoff, or remediation event. Without a stable work unit, evidence remains scattered across logs, tickets, approvals, messages, policies, and memories.

The second requirement is an authority boundary. Agentic systems often combine technical permission with business action. Technical permission answers whether a service can call a tool. Business authority answers whether that action was allowed under scope, condition, risk class, and responsibility. Auditability requires evidence of the authority boundary, not merely evidence that the tool call succeeded.

The third requirement is role mapping. Human role responsibility and agent role responsibility surface must be separated. The human role owns intent, approval, review, acceptance, escalation, or closure. The agent role describes bounded execution capability, constraints, tool access, and evidence obligations. Confusing these roles weakens accountability and makes review harder.

The fourth requirement is tool-action evidence. A tool action is often where AI-generated output becomes operational consequence. The evidence object should identify the tool, action, target system, affected data or

process, initiator, authority basis, reversibility, rollback path, and owner. This does not assign legal liability. It creates reviewability.

The fifth requirement is an evidence pointer. Auditability does not always require raw evidence to be duplicated into a central repository. It may require pointers to evidence, integrity hashes, redaction profiles, retention rules, and partition keys. The pointer must be strong enough that authorized reviewers can reconstruct the chain.

The sixth requirement is an accepted outcome state. An output is not the same as an accepted outcome. The model should distinguish produced, reviewed, accepted, rejected, disputed, remediated, and closed states. This is especially important when a downstream process treats agent work as completed before a responsible human or governance workflow has accepted it.

The seventh requirement is exception and remediation closure. Auditability is weak if exceptions disappear into operational tickets without lifecycle linkage. Exceptions should connect to authority, role, affected outcome, corrective action, recheck, closure owner, and reopen criteria.

The eighth requirement is privacy treatment. Evidence is not improved by unlimited collection. Reviewability must be balanced with minimization, selective disclosure, redaction, retention, and access control. The object model should identify which evidence can be disclosed, to whom, for what review purpose, and under what boundary.

Taken together, these fields create the Audit Evidence Chain. The chain does not replace professional judgment about sufficiency. It makes judgment possible by giving reviewers a reconstructable path.

Agentic Audit Object Model Table

Agentic Audit Object Model Table

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Lifecycle work unit ID

PURPOSE	Names the work under review
EXAMPLE EVIDENCE	Workflow/run/task ID, scope, intent
RELATED MRO	MRO-08
SOURCE / SYNTHESIS STATUS	author synthesis in this paper grounded in EVID-01
BOUNDARY	Not mandatory schema

Authority boundary

PURPOSE	Records delegated scope and limits
EXAMPLE EVIDENCE	Delegation record, condition, expiry, escalation rule
RELATED MRO	MRO-02, MRO-07
SOURCE / SYNTHESIS STATUS	GAIC-derived + AI-01/AI-08
BOUNDARY	Not legal delegation proof

Human role responsibility

PURPOSE	Maps human ownership
EXAMPLE EVIDENCE	Intent owner, review owner, acceptance owner, closure owner
RELATED MRO	MRO-01
SOURCE / SYNTHESIS STATUS	GAIC-derived + AUD-06/AUD-07
BOUNDARY	Not legal liability assignment

Agent role responsibility surface

PURPOSE	Defines bounded agent function
EXAMPLE EVIDENCE	Agent role, constraints, capability boundary, evidence duties
RELATED MRO	MRO-03, MRO-06
SOURCE / SYNTHESIS STATUS	GAIC-derived author synthesis + EVID-01
BOUNDARY	Agent is not human role

Tool-action evidence

PURPOSE	Records consequential action
EXAMPLE EVIDENCE	Tool ID, target system, affected data, rollback path
RELATED MRO	MRO-05
SOURCE / SYNTHESIS STATUS	GAIC-derived + EVID-02/EVID-03
BOUNDARY	Not liability conclusion

Evidence pointer

PURPOSE	Links to review evidence
EXAMPLE EVIDENCE	URI, hash, evidence pack reference, partition key
RELATED MRO	MRO-08, MRO-12
SOURCE / SYNTHESIS STATUS	Mixed
BOUNDARY	Not blanket disclosure

Accepted outcome

PURPOSE	Records outcome governance state
EXAMPLE EVIDENCE	Accepted/rejected/disputed/remediated status
RELATED MRO	MRO-04
SOURCE / SYNTHESIS STATUS	GAIC-derived + AUD-01/AUD-04
BOUNDARY	Not compliance proof

Exception state

PURPOSE	Records deviation or dispute
EXAMPLE EVIDENCE	Exception ID, trigger, owner, escalation path
RELATED MRO	MRO-07, MRO-16
SOURCE / SYNTHESIS STATUS	GAIC-derived + EVID-04
BOUNDARY	Not legal breach finding

Remediation closure

PURPOSE	Records corrective action and closure
EXAMPLE EVIDENCE	Corrective action, recheck, closure reviewer, reopen criteria
RELATED MRO	MRO-16
SOURCE / SYNTHESIS STATUS	GAIC-derived + AUD-04/AUD-05
BOUNDARY	Not settlement or regulator closure

Privacy treatment	
PURPOSE	Defines minimization and disclosure
EXAMPLE EVIDENCE	Redaction profile, retention rule, access scope
RELATED MRO	MRO-10 to MRO-13
SOURCE / SYNTHESIS STATUS	PRIV sources + GAIC synthesis
BOUNDARY	Not legal advice

Table note: This table is mixed source-grounded, GAIC-derived, and author-synthesis. It is a drafting model, not an audit standard, certification checklist, or legal template.

Cross-Links

- Chapter 6 maps the model to all MROs.
- Appendix A converts fields into evidence requests.
- Appendix B turns fields into a walkthrough template.

Boundary Note

The Agentic Audit Object Model should be read as a review architecture. It does not require MPLP, Cognitive OS, Validation Lab, or any particular vendor system.

6. MRO-to-Audit-Evidence Mapping

Chapter purpose: Translate GAIC MROs into audit evidence objects and evidence requests. **Reader question:** How does each MRO become evidence?

Key Claims

- GAIC MROs can be mapped into reviewable audit evidence objects.
- The mapping turns lifecycle governance objects into evidence request categories.
- MROs remain GAIC-derived governance objects; they are not legal mandates, audit standards, or certification criteria.

Source Grounding Note

MRO names and meanings use GAIC source truth. Audit evidence language uses AUD-01. AI audit/control context uses AUD-05 and AI-01.

Author Synthesis Note

The MRO-to-audit-evidence mapping is author synthesis in this paper derived from GAIC objects and externally sourced audit/control terminology.

Main Text

The Global AI Compliance White Paper 2026 introduced the sixteen Missing Regulatory Objects for agentic and multi-agent system governance. This paper reads those objects through an auditability lens. Each MRO asks a governance question. This paper adds the evidence question: what object would a reviewer need to inspect to reconstruct the lifecycle work?

This translation is not a claim that MROs are law. It is not a claim that the presence of an MRO proves compliance. It is not a claim that absence proves non-compliance. MROs are GAIC-derived lifecycle responsibility objects. Their value for this paper is that they give auditability a vocabulary for the work that logs alone do not capture.

The mapping also prevents a common failure in agentic governance: treating "evidence" as a flat export. Agentic systems do not need a pile of logs; they need evidence structured around review questions. Who owned the action? What authority applied? Which agent role acted? What tool created consequence? Was the outcome accepted? What exception occurred? What data was exposed? What changed after vendor, model, or runtime substitution? How did remediation close?

The sixteen MROs can be grouped into six auditability clusters:

1. Responsibility and role mapping: MRO-01, MRO-03, MRO-06.
2. Authority and action boundaries: MRO-02, MRO-05, MRO-07.
3. Outcome and closure: MRO-04, MRO-16.
4. Evidence partitioning and reuse: MRO-08, MRO-09.
5. Privacy and disclosure: MRO-10, MRO-11, MRO-12, MRO-13.

6. Vendor, processor, and substitution chains: MRO-14, MRO-15.

This cluster view helps readers see that auditability is not one control. It is the ability to reconstruct a lifecycle. A work unit can fail auditability at many points: no role map, no authority record, no accepted outcome state, no evidence partition, no privacy treatment, no substitution record, or no remediation closure.

Table 3: MRO-to-Audit-Evidence Mapping

Table 3: MRO-to-Audit-Evidence Mapping

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

MRO-01 Human Role to MAS Responsibility Mapping

AUDIT EVIDENCE OBJECT	Human-role responsibility map
SAMPLE EVIDENCE REQUEST	Show accountable role for intent, authorization, review, acceptance, exception, and remediation
SOURCE SUPPORT	AUD-06, AUD-07, AI-01
GUIDE 1 RELATION	Role/evidence architecture
GUIDE 2 RELATION	Responsibility governance
BOUNDARY	No legal liability assignment

MRO-02 Delegated Authority Boundary

AUDIT EVIDENCE OBJECT	Delegated authority record
SAMPLE EVIDENCE REQUEST	Show scope, condition, expiry, escalation, and reauthorization evidence
SOURCE SUPPORT	AI-01, AI-06, AI-08, AUD-06
GUIDE 1 RELATION	Authority control implementation
GUIDE 2 RELATION	Delegation policy-to-evidence
BOUNDARY	Not equal to IAM permission

MRO-03 Agent Role is not Human Role

AUDIT EVIDENCE OBJECT	Agent role boundary object
SAMPLE EVIDENCE REQUEST	Show agent role, capability limits, constraints, owner, and escalation boundary
SOURCE SUPPORT	AI-08, EVID-01, AI-01
GUIDE 1 RELATION	Agent role design
GUIDE 2 RELATION	Accountability model
BOUNDARY	Do not anthropomorphize agents

MRO-04 Accepted Outcome Compliance

AUDIT EVIDENCE OBJECT	Accepted outcome record
SAMPLE EVIDENCE REQUEST	Show acceptance criteria, reviewer, acceptance/rejection/dispute state, and evidence link
SOURCE SUPPORT	AUD-01, AUD-02, AUD-04
GUIDE 1 RELATION	Outcome state capture
GUIDE 2 RELATION	Review/acceptance governance
BOUNDARY	Acceptance is not compliance proof

MRO-05 Tool-Action Liability Boundary

AUDIT EVIDENCE OBJECT	Tool-action evidence object
SAMPLE EVIDENCE REQUEST	Show tool identity, action, affected system/data, authority basis, reversibility, rollback path, owner
SOURCE SUPPORT	EVID-02, EVID-03, EVID-04, AUD-05
GUIDE 1 RELATION	Tool evidence capture
GUIDE 2 RELATION	Tool-action approval policy
BOUNDARY	No legal liability conclusion

MRO-06 Responsibility Transfer Across Agents

AUDIT EVIDENCE OBJECT	Responsibility transfer record
SAMPLE EVIDENCE REQUEST	Show source/target agent, transferred scope, retained scope, inherited constraints, acceptance/rejection
SOURCE SUPPORT	EVID-01, AI-08, AUD-06
GUIDE 1 RELATION	Handoff record design
GUIDE 2 RELATION	Transfer governance
BOUNDARY	Not legal transfer

MRO-07 Authority Drift

AUDIT EVIDENCE OBJECT	Authority drift exception record
SAMPLE EVIDENCE REQUEST	Show baseline authority, observed action, trigger, escalation, stop/downgrade/reauthorization
SOURCE SUPPORT	AI-01, AI-02, EVID-04, AUD-04
GUIDE 1 RELATION	Drift monitoring
GUIDE 2 RELATION	Exception governance
BOUNDARY	Not every drift is a legal breach

MRO-08 MAS Evidence Partitioning

AUDIT EVIDENCE OBJECT	Partitioned evidence chain
SAMPLE EVIDENCE REQUEST	Show partition keys, evidence pointers, link identifiers, integrity hashes, access rules
SOURCE SUPPORT	EVID-01, EVID-02, EVID-03, AUD-01
GUIDE 1 RELATION	Evidence partitioning architecture
GUIDE 2 RELATION	Evidence access governance
BOUNDARY	No blanket retention/disclosure

MRO-09 Cross-Project Reuse Compliance

AUDIT EVIDENCE OBJECT	Reuse context validation record
SAMPLE EVIDENCE REQUEST	Show source context, target context, reset validation, privacy review, authorization refresh
SOURCE SUPPORT	AI-01, AI-05, PRIV-02, AUD-04
GUIDE 1 RELATION	Reuse validation workflow
GUIDE 2 RELATION	Reuse governance
BOUNDARY	No legal reuse clearance

MRO-10 Privacy / GDPR Lifecycle Mapping

AUDIT EVIDENCE OBJECT	Privacy lifecycle evidence object
SAMPLE EVIDENCE REQUEST	Show data flow, processing purpose, retention rule, rights workflow, minimization treatment
SOURCE SUPPORT	PRIV-01 to PRIV-05, AI-09
GUIDE 1 RELATION	Privacy-aware evidence capture
GUIDE 2 RELATION	Privacy governance
BOUNDARY	No legal advice or GDPR proof

MRO-11 Privacy-Preserving Third-Party Validation

AUDIT EVIDENCE OBJECT	Validation disclosure profile and verdict record
SAMPLE EVIDENCE REQUEST	Show ruleset identity, redaction profile, evidence pointer, verdict hash, reviewer boundary
SOURCE SUPPORT	PRIV-02, PRIV-03, BOUND-01, BOUND-02, AUD-03
GUIDE 1 RELATION	Evidence export/replay
GUIDE 2 RELATION	Third-party review governance
BOUNDARY	No certification claim

MRO-12 Evidence Minimization and Selective Disclosure

AUDIT EVIDENCE OBJECT	Evidence minimization and disclosure object
SAMPLE EVIDENCE REQUEST	Show minimization rule, disclosure justification, redaction profile, access log, expiry
SOURCE SUPPORT	PRIV-01 to PRIV-05
GUIDE 1 RELATION	Selective disclosure design
GUIDE 2 RELATION	Disclosure policy
BOUNDARY	No privacy compliance proof

MRO-13 Data Subject Rights vs Evidence Retention

AUDIT EVIDENCE OBJECT	Rights-retention reconciliation record
SAMPLE EVIDENCE REQUEST	Show request type, retention basis, hold status, deletion/redaction action, owner
SOURCE SUPPORT	PRIV-01, PRIV-02, PRIV-05, AI-09
GUIDE 1 RELATION	Rights-aware retention
GUIDE 2 RELATION	Legal/privacy governance
BOUNDARY	Requires legal review

MRO-14 Third-Party Processor / Subprocessor Chain

AUDIT EVIDENCE OBJECT	Processor/subprocessor chain evidence
SAMPLE EVIDENCE REQUEST	Show processor role, subprocessor chain, data agreement pointer, responsibility owner
SOURCE SUPPORT	AI-09, PRIV-01, PRIV-02, AUD-04
GUIDE 1 RELATION	Vendor evidence partitioning
GUIDE 2 RELATION	Processor governance
BOUNDARY	No vendor ranking

MRO-15 Vendor / Model / Runtime Substitution Conformance

AUDIT EVIDENCE OBJECT	Substitution conformance record
SAMPLE EVIDENCE REQUEST	Show prior/new component, control evidence, evidence-integrity check, reauthorization, revalidation
SOURCE SUPPORT	AI-01, AI-04, AI-05, AUD-04, BOUND-02
GUIDE 1 RELATION	Runtime/model change evidence
GUIDE 2 RELATION	Change governance
BOUNDARY	No conformance certification

MRO-16 Incident, Dispute, and Remediation Closure

AUDIT EVIDENCE OBJECT	Incident/remediation closure record
SAMPLE EVIDENCE REQUEST	Show event, affected outcome, owner, corrective action, recheck, closure acceptance, reopen condition
SOURCE SUPPORT	EVID-04, AI-02, AUD-04, AUD-05, AUD-06
GUIDE 1 RELATION	Remediation workflow evidence
GUIDE 2 RELATION	Closure governance
BOUNDARY	No legal closure or regulator acceptance

Table note: GAIC-derived and author-synthesis mapping grounded in audit/control/governance source language. It is not a legal checklist, audit procedure, certification criterion, or procurement tool.

Cross-Links

- Appendix C contains the expanded mapping.
- Chapter 13 maps the same objects into AARM dimensions.
- Guide 1 will later translate this table into implementation planning.
- Guide 2 will later translate it into compliance operating-model planning.

Boundary Note

This chapter does not invent new MRO numbering. It uses current GAIC source truth for MRO-01 through MRO-16.

7. Evidence Request List for Agentic Systems

Chapter purpose: Define the logic of evidence requests for agentic systems without turning them into formal audit procedures. **Reader question:** What should reviewers ask to reconstruct agentic lifecycle work?

Key Claims

- Evidence requests should be organized by lifecycle stage and review objective.
- Agentic evidence requests must cover role, authority, agent/tool action, outcomes, exceptions, privacy, partitioning, and closure.
- Evidence requests are readiness architecture, not universal legal demands or formal audit procedures.

Source Grounding Note

Audit evidence and evidence sufficiency language is grounded in AUD-01 and AUD-02. Control and AI audit practice context comes from AUD-04, AUD-05, and AUD-06. AI governance/logging context may use AI-06.

