

AIEDSM: AI EVIDENCE DATA SCIENCE METHODOLOGY V1.0

GOVERNING MACHINE-READABLE EVIDENCE INTELLIGENCE
AND TRACEABLE INSTITUTIONAL ANALYSIS

MWP04-AIEDSM • METHODOLOGY IV • 14 SEPTEMBER 2026

Publisher: EMJ LIFE HOLDINGS PTE. LTD. (Singapore)

DOI: [10.64969/emj.mwp04.aiedsm.2026.v1.0](https://doi.org/10.64969/emj.mwp04.aiedsm.2026.v1.0)

EMJ.LIFE
PRE DISCLOSURE INFRASTRUCTURE

Author: Anderson Yu

Founder & Chief Executive Officer
EMJ LIFE HOLDINGS PTE. LTD.

ORCID: [0009-0002-2161-5808](https://orcid.org/0009-0002-2161-5808)

METADATA PAGE

TITLE

AIEDSM: AI Evidence Data Science Methodology

Governing Machine-Readable Evidence Intelligence and Traceable Institutional Analysis

PUBLISHER

EMJ LIFE HOLDINGS PTE. LTD. (Singapore)

PUBLICATION SERIES

EMJ.LIFE Methodology White Paper Series

Methodology White Paper No. 04 (MWP04-AIEDSM)

METHODOLOGY OPERATOR

AI Evidence Data Science Methodology (AIEDSM)

Machine Evidence Intelligence Methodology

VERSION

V1.0 • 14 September 2026

PUBLICATION STATUS

Publication

IDENTIFIERS

DOI: [10.64969/emj.mwp04.aiedsm.2026.v1.0](https://doi.org/10.64969/emj.mwp04.aiedsm.2026.v1.0)

Canonical publication record: <https://www.emj.life/publication-aiedsm>

ORCID (Author): [0009-0002-2161-5808](https://orcid.org/0009-0002-2161-5808)

CORRESPONDING AUTHOR

Anderson Yu

Founder & Chief Executive Officer

EMJ LIFE HOLDINGS PTE. LTD.

Email: anderson@emj.life

ORCID: [0009-0002-2161-5808](https://orcid.org/0009-0002-2161-5808)

COPYRIGHT & LICENSE

© 2026 EMJ LIFE HOLDINGS PTE. LTD.

[Research Program: Evidence Infrastructure Research Series](#)

Released under the Creative Commons Attribution-NoDerivatives 4.0 International License (CC BY-ND 4.0)

<https://creativecommons.org/licenses/by-nd/4.0/>

PLACE OF PUBLICATION

Singapore

METHODOLOGY POSITION

Methodology IV • MWP04-AIEDSM

AIEDSM is an independently developed methodology governing how machine systems analyse ISA-defined Evidence Objects while preserving institutional context, authority boundaries, temporal applicability and analytical traceability.

AIEDSM does not redefine Evidence Object identity, establish external standard requirements, or transfer institutional authority to a machine-generated output.

Within the EMJ.LIFE Evidence Infrastructure research architecture, MWP01-PADV is the sole methodology route for forming participation-derived records eligible for MWP03-ISA Evidence Object admission. MWP02-NTCC is a conditional non-tradable recognition output of PADV, not an intervening process stage. AIEDSM functions as the Machine Evidence Intelligence layer after ISA governance.

METHODOLOGY OBJECTIVES

1. Preserve evidence identity, provenance, version, status, time, rights and authority during machine analysis.
2. Define context-preserving retrieval and relationship analysis for ISA-defined Evidence Objects.
3. Separate Directly-Supported claims, Deterministically-Derived results, Machine-Inferred propositions and unresolved claims through the multi-axis claim model.
4. Expose unresolved evidence divergence, conflict and qualified coverage gaps.
5. Produce traceable analytical outputs without conferring institutional decision authority on AI systems.

GUIDING PRINCIPLE

Machine Analysis. Preserved Evidence Meaning.

CORE METHODOLOGY

MWP01-PADV Evidence Formation; MWP03-ISA Evidence Object Governance; MWP04-AIEDSM Machine Analysis; Candidate TIIR; Institutional Review; Institutionally Accountable Decision. Conditional PADV output: MWP02-NTCC, which is not a process stage.

KEYWORDS

AI Evidence Data Science Methodology • AIEDSM • Machine Evidence Intelligence • Evidence Object • Evidence Context Envelope • Evidence Relationship Analysis • Institutional Equivalence • Temporal Evidence Science • Evidence-to-Inference Boundary • Traceable Institutional Intelligence

INDEPENDENT METHODOLOGY DEVELOPMENT

AIEDSM has been developed through independent research into evidence infrastructure, institutional information architecture, machine-mediated retrieval, provenance, version governance, knowledge representation and AI-assisted institutional analysis.

References to external standards, specifications, frameworks and technical implementations are used for comparative, descriptive or implementation-oriented analysis. Such references do not constitute endorsement, delegated authority, formal interpretation or a claim of conformance.

DISCLAIMER

This publication presents a methodology for context-preserving machine analysis of institutional evidence. It does not provide legal, audit, assurance, certification, regulatory, financial or investment advice.

AIEDSM does not establish factual truth, eliminate model error, determine compliance, determine materiality, issue assurance conclusions, certify evidence, replace professional judgement or create authority that is absent from the underlying Evidence Objects.

CONTENTS

Publication architecture and principal analytical components

DEFINITION STATEMENT	5
VALUE STATEMENT	6
ABSTRACT	6
PREFACE	9
INSTITUTIONAL BOUNDARY STATEMENT	11
EVIDENCE SOVEREIGNTY STATEMENT	12
CHAPTER 1: WHY AI EVIDENCE DATA SCIENCE METHODOLOGY?	15
CHAPTER 2: FROM DATA SCIENCE TO EVIDENCE DATA SCIENCE	21
CHAPTER 3: METHODOLOGY SCOPE, BOUNDARY AND PRINCIPLES	26
CHAPTER 4: EVIDENCE OBJECT AND EVIDENCE CONTEXT ARCHITECTURE	34
CHAPTER 5: EVIDENCE RELATIONSHIP ANALYSIS	44
CHAPTER 6: SEMANTIC SIMILARITY AND INSTITUTIONAL EQUIVALENCE	53
CHAPTER 7: TEMPORAL EVIDENCE SCIENCE	60
CHAPTER 8: BOUNDED MACHINE ANALYSIS AND CLAIM-LEVEL ANALYTICAL LINEAGE	68
CHAPTER 9: AIEDSM REFERENCE ARCHITECTURE AND IMPLEMENTATION BOUNDARY	78
CHAPTER 10: TRACEABLE INSTITUTIONAL INTELLIGENCE AND INSTITUTIONAL ACCOUNTABILITY	89
APPENDIX A: CONTROLLED TERMINOLOGY AND FOUNDATIONAL PROPOSITIONS	99
APPENDIX B: INSTITUTIONAL BOUNDARY AND EVIDENCE SOVEREIGNTY STATEMENT	105
APPENDIX C: EVIDENCE CONTEXT ENVELOPE MINIMUM FIELD FRAMEWORK	107
APPENDIX D: TAXONOMY, STATE, LIFECYCLE AND TESTING FRAMEWORK	115
APPENDIX E: REFERENCES	126

DEFINITION STATEMENT

AI Evidence Data Science Methodology, abbreviated as AIEDSM, is a proposed EMJ.LIFE methodology for governing how machine systems retrieve, contextualise, relate, compare, transform and analyse institutional evidence while preserving the conditions that make that evidence institutionally meaningful.

AIEDSM operates on Evidence Objects admitted and governed through MWP03-ISA: Institutional Standards Architecture. Within this methodology architecture, MWP01-PADV is the sole formation route for participation-derived records entering ISA consideration. AIEDSM does not redefine the Evidence Object, create an alternative admission route, or reduce the identity, provenance, classification, version, relationship or authority requirements established through ISA.

The methodological concern of AIEDSM begins when an ISA-governed Evidence Object enters a machine-mediated analytical environment. That environment may include retrieval systems, knowledge graphs, analytical models, generative systems, AI agents, institutional research tools or other computational processes. AIEDSM specifies the controls that should accompany the Evidence Object through selection, retrieval, contextualisation, comparison, transformation, inference and analytical output.

Its central research question is:

Under what conditions can AI systems analyse ISA-defined Evidence Objects while preserving claim-level traceability, temporal applicability, institutional authority and the boundary between source support and machine inference?

AIEDSM treats the preservation of institutional context as a condition of responsible machine analysis. Identity alone is insufficient. Provenance alone is insufficient. Semantic similarity alone is insufficient. An analytical system must also distinguish source type, version, status, effective period, jurisdiction, authority, declared relationships, machine-proposed relationships, permitted use and the epistemic status of each output proposition.

AIEDSM therefore defines an institutional analysis lifecycle rather than a model architecture. It is compatible with multiple technical approaches, but it does not prescribe a particular model, database, graph technology, retrieval method, software development kit, agent protocol or platform implementation.

The methodological outcome is Traceable Institutional Intelligence: machine-assisted analysis that remains inspectably connected to the Evidence Objects, sources, versions, temporal conditions, relationship assertions, analytical transformations, conflicts, gaps and limitations on which the analysis depends.

Traceable Institutional Intelligence supports an Institutionally Accountable Decision. It does not replace that decision and does not acquire the authority of the evidence from which it was generated.

VALUE STATEMENT

Institutional evidence is increasingly read, selected, summarised and compared by machines before it is reviewed by a human decision-maker. The analytical capacity created by this transition is substantial. The governance risk is equally substantial.

An AI system can retrieve an accurate sentence while selecting an obsolete version. It can identify strong semantic similarity while collapsing different jurisdictions, document types or levels of authority. It can produce a plausible synthesis while obscuring which statements were directly supported, which were deterministically calculated, and which were inferred by the model. It can retain a source link while losing the institutional conditions under which the source is applicable.

AIEDSM addresses this problem by making institutional context an active part of machine analysis rather than a passive annotation attached to a document.

Its value lies in six connected controls:

1. preserving the identity and provenance of each Evidence Object used;
2. binding each analytical use to an Evidence Context Envelope;
3. preserving version, status, temporal applicability and jurisdiction;
4. distinguishing semantic similarity from institutional equivalence;
5. classifying claim basis, resolution state, validation status and applicable temporal state independently;
6. recording the analytical basis of each output in a traceable and reviewable form.

These controls do not guarantee factual correctness, prevent every model error or determine institutional truth. They make the evidentiary basis, transformation history and inference boundary of machine-assisted analysis more inspectable.

AIEDSM converts structurally interoperable Evidence Objects into machine-readable analytical capacity without converting machine-generated outputs into institutional authority.

ABSTRACT

Artificial intelligence increasingly operates across institutional evidence that differs in source, authority, document type, version, effective period, jurisdiction, relationship and governance context. This evidence may be technically accessible and semantically intelligible while remaining institutionally ambiguous to a machine system.

The relevant problem is not simply whether an AI system can find information. It is whether the system can analyse that information without severing it from the institutional conditions that determine what the information is, where it came from, when it applies, what it relates to, what authority it carries and what uses are permitted.

AIEDSM responds to this machine interpretation problem. It establishes a proposed methodology for the context-preserving machine analysis of Evidence Objects governed through MWP03-ISA. It treats the Evidence Object as the institutional analytical unit and introduces the Evidence Context Envelope as a task-bound control structure for carrying the relevant institutional conditions into a specified analytical run.

The Evidence Context Envelope does not replace the Evidence Object or create a competing source record. It binds an analytical task to the selected Evidence Object identity, source, version, status, time, jurisdiction, authority, rights, relationship state, selection rationale and processing event. It enables contextual conditions to remain available when evidence crosses from a publication or repository environment into retrieval, graph, model, agent or output environments.

AIEDSM further distinguishes semantic similarity from institutional equivalence. Two sources may discuss the same subject and contain similar language while performing different institutional functions. A regulatory requirement, an entity disclosure, a professional interpretation, an academic paper and an editorial analysis may be semantically related without being institutionally equivalent. A machine may identify a candidate similarity. It may not treat that similarity as institutional equivalence without an explicit and governed basis.