Author Synthesis Note

The evidence request taxonomy is author synthesis in this paper derived from GAIC MROs and professional audit/control language.

Main Text

Agentic auditability becomes practical when it is expressed as evidence requests. A high-level statement that "the system is logged" is not enough. A reviewer needs to know what evidence should exist, where it lives, which lifecycle stage it supports, what role owns it, how sensitive fields are treated, and what boundary applies.

The evidence request list should begin with the lifecycle work unit. The reviewer should be able to ask: What work was initiated? What purpose and scope did it have? Who initiated it? Which agent or workflow received it? Which policy, authority, or business context applied? Without a work unit, later evidence cannot be connected.

The second request category is authority. What delegated authority allowed the agent to act? Was the action within scope? Did the authority have conditions, risk class, expiry, confirmation, or escalation requirements? If authority changed, was the change recorded? If authority drift occurred, was it escalated and closed?

The third category is responsibility. Which human role owned intent, review, acceptance, exception handling, and remediation? Which agent role executed the work? Was a responsibility transfer recorded across agents? Was a tool action tied back to a human or organizational owner? This category distinguishes accountability from technical execution.

The fourth category is tool-action evidence. Tool calls are where agent behavior often becomes operational consequence. Evidence should identify the tool, action, target system, affected data or process, response, reversibility, rollback path, and owner. The request is not asking for legal liability conclusions. It is asking for reconstructability.

The fifth category is outcome evidence. Reviewers should distinguish produced output from accepted outcome. Was the output accepted, rejected, disputed, escalated, or remediated? Who reviewed it? Which criteria or policy were applied? What evidence supports the final state?

The sixth category is exception and remediation. What event occurred? What lifecycle object was affected? Who owned the exception? What corrective action was taken? Was the action rechecked? Who accepted closure? What would reopen the issue?

The seventh category is privacy and selective disclosure. Evidence should not be over-collected or over-disclosed. Reviewers should ask for data categories, redaction profiles, retention rules, disclosure scope, access logs, and minimization rationale. Legal interpretation remains outside this paper, but the evidence architecture should make privacy-aware review possible.

Finally, evidence requests should include boundary statements. Each request should clarify whether it supports readiness, reconstruction, control review, third-party review, or assurance planning. It should also clarify what it does not prove.

Table 4: Lifecycle Stage -> Audit Evidence Request

Table 4: Lifecycle Stage → Audit Evidence Request

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Intent

CORE REVIEW QUESTION	What work was initiated and why?
EVIDENCE REQUESTED	Work unit ID, purpose, scope, initiating context
HUMAN ROLE	Request owner
AGENT/TOOL RECORD	Initial agent/workflow assignment
OUTCOME/CLOSURE	No accepted outcome yet
PRIVACY TREATMENT	Purpose and data categories

Delegation

CORE REVIEW QUESTION	Was action authorized within scope?
EVIDENCE REQUESTED	Authority grant, conditions, expiry, confirmation boundary
HUMAN ROLE	Delegating/approving role
AGENT/TOOL RECORD	Agent permission and tool scope
OUTCOME/CLOSURE	Escalation path
PRIVACY TREATMENT	Minimum necessary authority evidence

Planning

CORE REVIEW QUESTION	What plan or route was selected?
EVIDENCE REQUESTED	Plan version, constraints, review flags
HUMAN ROLE	Plan reviewer if required
AGENT/TOOL RECORD	Agent plan trace
OUTCOME/CLOSURE	Plan accepted/rejected state
PRIVACY TREATMENT	Redacted sensitive prompts if needed

Tool action

CORE REVIEW QUESTION	What external consequence occurred?
EVIDENCE REQUESTED	Tool identity, action, target, response, affected data, reversibility
HUMAN ROLE	Tool-action owner
AGENT/TOOL RECORD	Tool-call record, service trace
OUTCOME/CLOSURE	Action status and rollback path
PRIVACY TREATMENT	Data minimization and access scope

Handoff

CORE REVIEW QUESTION	Was responsibility transferred?
EVIDENCE REQUESTED	Source/target agent, transferred scope, retained scope, acceptance/rejection
HUMAN ROLE	Handoff owner
AGENT/TOOL RECORD	Agent transfer record
OUTCOME/CLOSURE	Transfer state
PRIVACY TREATMENT	Partition and disclosure rule

Review

CORE REVIEW QUESTION	Was work reviewed before acceptance?
EVIDENCE REQUESTED	Review record, criteria, evidence pointer
HUMAN ROLE	Reviewer role
AGENT/TOOL RECORD	Supporting traces/outputs
OUTCOME/CLOSURE	Accepted/rejected/disputed
PRIVACY TREATMENT	Redaction profile

Exception	
CORE REVIEW QUESTION	What deviation occurred?
EVIDENCE REQUESTED	Exception record, baseline, trigger, impact
HUMAN ROLE	Escalation owner
AGENT/TOOL RECORD	Incident trace or alert
OUTCOME/CLOSURE	Open/remediated/closed
PRIVACY TREATMENT	Sensitive-data handling

Closure	
CORE REVIEW QUESTION	Was remediation complete?
EVIDENCE REQUESTED	Corrective action, recheck evidence, closure owner, reopen criteria
HUMAN ROLE	Closure owner
AGENT/TOOL RECORD	Remediation workflow trace
OUTCOME/CLOSURE	Closure state
PRIVACY TREATMENT	Retention/disclosure rule

Table note: Mixed source-grounded and author-synthesis table. It is a readiness evidence architecture, not a formal audit procedure, legal demand, or certification criterion.

Cross-Links

- Appendix A expands the evidence request list.
- Appendix B turns these categories into a walkthrough template.
- Appendix E provides exception/remediation checklist detail.

Boundary Note

Evidence requests do not prove legal compliance or audit sufficiency. A qualified professional must evaluate sufficiency in context.

8. Lifecycle Walkthrough for AI Agent / MAS Work

Chapter purpose: Provide a reconstruction pattern for one agentic lifecycle run or work unit. **Reader question:** How is one agentic workflow walked from intent to accepted outcome?

Key Claims

- A lifecycle walkthrough should connect intent, authority, planning, execution, handoff, evidence, review, outcome, exception, remediation, and closure.
- The walkthrough is illustrative and readiness-oriented, not a formal audit procedure.
- Walkthroughs should preserve human role and agent role separation.

Source Grounding Note

Control and walkthrough-adjacent language uses AUD-04, AUD-05, and AUD-06. Provenance and trace context uses EVID-01 and EVID-02. GAIC source truth provides MRO lifecycle objects.

Author Synthesis Note

The agentic lifecycle walkthrough is author synthesis in this paper. It translates GAIC MROs and evidence request categories into a reconstruction pattern.

Main Text

A lifecycle walkthrough tests whether a single unit of agentic work can be reconstructed. It begins with a concrete work unit: a task, run, agent workflow, multi-agent handoff, tool action, or remediation event. The goal is not to prove compliance. The goal is to determine whether the evidence chain can explain the work from intent to closure.

The walkthrough should begin with intent and scope. What was the system asked to do? Who initiated the work? What business process, policy, or governance context applied? What was out of scope? The purpose of this step is to avoid reviewing an isolated output without understanding the work it belongs to.

The second step is authority. What authority allowed the agent or workflow to act? Was authority delegated by role, policy, system configuration, human confirmation, or another control? Did that authority include limits, expiry, risk class, confirmation requirements, or escalation paths? If the tool action exceeded scope, was that captured as exception or authority drift?

The third step is agent and tool identification. Which agent role acted? What constraints applied to that role? Which tools were available? Which tools were used? Which external systems, data stores, APIs, or processors were touched? The walkthrough should separate the agent role from the human role and the tool surface.

The fourth step is evidence pointer review. Where are logs, traces, workflow records, approvals, policies, tickets, evidence packs, or hash manifests stored? Are they linked by a stable work unit ID? Are there partition

keys by agent, tool, role, vendor, project, data class, and lifecycle stage? Can an authorized reviewer reconstruct the chain without overexposing sensitive data?

The fifth step is outcome review. What output or action occurred? Was the outcome accepted, rejected, disputed, escalated, remediated, or closed? Who had authority to accept it? What evidence supports that state? An accepted outcome is a governance state, not just a completed workflow event.

The sixth step is exception and remediation review. If something deviated, did the system record the baseline, trigger, owner, corrective action, recheck, closure state, and reopen criteria? A lifecycle walkthrough is incomplete if it stops at detection and never reaches closure.

The seventh step is boundary review. What does the walkthrough support? It may support readiness discussion, internal review, evidence request scoping, or future assurance planning. It does not itself produce an audit opinion, certification, legal conclusion, or regulator-approved result.

Table 5: Human Role / Agent / Tool Responsibility Matrix

Table 5: Human Role / Agent / Tool Responsibility Matrix

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Initiate work

HUMAN ROLE	Request owner
AGENT ROLE	Intake agent
TOOL/SYSTEM	Workflow system
AUTHORITY BASIS	Business request policy
EVIDENCE POINTER	Work unit record
ESCALATION PATH	Reject or request clarification

Approve external action

HUMAN ROLE	Business owner or delegated approver
AGENT ROLE	Execution agent
TOOL/SYSTEM	CRM/API/configuration tool
AUTHORITY BASIS	Delegated authority record
EVIDENCE POINTER	Authority evidence pointer and tool log
ESCALATION PATH	Human confirmation or stop

Transfer task	
HUMAN ROLE	Process owner
AGENT ROLE	Source and target agents
TOOL/SYSTEM	Orchestration runtime
AUTHORITY BASIS	Handoff rule and inherited constraints
EVIDENCE POINTER	Responsibility transfer record
ESCALATION PATH	Return/reject/escalate

Review output	
HUMAN ROLE	Reviewer role
AGENT ROLE	Drafting or analysis agent
TOOL/SYSTEM	Document/review system
AUTHORITY BASIS	Review policy
EVIDENCE POINTER	Output, review note, evidence pointer
ESCALATION PATH	Dispute or remediation

Handle exception	
HUMAN ROLE	Incident/remediation owner
AGENT ROLE	Monitoring or remediation agent
TOOL/SYSTEM	Ticketing/incident system
AUTHORITY BASIS	Escalation rule
EVIDENCE POINTER	Exception record and remediation evidence
ESCALATION PATH	Reopen, escalate, or close

Table note: GAIC-derived and author-synthesis table grounded in internal audit, control, provenance, and observability context. It assigns review responsibility for auditability purposes only and does not assign legal liability.

Walkthrough Output

A completed walkthrough should produce:

- Work unit identity and scope.
- Authority boundary and confirmation history.
- Human role and agent role map.
- Tool-action evidence and side-effect summary.
- Evidence chain with pointers and partition keys.
- Accepted outcome or dispute state.
- Exception/remediation closure state if applicable.
- Privacy treatment and disclosure profile.
- Boundary statement describing what the walkthrough does and does not prove.

Cross-Links

- Appendix B provides the walkthrough template.
- Chapter 9 addresses partitioning across agents, tools, vendors, and projects.
- Chapter 11 addresses exceptions and remediation closure.

Boundary Note

The walkthrough is a reconstruction pattern for readiness and evidence design. It is not a formal audit procedure, assurance engagement step, or legal compliance test.

9. Evidence Partitioning Across Agents, Tools, Roles, Vendors, and Projects

Chapter purpose: Define how evidence is separated and linked across lifecycle boundaries. **Reader question:** How do reviewers avoid flattened responsibility in multi-agent and tool-mediated work?

Key Claims

- Evidence must be partitioned and linkable across agent, tool, role, vendor, project, data class, jurisdiction, and lifecycle state where relevant.
- Partitioning enables reconstruction without blanket disclosure.
- Cross-project reuse, third-party processor chains, and vendor/model/runtime substitution require explicit evidence boundaries.

Source Grounding Note

Provenance and observability/logging context uses EVID-01, EVID-02, and EVID-03. Privacy context uses PRIV-01, PRIV-02, PRIV-03, and AI-09. GAIC source truth provides MRO-08, MRO-09, MRO-10, MRO-14, and MRO-15.

Author Synthesis Note

The evidence partitioning model is author synthesis in this paper grounded in provenance, observability, privacy, and GAIC MROs.

Main Text

Agentic evidence becomes difficult to review when everything is flattened into one trace, one log export, one transcript, or one vendor report. Multi-agent systems distribute work across agents, tools, human roles, vendors, processors, projects, and data boundaries. Auditability requires those boundaries to be visible.

Partitioning is the discipline of separating evidence by review-relevant boundary while preserving the links needed for reconstruction. It is not the same as hiding evidence. It is not the same as retaining everything. It is a way to make evidence reviewable, scoped, and privacy-aware.

The first partition is the work unit. Every evidence item should connect to a lifecycle work unit or reconstruction path. Without a work unit, reviewers may know that events happened but not which lifecycle task they belong to.

The second partition is the agent. Agent roles may differ in authority, capability, tool access, evidence duties, and escalation requirements. If evidence only shows that "the system" acted, responsibility surfaces become blurred. Reviewers should be able to distinguish planning agents, execution agents, review agents, monitoring agents, and remediation agents where those roles exist.

The third partition is the tool. Tool actions often create external consequences. Evidence should separate model output from tool execution, tool execution from downstream system state, and reversible actions from

irreversible or high-impact actions. Tool-specific evidence can include request/response logs, target system records, affected data categories, rollback records, and owner information.

The fourth partition is the human role. A user account is not enough. Auditability requires role responsibility: request owner, delegating role, reviewer, acceptor, exception owner, remediation owner, privacy owner, or internal audit reviewer. These roles should connect to evidence without implying legal liability.

The fifth partition is vendor, processor, and runtime boundary. Agentic systems may call third-party tools, external APIs, model providers, orchestration runtimes, vector stores, workflow systems, or processors/subprocessors. Evidence should identify where responsibility, data processing, and review access shift. This is not vendor ranking. It is boundary mapping.

The sixth partition is project and reuse context. Cross-project reuse can move prompts, agents, workflows, memory, tools, or policies into a new context. Reuse may be efficient, but auditability requires evidence that scope, authority, privacy treatment, and evidence obligations were reset or revalidated for the new context.

The seventh partition is privacy and disclosure class. Some evidence can be disclosed directly. Some should be redacted. Some may require evidence pointers, hashes, summaries, or reviewer-specific access. The partitioning model should make it possible to review what is necessary without exposing more than the review requires.

Table 6: Evidence Partitioning Matrix

Table 6: Evidence Partitioning Matrix

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Work unit

WHY IT MATTERS	Anchors reconstruction
EVIDENCE BOUNDARY	Work unit ID and lifecycle stage
ACCESS/DISCLOSURE RULE	Reviewer access tied to scope
FAILURE IF MISSING	Evidence cannot be connected
RELATED MRO	MRO-08

Agent

WHY IT MATTERS	Separates responsibility surfaces
EVIDENCE BOUNDARY	Agent ID, role, constraints, work segment
ACCESS/DISCLOSURE RULE	Role-based evidence disclosure
FAILURE IF MISSING	Flattened agent responsibility
RELATED MRO	MRO-03, MRO-06, MRO-08

Tool

WHY IT MATTERS	Tracks external consequence
EVIDENCE BOUNDARY	Tool ID, target system, action type, reversibility
ACCESS/DISCLOSURE RULE	Tool-specific disclosure and redaction
FAILURE IF MISSING	Unclear side effects
RELATED MRO	MRO-05, MRO-08

Human role

WHY IT MATTERS	Connects governance ownership
EVIDENCE BOUNDARY	Role map, approval, review, acceptance, closure
ACCESS/DISCLOSURE RULE	Need-to-know reviewer access
FAILURE IF MISSING	User identity mistaken for responsibility
RELATED MRO	MRO-01

Vendor / processor

WHY IT MATTERS	Maps third-party boundary
EVIDENCE BOUNDARY	Vendor/processor role, subprocessor chain, data path
ACCESS/DISCLOSURE RULE	Contract/privacy-aware disclosure
FAILURE IF MISSING	Responsibility chain disappears
RELATED MRO	MRO-14

Project / reuse context

WHY IT MATTERS	Prevents context drift
EVIDENCE BOUNDARY	Source context, target context, reset validation
ACCESS/DISCLOSURE RULE	Reuse-specific review access
FAILURE IF MISSING	Old authority/policy reused incorrectly
RELATED MRO	MRO-09

Runtime/model substitution

WHY IT MATTERS	Preserves evidence continuity after change
EVIDENCE BOUNDARY	Prior/new component, revalidation evidence
ACCESS/DISCLOSURE RULE	Change-scoped disclosure
FAILURE IF MISSING	Evidence chain breaks after substitution
RELATED MRO	MRO-15

Privacy class

WHY IT MATTERS	Limits overexposure
EVIDENCE BOUNDARY	Data category, redaction profile, retention rule
ACCESS/DISCLOSURE RULE	Minimized disclosure
FAILURE IF MISSING	Privacy leakage or evidence hoarding
RELATED MRO	MRO-10, MRO-12, MRO-13

Table note: Mixed source-grounded and author-synthesis table. It does not create universal disclosure or retention rules, vendor rankings, or procurement recommendations.

Cross-Links

- Chapter 10 covers privacy and selective disclosure.
- Appendix A includes partition fields for evidence requests.
- Appendix C maps partitioning to MRO-08, MRO-09, MRO-14, and MRO-15.

Boundary Note

Partitioning is a review architecture. It does not decide legal retention, data subject rights, processor obligations, or vendor liability.

10. Privacy, Selective Disclosure, and Audit Evidence Minimization

Chapter purpose: Balance auditability with privacy, confidentiality, minimization, and retention discipline.
Reader question: How can evidence be reviewable without overexposure?

Key Claims

- Auditability should not become evidence hoarding.
- Privacy-aware auditability uses minimization, redaction, retention discipline, evidence pointers, partitioning, and selective disclosure.
- this paper frames privacy/evidence tension; it does not provide legal advice or GDPR compliance proof.

Source Grounding Note

Privacy and minimization context uses PRIV-01, PRIV-02, PRIV-03, PRIV-04, PRIV-05, and AI-09. GAIC source truth provides MRO-10 through MRO-13 and the Validation Lab selective disclosure boundary.

Author Synthesis Note

The privacy-preserving audit evidence controls are author synthesis in this paper grounded in public privacy guidance and GAIC MROs.

Main Text

Auditability can create a privacy problem if it is misunderstood. The answer to weak evidence is not to collect everything forever. Agentic systems may process personal data, confidential business data, privileged material, regulated data, customer records, employee records, prompts, outputs, tool responses, logs, memory, and third-party processor data. If auditability becomes indiscriminate evidence capture, it may increase exposure and reduce trust.

The better question is: what evidence is necessary for the review objective, and how can it be disclosed safely? This requires minimization, partitioning, retention discipline, access control, redaction, evidence pointers, and reviewer-specific disclosure profiles.

Privacy-aware auditability begins with purpose. Evidence should be collected and retained for defined review purposes: reconstruction, control review, exception handling, third-party validation, remediation, or assurance planning. A generic desire to "keep all logs just in case" is not a mature evidence strategy.

The second principle is minimization. Evidence objects should identify what fields are needed for reconstruction and which fields can be redacted, tokenized, summarized, hashed, or replaced with pointers. A reviewer may need to know that a tool action affected a customer data category without seeing every raw value. Another reviewer may need access under a narrower approved scope. The evidence architecture should support both.

The third principle is selective disclosure. Different reviewers need different levels of evidence. Internal engineering may need logs and traces. Internal audit may need work unit chains and evidence pointers. Privacy teams may need data category and retention records. Third-party reviewers may need redacted evidence packs, ruleset identity, verdict hashes, or replay packages. The system should make disclosure explicit, scoped, and recorded.

The fourth principle is retention discipline. Evidence needs may conflict with storage limitation and deletion expectations. This paper does not decide legal retention periods. It asks that systems record retention rules, expiry, holds, deletion/redaction actions, and rights-retention reconciliation where relevant. A reviewable retention decision is stronger than a silent default.

The fifth principle is boundary clarity. Privacy sources provide context, not legal conclusions. A privacy-preserving evidence architecture does not prove GDPR compliance, UK GDPR compliance, EU AI Act compliance, or any jurisdiction-specific legal compliance. It creates the object layer that legal, privacy, audit, and governance professionals can review.

Table 7: Privacy-Preserving Audit Evidence Controls

Table 7: Privacy-Preserving Audit Evidence Controls

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Evidence minimization

PURPOSE	Avoid unnecessary collection
EVIDENCE ARTIFACT	Minimization rule, field inventory
SELECTIVE DISCLOSURE MECHANISM	Request only review-relevant fields
PRIVACY RISK REDUCED	Overcollection
RESIDUAL REVIEW NEED	Legal/privacy review for purpose and necessity

Redaction profile

PURPOSE	Remove or mask sensitive fields
EVIDENCE ARTIFACT	Redaction rule and redaction log
SELECTIVE DISCLOSURE MECHANISM	Redacted evidence pointer
PRIVACY RISK REDUCED	Overexposure to reviewer
RESIDUAL REVIEW NEED	Sufficiency check after redaction

Evidence pointer

PURPOSE	Reference evidence without copying raw data
EVIDENCE ARTIFACT	Pointer, hash, manifest entry
SELECTIVE DISCLOSURE MECHANISM	Reviewer-scoped access or replay
PRIVACY RISK REDUCED	Data duplication
RESIDUAL REVIEW NEED	Integrity and access review

Retention rule

PURPOSE	Limit storage duration or define hold
EVIDENCE ARTIFACT	Retention schedule, expiry, hold indicator
SELECTIVE DISCLOSURE MECHANISM	Time-bound access
PRIVACY RISK REDUCED	Excessive retention
RESIDUAL REVIEW NEED	Legal/privacy retention analysis

Disclosure profile

PURPOSE	Define who can see what
EVIDENCE ARTIFACT	Disclosure matrix, access log
SELECTIVE DISCLOSURE MECHANISM	Role-based disclosure package
PRIVACY RISK REDUCED	Unscoped sharing
RESIDUAL REVIEW NEED	Reviewer authorization review

Rights-retention reconciliation

PURPOSE	Address tension between rights requests and evidence retention
EVIDENCE ARTIFACT	Request record, action, hold status, rationale
SELECTIVE DISCLOSURE MECHANISM	Redaction/deletion/hold decision path
PRIVACY RISK REDUCED	Silent conflict between evidence and rights
RESIDUAL REVIEW NEED	Legal review

Third-party validation package	
PURPOSE	Support review without raw exposure
EVIDENCE ARTIFACT	Redacted evidence pack, ruleset ID, verdict hash
SELECTIVE DISCLOSURE MECHANISM	Non-raw evidence adjudication
PRIVACY RISK REDUCED	Third-party overexposure
RESIDUAL REVIEW NEED	Boundary and conflict review

Table note: Source-grounded and GAIC-derived table. It is not legal advice, privacy compliance proof, or universal retention guidance.

Cross-Links

- Chapter 9 addresses evidence partitioning.
- Chapter 12 addresses third-party validation boundaries.
- Appendix F provides privacy/legal boundary language.

Boundary Note

This chapter does not provide legal advice, interpret specific legal obligations, or claim that any architecture proves compliance with GDPR, UK GDPR, EU AI Act, or other privacy laws.

11. Exception, Dispute, and Remediation Closure

Chapter purpose: Define closure evidence for failed, disputed, drifted, or remediated agentic work. **Reader question:** How does auditability handle exceptions and closure?

Key Claims

- Exceptions, disputes, authority drift, outcome rejection, and remediation must close as evidence-backed lifecycle states.
- Incident logs alone do not prove closure.
- Closure evidence should identify owner, corrective action, recheck, accepted closure state, and reopen criteria.

Source Grounding Note

Incident/governance context uses EVID-04 and AI-02. Control and AI audit practice context uses AUD-04, AUD-05, and AUD-06. GAIC source truth provides MRO-04, MRO-07, and MRO-16.

Author Synthesis Note

Remediation closure as an audit evidence state is GAIC-derived author synthesis. It is not a legal settlement, regulator closure, or assurance conclusion.

Main Text

Agentic auditability is incomplete if it only reconstructs successful work. Many of the most important review questions arise when an agent acts outside expected authority, a tool action creates unintended consequence, an outcome is disputed, a privacy issue appears, a handoff fails, or remediation is required.

In ordinary operations, exceptions often become tickets. Tickets are useful, but they are not automatically lifecycle evidence. A ticket may describe an issue, assign an owner, record comments, and mark a status. For auditability, the ticket must connect to the agentic work unit, authority boundary, role map, tool action, evidence pointer, accepted outcome state, privacy treatment, corrective action, and closure state.

Closure is the critical concept. An incident can be detected without being remediated. A remediation can be attempted without being rechecked. A recheck can occur without being accepted. An accepted closure can later be reopened. Auditability requires those state transitions to be evidence-backed.

For agentic systems, exceptions should be categorized by lifecycle object. Was the issue an authority exception, where observed behavior exceeded delegated scope? Was it a responsibility exception, where no human owner could be identified? Was it a tool-action exception, where a tool created unexpected side effects? Was it an accepted-outcome dispute, where a result was produced but not accepted? Was it a privacy exception, where evidence collection or disclosure exceeded the intended boundary? Was it a substitution exception, where evidence integrity broke after runtime/model/tool change?

The closure record should answer five questions:

1. What happened and which lifecycle object was affected?
2. Who owned the exception and remediation?
3. What corrective action was taken?
4. What evidence shows recheck or review?
5. Who accepted closure, and what would reopen the issue?

None of these questions determines legal fault. None proves regulatory acceptance. None settles liability. They make the remediation lifecycle reconstructable.

Table 10: Exception / Dispute / Remediation Evidence Checklist

Table 10: Exception / Dispute / Remediation Evidence Checklist

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Authority exception

EVIDENCE REQUIRED	Baseline authority, observed action, trigger, affected work unit
OWNER	Escalation owner
ESCALATION PATH	Stop, downgrade, reauthorize, or human confirmation
REMEDIATION RECORD	Reauthorization or stop-action record
CLOSURE EVIDENCE	Closure acceptance and new boundary
BOUNDARY	Not legal determination

Outcome dispute

EVIDENCE REQUIRED	Output, acceptance criteria, dispute reason, reviewer notes
OWNER	Review owner
ESCALATION PATH	Secondary review or remediation
REMEDIATION RECORD	Correction/review path
CLOSURE EVIDENCE	Accepted, rejected, or remediated state
BOUNDARY	Not settlement proof

Tool-action side effect

EVIDENCE REQUIRED	Tool call, target system, affected data/process, reversibility
OWNER	Tool-action owner
ESCALATION PATH	Rollback, disable, escalate
REMEDIATION RECORD	Rollback or corrective action
CLOSURE EVIDENCE	Recheck and owner acceptance
BOUNDARY	Not liability finding

Privacy event

EVIDENCE REQUIRED	Data category, disclosure path, redaction/retention record
OWNER	Privacy owner
ESCALATION PATH	Privacy/legal review
REMEDIATION RECORD	Redaction, deletion, hold, or access restriction
CLOSURE EVIDENCE	Review note and disclosure update
BOUNDARY	Not legal advice

Handoff failure

EVIDENCE REQUIRED	Source/target agent, transferred scope, retained scope
OWNER	Process owner
ESCALATION PATH	Return, reassign, or escalate
REMEDIATION RECORD	Transfer correction
CLOSURE EVIDENCE	Accepted transfer or closed exception
BOUNDARY	Not legal responsibility transfer

Remediation defect

EVIDENCE REQUIRED	Prior closure evidence, reopened issue, failed recheck
OWNER	Remediation owner
ESCALATION PATH	Reopen and reassign
REMEDIATION RECORD	Updated corrective action
CLOSURE EVIDENCE	Recheck and closure reviewer
BOUNDARY	Not regulator closure

Table note: Mixed source-grounded, GAIC-derived, and author-synthesis checklist. It is not an audit procedure, legal remedy checklist, or certification criterion.

Cross-Links

- Appendix E provides an expanded closure checklist.
- Chapter 13 includes exception traceability and remediation closure as AARM dimensions.
- Chapter 16 connects closure to future Guide 1, Guide 2, and the later insurability white paper.

Boundary Note

Closure in this chapter means evidence-backed governance closure. It does not mean legal closure, settlement, regulator acceptance, assurance opinion, or certification.

12. Third-Party Validation Without Certification Claims

Chapter purpose: Define how evidence review can involve third parties without becoming certification, assurance opinion, conformity assessment, legal proof, or regulator approval. **Reader question:** What can third-party validation support without becoming certification?

Key Claims

- Third-party validation can support evidence review, replay, adjudication, and readiness discussion.
- Third-party validation does not certify compliance or issue assurance opinions unless a separate authorized engagement or conformity-assessment framework exists.
- Validation Lab is only a non-certifying evidence adjudication example.

Source Grounding Note

Assurance and attestation boundaries are grounded in AUD-03 and BOUND-03. Conformity-assessment and validation/verification boundaries are grounded in BOUND-01 and BOUND-02. Validation Lab boundary is GAIC source truth.

Author Synthesis Note

this paper's third-party validation boundary is author synthesis that uses professional boundary sources to avoid overclaim. It does not define a conformity assessment program.

Main Text

Agentic auditability often benefits from independent review. A third party, internal audit team, governance function, or specialized review group may inspect evidence packs, replay work units, test whether evidence objects exist, adjudicate ruleset conformance, or evaluate whether evidence chains are reconstructable. These activities can be valuable.

They are also easy to overstate. Evidence review is not automatically certification. Validation is not automatically an assurance opinion. A verdict hash is not regulator approval. A ruleset check is not legal compliance proof. A third-party review note is not an audit report unless it is performed under an appropriate professional engagement, criteria, scope, independence framework, and reporting structure.

This chapter exists to keep that boundary visible. This paper uses "third-party validation" in a narrow sense: review of evidence against defined rulesets or reconstruction objectives. The review may support readiness, evidence quality, selective disclosure, replay, exception closure, or future assurance planning. It does not itself create certification or assurance.

Validation Lab is a useful boundary example because GAIC already frames it as non-certifying evidence adjudication. In this paper, Validation Lab should only be described as one example of how evidence-based validation might be operationalized. It must not be described as the only validation path, a certification body, a

regulator, an audit firm, an assurance provider, a conformity assessment body, or proof of enterprise readiness.

The distinction matters for enterprises. If an organization receives a third-party evidence review, it should know exactly what was reviewed, against what ruleset, under what scope, with what evidence, and with what limitations. A review may show that evidence exists and can be replayed. It may show that a ruleset passed. It may show that disclosure was minimized. But it does not determine legal compliance or professional audit sufficiency.

The same distinction matters for reviewers. A review team should state its role: internal readiness review, independent evidence review, validation, verification, assurance engagement, attestation engagement, certification, or legal analysis. Each role has different requirements. this paper only defines the auditability object layer that may support those roles.

Table 8: Third-Party Validation Boundary Table

Table 8: Third-Party Validation Boundary Table

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Evidence adjudication

CAN SUPPORT	Ruleset conformance review
CANNOT PROVE	Legal compliance or audit opinion
EVIDENCE ARTIFACT	Ruleset ID, evidence pointer, verdict/hash record
BOUNDARY LANGUAGE	Non-certifying evidence review
SOURCE IDS	GAIC-SOURCE, BOUND-01

Evidence replay

CAN SUPPORT	Reconstructability check
CANNOT PROVE	Operating effectiveness across all periods
EVIDENCE ARTIFACT	Replay package, work unit chain, integrity record
BOUNDARY LANGUAGE	Scope-limited reconstruction
SOURCE IDS	EVID-01, AUD-01

Selective disclosure review

CAN SUPPORT	Privacy-aware review process
CANNOT PROVE	Privacy law compliance
EVIDENCE ARTIFACT	Redaction profile, disclosure log, reviewer access scope
BOUNDARY LANGUAGE	Privacy-aware evidence review, not legal advice
SOURCE IDS	PRIV-01 to PRIV-05

Third-party readiness review

CAN SUPPORT	Independent input to readiness discussion
CANNOT PROVE	Assurance conclusion
EVIDENCE ARTIFACT	Scope statement, findings, limitations
BOUNDARY LANGUAGE	Readiness input only
SOURCE IDS	AUD-03, BOUND-03

Validation Lab example

CAN SUPPORT	Non-certifying evidence adjudication example
CANNOT PROVE	Certification, regulator approval, conformity assessment, legal proof
EVIDENCE ARTIFACT	Verdict hash, ruleset identity, evidence pointer
BOUNDARY LANGUAGE	One optional example, not exclusive path
SOURCE IDS	GAIC-SOURCE, BOUND-01, BOUND-02

Assurance engagement

CAN SUPPORT	Professional assurance under separate criteria and engagement terms
CANNOT PROVE	Anything outside engagement scope
EVIDENCE ARTIFACT	Practitioner report, criteria, scope, evidence basis
BOUNDARY LANGUAGE	Outside this paper unless separately authorized
SOURCE IDS	AUD-03, BOUND-03

Table note: Source-grounded boundary table with GAIC-derived example. It is not a certification model, audit procedure, or assurance engagement template.

Cross-Links

- Appendix F provides reusable boundary language.
- Appendix D uses third-party review boundary as an AARM dimension.
- Chapter 15 explains professional use by audit and assurance firms.

Boundary Note

This chapter must not be read as claiming that any third party, including Validation Lab, certifies agentic systems, issues legal opinions, provides assurance opinions, or has regulator approval.

13. Agentic Auditability Readiness Model

Chapter purpose: Present AARM dimensions and readiness levels as a proposed readiness model. **Reader question:** How can auditability readiness be described without claiming assurance?