The methodology implements the Evidence-to-Inference Boundary through a multi-axis claim model. A Directly-Supported claim, Deterministically-Derived result and Machine-Inferred proposition retain distinct claim bases; each also carries an independent resolution state, validation status and applicable temporal state. A proposition advanced for investigation is represented as UNRESOLVED, with an optional investigative intended use, rather than as a separate evidentiary basis. Model confidence does not remove these distinctions or transfer source authority.

Temporal analysis is treated as a necessary part of evidence analysis. Publication date, effective date, reporting period, version state, supersession relationship, analytical cut-off time and retrieval time may each affect the responsible use of an Evidence Object. AIEDSM therefore requires analytical outputs to preserve the temporal conditions relevant to their evidentiary basis rather than treating the most recently retrieved fragment as automatically current or applicable.

Conflict and gap analysis are similarly bounded. When sources disagree, the machine should preserve each source, identify the type of conflict and avoid silently transferring resolution authority to the model. When expected evidence is missing, the machine may identify a qualified gap only against a declared

Expected Evidence Profile. Failure to retrieve an item is not, by itself, proof that the evidence does not exist or that an institution has failed to comply with a requirement.

The principal output of AIEDSM is Traceable Institutional Intelligence. Such an output records its analytical task, cut-off date, Evidence Object identities and versions, selection and exclusion basis, deterministic transformations, claim-level support, epistemic status, conflicts, gaps, limitations and review state. This record provides analytical lineage. It does not disclose or require a model's private chain of thought, and it does not guarantee that the resulting analysis is correct.

AIEDSM is implementation-neutral. It may be applied through established identifiers, metadata systems, provenance mechanisms, graph technologies, retrieval models, AI agents or other computational infrastructure where appropriate. It does not introduce a new persistent identifier, provenance ontology, catalogue vocabulary, content-authenticity mechanism, knowledge-graph technology or retrieval model. Its contribution lies in integrating relevant capabilities into a controlled institutional analysis lifecycle.

Within the EMJ.LIFE methodology architecture, MWP01-PADV is the sole methodology route through which eligible participation-derived activity is structured for admission as an MWP03-ISA-governed Evidence Object. MWP02-NTCC is not a process stage in that route: where separate NTCC conditions are met, it is a governed, non-tradable recognition output of PADV. MWP03-ISA governs the resultant Evidence Object, and MWP04-AIEDSM governs its machine analysis. NTCC does not create, replace or bypass ISA governance.

The governing doctrine of AIEDSM is:

AI may expand analytical capacity, but it must not erase the institutional conditions that make evidence meaningful.

PREFACE




Figure 1. Evidence Infrastructure method relationship and the AIEDSM analytical layer

Institutional evidence has traditionally been published for a human reader. A reader could recognise a cover page, distinguish a standard from an interpretation, identify a draft, notice a superseded edition, understand a citation hierarchy and infer the significance of an issuing institution. Much of this understanding depended on context that was visible around the content but was not encoded in the content itself.

Machine-mediated reading changes this condition.

An AI system may encounter institutional evidence as a search result, an extracted paragraph, a graph node, a vector representation, an application programming interface response, an agent message or a fragment within a retrieval context. The surrounding publication environment may not travel with the fragment. The machine may receive the words while losing the document identity, version status, institutional relationship, applicable jurisdiction or source authority that gave those words meaning.

The result is not necessarily false information. It may be correct information in an institutionally incomplete state.

This distinction matters. A statement can be textually accurate but temporally obsolete. It can be faithfully quoted but detached from the status of the source. It can be semantically similar to another

statement but institutionally non-equivalent. It can be derived from authoritative evidence without itself carrying the authority of that evidence.

MWP03-ISA addresses the structural interoperability of Evidence Objects across institutional systems and external frameworks. It provides the preceding methodological foundation for evidence identity, classification, provenance, version and relationship governance. AIEDSM begins from that governed object and asks a different question: what must remain true when a machine system uses the object for analysis?

The answer cannot be reduced to retrieval accuracy. Retrieval is only one event in a longer analytical lifecycle. Selection, exclusion, transformation, comparison, temporal interpretation, relationship construction, inference, summarisation and output presentation can each alter how evidence is understood.

AIEDSM therefore focuses on context survival across the full analytical lifecycle.

The conceptual problem was articulated in Evidence Infrastructure Analysis 011, *When the Reader Is No Longer Human*. That analysis identified Evidence Ambiguity and the Retrieval Boundary as emerging problems for institutional information. AIEDSM develops a methodological response. The analysis is a conceptual precursor, not a normative dependency. AIEDSM must remain understandable and usable without prior access to that publication.

Knowledge graphs, machine-readable publication environments, agent-to-agent exchanges and Model Context Protocol implementations may provide useful application settings for AIEDSM. They do not define the methodology. A reference implementation can demonstrate how a control is operationalised, but it cannot expand the authority of AIEDSM, establish external adoption or convert a platform design choice into a methodological requirement.

The same boundary applies to the relationship between AIEDSM and EMJ.NEXUS or Sustainability News Network environments. These environments may support later demonstrations or applications. AIEDSM remains platform-neutral and should be evaluated through the clarity of its concepts, controls, analytical records and stated limitations.

AIEDSM V1.0 uses one controlled methodological hierarchy: one Core Doctrine, five Operational Pillars, twelve Foundational Propositions and thirteen Control Principles.

Operational Pillar 1: Context Survival. Institutional context required for responsible use remains resolvable through discovery, selection, transformation, analysis, output, review and revalidation.

Operational Pillar 2: Authority Non-Transfer. Machine processing must not create, increase, confer or transfer institutional authority.

Operational Pillar 3: Institutional Equivalence Gate. Similarity, graph proximity and analytical convenience may identify a candidate correspondence but cannot establish equivalence or authorised substitution.

Operational Pillar 4: Evidence-to-Inference Separation. Each material analytical claim retains separate machine-readable claim-basis, resolution, validation and temporal states.

Operational Pillar 5: Institutionally Accountable Decision. Interpretation, reliance and consequential decision authority remain attributable to an authorised human or institutional function.

The Preface summarises that hierarchy; Appendix A is the authoritative consolidated statement of the Foundational Propositions. The hierarchy limits the methodology's claims: AIEDSM does not present AI as an institutional decision-maker, treat metadata as a cure for ambiguity, treat provenance as truth, or claim that traceability prevents hallucination.

The intended progression is therefore:



This progression defines the role of AIEDSM within Evidence Infrastructure. It extends analytical capacity while retaining institutional control.

INSTITUTIONAL BOUNDARY STATEMENT

AIEDSM is a proposed methodology developed and published by EMJ.LIFE. It is not a law, regulation, auditing framework, assurance engagement, certification programme, accreditation scheme, reporting requirement or professional opinion.

Reference to an external institution, framework, technology or publication does not imply endorsement, participation, adoption or approval by that institution. AIEDSM should not be represented as the official interpretation of an external standard, regulatory instrument or institutional publication.

AIEDSM does not determine:

1. legal or regulatory compliance;
2. whether evidence is legally admissible;

3. whether an institution has satisfied a disclosure requirement;
4. materiality under any reporting or decision framework;
5. audit or assurance conclusions;
6. the substantive truth of a source statement;
7. the legal effect of a document, version or effective date;
8. causal relationships merely from observed correlation;
9. capital allocation, credit approval or other institutional decisions;
10. the validity of professional judgement.

AIEDSM does not certify an Evidence Object. Admission and governance of Evidence Objects remain subject to ISA and any applicable source, organisational or legal controls.

AIEDSM does not convert an AI output into primary evidence. A summary, comparison, classification, inferred relationship, Candidate TIIR or other machine-generated proposition remains an Analytical Output. Within this methodology architecture, only the PADV-to-ISA route can form and admit an Evidence Object; source linkage, Institutional Review or publication of an output does not create a second admission route.

AIEDSM does not require access to a model's private chain of thought. Traceability is established through an Analytical Basis Record consisting of inspectable inputs, selection rules, transformations, support links, epistemic states, conflicts, gaps, tool or model versions, limitations and review status.

AIEDSM does not prescribe a technical product architecture. It may be implemented with different technologies where they can satisfy the applicable methodological controls. Use of a technology does not establish that an implementation satisfies AIEDSM, and this white paper does not make claims of external verification or implementation conformity.

AIEDSM supports institutional analysis. Responsibility for interpretation, review, use and decision remains with the authorised human or institutional process.

EVIDENCE SOVEREIGNTY STATEMENT

Evidence sovereignty means that machine access, processing or transformation does not transfer the identity, authority, ownership, rights, institutional role or decision power associated with underlying evidence.

An Evidence Object remains attributable to its identified source and governed by its applicable rights, version, status and institutional context. An AI system that retrieves or analyses the object does not become its issuer. An analytical output based on the object does not inherit the issuer's authority.

AIEDSM preserves four separations:

[Research Program: Evidence Infrastructure Research Series](#)

Persistent identity is not content integrity.

Content integrity is not source authority.

Source authority is not substantive truth.

Traceability does not establish correctness.

These separations prevent identity, provenance, authenticity, authority and truth from being collapsed into a single trust measure.

Evidence sovereignty also requires respect for rights and permitted use. Machine readability does not imply unrestricted permission to copy, transform, train on, redistribute or publicly disclose content. The Evidence Context Envelope must therefore preserve relevant rights, access restrictions and use conditions where they are available and applicable.

Where an Evidence Object is transformed, summarised or represented through a derived analytical structure, the transformation must remain distinguishable from the source. The source should remain canonically resolvable where permitted, and the analytical record should identify the transformation performed.

No machine-generated relationship, confidence score, summary or inference may silently overwrite a source-declared status or relationship. Machine proposals must remain distinguishable from source-declared, institutionally mapped, deterministically derived or human-proposed assertions and their independent validation statuses.

Evidence sovereignty is preserved when the conditions governing the evidence remain inspectable throughout the machine analytical lifecycle and when final institutional authority remains outside the model.

PART I: METHODOLOGY FOUNDATION

CHAPTER 1: WHY AI EVIDENCE DATA SCIENCE METHODOLOGY?

1.1 AI AS AN INSTITUTIONAL EVIDENCE READER

AI systems are becoming readers of institutional evidence before many human users encounter the underlying publication. Search engines generate answers from selected passages. Enterprise tools retrieve clauses and compare disclosures. Research systems cluster publications. AI agents assemble summaries, identify relationships and propose actions across multiple repositories.

The machine is therefore not merely a tool used after a human has understood the evidence. It increasingly participates in the conditions through which evidence is discovered, selected and framed.

This change affects institutional meaning.

A human reader may recognise that a document is an exposure draft rather than a final requirement. The reader may distinguish an issuer's statement from a third-party commentary. The reader may notice that a frequently cited paragraph belongs to a superseded edition. These distinctions are often carried by document architecture, publication history, institutional familiarity or surrounding context.

A machine system may receive only the selected text and a minimal source reference. If the relevant context is not represented and carried into the analytical task, the system can produce a fluent output that collapses those distinctions.

AIEDSM treats the machine as a context-dependent institutional evidence reader. The objective is not to make the machine an institution. The objective is to ensure that machine analysis remains bounded by inspectable institutional conditions.

The relevant analytical unit is not the isolated token, sentence, row or document fragment. For institutional analysis under AIEDSM, the analytical unit is the ISA-governed Evidence Object together with the task-bound context required for its responsible use.

1.2 FROM DATA ABUNDANCE TO EVIDENCE AMBIGUITY

Many AI deployments begin from an assumption of abundance. More documents, larger retrieval indexes and broader model context are expected to improve analytical quality. Abundance can improve coverage, but it can also increase institutional ambiguity.

Multiple versions may appear together. Draft and final publications may use similar wording. A regulator, reporting entity, professional adviser, academic researcher and journalist may all

discuss the same concept. Their statements may be related while carrying different authority, purpose and evidentiary roles.