Key Claims

- AARM describes readiness for reconstructability, testability, evidence structure, privacy-aware review, and third-party review boundaries.
- AARM is proposed author synthesis in this paper.
- AARM is not a score, audit standard, certification, assurance opinion, legal compliance proof, regulator approval, procurement tool, or vendor ranking.

Source Grounding Note

Audit evidence language uses AUD-01. AI audit practice context uses AUD-05 and AUD-06. AI governance context uses AI-01. Attestation boundary context uses BOUND-03. R4B/R1 AARM baseline and GAIC source truth provide the model structure.

Author Synthesis Note

AARM is author synthesis in this paper. It is a readiness vocabulary, not an external maturity model or professional assurance framework.

Main Text

Once evidence objects are defined, organizations need a way to discuss readiness. AARM provides that vocabulary. It asks whether agentic lifecycle work can be reconstructed, tested, evidenced, and reviewed across the object areas this paper has defined.

AARM deliberately avoids scoring. Scores invite overprecision and can be mistaken for certification, compliance grades, vendor rankings, or procurement guidance. AARM instead uses dimensions and levels. The dimensions ask what must be reviewable. The levels describe the degree to which the lifecycle work is observable, trace-linked, evidence-structured, auditability-ready, or assurance-ready for planning.

The word "Assurance-Ready" is sensitive. In AARM, it does not mean assurance has been issued. It means the evidence architecture may be mature enough to support qualified assurance planning or review within a defined scope. It still depends on professional scope, criteria, independence, engagement acceptance, and judgment, and it does not guarantee a positive result.

The ten AARM dimensions are:

1. Audit object clarity.
2. Authority traceability.
3. Human / agent responsibility mapping.
4. Evidence sufficiency.
5. Agent/tool partitioning.

6. Exception traceability.
7. Accepted outcome evidence.
8. Remediation closure.
9. Privacy and selective disclosure.
10. Third-party review boundary.

The levels describe readiness states. L0 Unobservable means lifecycle work cannot be reconstructed beyond informal narrative or final output. L1 Log-Visible means logs exist but responsibility linkage is weak. L2 Trace-Linked means traces connect some actions to agents, tools, or workflows. L3 Evidence-Structured means lifecycle evidence objects exist. L4 Auditability-Ready means defined-scope evidence chains can reconstruct authority, responsibility, tools, outcomes, exceptions, privacy, and closure. L5 Assurance-Ready means evidence architecture may support professional assurance planning or review within a defined scope.

The main use of AARM is conversation discipline. It helps technology leaders avoid saying "we have logs, so we are audit-ready." It helps governance teams ask whether accepted outcomes and remediation closure are evidence states. It helps audit teams separate raw observability from reconstructable evidence. It helps privacy teams see whether selective disclosure is built in. It helps third-party reviewers clarify what their review does and does not mean.

Table 9: Auditability Readiness Levels

Table 9: Auditability Readiness Levels

Readiness-level rubric rendered as split matrices in PDF mode; the level key is repeated to preserve level-to-level comparison.

Panel 1 of 2

Level	Definition	Observable Traits	Minimum Evidence
L0 Unobservable	Lifecycle work cannot be reconstructed beyond informal narrative or isolated outputs	Missing lifecycle IDs, no authority linkage, no role map, no accepted outcome state	None beyond ad hoc notes or final output
L1 Log-Visible	Activity logs exist without responsibility-linked evidence chains	Logs, traces, timestamps, tool events	Raw logs or trace exports
L2 Trace-Linked	Logs/traces link some actions to agents, tools, users, or workflows	Workflow IDs, agent/tool traces, partial actor linkage	Trace-linked run records and partial evidence map
L3 Evidence-Structured	Lifecycle evidence is organized into objects that support reconstruction	Role maps, authority objects, accepted outcome records, partitioning, exception records	MRO-aligned evidence objects and sampled walkthroughs
L4 Auditability-Ready	Evidence chains can reconstruct defined-scope agentic lifecycle work	Walkthroughs, evidence request response, privacy controls, closure evidence	Complete evidence chain for defined scope plus boundary statement
L5 Assurance-Ready	Evidence architecture may support professional assurance planning or review within a defined scope	Repeatable evidence packs, ruleset identity, reviewer scope, remediation/recheck procedure	Reviewable evidence packs, integrity records, boundary statement

Panel 2 of 2

Level	What It Does Not Prove	Source / Synthesis Note
L0 Unobservable	Absence of risk, absence of non-compliance, or system safety	Author synthesis grounded in AUD-01/AUD-02 expectations
L1 Log-Visible	Authority, responsibility, acceptance, exception closure, audit readiness	EVID-02/EVID-03 + author synthesis in this paper
L2 Trace-Linked	Complete lifecycle responsibility or sufficient evidence	EVID-01/EVID-02 + author synthesis in this paper
L3 Evidence-Structured	Assurance readiness, legal compliance, operating effectiveness	GAIC-derived + AUD-01/AUD-04/AUD-05
L4 Auditability-Ready	Audit opinion, certification, regulator approval, legal compliance	author synthesis in this paper grounded in audit/provenance/privacy sources
L5 Assurance-Ready	Actual assurance opinion, certification, guaranteed compliance, regulator acceptance	author synthesis in this paper bounded by AUD-03/BOUND-03

Table note: AARM is author synthesis grounded in source language. It is not a score, standard, certification, assurance conclusion, procurement benchmark, or vendor ranking.

AARM Dimension Summary

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Audit object clarity

AUDIT QUESTION	Can reviewers name the lifecycle unit under review?
MINIMUM EVIDENCE FAMILY	Work unit ID, scope, object inventory
PRIMARY MRO LINKS	MRO-01, MRO-03, MRO-04
BOUNDARY	Not legal category

Authority traceability

AUDIT QUESTION	Was the action authorized under scope?
MINIMUM EVIDENCE FAMILY	Authority grant, expiry, escalation
PRIMARY MRO LINKS	MRO-02, MRO-07
BOUNDARY	Not legal delegation proof

Human / agent responsibility mapping

AUDIT QUESTION	Who owned intent, execution, review, acceptance, remediation?
MINIMUM EVIDENCE FAMILY	Role map and transfer record
PRIMARY MRO LINKS	MRO-01, MRO-03, MRO-06
BOUNDARY	No anthropomorphism

Evidence sufficiency

AUDIT QUESTION	Can reviewers reconstruct work without narrative memory?
MINIMUM EVIDENCE FAMILY	Evidence chain and pointers
PRIMARY MRO LINKS	MRO-04, MRO-08, MRO-16
BOUNDARY	Not universal sufficiency

Agent/tool partitioning

AUDIT QUESTION	Can actions be separated across agents/tools/vendors/projects?
MINIMUM EVIDENCE FAMILY	Partition keys, tool evidence, substitution record
PRIMARY MRO LINKS	MRO-05, MRO-08, MRO-14, MRO-15
BOUNDARY	No vendor ranking

Exception traceability

AUDIT QUESTION	Can deviation be linked to owner and next action?
MINIMUM EVIDENCE FAMILY	Exception record and impact scope
PRIMARY MRO LINKS	MRO-07, MRO-16
BOUNDARY	Not legal violation

Accepted outcome evidence

AUDIT QUESTION	Who accepted or disputed the result?
MINIMUM EVIDENCE FAMILY	Outcome state, reviewer role, evidence link
PRIMARY MRO LINKS	MRO-04
BOUNDARY	Not compliance proof

Remediation closure

AUDIT QUESTION	What was remediated and who accepted closure?
MINIMUM EVIDENCE FAMILY	Corrective action, recheck, closure state
PRIMARY MRO LINKS	MRO-16
BOUNDARY	Not settlement

Privacy and selective disclosure

AUDIT QUESTION	Can review occur without unnecessary exposure?
MINIMUM EVIDENCE FAMILY	Redaction, retention, disclosure profile
PRIMARY MRO LINKS	MRO-10 to MRO-13
BOUNDARY	Not legal advice

Third-party review boundary

AUDIT QUESTION	What can a reviewer inspect and claim?
MINIMUM EVIDENCE FAMILY	Scope, ruleset, verdict, boundary statement
PRIMARY MRO LINKS	MRO-11, MRO-12, MRO-16
BOUNDARY	No certification/assurance opinion

Cross-Links

- Appendix D provides the expanded readiness matrix.
- Chapter 14 connects AARM to enterprise readiness.
- Chapter 15 connects AARM to professional use boundaries.

Boundary Note

AARM is a readiness model only. It does not prove legal compliance, issue assurance, certify systems, rank vendors, bind regulators, or require MPLP.

14. Enterprise Implementation: CIO / CTO / CCO Readiness

Chapter purpose: Connect this paper to enterprise readiness without replacing Guide 1 or Guide 2. **Reader question:** What must leaders prepare before auditability or assurance conversations become credible?

Key Claims

- Audit-ready agentic systems require architecture, governance, evidence ownership, privacy controls, exception closure, and role clarity.
- CIO/CTO readiness focuses on technical evidence architecture.
- CCO readiness focuses on policy-to-evidence governance.
- This paper bridges to Guide 1 and Guide 2 but does not replace them.

Source Grounding Note

Enterprise control context uses AUD-04 and AUD-07. AI governance context uses AI-01 and PRIV-03. Big Four sources BF-02 and BF-04 provide market context only. GAIC source truth provides Guide 1/Guide 2 boundaries.

Author Synthesis Note

The CIO/CTO/CCO readiness split is author synthesis in this paper derived from R4C Guide planning and audit evidence architecture in this paper.

Main Text

Enterprises do not become audit-ready by exporting logs at the end of an agentic workflow. Auditability must be designed into architecture, governance, review routines, and evidence ownership. This paper is not an implementation guide, but it defines the readiness questions leaders should ask before Guide 1 and Guide 2 translate the framework into practice.

For CIOs and CTOs, the first question is whether the system can produce evidence by design. Agent runtimes, orchestration layers, tool integrations, workflow engines, ticketing systems, approval systems, data stores, and logging platforms must be able to preserve the relationships that matter: work unit, authority, role, tool action, evidence pointer, outcome, exception, privacy treatment, and closure. If these relationships are not captured at runtime, auditability becomes after-the-fact reconstruction.

The second technology question is partitioning. Evidence should be separable by agent, tool, role, vendor, project, data class, lifecycle state, and review purpose. This requires stable identifiers, metadata, access controls, retention logic, evidence export, and selective disclosure. It may also require integration with existing observability, security, data governance, and workflow systems.

The third technology question is change. Agentic systems change as models, prompts, tools, vendors, runtimes, policies, workflows, and memory stores change. Auditability requires substitution records, evidence

continuity checks, regression evidence, reauthorization triggers, and revalidation boundaries. A system that loses evidence continuity after a runtime change is not auditability-ready for that scope.

For CCOs and governance leaders, the first question is policy-to-evidence translation. Policies should not only describe desired behavior; they should identify the evidence object that proves the policy was operationalized. A delegated authority policy should map to authority records. A human oversight policy should map to role and confirmation records. A remediation policy should map to closure evidence.

The second governance question is role operating model. The IIA Three Lines context is useful here as governance vocabulary, not as a this paper mandate. Management, risk/compliance, internal audit, privacy, security, legal, and technology teams need different evidence views and responsibilities. this paper helps name the evidence layer they must share.

The third governance question is review cadence. Agentic auditability is not a one-time readiness exercise. Evidence objects should be sampled, walked through, rechecked after change, reviewed after exceptions, and re-evaluated when authority, vendor, model, runtime, or privacy conditions change.

Guide 1 will later turn these questions into technical architecture planning. Guide 2 will later turn them into compliance operating-model routines. This paper should not absorb those details. It should define the evidence object layer that both guides will implement.

Enterprise Readiness Crosswalk

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Work unit identity

CIO / CTO QUESTION	Can systems assign stable lifecycle IDs?
CCO / GOVERNANCE QUESTION	Which work types require evidence?
EVIDENCE OBJECT	Work unit record
GUIDE RELATION	Guide 1 architecture; Guide 2 policy scope
BOUNDARY	Not audit procedure

Authority

CIO / CTO QUESTION	Can delegated authority be enforced and logged as business scope?
CCO / GOVERNANCE QUESTION	Which policies define authority, expiry, confirmation, and escalation?
EVIDENCE OBJECT	Delegated authority record
GUIDE RELATION	Guide 1 controls; Guide 2 governance
BOUNDARY	Not legal delegation proof

Role mapping

CIO / CTO QUESTION	Can human/agent roles be captured separately?
CCO / GOVERNANCE QUESTION	Who owns intent, review, acceptance, and closure?
EVIDENCE OBJECT	Responsibility map
GUIDE RELATION	Both guides
BOUNDARY	Not legal liability assignment

Tool actions

CIO / CTO QUESTION	Can consequential tool actions generate evidence?
CCO / GOVERNANCE QUESTION	Which tool actions require review or confirmation?
EVIDENCE OBJECT	Tool-action evidence
GUIDE RELATION	Guide 1 high relevance
BOUNDARY	No liability conclusion

Privacy

CIO / CTO QUESTION	Can evidence be minimized, redacted, retained, and selectively disclosed?
CCO / GOVERNANCE QUESTION	Which privacy rules and approvals govern review?
EVIDENCE OBJECT	Disclosure and retention profile
GUIDE RELATION	Both guides
BOUNDARY	No legal advice

Exceptions

CIO / CTO QUESTION	Can exceptions link to lifecycle objects and closure?
CCO / GOVERNANCE QUESTION	Who governs dispute/remediation workflow?
EVIDENCE OBJECT	Exception/remediation record
GUIDE RELATION	Both guides
BOUNDARY	No legal remedy claim

Third-party review	
CIO / CTO QUESTION	Can evidence be exported/replayed safely?
CCO / GOVERNANCE QUESTION	What can third parties claim or not claim?
EVIDENCE OBJECT	Validation scope and boundary statement
GUIDE RELATION	Guide 1 export; Guide 2 review boundary
BOUNDARY	No certification

Table note: Author-synthesis crosswalk grounded in controls, governance, privacy, and GAIC Guide planning. It is not procurement guidance or readiness guarantee.

Cross-Links

- Guide 1 will own technical implementation details.
- Guide 2 will own compliance operating-model details.
- Appendix D can be used as a readiness conversation input.

Boundary Note

This chapter does not claim that following this paper makes an enterprise compliant, audit-ready in a professional engagement, certified, regulator-approved, or procurement-preferred.

15. How Audit and Assurance Firms Can Use This Framework

Chapter purpose: Define cautious professional use without claiming endorsement, replacement methodology, assurance opinion, or standard-setting status. **Reader question:** How might audit and assurance firms use this paper responsibly?

Key Claims

- Audit and assurance firms may use this paper as a conceptual object/evidence request and readiness framing tool.
- This framework can support discussion, scoping, evidence request design, walkthrough planning, and methodology exploration.
- This paper does not replace professional methodology, engagement criteria, independence, judgment, standards, or reporting.

Source Grounding Note

Big Four context uses BF-02 and BF-03 as market context only. Assurance boundary uses AUD-03 and BOUND-03. AI audit practice context uses AUD-05 and AUD-06.

Author Synthesis Note

The professional-use framing is author synthesis in this paper. It is intentionally cautious and non-endorsement based.

Main Text

This framework is relevant to audit and assurance firms because agentic AI changes the evidence question. Firms already discuss AI assurance, trusted AI, audit transformation, technology-enabled assurance, governance, controls, and responsible AI. Agentic systems add another layer: lifecycle work that must be reconstructed across authority, responsibility, agents, tools, outcomes, exceptions, privacy, and remediation.

The safest professional use of this paper is as a framing tool. It can help teams ask what the review object is. It can help distinguish logs from audit evidence chains. It can help convert governance concerns into evidence requests. It can help structure walkthrough conversations. It can help identify whether an enterprise has evidence gaps before any formal assurance claim is discussed.

This paper may also help methodology teams think about agentic evidence architecture. A method team could use the Agentic Audit Object as a conceptual lens: What lifecycle work unit is under review? What authority record exists? What human role accepted the outcome? What tool-action evidence exists? What exception state or remediation closure exists? What privacy treatment governs disclosure?

Internal audit teams may use this paper differently. They may use it to scope readiness reviews, design evidence request lists, test walkthrough templates, ask whether human oversight is evidence-backed, or

evaluate whether remediation closure is reconstructable. Again, this use is readiness-oriented. It does not transform this paper into an internal audit standard.

Assurance teams must be especially careful. AARM's "Assurance-Ready" level does not mean assurance has been issued or is guaranteed. It means that an evidence architecture may be mature enough to support qualified assurance planning or review, subject to the firm's professional standards, criteria, engagement scope, independence, risk assessment, evidence evaluation, and reporting requirements.

This paper should never be represented as a Big Four methodology. It should never be described as endorsed by Deloitte, PwC, EY, KPMG, ISACA, IIA, IAASB, PCAOB, AICPA, ISO, NIST, OECD, EU institutions, Singapore agencies, ICO, or any source organization. Sources provide context and terminology. They do not approve this framework.