Evidence Ambiguity arises when content remains available but the machine cannot reliably distinguish the institutional conditions under which that content should be interpreted.

Typical sources of Evidence Ambiguity include:

1. missing or unstable identity;
2. incomplete provenance;
3. unclear document type;
4. hidden version and supersession relationships;
5. absent effective periods or analytical cut-off dates;
6. loss of jurisdiction and scope;
7. untyped relationships between sources and claims;
8. conflation of source authority with model confidence;
9. machine-proposed relationships presented as source-declared relationships;
10. transformed or editorial content presented as primary evidence.

Evidence Ambiguity does not require falsification. It can arise from accurate content used in the wrong temporal, institutional or analytical context.

The problem is therefore not solved by increasing retrieval volume. A system can retrieve more evidence and become less institutionally precise if the additional evidence arrives without governed context.

1.3 THE MACHINE INTERPRETATION PROBLEM

The Machine Interpretation Problem occurs when a computational system processes institutional evidence without preserving the conditions necessary to distinguish what the evidence is, how it relates to other evidence and what analytical claims it can support.

The problem appears across the analytical lifecycle.

During discovery, a system may rank a commentary above the primary publication.

During selection, it may choose a semantically relevant but superseded version.

During contextualisation, it may omit jurisdiction, effective period or source type.

During relationship construction, it may infer equivalence from textual similarity.

During transformation, it may summarise away a limitation or exception.

During inference, it may present a plausible connection as a sourced fact.

During output, it may cite several sources without showing which source supports which proposition.

These failures are not identical. They require separate controls. AIEDSM therefore does not treat grounding as a single yes-or-no property. It examines identity, selection, provenance, context, relationship, temporal state, transformation and epistemic status as connected but distinct dimensions.

The Machine Interpretation Problem also has an authority dimension. An AI system may analyse an authoritative source, but its output does not become authoritative through proximity. Authority remains with the source and with the institution empowered to make the relevant decision.

1.4 RESEARCH QUESTION

AIEDSM is organised around the following research question:

Under what conditions can AI systems analyse ISA-defined Evidence Objects while preserving claim-level traceability, temporal applicability, institutional authority and the boundary between source support and machine inference?

This question contains several subordinate questions:

1. How should an analytical task select and identify the Evidence Objects it uses?
2. What institutional context must accompany each selected object?
3. How should source-declared, mapped, derived and machine-proposed relationships be distinguished?
4. How should a system treat semantic similarity where institutional equivalence has not been established?
5. How should version, effective period and analytical cut-off affect selection and output?
6. How should claim basis, resolution state, validation status and applicable temporal state remain independently inspectable?
7. How should conflicts be preserved without silently transferring resolution authority to the model?

8. Under what declared expectation may a missing item be characterised as an evidence gap?
9. What information must be recorded so that an authorised reviewer can inspect the analytical basis?

AIEDSM does not assume that one technical answer will satisfy every institutional setting. It establishes methodological conditions that can be implemented in different environments.

1.5 METHODOLOGY OBJECTIVE

The objective of AIEDSM is to establish a controlled institutional analysis lifecycle for machine use of ISA-governed Evidence Objects.

The lifecycle should enable a machine system to:

1. discover and retrieve Evidence Objects through resolvable identity;
2. select the version and status appropriate to a declared analytical task;
3. preserve provenance, authority, jurisdiction, rights and temporal conditions;
4. identify and type relationships without silently promoting similarity into equivalence;
5. separate claim basis, resolution state, validation status and temporal state;
6. expose conflicts and qualify gaps;
7. record transformations and analytical limitations;
8. produce a Traceable Institutional Intelligence Record;
9. support an Institutionally Accountable Decision.

The methodology does not aim to eliminate all uncertainty. Institutional evidence can be incomplete, contested, conditional or evolving. A responsible machine system should be able to preserve these states rather than conceal them behind a single answer.

The objective is therefore not maximum automation. It is controlled analytical expansion.

1.6 INTENDED METHODOLOGICAL OUTCOMES

AIEDSM is intended to produce a coherent set of governance concepts and implementable control structures.

The principal outcomes are:

1. **Evidence Object Admission Boundary**

AIEDSM accepts Evidence Objects governed through ISA. It does not establish a competing qualification route.

2. Evidence Context Envelope

A task-bound control object that carries selected institutional context into a defined analytical run.

3. Typed Relationship and Assertion-Mode Framework

A structure for distinguishing the meaning of a relationship from the manner in which that relationship was asserted.

4. Temporal Evidence Controls

A method for preserving publication, version, effective, supersession and analytical cut-off conditions.

5. Evidence-to-Inference State Model

A model that separates directly supported claims, deterministically derived results, machine-inferred propositions and hypotheses.

6. Conflict Preservation and Gap Qualification

Controls that expose unresolved conflicts and limit gap claims to declared Expected Evidence Profiles.

7. Traceable Institutional Intelligence Record

An inspectable record of evidence selection, transformation, support, inference, conflict, gap, limitation and review status.

8. Institutionally Accountable Decision Boundary

A governance condition preserving decision authority outside the machine system.

These outcomes should be testable through whether required context survives an analytical lifecycle, whether authority remains non-transferable, whether conflicts remain visible, and whether machine inference remains distinguishable from source support.

1.7 CORE DOCTRINE

The Core Doctrine of AIEDSM is:

AI may expand analytical capacity, but it must not erase the institutional conditions that make evidence meaningful.

The doctrine is operationalised through five Operational Pillars. These pillars are not a second set of Foundational Propositions; Appendix A contains the authoritative twelve propositions.

Operational Pillar 1 - Context Survival. The institutional context required for responsible use must remain resolvable through discovery, selection, transformation, analysis, output, review and revalidation.

Operational Pillar 2 - Authority Non-Transfer. Machine processing must not create, increase, confer or transfer institutional authority. An output may preserve, cite or characterise source authority; it does not inherit that authority.

Operational Pillar 3 - Institutional Equivalence Gate. Semantic similarity, graph proximity, shared terminology or analytical convenience may identify a candidate correspondence but cannot establish institutional equivalence or authorised substitutability.

Operational Pillar 4 - Evidence-to-Inference Separation. Each material analytical claim must retain a machine-readable classification of its claim basis, resolution state, validation status and applicable temporal state.

Operational Pillar 5 - Institutionally Accountable Decision. AIEDSM supports analysis. Interpretation, reliance and consequential decision authority remain attributable to an authorised human or institutional function.

CHAPTER 2: FROM DATA SCIENCE TO EVIDENCE DATA SCIENCE

2.1 DATA

Data consists of recorded values, symbols, observations or representations that can be stored, transmitted or processed. Data may be structured or unstructured, quantitative or qualitative, human-generated or machine-generated.

Within AIEDSM, data is not treated as institutionally self-explanatory. A value does not reveal by itself who produced it, why it was produced, which definition was applied, what period it represents, whether it has been revised or what claim it can support.

The same recorded value may perform different functions in different contexts. It may be an operational measurement, a disclosure input, a calculation output, a model feature or an item quoted in an analysis. Its existence as data does not establish its evidentiary role.

AIEDSM therefore adopts the proposition:

Data is not automatically evidence.

This proposition does not diminish the importance of data. It establishes that evidentiary use requires additional institutional conditions.

2.2 INFORMATION

Information is data organised or expressed so that it communicates meaning within a context. Labels, definitions, units, dates and narrative explanation can make data intelligible to a reader or system.

Information remains broader than evidence. A description may be informative without supporting a specified claim. A generated summary may communicate meaning without carrying the provenance, authority or temporal conditions of the underlying source.

Machine systems frequently transform data into information. They classify, summarise, translate and relate content. These transformations can improve accessibility while also introducing interpretation.

AIEDSM requires the transformation to remain visible. A machine summary should not be presented as though it were the source publication. A translation should remain linked to the source language. A classification should identify whether it was source-declared, institutionally

mapped, deterministically derived, machine-proposed or human-proposed; its validation status is recorded independently.

Information becomes analytically safer when its transformation history remains inspectable.

2.3 EVIDENCE

Evidence is data or information used to support, challenge, contextualise or assess a specified proposition, requirement, event or institutional question.

Evidentiary relevance is relational. A source is evidence for something. Its relevance depends on the analytical task, the proposition being considered and the conditions under which the source applies.

A publication may provide direct support for one claim, contextual support for another and no support for a third. A source citation alone does not establish the scope of support.

AIEDSM therefore requires claim-level attention. The analytical record should indicate which Evidence Object supports which output proposition and whether the relationship is direct, derived, inferred, contested or unresolved.

Evidence can be authoritative without being conclusive. Evidence can also be relevant while carrying limited authority. AIEDSM preserves these distinctions rather than reducing them to a single score.

2.4 INSTITUTIONAL EVIDENCE

Institutional evidence is evidence whose interpretation or use depends materially on institutional conditions such as issuer, mandate, document type, jurisdiction, version, effective period, governance status, relationship and permitted use.

Examples may include regulatory publications, standards-related materials, organisational records, consultation documents, responses, disclosures, verification records, controlled research outputs and other governed sources.

The term does not imply that every institutional publication is true, final or authoritative for every purpose. It indicates that institutional attributes are necessary to understand the evidentiary role of the source.

Institutional evidence often exists within a relationship network. A consultation notice relates to an exposure draft. A response relates to consultation questions. A final publication may replace a

draft. Guidance may interpret a requirement without replacing it. An entity disclosure may respond to a reporting requirement without becoming the requirement itself.

Machine analysis must preserve these relationships if it is to preserve institutional meaning.

2.5 EVIDENCE OBJECT

For AIEDSM, Evidence Object has the meaning assigned through MWP03-ISA. AIEDSM does not reproduce a competing definition.

The Evidence Object is inherited as the institutional analytical unit. Its identity, provenance, type, version, status, relationship and other governed attributes remain subject to ISA.

AIEDSM adds no lower threshold through which raw data, an extracted sentence, a model output or a source link automatically becomes an Evidence Object.

When an Evidence Object enters an analytical run, AIEDSM binds it to an Evidence Context Envelope. The envelope records the task-specific conditions relevant to the use of that object. It does not alter the source object and does not become a substitute for ISA governance.

An AIEDSM output is also not automatically an Evidence Object. If an analytical output is later intended to function as evidence, it must enter the applicable evidence governance process separately.

This inheritance boundary prevents circular elevation, in which AI analysis is treated as evidence merely because it analysed evidence.

2.6 EVIDENCE INTELLIGENCE

Evidence Intelligence is the capacity to discover, identify, contextualise, relate, compare and analyse Evidence Objects while preserving the basis on which each analytical proposition is made.

It differs from unrestricted information synthesis. Its value depends on the continued visibility of source, version, time, authority, relationship, transformation and inference status.

Evidence Intelligence includes capabilities such as:

1. locating canonical Evidence Objects;
2. selecting a task-relevant version;
3. distinguishing current, historical, proposed and superseded states;
4. comparing Evidence Objects without assuming equivalence;
5. identifying supporting, contradicting, dependent and superseding relationships;

6. detecting conflicts without silently resolving them;
7. detecting qualified gaps against a declared expectation;
8. generating claim-level traceable analysis;
9. preserving limitations and unresolved states.

Evidence Intelligence does not mean that the system knows institutional truth. It means that the system performs analysis within explicit evidentiary and institutional controls.

2.7 FROM PREDICTION TO TRACEABLE INSTITUTIONAL ANALYSIS

Conventional data science is often represented as a sequence from data to model to prediction. This representation is useful for some analytical tasks, but it is incomplete for institutional evidence.

AIEDSM adopts a different sequence:

```

ISA-governed Evidence Object
↓
Evidence Context Envelope
↓
Context-preserving retrieval and selection
↓
Relationship and temporal analysis
↓
Bounded transformation and inference
↓
Traceable Institutional Intelligence Record
↓
Institutionally Accountable Decision

```

This sequence does not prohibit prediction. It requires predictive or generative analysis to remain connected to the institutional basis from which it was produced.

The distinction is especially important when analysis crosses source types. A machine may compare a standard-related publication with an entity disclosure or an editorial analysis. The comparison may be useful, but it must not erase their different institutional roles.