Professional Use Matrix

Professional Use Matrix

Use case	This framework can support	this paper cannot replace	Useful objects	Boundary
Readiness conversation	Shared vocabulary for auditability gaps	Professional judgment or engagement acceptance	AARM, evidence chain, object model	No assurance conclusion
Evidence request design	Categories for authority, roles, tools, outcomes, exceptions, privacy, closure	Audit procedures or legal demands	Appendix A request list	Not formal procedure
Walkthrough scoping	Lifecycle reconstruction pattern	Sampling methodology or engagement workplan	Appendix B walkthrough	Illustrative only
Internal audit review	Auditability gap identification	Internal audit standards or final reports	MRO mapping, closure checklist	No standard-setting
Methodology exploration	Object model for agentic work	Firm methodology or regulatory requirements	Agentic Audit Object, Audit Evidence Chain	No Big Four endorsement
Assurance planning	Evidence architecture discussion	Assurance criteria, independence, evidence evaluation, opinion	AARM L4/L5 boundary	No assurance opinion

Table note: Professional-use matrix is author synthesis grounded in assurance and AI audit practice sources. It is not a firm methodology, engagement template, or endorsement claim.

Cross-Links

- Chapter 1 defines scope and boundary.
- Chapter 12 defines third-party validation boundaries.
- Appendix F provides safe language.

Boundary Note

This chapter does not claim that audit or assurance firms need this paper, endorse this paper, have adopted this paper, or should replace their methods with this paper.

16. Conclusion

Chapter purpose: Close this paper's argument and bridge to Guide 1, Guide 2, and the later insurability white paper without publication or insurability overclaim. **Reader question:** What must exist before agentic assurance and insurability conversations can mature?

Key Claims

- Auditability begins when agentic lifecycle work becomes reconstructable.
- Logs record activity; evidence chains support responsibility review.
- This paper prepares the path for Guide 1, Guide 2, and the later insurability white paper.

Source Grounding Note

Audit evidence and assurance boundary language uses AUD-01 and AUD-03. AI governance context uses AI-01. GAIC source truth supports the Global AI Compliance White Paper 2026 / this paper / later insurability white paper object-chain relationship and companion-paper boundary.

Author Synthesis Note

The Compliance Object → Audit Evidence Object → Insurable Risk Object chain is GAIC/author synthesis in this paper. The later insurability white paper source research is deferred.

Main Text

Agentic AI changes the auditability problem because agentic systems do not merely produce outputs. They perform lifecycle work. They act through roles, tools, delegated authority, handoffs, privacy boundaries, exceptions, and remediation workflows. If that work cannot be reconstructed, the system may be impressive, observable, and even well-governed in parts, but it is not audit-ready in the sense this paper defines.

The core thesis therefore remains simple: AI agent auditability cannot be built on logs alone. Logs record activity; evidence chains support responsibility review. The difference is the lifecycle responsibility layer. A log may show that an event occurred. An Audit Evidence Chain shows how that event relates to authority, human role, agent role, tool action, evidence pointer, accepted outcome, exception, privacy treatment, and closure.

This paper's contribution is the Audit Evidence Object layer. The Global AI Compliance White Paper 2026 defined the compliance object layer through Agentic Lifecycle Governance and MROs. This paper translates that layer into evidence requests, walkthroughs, partitioning rules, privacy-aware disclosure, validation boundaries, and readiness levels. It defines Agentic AI Auditability as the ability to reconstruct, test, and evidence agentic lifecycle work across authority, responsibility, tools, outcomes, exceptions, and remediation.

The paper also sets limits. It does not issue assurance. It does not certify systems. It does not prove legal compliance. It does not bind regulators. It does not replace professional audit or assurance methodology. It does not rank vendors. It does not require MPLP. It does not make Validation Lab a certification authority.

These limits are not disclaimers at the edge of the work; they are part of the architecture. Auditability is strongest when the evidence object and the claim boundary are both explicit.

For enterprises, the next step is design. CIOs and CTOs need evidence architecture that captures lifecycle responsibility by design, not after-the-fact logs by accident. CCOs and governance leaders need policy-to-evidence operating models that turn authority, oversight, privacy, exception, and remediation policies into reviewable objects. Internal audit and risk teams need walkthrough structures that test reconstructability before assurance claims appear.

For audit and assurance firms, the next step is method dialogue. This framework can support object-model discussions, readiness scoping, evidence request framing, and lifecycle walkthrough design. It cannot replace professional standards, engagement criteria, independence, evidence evaluation, or reporting.

For the series, this paper completes the second layer of the 3+3 front half. The Global AI Compliance White Paper 2026 established the compliance object layer. This paper establishes the audit evidence object layer. Guide 1 will translate auditability into technical architecture. Guide 2 will translate auditability into compliance operating model. The later insurability white paper will later ask what evidence and governance structures are needed before insurable risk can be evaluated.

The object chain is:

Series layer	Object layer	Question	Boundary
Global AI Compliance White Paper 2026	Compliance Object	What lifecycle responsibility objects does agentic governance require?	Not legal compliance proof
Agentic AI Auditability & Assurance White Paper 2026	Audit Evidence Object	Can agentic lifecycle work be reconstructed through responsibility-linked evidence chains?	Not audit standard, certification, or assurance opinion
Later insurability white paper	Insurable Risk Object	What risk evidence may later support insurability analysis?	Deferred; no insurance guarantee
Guide 1	Technical implementation path	How can systems capture audit-ready evidence by design?	Future guide, not drafted here
Guide 2	Compliance operating model path	How can governance translate policy into evidence?	Future guide, not drafted here

Table note: GAIC/author-synthesis table. It preserves companion-paper sequencing and does not claim conclusions for the later insurability white paper.

Auditability begins when agentic lifecycle work becomes reconstructable. Logs record activity; evidence chains support responsibility review. This paper prepares the path for Guide 1, Guide 2, and the later insurability white paper by defining the evidence object layer that must exist before agentic assurance or insurability conversations can be serious.

Cross-Links

- Appendix F preserves conclusion-safe boundary language.
- Guide 1 and Guide 2 remain future practitioner guides.
- The later insurability white paper remains a future white paper with separate source research.

Boundary Note

This chapter does not claim this paper is published, final, sealed, regulator-approved, adopted by firms, or sufficient for assurance, certification, legal compliance, procurement, or insurance.

Appendix A — Agentic Audit Evidence Request List

Purpose: Provide a structured evidence request catalog for agentic systems. **Intended reader:** Audit/assurance teams, internal audit, AI governance, CIO/CTO/CCO teams, technology risk teams, and evidence architecture owners.

Source Grounding

Audit evidence and control vocabulary: AUD-01, AUD-02, AUD-04, AUD-05, AUD-06. GAIC dependencies: MRO-01 through MRO-16, Evidence-Based Validation Pattern, Validation Lab boundary.

Boundary

This appendix is an evidence request architecture. It is not a formal audit procedure, legal demand, certification criterion, procurement checklist, assurance engagement plan, or proof of legal compliance.

Evidence Request Catalog

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

ER-01

LIFECYCLE STAGE	Work initiation
EVIDENCE REQUESTED	Work unit ID, task intent, request source, scope, excluded scope
RELATED MRO	MRO-08
OWNER	Request owner
PRIVACY TREATMENT	Record only necessary purpose/context fields
BOUNDARY NOTE	Does not prove business appropriateness

ER-02

LIFECYCLE STAGE	Role mapping
EVIDENCE REQUESTED	Human role for intent, authorization, review, acceptance, exception, remediation
RELATED MRO	MRO-01
OWNER	Governance/control owner
PRIVACY TREATMENT	Role data minimized to function and authority
BOUNDARY NOTE	Does not assign legal liability

ER-03

LIFECYCLE STAGE	Delegated authority
EVIDENCE REQUESTED	Authority grant, conditions, expiry, risk class, confirmation and escalation rules
RELATED MRO	MRO-02
OWNER	Delegating role
PRIVACY TREATMENT	Restrict authority evidence to review scope
BOUNDARY NOTE	Not legal delegation proof

ER-04

LIFECYCLE STAGE	Agent role
EVIDENCE REQUESTED	Agent role, capability boundary, constraints, evidence duties, escalation route
RELATED MRO	MRO-03
OWNER	Platform owner
PRIVACY TREATMENT	Avoid storing unnecessary prompt/user data
BOUNDARY NOTE	Agent is not a human role

ER-05

LIFECYCLE STAGE	Accepted outcome
EVIDENCE REQUESTED	Acceptance criteria, reviewer role, accepted/rejected/disputed/remediated state
RELATED MRO	MRO-04
OWNER	Review owner
PRIVACY TREATMENT	Redact sensitive output fields where possible
BOUNDARY NOTE	Acceptance is not compliance proof

ER-06

LIFECYCLE STAGE	Tool action
EVIDENCE REQUESTED	Tool ID, action type, target system, affected data/process, reversibility, rollback path
RELATED MRO	MRO-05
OWNER	Tool-action owner
PRIVACY TREATMENT	Classify affected data and redact as needed
BOUNDARY NOTE	Not legal liability finding

ER-07

LIFECYCLE STAGE	Responsibility transfer
EVIDENCE REQUESTED	Source agent, target agent, transferred scope, retained scope, inherited constraints
RELATED MRO	MRO-06
OWNER	Process owner
PRIVACY TREATMENT	Partition by work unit and agent role
BOUNDARY NOTE	Not legal responsibility transfer

ER-08

LIFECYCLE STAGE	Authority drift
EVIDENCE REQUESTED	Baseline authority, observed action, drift trigger, escalation, stop/downgrade/reauthorization
RELATED MRO	MRO-07
OWNER	Escalation owner
PRIVACY TREATMENT	Limit exposure to event-relevant fields
BOUNDARY NOTE	Not every drift is legal breach

ER-09

LIFECYCLE STAGE	Evidence partitioning
EVIDENCE REQUESTED	Partition keys, links, evidence pointers, integrity hashes, access/disclosure rules
RELATED MRO	MRO-08
OWNER	Evidence owner
PRIVACY TREATMENT	Use scoped access and redaction profiles
BOUNDARY NOTE	Not blanket retention

ER-10

LIFECYCLE STAGE	Cross-project reuse
EVIDENCE REQUESTED	Source context, target context, reset validation, privacy review, authorization refresh
RELATED MRO	MRO-09
OWNER	Reuse owner
PRIVACY TREATMENT	Revalidate data categories and retention
BOUNDARY NOTE	Not legal reuse clearance

ER-11

LIFECYCLE STAGE	Privacy lifecycle
EVIDENCE REQUESTED	Data flow, processing purpose, retention rule, rights workflow, minimization treatment
RELATED MRO	MRO-10
OWNER	Privacy owner
PRIVACY TREATMENT	Apply privacy-by-design review
BOUNDARY NOTE	Not legal advice

ER-12

LIFECYCLE STAGE	Third-party validation
EVIDENCE REQUESTED	Ruleset identity, evidence pointer, disclosure/redaction profile, verdict hash, reviewer boundary
RELATED MRO	MRO-11
OWNER	Validation/review owner
PRIVACY TREATMENT	Use selective disclosure package
BOUNDARY NOTE	Not certification

ER-13

LIFECYCLE STAGE	Evidence minimization
EVIDENCE REQUESTED	Minimization rule, disclosure justification, redaction profile, access log, expiry
RELATED MRO	MRO-12
OWNER	Evidence/privacy owner
PRIVACY TREATMENT	Minimize and document residual exposure
BOUNDARY NOTE	Not privacy compliance proof

ER-14

LIFECYCLE STAGE	Rights vs retention
EVIDENCE REQUESTED	Request type, retention basis, hold status, deletion/redaction action, review owner
RELATED MRO	MRO-13
OWNER	Privacy/legal owner
PRIVACY TREATMENT	Requires jurisdiction-specific review
BOUNDARY NOTE	Not legal interpretation

ER-15

LIFECYCLE STAGE	Processor chain
EVIDENCE REQUESTED	Processor role, subprocessor chain, data processing pointer, responsibility owner
RELATED MRO	MRO-14
OWNER	Vendor/privacy owner
PRIVACY TREATMENT	Restrict processor evidence by need-to-know
BOUNDARY NOTE	Not vendor ranking

ER-16

LIFECYCLE STAGE	Substitution conformance
EVIDENCE REQUESTED	Prior/new model, vendor, runtime, tool, evidence-integrity check, reauthorization
RELATED MRO	MRO-15
OWNER	Change owner
PRIVACY TREATMENT	Preserve evidence continuity without over-retention
BOUNDARY NOTE	Not conformance certification

ER-17

LIFECYCLE STAGE	Incident/remediation closure
EVIDENCE REQUESTED	Event, affected outcome, owner, corrective action, recheck, closure acceptance, reopen criteria
RELATED MRO	MRO-16
OWNER	Remediation owner
PRIVACY TREATMENT	Redact sensitive incident details as needed
BOUNDARY NOTE	Not legal closure

Guide Relationship

Guide 1 should translate this appendix into system output requirements, evidence export design, runtime metadata, and workflow architecture. Guide 2 should translate it into policy-to-evidence governance, audit preparation, review cadence, and closure routines.

Placement Recommendation

Retain a condensed version in the main artifact. A fuller operational checklist can be deferred to Guide 1 and Guide 2.

Appendix B — Agentic System Walkthrough Template

Purpose: Provide a repeatable walkthrough structure for a single agentic lifecycle run. **Intended reader:** Internal audit, technology risk, AI governance committees, architecture review teams, platform owners, and compliance engineering teams.

Source Grounding

Control and audit practice context: AUD-04, AUD-05, AUD-06. Provenance and observability context: EVID-01, EVID-02. GAIC dependencies: MRO-01 through MRO-08 and MRO-16.

Boundary

This template is illustrative. It is not a formal audit procedure, assurance engagement program, legal compliance test, certification path, or regulator-approved walkthrough.

Walkthrough Fields

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Work unit ID

DESCRIPTION	Stable identifier for the lifecycle work under review
EVIDENCE POINTER	Workflow/run/task record
RELATED MRO	MRO-08
BOUNDARY NOTE	Not mandatory schema

Intent and scope

DESCRIPTION	Purpose, initiating request, in-scope and out-of-scope actions
EVIDENCE POINTER	Request record, policy reference
RELATED MRO	MRO-01, MRO-08
BOUNDARY NOTE	Does not prove appropriateness

Initiator role

DESCRIPTION	Human or system role that initiated the work
EVIDENCE POINTER	Role map
RELATED MRO	MRO-01
BOUNDARY NOTE	Not legal liability assignment

Delegated authority

DESCRIPTION	Authority grant, conditions, expiry, confirmation, escalation
EVIDENCE POINTER	Authority record
RELATED MRO	MRO-02
BOUNDARY NOTE	Not legal delegation proof

Agent role

DESCRIPTION	Agent identity, role boundary, constraints, capability limits
EVIDENCE POINTER	Agent registry/config record
RELATED MRO	MRO-03
BOUNDARY NOTE	Agent role is not human role

Tool/system used

DESCRIPTION	Tool, external system, API, processor, runtime
EVIDENCE POINTER	Tool-action evidence
RELATED MRO	MRO-05, MRO-14
BOUNDARY NOTE	Not vendor ranking

Evidence pointers

DESCRIPTION	Logs, traces, approvals, tickets, records, hash manifests
EVIDENCE POINTER	Evidence chain references
RELATED MRO	MRO-08, MRO-12
BOUNDARY NOTE	Not blanket disclosure

Human review

DESCRIPTION	Reviewer role, review criteria, review result
EVIDENCE POINTER	Review record
RELATED MRO	MRO-01, MRO-04
BOUNDARY NOTE	Not assurance opinion

Accepted outcome

DESCRIPTION	Produced/accepted/rejected/disputed/remediated state
EVIDENCE POINTER	Outcome record
RELATED MRO	MRO-04
BOUNDARY NOTE	Not compliance proof

Exception state

DESCRIPTION	Drift, dispute, failure, privacy event, tool side effect
EVIDENCE POINTER	Exception record
RELATED MRO	MRO-07, MRO-16
BOUNDARY NOTE	Not legal breach finding

Remediation closure

DESCRIPTION	Corrective action, recheck, closure acceptance, reopen criteria
EVIDENCE POINTER	Remediation record
RELATED MRO	MRO-16
BOUNDARY NOTE	Not legal settlement

Privacy treatment

DESCRIPTION	Data category, redaction, retention, disclosure profile
EVIDENCE POINTER	Privacy/disclosure record
RELATED MRO	MRO-10 to MRO-13
BOUNDARY NOTE	Not legal advice

Third-party review boundary

DESCRIPTION	Reviewer role, ruleset, scope, result, limitation
EVIDENCE POINTER	Validation/review record
RELATED MRO	MRO-11
BOUNDARY NOTE	Not certification

Walkthrough Sequence

1. Identify the work unit and review scope.
2. Confirm the authority boundary and role ownership.
3. Reconstruct agent actions, tool actions, and handoffs.
4. Link logs/traces to evidence pointers and lifecycle objects.
5. Verify accepted outcome, dispute, or rejection state.
6. Review exceptions and remediation closure.

7. Confirm privacy/selective disclosure treatment.
8. Record what the walkthrough supports and what it does not prove.

Guide Relationship

Guide 1 should use this template as a technical evidence export and walkthrough architecture pattern. Guide 2 should use it as an audit preparation and governance review workflow.

Placement Recommendation

Include in the main artifact if space allows. Expand implementation details in Guide 1 and governance workflow details in Guide 2.

Appendix C — MRO-to-Audit-Evidence Mapping

Purpose: Preserve the full mapping from GAIC MROs to audit evidence objects in this paper. **Intended reader:** Audit/assurance method teams, compliance engineering, internal audit, AI governance, Guide 1/2 authors.

Source Grounding

GAIC source truth defines MRO-01 through MRO-16. Audit evidence and control language is grounded in AUD-01, AUD-05, and AI-01. This appendix is GAIC-derived plus author synthesis in this paper.

Boundary

MROs are proposed GAIC governance objects. They are not legal mandates, audit standards, certification criteria, procurement criteria, or legal compliance proof.

Full Mapping

Full Mapping

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

MRO-01 Human Role to MAS Responsibility Mapping

EVIDENCE OBJECT	Human-role responsibility map
AUDIT QUESTION	Who owned intent, authorization, review, acceptance, exception, and remediation?
EVIDENCE REQUEST	Role map with lifecycle stage and evidence duty
AARM DIMENSION	Human / agent responsibility mapping
GUIDE 1 RELATION	Capture roles in workflows and evidence metadata
GUIDE 2 RELATION	Define role accountability routines
BOUNDARY	No legal liability assignment

MRO-02 Delegated Authority Boundary

EVIDENCE OBJECT	Delegated authority record
AUDIT QUESTION	Was the agent authorized for the action under scope and condition?
EVIDENCE REQUEST	Authority grant, expiry, risk class, escalation, reauthorization
AARM DIMENSION	Authority traceability
GUIDE 1 RELATION	Implement authority checks and evidence records
GUIDE 2 RELATION	Govern delegation policy and approval
BOUNDARY	Not IAM-only and not legal proof

MRO-03 Agent Role is not Human Role

EVIDENCE OBJECT	Agent role boundary object
AUDIT QUESTION	Was the agent role separated from human accountability?
EVIDENCE REQUEST	Agent role, constraints, owner, escalation
AARM DIMENSION	Human / agent responsibility mapping
GUIDE 1 RELATION	Define agent role registry and capability boundaries
GUIDE 2 RELATION	Govern agent role ownership
BOUNDARY	Do not anthropomorphize

MRO-04 Accepted Outcome Compliance

EVIDENCE OBJECT	Accepted outcome record
AUDIT QUESTION	Was the output accepted, rejected, disputed, or remediated?
EVIDENCE REQUEST	Acceptance criteria, reviewer, outcome state, evidence link
AARM DIMENSION	Accepted outcome evidence
GUIDE 1 RELATION	Capture outcome state transitions
GUIDE 2 RELATION	Govern review and acceptance
BOUNDARY	Not compliance proof

MRO-05 Tool-Action Liability Boundary

EVIDENCE OBJECT	Tool-action evidence object
AUDIT QUESTION	What tool created consequence, under what authority, and with what reversibility?
EVIDENCE REQUEST	Tool, target, action, affected data, rollback, owner
AARM DIMENSION	Agent/tool partitioning
GUIDE 1 RELATION	Instrument tool-action evidence
GUIDE 2 RELATION	Define tool-action approval and review
BOUNDARY	No legal liability conclusion

MRO-06 Responsibility Transfer Across Agents

EVIDENCE OBJECT	Responsibility transfer record
AUDIT QUESTION	Did responsibility transfer or remain with source role?
EVIDENCE REQUEST	Source/target agent, transferred/retained scope, inherited constraints
AARM DIMENSION	Human / agent responsibility mapping
GUIDE 1 RELATION	Capture handoff records
GUIDE 2 RELATION	Govern transfer and escalation
BOUNDARY	Not legal transfer

MRO-07 Authority Drift

EVIDENCE OBJECT	Authority drift exception record
AUDIT QUESTION	Did observed action exceed delegated authority?
EVIDENCE REQUEST	Baseline, observed action, trigger, escalation, closure
AARM DIMENSION	Authority traceability; exception traceability
GUIDE 1 RELATION	Detect and record drift
GUIDE 2 RELATION	Govern drift escalation
BOUNDARY	Not automatic legal breach

MRO-08 MAS Evidence Partitioning

EVIDENCE OBJECT	Partitioned evidence chain
AUDIT QUESTION	Can evidence be separated and linked for review?
EVIDENCE REQUEST	Partition keys, links, hashes, evidence pointers, access rules
AARM DIMENSION	Evidence sufficiency; agent/tool partitioning
GUIDE 1 RELATION	Build partitioned evidence architecture
GUIDE 2 RELATION	Govern evidence access
BOUNDARY	No blanket retention

MRO-09 Cross-Project Reuse Compliance

EVIDENCE OBJECT	Reuse context validation record
AUDIT QUESTION	Was reused agent/workflow context reset and revalidated?
EVIDENCE REQUEST	Source/target context, reset validation, privacy review, authorization refresh
AARM DIMENSION	Agent/tool partitioning
GUIDE 1 RELATION	Implement reuse validation checks
GUIDE 2 RELATION	Govern reuse policy
BOUNDARY	No legal reuse clearance

MRO-10 Privacy / GDPR Lifecycle Mapping

EVIDENCE OBJECT	Privacy lifecycle evidence object
AUDIT QUESTION	How does lifecycle work map to data processing and retention?
EVIDENCE REQUEST	Data flow, purpose, retention rule, rights workflow, minimization
AARM DIMENSION	Privacy and selective disclosure
GUIDE 1 RELATION	Capture privacy metadata
GUIDE 2 RELATION	Govern privacy review
BOUNDARY	Not legal advice

MRO-11 Privacy-Preserving Third-Party Validation

EVIDENCE OBJECT	Validation disclosure profile and verdict record
AUDIT QUESTION	Can third-party review happen without raw over-disclosure?
EVIDENCE REQUEST	Ruleset, redaction/disclosure profile, evidence pointer, verdict hash
AARM DIMENSION	Third-party review boundary; privacy
GUIDE 1 RELATION	Support export/replay packages
GUIDE 2 RELATION	Govern third-party review
BOUNDARY	Not certification

MRO-12 Evidence Minimization and Selective Disclosure

EVIDENCE OBJECT	Evidence minimization and disclosure object
AUDIT QUESTION	Is disclosed evidence limited to review purpose?
EVIDENCE REQUEST	Minimization rule, disclosure justification, redaction profile, access log
AARM DIMENSION	Privacy and selective disclosure
GUIDE 1 RELATION	Build redaction/selective disclosure
GUIDE 2 RELATION	Govern disclosure approvals
BOUNDARY	Not privacy compliance proof

MRO-13 Data Subject Rights vs Evidence Retention

EVIDENCE OBJECT	Rights-retention reconciliation record
AUDIT QUESTION	How are rights requests reconciled with evidence retention?
EVIDENCE REQUEST	Request, retention basis, hold status, redaction/deletion action
AARM DIMENSION	Privacy and selective disclosure
GUIDE 1 RELATION	Support retention/rights workflows
GUIDE 2 RELATION	Govern legal/privacy review
BOUNDARY	Requires counsel

MRO-14 Third-Party Processor / Subprocessor Chain

EVIDENCE OBJECT	Processor/subprocessor chain evidence
AUDIT QUESTION	Which third parties processed or affected lifecycle work?
EVIDENCE REQUEST	Processor role, subprocessor chain, data agreement pointer, responsibility owner
AARM DIMENSION	Agent/tool partitioning
GUIDE 1 RELATION	Capture vendor/processor evidence
GUIDE 2 RELATION	Govern processor chain
BOUNDARY	No vendor ranking

MRO-15 Vendor / Model / Runtime Substitution Conformance

EVIDENCE OBJECT	Substitution conformance record
AUDIT QUESTION	Did evidence continuity survive component change?
EVIDENCE REQUEST	Prior/new component, integrity check, reauthorization, revalidation
AARM DIMENSION	Agent/tool partitioning; evidence sufficiency
GUIDE 1 RELATION	Implement substitution evidence controls
GUIDE 2 RELATION	Govern change review
BOUNDARY	Not conformance certification

MRO-16 Incident, Dispute, and Remediation Closure

EVIDENCE OBJECT	Incident/remediation closure record
AUDIT QUESTION	Did exception or dispute close with evidence?
EVIDENCE REQUEST	Event, affected outcome, owner, corrective action, recheck, closure
AARM DIMENSION	Exception traceability; remediation closure
GUIDE 1 RELATION	Build closure workflow evidence
GUIDE 2 RELATION	Govern remediation closure
BOUNDARY	Not legal closure

Guide Relationship

Guide 1 should treat this appendix as an implementation-neutral architecture checklist. Guide 2 should treat it as a policy-to-evidence and compliance operating model checklist.

Placement Recommendation

Include full mapping in the this paper appendix. Consider a shorter summary table in Chapter 6.

Appendix D — Auditability Readiness Matrix

Purpose: Provide AARM dimension and readiness-level crosswalk. **Intended reader:** Audit readiness teams, CIO/CTO/CCO leaders, internal audit, governance committees, platform teams.

Source Grounding

Audit evidence: AUD-01. AI audit practice: AUD-05, AUD-06. AI governance: AI-01. Boundary: BOUND-03. GAIC dependency: R4B/R1 AARM baseline, MROs, ALCS context, Validation Lab boundary.

Boundary

AARM is not an audit standard, certification, assurance opinion, legal compliance proof, regulator approval, procurement recommendation, vendor ranking, or score.

AARM Dimensions

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Audit object clarity

DEFINITION	System can identify lifecycle work under review
AUDIT QUESTION	Can reviewers name the lifecycle unit, scope, and outcome?
EVIDENCE EXPECTED	Lifecycle ID, workflow scope, object inventory
GAIC/MRO DEPENDENCY	MRO-01, MRO-03, MRO-04
AUTHOR-SYNTHESIS NOTE	AARM synthesis
BOUNDARY	Not legal category

Authority traceability

DEFINITION	Delegated authority can be traced to action
AUDIT QUESTION	Was the action authorized within scope?
EVIDENCE EXPECTED	Authority grant, expiry, escalation, reauthorization
GAIC/MRO DEPENDENCY	MRO-02, MRO-07
AUTHOR-SYNTHESIS NOTE	AARM synthesis
BOUNDARY	Not legal delegation proof

Human / agent responsibility mapping

DEFINITION	Human and agent roles are separated and mapped
AUDIT QUESTION	Who owned intent, execution, review, acceptance, remediation?
EVIDENCE EXPECTED	Role-responsibility map, transfer record
GAIC/MRO DEPENDENCY	MRO-01, MRO-03, MRO-06
AUTHOR-SYNTHESIS NOTE	AARM synthesis
BOUNDARY	No anthropomorphism

Evidence sufficiency

DEFINITION	Evidence supports reconstruction under review scope
AUDIT QUESTION	Can reviewers reconstruct without narrative memory?
EVIDENCE EXPECTED	Evidence chain, pointers, integrity records
GAIC/MRO DEPENDENCY	MRO-04, MRO-08, MRO-16
AUTHOR-SYNTHESIS NOTE	AARM synthesis
BOUNDARY	Not professional sufficiency conclusion

Agent/tool partitioning

DEFINITION	Evidence separates agent, tool, vendor, processor, project
AUDIT QUESTION	Can reviewers identify which surface produced each evidence item?
EVIDENCE EXPECTED	Agent/tool trace, partition keys, substitution record
GAIC/MRO DEPENDENCY	MRO-05, MRO-08, MRO-14, MRO-15
AUTHOR-SYNTHESIS NOTE	AARM synthesis
BOUNDARY	No vendor ranking

Exception traceability

DEFINITION	Exceptions link to lifecycle objects and owners
AUDIT QUESTION	Can deviation be linked to authority, impact, and next action?
EVIDENCE EXPECTED	Exception record, impact scope, escalation
GAIC/MRO DEPENDENCY	MRO-07, MRO-16
AUTHOR-SYNTHESIS NOTE	GAIC-derived synthesis
BOUNDARY	Not legal violation finding

Accepted outcome evidence

DEFINITION	Outcome states are recorded with role and evidence linkage
AUDIT QUESTION	Who accepted, rejected, disputed, or remediated the result?
EVIDENCE EXPECTED	Outcome state, reviewer role, evidence link
GAIC/MRO DEPENDENCY	MRO-04
AUTHOR-SYNTHESIS NOTE	GAIC-derived synthesis
BOUNDARY	Acceptance is not compliance proof

Remediation closure

DEFINITION	Corrective action and closure are evidence-backed
AUDIT QUESTION	What was remediated and who accepted closure?
EVIDENCE EXPECTED	Remediation plan, correction evidence, recheck, closure status
GAIC/MRO DEPENDENCY	MRO-16
AUTHOR-SYNTHESIS NOTE	GAIC-derived synthesis
BOUNDARY	Not legal settlement

Privacy and selective disclosure

DEFINITION	Evidence can be minimized and disclosed by scope
AUDIT QUESTION	Can review happen without exposing more than needed?
EVIDENCE EXPECTED	Redaction, retention, disclosure profile
GAIC/MRO DEPENDENCY	MRO-10 to MRO-13
AUTHOR-SYNTHESIS NOTE	author synthesis in this paper
BOUNDARY	Not legal advice

Third-party review boundary

DEFINITION	Review role and claim boundary are explicit
AUDIT QUESTION	What can reviewers inspect and not claim?
EVIDENCE EXPECTED	Scope, ruleset, verdict, boundary statement
GAIC/MRO DEPENDENCY	MRO-11, MRO-12, MRO-16
AUTHOR-SYNTHESIS NOTE	author synthesis in this paper
BOUNDARY	No certification/opinion

Readiness Levels

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

L0 Unobservable

DEFINITION	Lifecycle work cannot be reconstructed
OBSERVABLE TRAITS	Outputs only, informal memory, no lifecycle ID
MINIMUM EVIDENCE	None or ad hoc notes
WHAT IT DOES NOT PROVE	Absence of risk or non-compliance
SOURCE / SYNTHESIS CLASSIFICATION	author synthesis in this paper grounded in audit evidence concepts

L1 Log-Visible

DEFINITION	Activity logs exist but are not responsibility-linked
OBSERVABLE TRAITS	Logs, traces, timestamps, event streams
MINIMUM EVIDENCE	Raw logs or trace exports
WHAT IT DOES NOT PROVE	Authority, responsibility, accepted outcome, closure
SOURCE / SYNTHESIS CLASSIFICATION	EVID-02/EVID-03 + author synthesis in this paper

L2 Trace-Linked

DEFINITION	Logs/traces connect some actions to workflows or actors
OBSERVABLE TRAITS	Workflow IDs, agent/tool traces, partial actor link
MINIMUM EVIDENCE	Trace-linked run records
WHAT IT DOES NOT PROVE	Complete lifecycle responsibility or sufficient evidence
SOURCE / SYNTHESIS CLASSIFICATION	EVID-01/EVID-02 + author synthesis in this paper

L3 Evidence-Structured

DEFINITION	Lifecycle objects exist for reconstruction
OBSERVABLE TRAITS	Role maps, authority objects, outcome records, exception records
MINIMUM EVIDENCE	MRO-aligned evidence objects
WHAT IT DOES NOT PROVE	Assurance readiness, legal compliance, operating effectiveness
SOURCE / SYNTHESIS CLASSIFICATION	GAIC-derived + audit/control grounding

L4 Auditability-Ready

DEFINITION	Defined-scope evidence chains reconstruct lifecycle work
OBSERVABLE TRAITS	Walkthroughs, evidence requests, privacy controls, closure records
MINIMUM EVIDENCE	Complete evidence chain for review scope
WHAT IT DOES NOT PROVE	Audit opinion, certification, regulator approval
SOURCE / SYNTHESIS CLASSIFICATION	author synthesis in this paper

L5 Assurance-Ready

DEFINITION	Evidence architecture may support professional assurance planning or review within a defined scope
OBSERVABLE TRAITS	Repeatable evidence packs, reviewer scope, rulesets, recheck procedures
MINIMUM EVIDENCE	Reviewable evidence packs and boundary statement
WHAT IT DOES NOT PROVE	Actual assurance, certification, legal compliance, regulator acceptance
SOURCE / SYNTHESIS CLASSIFICATION	author synthesis in this paper bounded by AUD-03/BOUND-03

Minimum Evidence Crosswalk

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Audit object clarity

L0/L1 CONCERN	No stable work unit
L2/L3 TRANSITION	Work unit IDs appear in traces and evidence objects
L4/L5 TRANSITION	Work units support scoped walkthroughs and evidence packs

Authority traceability

L0/L1 CONCERN	Access logs only
L2/L3 TRANSITION	Partial authority metadata
L4/L5 TRANSITION	Delegated authority records are reconstructable

Responsibility mapping

L0/L1 CONCERN	User/service IDs only
L2/L3 TRANSITION	Partial role map
L4/L5 TRANSITION	Human/agent responsibility chain is explicit

Evidence sufficiency

L0/L1 CONCERN	Narrative memory
L2/L3 TRANSITION	Evidence pointers emerge
L4/L5 TRANSITION	Evidence chain supports review objective

Partitioning

L0/L1 CONCERN	Flattened logs
L2/L3 TRANSITION	Partial partitions
L4/L5 TRANSITION	Agent/tool/vendor/project/privacy partitions are linked

Exception traceability

L0/L1 CONCERN	Incident tickets only
L2/L3 TRANSITION	Exceptions linked to work units
L4/L5 TRANSITION	Exceptions include owner, action, recheck, closure

Accepted outcome

L0/L1 CONCERN	Output completion only
L2/L3 TRANSITION	Partial review records
L4/L5 TRANSITION	Accepted/rejected/disputed/remediated states are evidence-backed

Remediation closure

L0/L1 CONCERN	Closed ticket only
L2/L3 TRANSITION	Corrective action tracked
L4/L5 TRANSITION	Closure is accepted and reopen criteria exist

Privacy/selective disclosure

L0/L1 CONCERN	Collect everything or disclose ad hoc
L2/L3 TRANSITION	Redaction/retention rules appear
L4/L5 TRANSITION	Selective disclosure packages support review

Third-party boundary

L0/L1 CONCERN	No scope statement
L2/L3 TRANSITION	Limited reviewer scope
L4/L5 TRANSITION	Ruleset, scope, verdict, and non-certifying boundary are explicit

Appendix E — Exception / Remediation Closure Checklist

Purpose: Define closure evidence requirements for exceptions, disputes, and remediation. **Intended reader:** Internal audit, incident governance, compliance, risk, security, AI platform owners, remediation owners.