Traceable institutional analysis therefore requires more than a citation list. It requires an analytical basis showing selection, exclusion, transformation, claim-level support, epistemic status, conflicts, gaps, limitations and review state.

[*Research Program: Evidence Infrastructure Research Series*](#)

This is the transition from general machine use of data to governed machine analysis of institutional evidence.

CHAPTER 3: METHODOLOGY SCOPE, BOUNDARY AND PRINCIPLES

3.1 METHODOLOGY SCOPE

AIEDSM applies to machine-assisted analytical tasks that use ISA-governed Evidence Objects in institutional contexts.

The methodology begins after Evidence Object governance has been established through ISA. It addresses how those objects are selected, contextualised, related, transformed, analysed and represented in machine-generated outputs.

The scope includes:

1. retrieval and selection of Evidence Objects;
2. identity and canonical resolution;
3. task-bound contextualisation;
4. provenance and transformation lineage;
5. version, status and temporal analysis;
6. relationship typing and assertion-mode control;
7. semantic comparison;
8. conflict identification and preservation;
9. qualified evidence gap detection;
10. evidence-to-inference separation;
11. claim-level analytical traceability;
12. AI agent use of institutional evidence;
13. human review and independent decision boundaries.

AIEDSM can be applied to sustainability, due diligence, procurement, supply-chain analysis, finance, regulation, academic research, public policy, AI governance, institutional media and other fields where the institutional conditions of evidence matter.

Application in a field does not grant AIEDSM authority within that field. Applicable law, professional requirements, contractual controls and source-specific rights continue to govern.

3.2 PERMITTED ANALYTICAL FUNCTIONS

Within its boundary, AIEDSM supports machine systems that perform controlled analytical functions.

Permitted functions may include:

1. retrieving Evidence Objects through persistent or canonical references;
2. comparing source content, classifications and relationships;
3. identifying patterns across multiple Evidence Objects;
4. analysing evidence density within a defined scope;
5. identifying missing evidence against a declared Expected Evidence Profile;
6. detecting potential conflicts and classifying conflict type;
7. analysing temporal relationships and supersession;
8. proposing candidate relationships for review;
9. generating summaries with claim-level support;
10. producing traceable research and analytical outputs;
11. supporting authorised human review.

Permission under the methodology does not override legal rights, licensing terms, confidentiality, access restrictions or organisational controls.

The analytical function must remain proportionate to the quality and scope of the available evidence. A system should not present an institution-wide conclusion when its evidence basis is limited to a small, unrepresentative or temporally incomplete set.

3.3 EXCLUDED INSTITUTIONAL FUNCTIONS

AIEDSM does not authorise a machine system to perform institutional functions that require legal mandate, professional responsibility, assurance authority or independent human judgement.

Excluded functions include:

1. determining legal compliance;
2. issuing audit or assurance conclusions;
3. certifying or validating evidence as true;
4. determining materiality;

5. issuing an official interpretation of a regulation or external standard;
6. replacing professional judgement;
7. determining truth merely because a model predicts an answer;
8. creating institutional authority that does not exist in the underlying evidence;
9. transforming correlation into causation without an adequate basis;
10. transforming editorial analysis into primary evidence;
11. silently resolving institutional conflicts;
12. making final capital allocation, procurement, credit or regulatory decisions.

AIEDSM may support research relevant to these functions. Support is not substitution.

Where a machine output is used within an institutional process, the responsible institution should define who may review, approve, reject, override or rely on the output.

3.4 AUTHORITY NON-TRANSFER PRINCIPLE

The Authority Non-Transfer Principle states:

Machine processing must not create, increase, confer or transfer institutional authority.

A machine output may preserve, cite or characterise the authority of an identified source, but it does not inherit that authority. Institutional Review may authorise use under an existing mandate; it does not turn a machine output into a source-issued statement, official interpretation, assurance conclusion or external determination.

The principle applies to retrieval, contextualisation, transformation, graph representation, citation, fluent expression, scoring, review and publication. Repetition, reviewer acceptance or proximity to authoritative evidence must not rewrite assertion origin, claim basis or source authority.

Relationship type, assertion origin, validation status and lifecycle status perform separate functions. They must remain independently inspectable and must not be collapsed into a single standing or confidence measure.

3.5 INSTITUTIONALLY ACCOUNTABLE DECISION

AIEDSM concludes with an Institutionally Accountable Decision rather than an automated institutional decision.

Independence in this context means that the authorised decision-maker can inspect the analytical basis, consider evidence and limitations, apply relevant judgement and reach a decision that is not treated as predetermined by the model output.

The decision-maker should be able to identify:

1. the analytical question asked;
2. the evidence set selected;
3. the versions and temporal cut-off applied;
4. the selection and exclusion basis;
5. the claims directly supported;
6. the results deterministically derived;
7. the propositions inferred by the machine;
8. the conflicts and gaps identified;
9. the limitations and unresolved questions;
10. the model, tool and review status relevant to the output.

Human review should not be reduced to a ceremonial approval. The reviewer must have the authority and sufficient information to challenge the analytical result.

AIEDSM does not prescribe who the decision-maker must be. That role depends on the applicable governance structure, professional duty and legal environment.

3.6 IMPLEMENTATION-NEUTRAL BOUNDARY

AIEDSM is defined through analytical controls rather than a prescribed technical stack.

An implementation may use document repositories, metadata catalogues, relational databases, graph databases, vector retrieval, rules engines, language models, AI agents, provenance services, content credentials or other technologies.

No single technology establishes the methodology.

A persistent identifier can support resolution but does not prove content integrity. A cryptographic record can support integrity verification but does not establish source authority or substantive truth. A knowledge graph can express relationships but does not determine whether those relationships are institutionally valid. A retrieval system can improve grounding but does not

eliminate model inference. A rights field can record a restriction but does not itself enforce every permitted use.

Implementation neutrality also applies to interface protocols. Agent-to-agent exchange, Model Context Protocol or application programming interfaces may transport Evidence Objects and analytical context. Transport does not create authority, validate evidence or satisfy the methodology unless the required controls remain operational across the exchange.

AIEDSM therefore evaluates whether institutional meaning survives the analytical lifecycle, not whether a particular technology has been adopted.

3.7 AIEDSM CORE PRINCIPLES

AIEDSM is governed by the following principles.

3.7.1 IDENTITY ANCHORED

Every Evidence Object used in analysis must remain linked to its ISA-governed identity and canonical reference, where a canonical reference lawfully exists.

3.7.2 PROVENANCE PRESERVED

Source origin and material transformation history must remain inspectable at object, field, relationship and claim level.

3.7.3 CONTEXT RETAINED

Task-relevant institutional conditions must accompany the Evidence Object through selection, transformation, analysis, output and review.

3.7.4 TEMPORAL STATE EXPLICIT

Source publication status, version state, effective state, knowledge boundary and analytical cut-off must not be collapsed into a single label of current.

3.7.5 RELATIONSHIPS TYPED

Each material relationship must record its controlled type, direction, scope, assertion origin, validation status, provenance, temporal conditions and lifecycle status.

3.7.6 INSTITUTIONAL EQUIVALENCE GOVERNED

Similarity and correspondence must remain distinct from authorised substitutability and institutional equivalence.

3.7.7 AUTHORITY NON-TRANSFER

Machine processing and Institutional Review must not create or transfer source authority.

3.7.8 INFERENCE BOUNDED

Claim basis, resolution state, validation status and temporal state must remain separately inspectable. A score must not overwrite any of these states.

3.7.9 CONFLICT PRESERVED

Material conflicts must remain visible with their sources, scopes, dates, jurisdictions, authority contexts and dispositions.

3.7.10 GAPS QUALIFIED

A gap statement must identify the governed expectation and the search, access, rights and time boundary against which the gap was determined.

3.7.11 OUTPUTS TRACEABLE

Each material analytical claim must resolve to its Evidence Objects, transformations, relationship assertions, methods, states and limitations.

3.7.12 RIGHTS ENFORCED

Machine readability does not imply permission. Unresolved or prohibited use conditions must prevent unauthorised processing or release.

3.7.13 DECISION AUTHORITY INDEPENDENT

AIEDSM outputs may inform, but must not predetermine or impersonate, the authorised institutional decision process.

These principles establish the minimum governance logic of AIEDSM. The chapters that follow translate them into context, relationship, temporal, inference, conflict, gap, agent and output architectures.

3.8 STATUS AND NORMATIVE LANGUAGE

AIEDSM is a methodology publication, not an external standard, certification scheme, assurance procedure or conformity-assessment regime. It specifies controls required to preserve the institutional meaning of evidence during machine analysis. A claim of implementation or conformance requires a separately declared implementation profile, scope and test record.

In this publication, must expresses a requirement necessary to preserve methodological integrity; must not expresses a corresponding prohibition; should expresses a recommended control from which an implementation may depart only with documented rationale and risk treatment; may expresses permission or an optional capability; and can expresses possibility or capability and is not normative. must is not used as a separate normative category.

Passing an AIEDSM control test means only that the specified records and controls were inspectable within the declared test scope. It does not establish certification, assurance, legal compliance, substantive correctness or external endorsement.

MATERIAL ANALYTICAL CLAIM

A Material Analytical Claim is an output proposition whose error, omission, ambiguity, qualification or use could reasonably affect the declared institutional purpose, an Institutional Review or an Institutionally Accountable Decision. A machine system may apply institutionally approved significance criteria or declared escalation thresholds. It must not independently determine legal, accounting, assurance, financial, sustainability or regulatory materiality.

PART II: EVIDENCE INTELLIGENCE ARCHITECTURE

CHAPTER 4: EVIDENCE OBJECT AND EVIDENCE CONTEXT ARCHITECTURE

4.1 THE ANALYTICAL UNIT INHERITED FROM ISA

AIEDSM begins only after a participation-derived record has passed through the MWP01-PADV formation route and has been admitted and governed as an Evidence Object under MWP03-ISA. The definition, admission and structural governance of the Evidence Object remain within ISA and are not repeated or altered by AIEDSM. Chunking, embedding, extraction, summarisation, review or publication of an Analytical Output does not create another Evidence Object formation route.

The Evidence Object remains the persistent referential anchor for machine analysis. Context may be held in controlled registries, resolvers or related records, but the object must resolve to the institutional conditions required for the declared analytical use.