Source Grounding

Incident/governance context: EVID-04, AI-02. Controls and audit practice: AUD-04, AUD-05, AUD-06. GAIC dependencies: MRO-04, MRO-07, MRO-16.

Boundary

This checklist frames closure as evidence-backed governance state. It does not prove legal remedy, settlement, regulator acceptance, operational effectiveness, certification, or assurance.

Closure Checklist

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Authority exception	
EVIDENCE REQUIRED	Baseline authority, observed action, trigger, work unit, evidence pointer
OWNER	Escalation owner
CORRECTIVE ACTION	Stop, downgrade, reauthorize, confirm, or constrain
CLOSURE STATUS	Open / remediated / closed
REVIEWER	Authority reviewer
REOPEN CRITERIA	Repeat drift, scope expansion, failed confirmation
BOUNDARY NOTE	Not legal breach finding

Responsibility gap

EVIDENCE REQUIRED	Missing role owner, affected lifecycle stage, agent role, work unit
OWNER	Governance owner
CORRECTIVE ACTION	Assign owner, update role map, recheck workflow
CLOSURE STATUS	Open / assigned / closed
REVIEWER	Governance reviewer
REOPEN CRITERIA	Owner missing in later run
BOUNDARY NOTE	Not legal liability assignment

Tool-action side effect

EVIDENCE REQUIRED	Tool call, affected system/data, reversibility, rollback evidence
OWNER	Tool-action owner
CORRECTIVE ACTION	Rollback, correction, disable tool, add approval
CLOSURE STATUS	Open / corrected / closed
REVIEWER	System owner
REOPEN CRITERIA	New side effect or rollback failure
BOUNDARY NOTE	Not liability conclusion

Outcome dispute

EVIDENCE REQUIRED	Output, criteria, dispute reason, reviewer notes
OWNER	Review owner
CORRECTIVE ACTION	Re-review, correct, reject, accept with limitation
CLOSURE STATUS	Open / disputed / remediated / accepted
REVIEWER	Outcome reviewer
REOPEN CRITERIA	New evidence or recurring dispute
BOUNDARY NOTE	Not settlement proof

Privacy event

EVIDENCE REQUIRED	Data category, disclosure path, redaction/retention record
OWNER	Privacy owner
CORRECTIVE ACTION	Redact, restrict, delete, hold, notify internal process
CLOSURE STATUS	Open / contained / closed
REVIEWER	Privacy reviewer
REOPEN CRITERIA	Unauthorized exposure or rights request
BOUNDARY NOTE	Not legal advice

Handoff failure

EVIDENCE REQUIRED	Source/target agent, transferred scope, retained scope, rejection reason
OWNER	Process owner
CORRECTIVE ACTION	Return, reassign, clarify constraints
CLOSURE STATUS	Open / reassigned / closed
REVIEWER	Process reviewer
REOPEN CRITERIA	Repeated failed handoff
BOUNDARY NOTE	Not legal transfer

Cross-project reuse issue

EVIDENCE REQUIRED	Source context, target context, reuse decision, missing reset
OWNER	Reuse owner
CORRECTIVE ACTION	Revalidate, reset, withdraw, update evidence
CLOSURE STATUS	Open / revalidated / closed
REVIEWER	Reuse/governance reviewer
REOPEN CRITERIA	Reuse in new context
BOUNDARY NOTE	Not legal reuse clearance

Substitution evidence gap

EVIDENCE REQUIRED	Prior/new component, evidence break, regression or revalidation status
OWNER	Change owner
CORRECTIVE ACTION	Reauthorize, retest, revalidate evidence chain
CLOSURE STATUS	Open / revalidated / closed
REVIEWER	Change reviewer
REOPEN CRITERIA	Component change or failed evidence continuity
BOUNDARY NOTE	Not conformance certification

Remediation defect

EVIDENCE REQUIRED	Prior remediation, failed recheck, affected outcome
OWNER	Remediation owner
CORRECTIVE ACTION	Reopen, assign corrective action, recheck
CLOSURE STATUS	Reopened / corrected / closed
REVIEWER	Closure reviewer
REOPEN CRITERIA	Failed recheck or recurring exception
BOUNDARY NOTE	Not regulator closure

Closure Evidence Fields

- Event ID.
- Work unit ID.
- Affected MRO.
- Affected outcome.
- Owner.
- Evidence pointer.
- Corrective action.
- Recheck evidence.
- Closure reviewer.
- Closure status.
- Reopen criteria.
- Privacy treatment.
- Boundary statement.

Guide Relationship

Guide 1 should translate closure fields into workflow, ticketing, runtime, and evidence-export architecture.
Guide 2 should translate closure fields into incident/dispute/remediation governance routines.

Appendix F — Boundary Language

Purpose: Provide safe language for auditability, assurance, validation, MRO, AARM, MPLP, privacy, and professional-use boundaries. **Intended reader:** Author, editors, reviewers, publication QA, Guide 1/2 teams, source QA.

Source Grounding

Assurance and attestation boundary: AUD-03, BOUND-03. Conformity/certification boundary: BOUND-01, BOUND-02. GAIC dependencies: Validation Lab boundary, MPLP boundary, companion-paper boundary, R4A/R4B/R4C boundary QA.

Allowed / Forbidden / Replacement Language

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

Paper status

ALLOWED WORDING	"public research edition"; "public HTML artifact"; "public PDF artifact"
FORBIDDEN WORDING	"published"; "final"; "sealed"; "live public release"
SAFE REPLACEMENT	"public research edition, not an audit standard, certification, assurance opinion, or endorsement"

Auditability

ALLOWED WORDING	"proposed auditability framework"; "lifecycle evidence guide"
FORBIDDEN WORDING	"audit standard"; "official audit methodology"
SAFE REPLACEMENT	"auditability readiness framework"

Evidence

ALLOWED WORDING	"evidence chain supports reconstruction"
FORBIDDEN WORDING	"evidence chain proves compliance"
SAFE REPLACEMENT	"evidence chain supports responsibility review"

Agentic Audit Object

ALLOWED WORDING	"proposed object model in this paper"
FORBIDDEN WORDING	"required schema"; "legal category"
SAFE REPLACEMENT	"conceptual review object"

MRO

ALLOWED WORDING	"GAIC-derived governance object"
FORBIDDEN WORDING	"legal requirement"; "regulator-mandated object"
SAFE REPLACEMENT	"proposed lifecycle responsibility object"

AARM

ALLOWED WORDING	"proposed readiness model"
FORBIDDEN WORDING	"certification level"; "maturity score"; "assurance grade"
SAFE REPLACEMENT	"readiness vocabulary"

Assurance

ALLOWED WORDING	"may support assurance planning subject to professional scope"
FORBIDDEN WORDING	"assures"; "guarantees assurance"; "issues assurance opinion"
SAFE REPLACEMENT	"may support readiness discussion"

Certification

ALLOWED WORDING	"non-certifying review"; "not certification"
FORBIDDEN WORDING	"certified"; "certifies systems"; "certification-ready"
SAFE REPLACEMENT	"reviewable under defined scope"

Legal/privacy

ALLOWED WORDING	"privacy/legal context"; "requires legal review"
FORBIDDEN WORDING	"GDPR compliant"; "legal compliance proof"; "legal advice"
SAFE REPLACEMENT	"privacy-aware evidence architecture"

Big Four

ALLOWED WORDING	"reviewed sources create market context"
FORBIDDEN WORDING	"Big Four endorse this paper"; "Big Four need this paper"
SAFE REPLACEMENT	"Big Four sources discuss related AI assurance/governance topics"

Audit/professional bodies

ALLOWED WORDING	"supports terminology or boundary"
FORBIDDEN WORDING	"endorses this paper"; "adopts AARM"
SAFE REPLACEMENT	"provides external professional-language context"

Validation Lab

ALLOWED WORDING	"non-certifying evidence adjudication example"
FORBIDDEN WORDING	"certification authority"; "proves compliance"; "regulator-approved"
SAFE REPLACEMENT	"one optional example of evidence review boundary"

MPLP

ALLOWED WORDING	"one optional protocol path"
FORBIDDEN WORDING	"required"; "industry standard"; "only correct path"
SAFE REPLACEMENT	"one possible implementation path"

Cognitive OS

ALLOWED WORDING	"possible runtime path for evidence capture in related product-stack planning"
FORBIDDEN WORDING	"required for auditability"
SAFE REPLACEMENT	"one possible implementation environment, outside this paper requirement"

SoloCrew

ALLOWED WORDING	"product proof path for workflow patterns in planning notes"
FORBIDDEN WORDING	"proves enterprise readiness"
SAFE REPLACEMENT	"illustrative product proof path, not an enterprise-readiness claim"

SEO/GEO

ALLOWED WORDING	"future planning note"
FORBIDDEN WORDING	"uplift achieved"; "answer-engine recognized"
SAFE REPLACEMENT	"future publication/visibility work after artifact exists"

Procurement

ALLOWED WORDING	"readiness conversation input"
FORBIDDEN WORDING	"recommended vendor"; "procurement ranking"
SAFE REPLACEMENT	"non-procurement framework"

Later insurability white paper

ALLOWED WORDING	"future insurability paper"
FORBIDDEN WORDING	"insurance guarantee"; "underwriting standard"
SAFE REPLACEMENT	"future source research required"

Big Four Boundary

Use:

- "The reviewed Big Four sources discuss AI assurance, trusted AI, governance, risk, controls, audit transformation, and agentic AI adoption."
- "These sources create market context for this paper."

Do not use:

- "Big Four firms endorse this paper."
- "The Big Four need this paper."
- "this paper replaces Big Four methodology."

Audit / Professional Body Boundary

Use professional sources for terminology and boundary. Do not imply that PCAOB, IAASB, AICPA, ISACA, IIA, ISO, NIST, OECD, EU institutions, Singapore agencies, ICO, or any other organization endorses or adopts this paper.

Legal / Privacy Boundary

Use privacy and legal sources as context. State that jurisdiction-specific interpretation belongs to qualified legal/privacy professionals. Do not say any evidence object in this paper proves GDPR, EU AI Act, UK GDPR, or other legal compliance.

Validation Lab Boundary

Validation Lab may be described only as one non-certifying evidence adjudication example. It does not certify systems, issue compliance certificates, provide legal opinions, issue assurance opinions, act as regulator, act as conformity assessment body, prove enterprise readiness, or guarantee insurability.

MPLP Boundary

MPLP may be discussed only as one optional protocol path. This paper does not require MPLP, does not claim MPLP is industry-standard, and does not claim MPLP proves compliance, auditability, assurance, or enterprise readiness.

AARM Boundary

AARM is a proposed readiness model. It does not score vendors, certify systems, issue assurance, prove legal compliance, bind regulators, or provide procurement guidance. L5 "Assurance-Ready" means evidence

architecture may support qualified assurance planning or review within a defined scope, subject to professional scope and judgment. It does not mean assurance has been issued.

Source Note and Citation Register

This paper uses source-ID citation notes rather than public footnotes. Source IDs appear in chapter source notes and table source-support columns. The candidate package keeps the full source register and chapter-level citation map as internal package records.

Citation Style Decision

- Internal style for R4: source IDs in text and tables, plus a package-level source register.
- Public style for a later wave: decide whether to convert source IDs into endnotes, footnotes, or an appendix source register.
- Quote guidance: paraphrase by default; avoid direct quotation unless a later citation QA wave approves a short, necessary excerpt.
- Boundary guidance: source presence does not imply endorsement, standard-setting, certification, assurance, legal advice, procurement approval, or regulator approval.

Source Role Summary

- BF sources: market context only.
- AUD sources: audit/assurance terminology and boundary context.
- AI sources: governance context, not legal compliance proof.
- EVID sources: observability, provenance, logs, and evidence-chain distinction.
- PRIV sources: minimization, retention, selective disclosure, and privacy/evidence tension.
- BOUND sources: validation, verification, attestation, certification, and conformity-assessment boundaries.
- GAIC-SOURCE: MROs, RCCS-M/ALCS context, Validation Lab boundary, and companion-paper sequencing.

Artifact Integrity Note

The public HTML and PDF artifacts are published with manifest and checksum records for integrity verification. Source-register and citation-map material remains included for transparency, without creating certification, endorsement, or assurance opinion.

Source Note and Citation Register

The source register and citation map below support source traceability. Source IDs remain citation handles for review and future citation styling.

AIAAWP Source Register

Status: public research edition source register **Boundary:** Sources listed here are used for context, terminology, grounding, or boundary discipline. Inclusion does not imply endorsement of this paper, GAIC, MPLP, AARM, Validation Lab, or Jearon Wong.

Source Use Rules

- Big Four sources are market context only.
- Audit/professional sources support audit evidence, assurance, controls, internal audit, and boundary language.
- AI governance/public guidance sources support governance context, risk management, accountability, documentation, monitoring, and human oversight.
- Evidence/provenance/observability/logging sources support the distinction between logs/traces and responsibility-linked audit evidence chains.
- Privacy sources support minimization, retention, selective disclosure, and evidence/privacy tension.
- Boundary sources support distinctions among validation, assurance, attestation, certification, conformity assessment, and audit opinion.
- GAIC source truth supports MROs, RCCS-M/ALCS context, Evidence-Based Validation Pattern, Validation Lab boundary, and companion-paper boundary.

Register

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review surfaces.

BF-01

TITLE	AI agents: Scaling faster than expected
ORGANIZATION	Deloitte
SOURCE CLASS	Big Four market context
USE ROLE	Agentic AI adoption and enterprise readiness context
CHAPTERS USED	0, 2, 14, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No endorsement claim.

BF-02

TITLE	Assurance for AI
ORGANIZATION	PwC
SOURCE CLASS	Big Four market context
USE ROLE	AI assurance market context
CHAPTERS USED	0, 1, 14, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Context only; not validation of this paper.

BF-03

TITLE	EY Assurance releases new technology capabilities strengthening confidence and trust
ORGANIZATION	EY
SOURCE CLASS	Big Four market context
USE ROLE	Audit transformation and assurance technology context
CHAPTERS USED	1, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No methodology replacement claim.

BF-04

TITLE	KPMG Trusted AI
ORGANIZATION	KPMG
SOURCE CLASS	Big Four market context
USE ROLE	Trusted AI, governance, privacy, and accountability context
CHAPTERS USED	2, 10, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Context only.

BF-05

TITLE	AI Agents Move Beyond Experimentation as Leaders Prepare for Competitive Transformation Within 24 Months
ORGANIZATION	KPMG
SOURCE CLASS	Big Four market context
USE ROLE	Agentic AI adoption and risk context
CHAPTERS USED	0, 2, 14, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No adoption or endorsement claim.

AUD-01

TITLE	AS 1105: Audit Evidence
ORGANIZATION	PCAOB
SOURCE CLASS	Audit/professional body
USE ROLE	Audit evidence terminology and sufficiency/appropriateness framing
CHAPTERS USED	0, 3, 4, 5, 6, 7, 13, 16
QUOTE/PARAPHRASE NOTE	Paraphrase; short definition quote only if later approved
BOUNDARY NOTE	This paper is not a PCAOB audit standard.

AUD-02

TITLE	ISA 500 (Revised), Audit Evidence
ORGANIZATION	IAASB
SOURCE CLASS	Audit/professional body
USE ROLE	Audit evidence and technology-enabled evidence context
CHAPTERS USED	3, 4, 7, 13
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	This paper is not ISA methodology.