CANONICAL MWP03 DEPENDENCY RECORD

Field	Canonical MWP03 information
Document code and series	MWP03; EMJ.LIFE Methodology White Paper Series; Evidence Infrastructure Research Series
Title	ISA: Institutional Standards Architecture
Subtitle	A Structural Interoperability Methodology for Reusable Evidence Ecosystems
Author	Anderson Yu; ORCID 0009-0002-2161-5808
Publisher and place	EMJ LIFE HOLDINGS PTE. LTD.; Singapore
Version, status, date and extent	V3.0; Version of Record; published 10 September 2026; 271 pages
DOI	https://doi.org/10.64969/padv.isa.2026.v3
Canonical landing page	https://www.emj.life/publication-isa
Licence	Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0)
Official PDF SHA-256	ea32924428d8c9339027e9f9876f24950703d1773a41eceff5dc02fb8ae2848

Canonical citation: Yu, Anderson. (2026). ISA: Institutional Standards Architecture: A Structural Interoperability Methodology for Reusable Evidence Ecosystems (MWP03, Version 3.0). EMJ LIFE HOLDINGS PTE. LTD. <https://doi.org/10.64969/padv.isa.2026.v3>

4.2 FORMAL DEFINITION OF THE EVIDENCE CONTEXT ENVELOPE

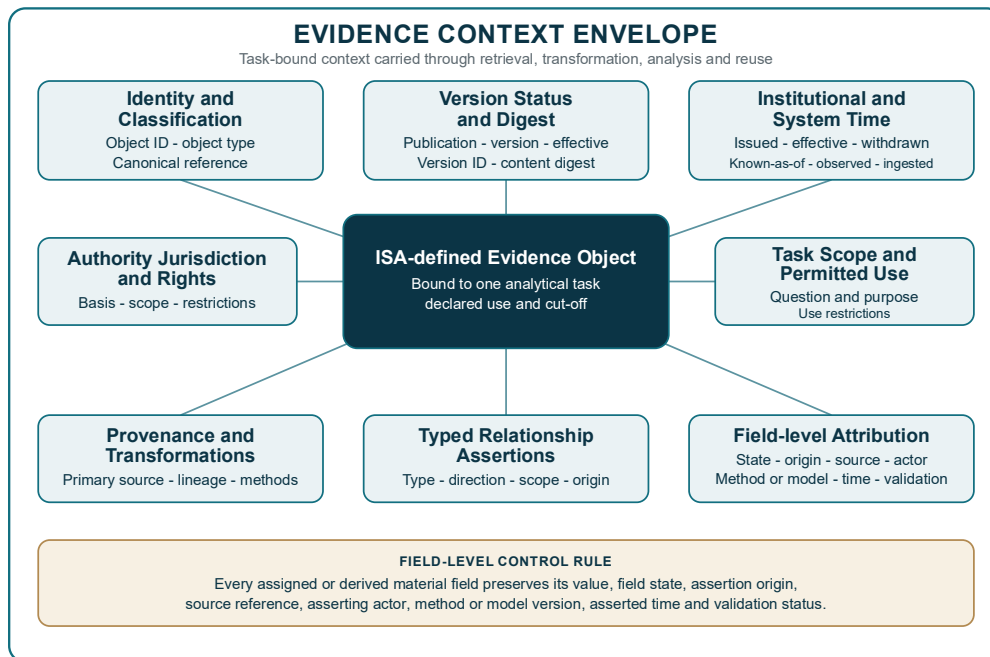


Figure 2. Evidence Context Envelope bound to an ISA-defined Evidence Object

Formal Definition

An Evidence Context Envelope is a versioned, machine-readable context layer logically bound to an Evidence Object that preserves and exposes the institutional conditions necessary for its interpretation, retrieval, comparison, transformation and reuse.

The Evidence Context Envelope, abbreviated as ECE, does not replace the Evidence Object. It makes the object's relevant institutional conditions available to a machine system. Those conditions include identity, provenance, version, status, effective time, authority, jurisdiction, relationships, rights and transformation history.

Logical binding does not require one storage arrangement. The ECE may be packaged with the object or resolved from a governed registry, provided that the association and applicable envelope version can be verified.

The ECE is not an embedding vector. It may reference machine representations, but semantic coordinates cannot independently express authority, status, provenance, effective period, rights or applicability and cannot substitute for the ECE.

Boundary Statement

An Evidence Context Envelope does not modify the underlying Evidence Object, cure absent provenance, establish factual truth, confer institutional authority or convert an analytical annotation into primary evidence.

Where context is unknown, disputed, unavailable, restricted or invalid, the ECE must expose the applicable field state. AIEDSM does not permit a machine system to fill an institutional field through unlabelled prediction. A predicted jurisdiction, inferred source type or estimated effective period may be stored only through a Field Assertion Record with immutable assertion origin, attribution basis, method version and validation status.

4.3 THE THREE CONTEXT LAYERS

An ECE separates context into three layers. This separation prevents downstream users from mistaking machine enrichment for source-originating information.

4.3.1 SOURCE-NATIVE CONTEXT

Source-native context is stated by, embedded in or authoritatively associated with the originating record. It includes attributes such as title, issuing body, publication date, declared status, identifier, version and effective date. Source-native describes origin, not absolute correctness. The ECE records both the value and its attribution basis.

4.3.2 INSTITUTIONALLY ASSIGNED CONTEXT

Institutionally assigned context is added through an authorised governance process external to the originating source, including internal classifications, access categories, approved mappings, retention periods and applicability determinations. It must identify the assigning authority, method, date, scope and status and must not be attributed to the originating publisher.

4.3.3 MACHINE-GENERATED CONTEXT

Machine-generated context is produced by extraction, classification, similarity analysis, entity recognition, graph inference, summarisation or another automated process. It must preserve the method, model or rule version, generation time, input version, assertion origin, validation status and Field Assertion Record. It cannot silently overwrite source-native or institutionally assigned context. Where values differ, the ECE records the conflict.

This separation establishes a foundational AIEDSM rule:

Machine enrichment may expand the usable context of an Evidence Object, but it must not obscure where that context came from.

4.4 MINIMUM ECE FIELD FRAMEWORK

The following field groups constitute the implementation-neutral ECE minimum. Implementations may use different schemas or storage arrangements, but must preserve each institutional distinction and its machine-resolvable attribution basis.

Field group	Minimum content	Governance purpose
Envelope control	ECE identity, schema version, creation time, run, task, controller and record status	Separates the task-bound control object from the Evidence Object
Object binding	Evidence Object identifier, version, binding method and representation observation	Establishes the exact ISA-governed object and representation used
Canonical resolution	Canonical reference or controlled registry identifier, locator, resolver state and last resolution time	Supports resolution without inventing a public URI
Source identity	Originating institution, publisher or responsible actor and source locus	Preserves attribution and institutional origin
Evidence classification	Object type, category, scheme and scheme version	Enables controlled interpretation without redefining the object
Field attribution	Field state and Field Assertion Record for each assigned, mapped, derived or proposed value	Separates values from their origin and review disposition
Provenance	Source chain, acquisition, responsible actors, derivations and event references	Makes origin and transformation inspectable
Publication, version and effectivity	Native label plus separate publication, version and effective states	Prevents status collapse
Temporal context	Event, creation, publication, availability, access, ingestion, effective, retrieval and transformation times	Supports bounded and reproducible analysis
Authority context	Authority type, basis, scope and limitations	Prevents relevance from becoming authority
Jurisdiction and applicability	Jurisdiction, declared scope, applicability state, conditions and exclusions	Preserves the conditions of use

Field group	Minimum content	Governance purpose
Subject and entity links	Referenced entity, activity, claim, requirement or reporting boundary	Connects the object to the analytical subject
Relationship references	AIEDSM-RV 1.0 assertion identifiers and versions	Supports graph reconstruction without embedding ungoverned edges
Rights and access	Ownership or licence, access class, permitted and prohibited uses, disclosure limits	Separates technical access from permission
Transformation history	Derived representations, methods, outputs and invalidation state	Preserves lineage of machine artefacts
Expected Evidence Profile	Applicable EEP identity, version, validation and authority, where used	Prevents a system from self-creating a governed gap basis
Revalidation	Trigger rule, next review, dependencies and change-impact state	Prevents stale outputs from silent reuse
Limitations	Unknown, unavailable, restricted, conflicting and excluded conditions	Prevents incomplete context from appearing complete

Every required or applicable Conditional field must carry a field_state. Controlled values are KNOWN, UNKNOWN, NOT_APPLICABLE, UNAVAILABLE, RESTRICTED, CONFLICTING and INVALID. A blank value is not a substitute for a field state.

FIELD ASSERTION RECORD

Each material ECE field assignment must resolve to a Field Assertion Record. Global object-level provenance does not replace field-level attribution where a value is assigned, normalised, mapped, derived or proposed.

Field	Minimum requirement
field_assertion_id	Persistent identity for the field assertion
ece_id and ece_version	Envelope identity and applicable version
field_path and value	Exact field path and value, omitted only where field_state precludes a value
field_state	One controlled state from the ECE field-state model
assertion_origin	SOURCE_DECLARED, INSTITUTIONALLY_MAPPED, DETERMINISTICALLY_DERIVED, MACHINE_PROPOSED or HUMAN_PROPOSED

Field	Minimum requirement
source_reference	Exact object, location, rule, event or attribution basis
asserting_actor_id	Identified source, institution, person, system, method or model
method_or_model	Identifier and version; required for mapped, derived or machine-proposed values
asserted_at and valid_time	Transaction timestamp and applicable valid interval
validation_status	Current review disposition, separate from assertion origin
review_event_id	Required when validation is not UNREVIEWED
confidence_or_score	Optional technical signal; never a substitute for origin, authority or state
supersedes_field_assertion_id	Required when a new assertion replaces an earlier assertion

4.5 CANONICAL RESOLUTION

Machine systems encounter institutional content through publisher pages, files, repositories, indexes, archives, quotations and caches. These are retrieval routes, not necessarily distinct Evidence Objects, and the initially discovered route does not become the object's identity.

Canonical resolution is the controlled process of resolving an encountered representation to the persistent Evidence Object identity, canonical reference or governed registry identifier, applicable version, exact representation and current resolution state.

evidence_object_id is required. canonical_reference is required when a resolvable canonical reference lawfully exists. An internal or restricted Evidence Object may instead resolve through a controlled registry identifier and access-governed locator. The absence of a public URI must not cause invention of one.

When a web page, API response, database view or other dynamic representation can change without a new public locator, each analytical use must create or resolve to a representation observation. The observation records its identifier, observed time, source-declared last-modified value where available, media type, retrieval context and content digest. Where rights permit, a snapshot or preservation locator should be retained. A stable URL is not evidence of stable content.

A content digest may assist integrity and duplicate detection, but it is not a complete institutional identity. Identical files may carry different authority or rights conditions, while corrected files may remain versions of one continuing object.

When a canonical source cannot be reached, a preserved representation may be used only where rights and the declared use permit. The unresolved state and its consequences must remain visible. Different versions must remain distinct even where most content is identical.

4.6 CONTEXT SURVIVAL CONDITIONS

Context survives when the conditions needed to interpret an Evidence Object remain available or durably resolvable after retrieval or transformation.

An Evidence Object satisfies the AIEDSM context survival conditions for a declared analytical use only when all material conditions below are met.

1. **Identity survival:** the retrieved or transformed item resolves to the persistent Evidence Object identity and relevant object version.
2. **Provenance survival:** the originating source, acquisition route and material transformations remain inspectable.
3. **Classification survival:** the system preserves the object's evidence type and distinguishes source classification from assigned or inferred classification.
4. **Version survival:** the machine can distinguish the analysed version from earlier, later, amended or superseding versions.
5. **Status survival:** draft, proposed, issued, effective, amended, superseded, withdrawn or other controlled states do not disappear during retrieval.
6. **Temporal survival:** material event, publication, effective, observation and retrieval times remain associated with the relevant proposition.
7. **Authority survival:** the identity and scope of the originating or assigning authority remain visible and are not transferred to a downstream summary.
8. **Relationship survival:** material links to claims, requirements, entities, versions and related objects remain resolvable.
9. **Rights survival:** access, reuse, disclosure and transformation restrictions remain attached to the object and its derivatives.

10. **Transformation survival:** every material machine representation can be traced to its input object version and declared generation process.
11. **Limitation survival:** known gaps, disputes, exclusions and unresolved fields remain visible after transformation.
12. **Claim-level survival:** a material analytical proposition can be traced to the specific Evidence Object or derived method that supports it, rather than only to a general document collection.

Context survival is purpose-sensitive. The system must declare the analytical use and test whether the available context is sufficient. An object may be discoverable but ineligible for a particular analysis.

Boundary Statement

Retrieval success is not context survival. A passage has not retained its institutional meaning merely because its words can be found, embedded or quoted.

4.7 DERIVED MACHINE REPRESENTATIONS

Chunks, embeddings, extracts, summaries, classifications, scores and graph assertions are Derived Machine Representations. They are not independent Evidence Objects merely because they are stored or cited.

Every material derivative must resolve to the exact parent Evidence Object version and source segment, retain rights and permitted-use conditions, and preserve its immutable `assertion_origin` independently from `validation_status`.

Minimum field	Required function
Representation identifier and type	Distinguish the derivative and identify whether it is a chunk, embedding, extraction, summary, classification, score, graph assertion or other output
Parent object, version and source locus	Preserve lineage to the exact input and page, paragraph, field, table, timestamp or structured element
Representation observation	Record observation identity, time, media type, retrieval context and digest for dynamic inputs
Generation method and version	Identify the rule, model, parser, algorithm or workflow

Minimum field	Required function
Generation time, parameters and scope	Distinguish outputs created under different system states and boundaries
assertion_origin	Preserve how the representation or annotation entered the environment
validation_status	Record review disposition without overwriting origin
Rights inheritance and prohibited uses	Keep source and derivative-specific restrictions operative
Invalidation and revalidation rule	Identify source, model, schema, rights, status or use changes requiring action

4.8 OPERATIONAL TESTS FOR EVIDENCE CONTEXT ARCHITECTURE

An AIEDSM workflow should apply the following tests before relying on an Evidence Object for material machine analysis.

Test	Operational question
Object Binding	Does the input resolve to a specific ISA-defined Evidence Object and version?
Source Separation	Can source-native, institutionally assigned and machine-generated context be distinguished?
Canonical Resolution	Are the persistent object, canonical record, retrieved representation and version relationship identifiable?
Context Portability	Do material identity, provenance, version, status, time, authority, relationship, rights and limitations survive transfer?
Derivative Lineage	Can every material derivative be traced to its parent version, source segment and method?
Invalidation	Do material source, status, rights or method changes trigger review or regeneration?
Unknown Context	Are unknown fields exposed and use restricted rather than silently predicted?
Rights Survival	Do access and reuse restrictions remain effective after transformation or agent transfer?
Claim Trace	Can each material proposition be traced to qualifying evidence or a declared method-derived result?

Failure of one test does not necessarily require deletion of the object. It may restrict the object to discovery, require remediation, trigger review or exclude the object from a declared analysis. The response depends on the purpose, materiality of the missing context and applicable governance.

4.9 CHAPTER BOUNDARY

The ECE makes context available but does not determine relationships, equivalence or applicability. AIEDSM therefore carries forward the following rule:

The Evidence Object remains the analytical anchor. Every machine representation, relationship and conclusion must remain subordinate to its resolvable identity and institutional context.

CHAPTER 5: EVIDENCE RELATIONSHIP ANALYSIS

5.1 FROM OBJECT RETRIEVAL TO RELATIONSHIP ANALYSIS

Institutional analysis rarely depends on one Evidence Object in isolation. Requirements, guidance, procedures, records, amendments and analytical outputs form typed relationships. Semantic retrieval may discover related content, but provenance, direction, time, assertion origin and institutional scope determine how the relationship can be used.

AIEDSM addresses this requirement through the Evidence Relationship Graph and Evidence Relationship Analysis.

5.2 FORMAL DEFINITIONS

Evidence Relationship Graph

An Evidence Relationship Graph is a versioned, typed and provenance-bearing graph in which nodes represent identified Evidence Objects, claims, requirements, standards, entities, activities and analytical outputs, and edges represent scoped relationship assertions.

The ERG is a controlled analytical structure, not a visualisation alone. Its underlying records preserve the attributes needed to inspect each node and relationship.

Evidence Relationship Assertion

An Evidence Relationship Assertion is a directional, typed and scoped proposition that a defined relationship exists between identified nodes, together with the provenance, method, time, status and limitations of that proposition.

Every edge is an assertion. Its presence does not establish source origin, institutional confirmation or current validity.

Evidence Relationship Analysis

Evidence Relationship Analysis is the controlled process of constructing, qualifying, validating, traversing and comparing relationships among Evidence Objects and related institutional constructs.

Evidence Relationship Analysis, abbreviated as ERA, includes relationship discovery, but it is not limited to discovery. It also governs how a proposed relationship is classified, reviewed, applied, updated, disputed and retired.

Boundary Statement

A relationship edge is an inspectable assertion, not self-validating truth. Semantic proximity may generate a candidate relationship, but it must not establish institutional equivalence, authority or applicability.

5.3 ERG NODE ARCHITECTURE

AIEDSM Controlled Taxonomy v1.0 defines the following node classes. Implementations may add namespaced extensions, but must not redefine the meaning or governance boundary of a v1.0 class.

Code	Node class	Governance boundary
EVIDENCE_OBJECT	Evidence Object	ISA-governed anchor; graph inclusion does not alter identity or status
CLAIM	Claim	A proposition; it is not evidence merely because an Evidence Object points to it
REQUIREMENT	Requirement	An identified criterion; AIEDSM does not create or legally interpret it
INSTITUTIONAL_REFERENCE	Institutional Reference Object	A standard, law, policy, guidance item or FAQ that is not automatically an Evidence Object
ENTITY	Entity	An organisation, person, asset, product, jurisdiction or subject; resolution uncertainty remains visible
ACTIVITY	Activity	An observed, reported or inferred event; observation basis must be stated
VERSION_EVENT	Version or Status Event	An issuance, amendment, correction, effectiveness, supersession or withdrawal event
METHOD	Method or Rule	A declared transformation, mapping, calculation or analytical procedure; presence does not validate result
ANALYTICAL_OUTPUT	Analytical Output	A result, inference, Candidate TIIR or TIIR; it does not inherit authority and is not automatically evidence
REVIEW_EVENT	Review Event	Immutable record of mandate, scope, findings, disposition and conditions
DECISION_RECORD	Decision Record	Record attributable to the authorised decision-maker; distinct from the TIIR

5.4 MINIMUM RELATIONSHIP ASSERTION FIELDS

A relationship assertion is a versioned analytical record, not a self-validating truth. Every material assertion must record the dimensions below.

Field	Minimum requirement
Assertion identity	Persistent assertion identifier and assertion version or immutable change-event reference
Source and target	Node identifiers and applicable versions
relationship_type	Controlled directional predicate from AIEDSM-RV 1.0
Direction and scope	Claim, section, field, period, jurisdiction, entity, purpose or other bounded scope
assertion_origin	Immutable origin: source, institutional mapping, deterministic method, machine proposal or human proposal
Assertion provenance	Asserting source, person, institution, model, rule or workflow
Evidence basis	Evidence Object references, exact locations or reproducible method
Valid and transaction time	Applicable interval plus creation, change and review events
Authority basis and limit	Authority of the assertion source and boundaries of that authority
validation_status	Review disposition, independent of origin
relationship_lifecycle_status	ACTIVE, CONTESTED, SUPERSEDED, RETIRED or WITHDRAWN
Score	Optional method-specific signal, separate from origin, validation, authority and truth
Conditions and exclusions	Purpose, jurisdiction, version, direction, exceptions and prohibited expansion
Supersession reference	Later assertion or immutable event that modifies or replaces this assertion

5.5 ASSERTION ORIGINS

5.5 ASSERTION ORIGIN, VALIDATION AND LIFECYCLE

Relationship type states what is asserted. `assertion_origin` states how the assertion entered the analytical environment. `validation_status` states the review disposition. `relationship_lifecycle_status` states whether it remains operative. These dimensions must not be collapsed.

5.5.1 ASSERTION ORIGIN

Code	Display label	Meaning
SOURCE_DECLARED	Source-Declared	Expressly stated by an identified eligible source
INSTITUTIONALLY_MAPPED	Institutionally-Mapped	Established through an identified institution's authorised mapping process for a defined use
DETERMINISTICALLY_DERIVED	Deterministically-Derived	Follows from an explicit, versioned and reproducible rule applied to identified inputs
MACHINE_PROPOSED	Machine-Proposed	Proposed by a probabilistic, semantic, relational or other model-mediated process
HUMAN_PROPOSED	Human-Proposed	Proposed by an identified person outside a source declaration or authorised institutional mapping

5.5.2 VALIDATION STATUS

Code	Meaning
UNREVIEWED	No completed accountable review is recorded
REVIEW_IN_PROGRESS	Review has begun but no completed disposition exists
REVIEWED	Review is complete; this state alone does not authorise use
ACCEPTED_FOR_DECLARED_USE	An authorised reviewer accepted use for the recorded purpose, scope, period and conditions
ACCEPTED_WITH_QUALIFICATION	Use is accepted only with recorded limitations or conditions
DISPUTED	A material challenge remains recorded
REJECTED	The assertion is not accepted for the declared use

5.5.3 RELATIONSHIP LIFECYCLE STATUS

Code	Meaning
ACTIVE	Operative for its declared scope, subject to validation and temporal conditions
CONTESTED	A material challenge or competing assertion remains active
SUPERSEDED	Replaced by an identified later assertion for a declared scope and time
RETIRED	Not used for current analysis but preserved for reconstruction

Code	Meaning
WITHDRAWN	Withdrawn by the asserting authority or accountable controller under recorded conditions

5.5.4 IMMUTABLE ORIGIN RULE

Assertion origin is immutable. Review, acceptance, repetition or publication in a TIIR must not rewrite MACHINE_PROPOSED or HUMAN_PROPOSED as SOURCE_DECLARED or INSTITUTIONALLY_MAPPED. Validation can change permitted reliance for a declared use; it does not establish truth, source authority or assurance.

5.6 RELATIONSHIP VOCABULARY AND SCOPE

AIEDSM Relationship Vocabulary Profile, AIEDSM-RV 1.0, is the canonical predicate set for this publication. Implementations may add namespaced extensions only when the owner, namespace, version, definition, node constraints, direction, temporal behaviour and validation requirements remain explicit.

Relationship type	Locked meaning	Mandatory control
CITES	Source explicitly references target	Preserve citation locus and version; citation is not support
HAS_PRIMARY_SOURCE	Secondary object identifies a primary source	Preserve characterisation and independently resolve target where possible
SUPPORTS	Source materially supports an identified claim	Record exact claim, locus, context and limitations; not transitive
QUALIFIES	Source limits, narrows or conditions a claim or relationship	Preserve the qualified proposition and effect
CONTRADICTS	Source is materially incompatible after context checks	Preserve both positions, conflict type and disposition
DERIVED_FROM	Object or result was materially transformed from target	Preserve transformation, inputs, method, version, parameters and time
PROPOSES	Draft or consultation introduces content for consideration	Do not treat as final, adopted or effective

Relationship type	Locked meaning	Mandatory control
RESPONDS_TO	Response addresses a proposal, question or request	Preserve responder, date, scope and status
FINALISES	Final publication completes a declared proposal process	Preserve authority and identify what was and was not finalised
REVISION_OF	Newer object revises earlier object	Identify changed scope and continuing historical relevance
PRECEDES	Source is earlier in a declared sequence	Sequence does not imply revision, approval or causation
SUPERSEDES	Source authoritatively replaces target	Require authority basis, scope and effective time
AMENDS	Source authoritatively changes identified parts	Identify affected and unaffected provisions
CORRECTS	Source authoritatively corrects target content	Preserve both versions, notice and impact
CONSOLIDATES	Source combines identified earlier material	Do not erase component authority or history without separate basis
WITHDRAWS	Source records cessation of target	Preserve issuer, basis, date, scope and residual use
INTERPRETS	Source provides an interpretation of target	Preserve interpreting authority; do not transfer target authority
SUMMARISES	Source abridges target	Preserve omissions, version, actor or method and status
IMPLEMENTS	Source operationalises a requirement or policy	Preserve scope, jurisdiction, institution and effective conditions
APPLIES_TO	Connects object to entity, period, jurisdiction, purpose or use	Preserve basis, conditions and state
REQUIRES	Source declares an evidence element, action or condition	Preserve authority, scope, period and exceptions
DEPENDS_ON	Claim, result or decision materially depends on target	Preserve direction and conditions; not automatically transitive

Relationship type	Locked meaning	Mandatory control
TRANSLATION_OF	Source is presented as a translation of target	Preserve language, translator, status and authoritative-language rule
ALTERNATE_REPRESENTATION_OF	Nodes represent one governed institutional object	Require shared identity or authorised mapping; no equal legal effect implied
SEMANTICALLY_SIMILAR_TO	Declared method detects semantic similarity	Discovery only; preserve method, version, score and threshold
RELATES_TO	Weak association supports discovery	Cannot alone establish support, applicability, substitution or equivalence
CORRESPONDS_TO	Directional mapping for a declared analytical purpose	Preserve authority, method, scope, period and exclusions
AUTHORISED_SUBSTITUTE_FOR	Authority permits source to replace target for a defined use	Directional; preserve authority, jurisdiction, period and revocation
INSTITUTIONALLY_EQUIVALENT_TO	Authority establishes reciprocal equivalence within conditions	Require competent authority; never infer from similarity or correspondence

An unqualified equivalence predicate and the earlier suffixed support and contradiction predicates are retired aliases and must not appear in new records. The controlled v1.0 predicates are INSTITUTIONALLY_EQUIVALENT_TO, SUPPORTS and CONTRADICTS.