AUD-03

TITLE	ISAE 3000 (Revised)
ORGANIZATION	IAASB
SOURCE CLASS	Audit/professional body
USE ROLE	Assurance engagement boundary
CHAPTERS USED	1, 12, 13, 15, 16
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	This paper does not issue assurance conclusions.

AUD-04

TITLE	Trust Services Criteria
ORGANIZATION	AICPA
SOURCE CLASS	Audit/professional body
USE ROLE	Control vocabulary, monitoring, privacy/security, control activities
CHAPTERS USED	5, 7, 8, 10, 11, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	This paper is not SOC or attestation guidance.

AUD-05

TITLE	Audit and Assurance Guidance for the NIST Cybersecurity Framework 2.0 and Artificial Intelligence
ORGANIZATION	ISACA
SOURCE CLASS	Audit/professional body
USE ROLE	AI audit and assurance practice context
CHAPTERS USED	6, 7, 8, 11, 13, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No ISACA endorsement.

AUD-06

TITLE	Global Technology Audit Guide: Auditing Artificial Intelligence
ORGANIZATION	IIA
SOURCE CLASS	Audit/professional body
USE ROLE	Internal audit, AI governance, controls, human oversight
CHAPTERS USED	2, 7, 8, 11, 13, 14, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Internal audit language only.

AUD-07

TITLE	The IIA's Three Lines Model
ORGANIZATION	IIA
SOURCE CLASS	Audit/professional body
USE ROLE	Governance role separation
CHAPTERS USED	1, 14, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Not an enterprise mandate.

AI-01

TITLE	Artificial Intelligence Risk Management Framework 1.0
ORGANIZATION	NIST
SOURCE CLASS	Regulator/public guidance
USE ROLE	AI risk management, governance, monitoring, accountability
CHAPTERS USED	0, 2, 3, 5, 6, 13, 14, 16
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Guidance context, not legal proof.

AI-02

TITLE	NIST AI 600-1 Generative AI Profile
ORGANIZATION	NIST
SOURCE CLASS	Regulator/public guidance
USE ROLE	Generative AI risk, incidents, monitoring, documentation
CHAPTERS USED	2, 7, 11, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Guidance context only.

AI-03

TITLE	OECD AI Principles
ORGANIZATION	OECD
SOURCE CLASS	Public guidance
USE ROLE	Responsible AI accountability and transparency context
CHAPTERS USED	1, 2, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Policy context only.

AI-04

TITLE	ISO/IEC 42001:2023 AI management system
ORGANIZATION	ISO
SOURCE CLASS	Standards body
USE ROLE	AI management-system context
CHAPTERS USED	2, 13, 14
QUOTE/PARAPHRASE NOTE	Paraphrase public summary only
BOUNDARY NOTE	No compliance claim.

AI-05

TITLE	ISO/IEC 23894:2023 AI risk management
ORGANIZATION	ISO
SOURCE CLASS	Standards body
USE ROLE	AI risk-management context
CHAPTERS USED	2, 13, 14
QUOTE/PARAPHRASE NOTE	Paraphrase public summary only
BOUNDARY NOTE	No standard compliance claim.

AI-06

TITLE	Regulation (EU) 2024/1689 Artificial Intelligence Act
ORGANIZATION	European Union
SOURCE CLASS	Regulator/public guidance
USE ROLE	AI governance, logging, human oversight, documentation context
CHAPTERS USED	1, 2, 4, 7, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Legal context only; no legal advice.

AI-07

TITLE	Model AI Governance Framework for Generative AI
ORGANIZATION	AI Verify Foundation / IMDA
SOURCE CLASS	Public guidance
USE ROLE	Generative AI governance and testing context
CHAPTERS USED	1, 2, 7, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No endorsement.

AI-08

TITLE	Model AI Governance Framework for Agentic AI
ORGANIZATION	Singapore MDDI / AI Verify Foundation
SOURCE CLASS	Public guidance
USE ROLE	Agentic AI governance, autonomy, accountability, safeguards
CHAPTERS USED	0, 2, 3, 5, 6, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Context only.

AI-09

TITLE	Guidance on AI and Data Protection
ORGANIZATION	UK ICO
SOURCE CLASS	Regulator/public guidance
USE ROLE	AI/privacy context
CHAPTERS USED	9, 10, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No legal advice.

EVID-01

TITLE	PROV Overview
ORGANIZATION	W3C
SOURCE CLASS	Technical documentation
USE ROLE	Provenance vocabulary for entities, activities, agents, derivation, responsibility
CHAPTERS USED	3, 4, 5, 8, 9, 13
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Technical vocabulary, not audit standard.

EVID-02

TITLE	OpenTelemetry Observability Primer
ORGANIZATION	OpenTelemetry
SOURCE CLASS	Technical documentation
USE ROLE	Logs, metrics, traces, observability context
CHAPTERS USED	0, 4, 8, 9, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Do not criticize vendors.

EVID-03

TITLE	SP 800-92 Guide to Computer Security Log Management
ORGANIZATION	NIST
SOURCE CLASS	Public guidance
USE ROLE	Log management context
CHAPTERS USED	4, 9, 10
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Older official logging source; not agentic auditability.

EVID-04

TITLE	Cybersecurity Framework 2.0
ORGANIZATION	NIST
SOURCE CLASS	Public guidance
USE ROLE	Governance, detection, response, recovery, incident context
CHAPTERS USED	4, 11, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Cybersecurity context only.

PRIV-01

TITLE	GDPR Article 5
ORGANIZATION	European Union
SOURCE CLASS	Regulator/public guidance
USE ROLE	Minimization, storage limitation, accountability context
CHAPTERS USED	9, 10
QUOTE/PARAPHRASE NOTE	Paraphrase; short quote only if later approved
BOUNDARY NOTE	Legal context, not advice.

PRIV-02

TITLE	EDPB Guidelines 4/2019 on Article 25
ORGANIZATION	EDPB
SOURCE CLASS	Regulator/public guidance
USE ROLE	Data protection by design/default and minimization context
CHAPTERS USED	9, 10
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No GDPR proof.

PRIV-03

TITLE	NIST Privacy Framework 1.0
ORGANIZATION	NIST
SOURCE CLASS	Public guidance
USE ROLE	Privacy risk management and governance context
CHAPTERS USED	10, 14
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	Guidance only.

PRIV-04

TITLE	Data minimisation
ORGANIZATION	UK ICO
SOURCE CLASS	Regulator/public guidance
USE ROLE	Data minimization context
CHAPTERS USED	10
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No legal advice.

PRIV-05

TITLE	Storage limitation
ORGANIZATION	UK ICO
SOURCE CLASS	Regulator/public guidance
USE ROLE	Retention and storage limitation context
CHAPTERS USED	10
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	No retention mandate.

BOUND-01

TITLE	CASCO - Committee on conformity assessment
ORGANIZATION	ISO
SOURCE CLASS	Standards body
USE ROLE	Conformity-assessment boundary
CHAPTERS USED	12, 13
QUOTE/PARAPHRASE NOTE	Paraphrase public page only
BOUNDARY NOTE	Validation Lab is not a conformity assessment body.

BOUND-02

TITLE	ISO/IEC 17029:2019 summary
ORGANIZATION	ISO
SOURCE CLASS	Standards body
USE ROLE	Validation/verification body boundary
CHAPTERS USED	12, 13
QUOTE/PARAPHRASE NOTE	Paraphrase public summary only
BOUNDARY NOTE	No conformance claim.

BOUND-03

TITLE	Attestation standards
ORGANIZATION	AICPA
SOURCE CLASS	Audit/professional body
USE ROLE	Attestation and practitioner-report boundary
CHAPTERS USED	12, 13, 15
QUOTE/PARAPHRASE NOTE	Paraphrase only
BOUNDARY NOTE	This paper is not an attestation engagement.

GAIC-SOURCE

TITLE	GAIC v0.3.2-FRC-R3 source truth
ORGANIZATION	Jearon Wong / GAIC source master
SOURCE CLASS	GAIC-derived source truth
USE ROLE	MROs, RCCS-M/ALCS context, Evidence-Based Validation Pattern, Validation Lab boundary, companion-paper boundary
CHAPTERS USED	1, 3, 5, 6, 9, 10, 11, 12, 13, 16; appendices
QUOTE/PARAPHRASE NOTE	Internal source reference
BOUNDARY NOTE	GAIC-derived constructs must be labeled as author framework, not external standard.

Source Register Result

PASS - all sources used in this public research edition are drawn from the source inventory or GAIC source truth.

Citation Style

This paper uses source IDs in text, tables, and source notes with this register as the package-level lookup. Public endnote or footnote formatting remains a future citation-style decision.

Package Citation Map

AIAAWP Citation Map

Status: public research edition citation map **Boundary:** This map supports source review. It does not create final citation formatting, quote source text, or claim endorsement by any source organization.

Chapter Citation Map

Table rendered as semantic row cards to preserve GAIC-style readability across HTML and PDF review

surfaces.

0 Executive Summary

SOURCE IDS USED	BF-01, BF-02, BF-05, AUD-01, AI-01, EVID-02, GAIC-SOURCE
SOURCE ROLE	Market context; audit evidence; governance; observability; GAIC source truth
USE PLACEMENT	Main argument and source note
CITATION RISK	Big Four endorsement overclaim
REVIEW NOTE	Keep Big Four as context only.

1 Scope, Audience, Boundary

SOURCE IDS USED	AUD-03, BOUND-01, BOUND-02, BOUND-03, AUD-07, BF-02, BF-03, GAIC-SOURCE
SOURCE ROLE	Boundary, audience, assurance/attestation context
USE PLACEMENT	Boundary section and table notes
CITATION RISK	Treating boundary sources as endorsement
REVIEW NOTE	Repeat non-standard and non-certifying boundary.

2 Difference from Model Governance

SOURCE IDS USED	BF-01, BF-04, BF-05, AI-01, AI-02, AI-04, AI-05, AI-08, AUD-06
SOURCE ROLE	Market and AI governance context
USE PLACEMENT	Main text and Table 1
CITATION RISK	Dismissing model governance
REVIEW NOTE	Maintain "necessary but insufficient" framing.

3 Audit Object Shift

SOURCE IDS USED	AUD-01, AUD-02, AI-01, AI-08, EVID-01, GAIC-SOURCE
SOURCE ROLE	Audit evidence, governance, provenance, GAIC MROs
USE PLACEMENT	Main text
CITATION RISK	Treating Agentic Audit Object as externally sourced
REVIEW NOTE	Label as author synthesis in this paper.

4 Logs vs Evidence Chains

SOURCE IDS USED	AUD-01, AUD-02, EVID-01, EVID-02, EVID-03, EVID-04, AI-06
SOURCE ROLE	Audit evidence, observability, logging, incident/governance context
USE PLACEMENT	Main text and Table 2
CITATION RISK	Saying logs are useless
REVIEW NOTE	State logs are useful evidence ingredients.

5 Agentic Audit Object Model

SOURCE IDS USED	AUD-01, AUD-04, AI-01, EVID-01, GAIC-SOURCE
SOURCE ROLE	Evidence, controls, governance, provenance, GAIC objects
USE PLACEMENT	Main text and object table
CITATION RISK	Mandatory schema overclaim
REVIEW NOTE	Keep model as proposed architecture.

6 MRO Mapping

SOURCE IDS USED	GAIC-SOURCE, AUD-01, AUD-05, AI-01
SOURCE ROLE	GAIC MRO source truth; evidence/control language
USE PLACEMENT	Main text and Table 3
CITATION RISK	MRO legal mandate overclaim
REVIEW NOTE	Keep MROs as GAIC-derived governance objects.

7 Evidence Request List

SOURCE IDS USED	AUD-01, AUD-02, AUD-04, AUD-05, AUD-06, AI-06
SOURCE ROLE	Audit evidence, controls, AI audit practice
USE PLACEMENT	Main text and Table 4
CITATION RISK	Evidence request as formal audit procedure
REVIEW NOTE	Mark as readiness/evidence architecture.

8 Lifecycle Walkthrough

SOURCE IDS USED	AUD-04, AUD-05, AUD-06, EVID-01, EVID-02, GAIC-SOURCE
SOURCE ROLE	Controls, audit practice, provenance, traces
USE PLACEMENT	Main text and Table 5
CITATION RISK	Formal audit procedure implication
REVIEW NOTE	Keep walkthrough illustrative.

9 Evidence Partitioning

SOURCE IDS USED	EVID-01, EVID-02, EVID-03, PRIV-01, PRIV-02, PRIV-03, AI-09, GAIC-SOURCE
SOURCE ROLE	Provenance/logging/privacy/context partitioning
USE PLACEMENT	Main text and Table 6
CITATION RISK	Universal disclosure or retention
REVIEW NOTE	Stress scope, access, and minimization.

10 Privacy and Minimization

SOURCE IDS USED	PRIV-01, PRIV-02, PRIV-03, PRIV-04, PRIV-05, AI-09, GAIC-SOURCE
SOURCE ROLE	Privacy and retention context
USE PLACEMENT	Main text and Table 7
CITATION RISK	Legal advice or GDPR proof
REVIEW NOTE	Require legal/privacy review in R3.

11 Exception and Closure

SOURCE IDS USED	EVID-04, AI-02, AUD-04, AUD-05, AUD-06, GAIC-SOURCE
SOURCE ROLE	Incident, control, remediation, audit practice
USE PLACEMENT	Main text and Table 10
CITATION RISK	Legal remedy/settlement overclaim
REVIEW NOTE	Closure is evidence state only.

12 Third-Party Validation Boundary

SOURCE IDS USED	AUD-03, BOUND-01, BOUND-02, BOUND-03, GAIC-SOURCE
SOURCE ROLE	Assurance, attestation, conformity, validation boundary
USE PLACEMENT	Main text and Table 8
CITATION RISK	Certification or assurance opinion overclaim
REVIEW NOTE	Review every Validation Lab paragraph.

13 AARM

SOURCE IDS USED	AUD-01, AUD-05, AUD-06, AI-01, BOUND-03, GAIC-SOURCE
SOURCE ROLE	Evidence, AI audit practice, governance, attestation boundary, AARM source truth
USE PLACEMENT	Main text and Table 9
CITATION RISK	Maturity/certification overclaim
REVIEW NOTE	Keep "what it does not prove" column.

14 Enterprise Readiness

SOURCE IDS USED	BF-02, BF-04, AUD-04, AUD-07, AI-01, PRIV-03, GAIC-SOURCE
SOURCE ROLE	Enterprise governance, controls, three lines, privacy, market context
USE PLACEMENT	Main text
CITATION RISK	Procurement/readiness guarantee
REVIEW NOTE	Bridge only to Guide 1/Guide 2.

15 Audit and Assurance Firm Use

SOURCE IDS USED	BF-02, BF-03, AUD-03, AUD-05, AUD-06, BOUND-03
SOURCE ROLE	Market context, professional practice, assurance/attestation boundary
USE PLACEMENT	Main text
CITATION RISK	Big Four endorsement or methodology replacement
REVIEW NOTE	Use "may use as discussion/readiness structure."

16 Conclusion

SOURCE IDS USED	GAIC-SOURCE, AUD-01, AUD-03, AI-01
SOURCE ROLE	Object chain, audit evidence, assurance boundary, governance context
USE PLACEMENT	Main text
CITATION RISK	Insurability overclaim before the insurability and risk transfer white paper
REVIEW NOTE	Keep the bridge to the insurability and risk transfer white paper high level.

Appendix Citation Map

Appendix	Source IDs used	Source role	Citation risk	Review note
A Evidence Request List	AUD-01, AUD-02, AUD-04, AUD-05, AUD-06, GAIC-SOURCE	Evidence/control vocabulary and MRO mapping	Formal audit-procedure implication	Keep request list as readiness architecture.
B Walkthrough Template	AUD-04, AUD-05, AUD-06, EVID-01, EVID-02, GAIC-SOURCE	Controls, provenance, traces, MRO lifecycle	Audit procedure implication	Keep illustrative.
C MRO Mapping	GAIC-SOURCE, AUD-01, AUD-05, AI-01	GAIC MRO source truth and audit/control language	Legal mandate overclaim	Verify MRO names exactly.
D Readiness Matrix	AUD-01, AUD-05, AUD-06, AI-01, BOUND-03, GAIC-SOURCE	Evidence, audit practice, governance, boundary	Certification/maturity-score overclaim	Keep no scores.
E Closure Checklist	EVID-04, AI-02, AUD-04, AUD-05, AUD-06, GAIC-SOURCE	Incident, control, remediation, MRO closure	Legal closure overclaim	Closure means evidence state only.
F Boundary Language	AUD-03, BOUND-01, BOUND-02, BOUND-03, GAIC-SOURCE	Assurance, conformity, attestation, Validation Lab and companion-paper boundary	Boundary wording becoming legal disclaimer overreach	Keep as editorial guardrail.

Citation Map Result

PASS - this public research edition uses source IDs and GAIC source truth with explicit source-role and citation-risk notes.