5.7 DIRECTION, TRANSITIVITY AND GRAPH PATHS

Relationship properties are controlled independently of their labels.

No controlled relationship is presumed transitive unless a versioned profile expressly permits transitivity for defined node classes and scope. No inverse assertion may be materialised without an explicit inverse rule and provenance record. Symmetry must not be inferred from wording.

INSTITUTIONALLY_EQUIVALENT_TO is reciprocal only within its authorised purpose, scope, jurisdiction and period. AUTHORISED_SUBSTITUTE_FOR and CORRESPONDS_TO remain directional.

A multi-edge path is a Derived Machine Representation. It must preserve every edge, assertion origin, validation status, version, temporal condition and path rule. A material change creates a new assertion version or immutable event; historical TIIR dependencies remain reconstructable subject to lawful retention and rights.

5.8 RELATIONSHIP AUTHORITY AND EVIDENCE WEIGHT

Node presence, density or centrality does not determine authority, evidentiary weight or materiality. A widely cited commentary may be more central than a primary source.

AIEDSM keeps at least three dimensions separate:

1. **Relationship existence:** whether a typed relationship has been asserted.
2. **Assertion standing:** how the assertion was created and reviewed.
3. **Institutional authority:** what authority the asserting source holds for the defined purpose.

An authoritative source may be outdated or outside the relevant scope. A valuable machine proposal may carry no authority, and a high score may remain a candidate. Keeping these dimensions separate permits discovery without claiming institutional resolution.

5.9 ERA OPERATIONAL TESTS

Test	Operational question
Edge inspection	Are identities, type, direction, scope, origin, provenance, time, validation and limitations inspectable?
Origin and validation separation	Does review preserve the original SOURCE_DECLARED, mapped, derived, machine-proposed or human-proposed origin?
Directionality	Does the system block an invalid inverse when nodes are reversed?
Transitivity	Does the engine block unsupported propagation across paths?
Authority Non-Transfer	Does commentary or machine output remain distinct from source authority?
Temporal assertion	Does retrieval apply the assertion version and valid period relevant to the query?
Contestation	Can competing assertions remain visible without forced resolution?
Lifecycle	Are retired, superseded and withdrawn assertions excluded or qualified for current use but preserved historically?
Weak relation	Do SEMANTICALLY_SIMILAR_TO and RELATES_TO remain discovery relations?
Path reconstruction	Can every material multi-edge inference be reconstructed from versioned assertions and path rules?

5.10 CHAPTER BOUNDARY

ERA represents institutional relationships without treating similarity as equivalence, proximity as authority or a graph path as a conclusion. The governing rule is:

A machine may propose how Evidence Objects relate, but the graph must preserve who or what asserted the relationship, under which conditions, for which period and with what institutional standing.

CHAPTER 6: SEMANTIC SIMILARITY AND INSTITUTIONAL EQUIVALENCE

6.1 THE EQUIVALENCE ERROR

Semantic similarity supports discovery, clustering, cross-reference and candidate mapping but cannot establish interchangeability. Requirements, guidance, organisational statements, reports, editorial analysis and machine summaries may use similar language while retaining different authority, purpose, scope, time and evidentiary roles. Conversely, institutionally related expressions may use different terminology. Low similarity does not prove that no relationship exists.

AIEDSM therefore treats semantic similarity as an analytical signal, not an institutional determination.

6.2 FORMAL DEFINITIONS

SEMANTIC SIMILARITY

A method-generated indication that representations are alike in language, concepts or embedding space. It supports discovery and candidate formation only.

INSTITUTIONAL COMPARABILITY

A determination that objects can be examined against declared dimensions without implying substitution, acceptance or equivalence.

INSTITUTIONAL CORRESPONDENCE

A bounded, normally directional relationship indicating that objects or concepts perform related functions for a declared analytical purpose. It is represented through CORRESPONDS_TO.

AUTHORISED SUBSTITUTABILITY

A directional determination by an authority competent for the declared use that one object may be used in place of another within a recorded purpose, scope, jurisdiction, period, conditions and revocation rule. It is represented through AUTHORISED_SUBSTITUTE_FOR.

INSTITUTIONAL EQUIVALENCE

A reciprocal determination by an authority competent for the declared equivalence that two objects or concepts may be treated as equivalent for a defined purpose, scope, jurisdiction and period. It is represented through INSTITUTIONALLY_EQUIVALENT_TO. Semantic similarity, internal convenience or review alone cannot establish external regulatory equivalence.

Authority classes remain explicit: an external regulator or legal authority, framework owner, contractual authority and internal governance role do not possess interchangeable mandates. An internal acceptance may govern internal use; it cannot establish equivalence under an external regime unless that regime authorises it.

6.3 THE INSTITUTIONAL COMPARISON DIMENSIONS

Before proposing correspondence or equivalence, an AIEDSM analysis should compare the following dimensions.

Dimension	Required question	Risk if omitted
Institutional function	What role does each object perform: requirement, guidance, evidence, interpretation, decision or analysis?	Objects with different functions may be treated as substitutes
Subject and unit of analysis	What entity, activity, claim, asset, disclosure or decision does each expression address?	Similar language may refer to different subjects
Definition	Are key terms defined in the same way and at the same level?	Lexical matches may conceal definitional differences
Scope and boundary	Which operations, entities, periods, value-chain boundaries or exclusions apply?	Partial overlap may be represented as full alignment
Measurement or assessment basis	Which method, unit, threshold, assumption or evidence condition applies?	Comparable topics may generate non-comparable results
Authority	Who issued, interpreted or assigned the expression, and within what authority?	Commentary or machine output may inherit source authority

Dimension	Required question	Risk if omitted
Jurisdiction	Under which legal, regulatory, organisational or contractual context does it operate?	Cross-jurisdictional similarity may be mistaken for applicability
Version and status	Is each item draft, proposed, issued, effective, amended, superseded or withdrawn?	An outdated or proposed item may replace a current one
Temporal applicability	For which period or event is each expression relevant?	Current text may be applied to historical facts, or vice versa
Intended purpose	For what analytical, reporting, operational or decision purpose was it designed?	Similar structures may not support the same use
Direction	Is the proposed relationship one-way or reciprocal?	A valid directional mapping may be reversed incorrectly
Rights and permitted use	Can the item be reproduced, transformed, compared or disclosed for the intended use?	Machine access may be treated as use permission
Exceptions and conditions	Which exclusions, transition provisions or qualifiers apply?	Conditional correspondence may become absolute equivalence

The comparison need not produce a binary outcome. An analysis may conclude that two objects are topically related, structurally comparable, partially corresponding, operationally substitutable for a narrow task or not comparable under the available context.

6.4 THE SIMILARITY-TO-RELATIONSHIP CONTROL SEQUENCE

Similarity is a candidate-generation signal. It becomes institutionally usable only through an inspectable sequence whose states are not automatically progressive.

Stage	Control
1. Similarity Detection	Record representation, model, input versions, segmentation, score and threshold
2. Candidate Formation	Create SEMANTICALLY_SIMILAR_TO or another Machine-Proposed candidate relationship

Stage	Control
3. Context Qualification	Compare identity, function, definition, scope, authority, jurisdiction, version, time, rights and direction
4. Classification	Assign UNRELATED, RELATED, COMPARABLE, CORRESPONDING, CANDIDATE_SUBSTITUTION, INSUFFICIENT_CONTEXT or REJECTED_NON_EQUIVALENT
5. Authority Gate	Require the competent authority and conditions before AUTHORISED_SUBSTITUTE_FOR or INSTITUTIONALLY_EQUIVALENT_TO
6. Controlled Use	Apply only within recorded purpose, direction, period, version, jurisdiction and scope
7. Review and Invalidation	Reassess on source, status, jurisdiction, evidence, relationship, authority or method change

6.5 EQUIVALENCE ASSERTION CONTROLS

An equivalence or substitution assertion must identify the exact source and target, relationship type, direction, declared purpose, subject, scope, jurisdiction, period, version, competent authorising authority, authority basis, validation status, exceptions, revocation rule and review date.

CORRESPONDS_TO is directional by default and never implies substitution. AUTHORISED_SUBSTITUTE_FOR is directional and limited to its declared conditions. INSTITUTIONALLY_EQUIVALENT_TO is reciprocal only within its expressly authorised conditions.

A Machine-Proposed or Human-Proposed candidate cannot be elevated by similarity score, graph centrality, repetition or internal review alone. If the authority basis is absent, disputed or out of scope, classification remains INSUFFICIENT_CONTEXT, CORRESPONDING or REJECTED_NON_EQUIVALENT as applicable.

C2PA can bind provenance and integrity assertions to digital content, but successful content-credential validation does not establish that a proposition is true, authoritative, applicable or sufficient for a declared institutional use (Coalition for Content Provenance and Authenticity, 2026).

6.6 AUTHORITY PRESERVATION IN INTERPRETATION AND TRANSLATION

Interpretations, translations and summaries can create authority confusion. AIEDSM records them as separate nodes linked to the source through typed relationships.

The source authority does not flow automatically through INTERPRETS, TRANSLATION_OF, SUMMARISES or CITES. A downstream analytical output must identify both the source and the intermediary when it relies on the intermediary's wording or judgement.

For translation, the ECE should preserve:

1. source and target language;
2. source object and version;
3. translator or generating system;
4. translation status;
5. authoritative-language conditions where stated;
6. review status;
7. known terminology decisions;
8. segment alignment where material;
9. generation and revision dates.

Similarity between a translation and source does not establish authority. The appropriate determination depends on review, context and intended use.

6.7 COMPARISON CLASSIFICATION AND VALIDATION

Comparison outputs use explicit classification, assertion-origin and validation fields rather than a single similarity or standing label.

Classification	Meaning	Permitted use
SIMILARITY_SIGNAL	Method detects content likeness	Discovery and ranking only
CANDIDATE_CORRESPONDENCE	Machine or human proposes a bounded relationship	Review, mapping preparation and exploratory analysis
CANDIDATE_DIVERGENCE	Method detects a material difference requiring context analysis	Conflict screening and review routing
CONTEXT_QUALIFIED_COMPARISON	Required dimensions have been compared for a declared purpose	Controlled analysis within stated limits
ACCEPTED_CORRESPONDENCE	Authorised role accepts correspondence for a defined use	Use only within approved purpose, period and scope

Classification	Meaning	Permitted use
AUTHORISED_SUBSTITUTION	Competent authority permits directional substitution	Substitution only in recorded direction and conditions
INSTITUTIONAL_EQUIVALENCE	Competent authority establishes reciprocal equivalence	Reciprocal use only within recorded purpose, jurisdiction, period and scope
CONTESTED	Material challenge remains unresolved	Preserve and expose competing positions
INSUFFICIENT_CONTEXT	Information is insufficient to classify	No equivalence or substitution claim
RETIRED	Relationship is no longer active	Historical reconstruction only

The comparison classification is not assertion origin. The original SOURCE_DECLARED, INSTITUTIONALLY_MAPPED, DETERMINISTICALLY_DERIVED, MACHINE_PROPOSED or HUMAN_PROPOSED origin remains immutable, while validation_status records review disposition.

6.8 SEMANTIC AND INSTITUTIONAL TEST CASES

The following tests expose common equivalence errors.

Test	Required behaviour
Authority Reversal	A similar commentary remains distinct from the authority of the primary source
Version Substitution	Similar versions retain different status and effective periods
Jurisdiction Substitution	Cross-jurisdictional similarity may support correspondence but not unqualified applicability
Purpose Substitution	Topical similarity between a procedure and requirement does not establish satisfaction
Evidence Role	Statements, reports and summaries retain distinct evidence roles
Translation Status	Authorised, unofficial and machine translations retain status and provenance
High Similarity	Context qualification remains required above a similarity threshold
Low Similarity	Externally established relationships survive low semantic similarity

Test	Required behaviour
Directionality	Partial correspondence cannot be reversed or expanded automatically
Exception Survival	Exclusions and transition conditions remain visible downstream

6.9 OPERATIONAL DECISION RULE

A machine may detect similarity, produce a comparison and propose a relationship. It must not assert authorised substitutability or institutional equivalence unless the competent authority, declared purpose, direction or reciprocity, scope, jurisdiction, applicable period, conditions and validation record are all present and resolvable.

Where context is incomplete, the correct result is INSUFFICIENT_CONTEXT or another qualified state, not an invented relationship. Where a candidate has been tested and fails the declared conditions, it may be REJECTED_NON_EQUIVALENT without implying that all other forms of relationship are absent.

6.10 CHAPTER BOUNDARY

AIEDSM supports semantic discovery and comparison without treating machine representations as institutional determinations. The governing rule is:

Two expressions may be semantically similar without being institutionally equivalent, and institutionally related without being semantically similar. AIEDSM preserves both possibilities.

CHAPTER 7: TEMPORAL EVIDENCE SCIENCE

7.1 TIME AS AN INSTITUTIONAL CONDITION

Evidence operates across event, creation, publication, effective, amendment and retrieval times. Machine retrieval may privilege the latest indexed item, although it may not be effective or applicable to the event. Later corrections change current knowledge without changing the earlier knowledge boundary, and superseded versions may remain necessary for historical reconstruction.

AIEDSM therefore treats temporal context as part of evidence identity and applicability, not as optional metadata.

7.2 FORMAL DEFINITION

Temporal Evidence Science is the disciplined analysis and control of how event time, record time, publication time, effective time, version, status and applicability affect the machine interpretation and institutional use of Evidence Objects.

Temporal Evidence Science does not create an external effective date or determine legal effect. It preserves the source-declared and institutionally assigned temporal conditions needed for traceable analysis.

Boundary Statement

Latest is a retrieval ordering. Current is a context-dependent status. Effective is an institutional condition. Applicable is a bounded determination. AIEDSM does not treat these terms as interchangeable.

7.3 MINIMUM TEMPORAL DIMENSIONS

Temporal evidence analysis preserves distinct clocks rather than compressing them into one date.

Temporal dimension	Meaning	Analytical significance
Event or observation time	When the underlying activity or condition occurred	Connects evidence to the fact pattern or reporting period
Record creation time	When the record was created	Distinguishes contemporaneous from retrospective evidence
Issuance or publication time	When the source was formally issued	Supports publication chronology

Temporal dimension	Meaning	Analytical significance
Institutionally available time	When the representation became available under applicable access conditions	Defines the institutional availability boundary
System access time	When the analytical environment first obtained access	Separates availability from system capability
System ingestion time	When the representation entered the analytical index or store	Supports reconstruction of system knowledge
Effective-from time	When a requirement, status or relationship begins to operate	Prevents premature application
Effective-to time	When stated effect ends or is replaced	Prevents use beyond the applicable period
Observation or coverage period	Period described by the Evidence Object	Prevents period mismatch
Amendment or correction time	When a source or relationship changed	Supports version history and impact analysis
Supersession or withdrawal time	When an object ceased to be current for a scope	Preserves historical use without current substitution
Retrieval time	When the representation was selected for a run	Reconstructs the analytical environment
Transformation time	When a derivative or inference was produced	Distinguishes outputs under different source or model states
Review or validation time	When context, relationship or output was reviewed	Prevents indefinite reliance on stale review
Decision time	When an accountable decision considered the record	Supports accountability and historical reconstruction

7.4 VERSION, STATUS AND APPLICABILITY

A generic status field is insufficient. AIEDSM preserves the native status label and separately records publication, version, effective, knowledge-boundary and task-applicability dimensions.

Field	Controlled values or rule
source_publication_status	DRAFT, PROPOSED, FINAL, WITHDRAWN, OTHER or UNKNOWN, plus native label and source
version_state	CURRENT, SUPERSEDED, HISTORICAL or UNKNOWN
effective_state	NOT_YET_EFFECTIVE, EFFECTIVE, EXPIRED, NOT_APPLICABLE or UNRESOLVED
knowledge_boundary_state	KNOWN_BY_CUTOFF, INSTITUTIONALLY_AVAILABLE_BUT_NOT_SYSTEM_AVAILABLE, INGESTED_AFTER_CUTOFF, NOT_YET_PUBLISHED_AT_CUTOFF or UNKNOWN
applicability_state	APPLICABLE, CONDITIONALLY_APPLICABLE, NOT_APPLICABLE, DISPUTED or UNRESOLVED
native_status	Unmodified source label, asserting source, date and exact reference
status_assertion	Assertion origin, validation status, valid time and supersession event
conditions and exclusions	Thresholds, transitions, exceptions, dependencies and residual obligations
review and revalidation	Review date, trigger, current record status and downstream dependencies

A newer item is not necessarily effective; a final item is not necessarily applicable; and system ingestion is not the same as institutional availability.

7.5 VERSION GRAPH CONTROLS

The ERG represents version relationships as explicit assertions. The following distinctions are material:

1. PRECEDES records sequence but does not establish replacement.
2. SUPERSEDES records replacement for a defined scope and effective period.
3. AMENDS changes identified parts while other parts may continue.
4. CORRECTS addresses an identified error and preserves the correction history.
5. CONSOLIDATES combines prior material into a new representation but does not necessarily erase the authority or historical relevance of component sources.
6. TRANSLATION_OF links language versions while preserving translation status and authoritative-language conditions.
7. DERIVED_FROM records transformation without claiming equivalence or replacement.
8. WITHDRAWALS records removal or cessation under the source's declared conditions.

Each version edge must identify affected scope. Partial amendment must not imply whole-object replacement. Version histories remain additive, with prior versions retained for historical reconstruction subject to rights and retention rules. Changed content at a stable locator must create or update a version assertion because locator stability does not prove content stability.

7.6 TEMPORAL QUERY MODES

Query mode	Selection rule
Latest Available	Most recently available eligible version under a declared ordering, without implying effectiveness
Currently Effective	Objects effective at query time for the declared context
Effective-As-Of	Version and status effective at a specified time under the declared authority basis
Known-As-Of	Apply both the institutional-availability and system-knowledge boundaries declared for the run; report which boundary excluded an item
Event-Time	Evidence connected to an event or observation period, including later-created records
Publication-Time	Material issued within a period, without presuming effectiveness
Proposed-State	Draft or proposed material for horizon analysis with non-final status preserved
Historical Reconstruction	Object, relationship, context, system and review state relevant to a past output or decision

7.7 INSTITUTIONAL TIME AND SYSTEM TIME

Institutional Time describes when a source, rule, relationship or determination existed or applied within its institutional environment. System Time describes when the analytical environment accessed, ingested, transformed or reviewed the relevant representation.

A Known-As-Of analysis must distinguish institutional availability from system acquisition. `institutionally_available_at` records when the representation was available under applicable access conditions. `system_accessed_at` and `system_ingested_at` record when the analytical environment first accessed and incorporated it.

An item first ingested after the analytical cut-off must not be silently included in a reconstruction of the system's earlier knowledge state, even if the source bears an earlier publication date. Conversely, absence from the system does not establish that the item was institutionally unavailable.

7.8 LATE-ARRIVING EVIDENCE, CORRECTIONS AND RESTATEMENT

Late-arriving evidence, corrections and restatements create new analytical events. They must not overwrite the object, relationship, ECE, TIIR or decision state that existed at an earlier cut-off.

A material correction or late arrival triggers a change-impact assessment. The system identifies affected Evidence Objects, Field Assertion Records, relationships, claims, TIIRs, reviews and decisions. Where a material dependency exists, an approved TIIR moves to STALE or PENDING_REVALIDATION until an authorised disposition is recorded.

Prior decisions remain historically attributable. New reliance must wait until the affected record is revalidated, replaced, expressly restricted to historical use or otherwise disposed of by an authorised role.

7.9 TEMPORAL APPLICABILITY CONTROLS

Applicability is a task-bound determination that must identify the target entity, activity, claim, jurisdiction, reporting period or decision purpose; the version and effective state; authority basis; assertion origin; validation status; conditions; and analytical cut-off.

Every material time interval must specify its time standard and precision, whether each endpoint is inclusive or exclusive, any open endpoint, the source of the date and any uncertainty. Timestamps should use ISO 8601 with a time zone or UTC offset. A date-only value must not be represented as a precise timestamp. Estimated or inferred dates must retain their assertion origin and uncertainty.

AIEDSM does not determine legal applicability, regulatory scope or materiality. A machine may apply institutionally approved significance criteria and route uncertain applicability for review; it must not invent the institutional determination.

7.10 TEMPORAL CONTEXT SURVIVAL IN MACHINE SYSTEMS

Temporal context often disappears during chunking. A passage about a proposal, historical requirement or expired status may be presented as current when retrieved alone.

Every material Derived Machine Representation should therefore inherit or resolve to:

1. parent object and version;
2. native and controlled status;
3. publication or issuance date;

4. relevant effective period;
5. observation or coverage period;
6. supersession or withdrawal state;
7. generation time;
8. temporal query mode;
9. invalidation condition.

Summaries and analytical outputs should state an as-of time where current status matters. A current-status answer should disclose material proposed, disputed or superseded sources when they affect interpretation. A historical answer should not be silently updated with later evidence unless the query asks for retrospective analysis.

Caches, indexes and agent memories require temporal invalidation. When an object changes version or status, the system should identify affected chunks, embeddings, relationship assertions, summaries and outputs. Regeneration is not always necessary, but review status and lineage must remain accurate.

7.11 TEMPORAL OPERATIONAL TESTS

Test	Operational question
As-Of reproducibility	Can eligible object, assertion, ECE and system states be reproduced for the declared cut-off?
Latest versus effective	Is a newer non-effective item distinct from the effective item?
Publication-status separation	Do draft, proposed, final and withdrawn states remain distinct from version and effectivity?
Version chain	Can proposal, revision, amendment, correction, consolidation, supersession and withdrawal be reconstructed?
Partial amendment	Does a partial change avoid whole-object replacement?
Applicability separation	Is general source status separate from task-specific applicability?
Institutional availability	Can the record show when the source was institutionally available under applicable access conditions?
System knowledge	Can access and ingestion after the cut-off be excluded from a historical system reconstruction?

Test	Operational question
Late arrival	Are event, availability, access and ingestion times preserved without rewriting prior outputs?
Correction propagation	Do corrections trigger recorded change-impact and revalidation rules?
Stable-locator change	Does changed content at one locator create a distinct observation and digest?
Interval precision	Are endpoint inclusion, precision, time zone, open bounds and uncertainty explicit?
Temporal claim trace	Can each time-sensitive claim resolve to the applicable states, intervals and query mode?

7.12 CHAPTER BOUNDARY

Temporal Evidence Science enables AIEDSM to preserve what applied, what was known and what changed. It does not determine external legal effect or remove the need for authorised institutional judgement.

The governing rule is:

Machine analysis must identify not only which Evidence Object supports a proposition, but which version, status, period and applicability determination made that support relevant at the time of analysis.

Chapters 4 to 7 establish the AIEDSM analytical substrate: the inherited Evidence Object, context-preserving ECE, governed relationship assertions, equivalence controls and temporal controls. Later chapters distinguish evidence support from inference, expose conflicts and gaps, and produce Traceable Institutional Intelligence without transferring decision authority to the machine.

PART III: BOUNDED MACHINE ANALYSIS

CHAPTER 8: BOUNDED MACHINE ANALYSIS AND CLAIM-LEVEL ANALYTICAL LINEAGE

8.1 FROM EVIDENCE AVAILABILITY TO GOVERNED ANALYTICAL USE

MULTI-AXIS EVIDENCE-TO-INFERENCE CLASSIFICATION

Classify every claim on every axis. No state substitutes for, or silently changes, the claim basis.




Figure 3. Multi-axis Evidence-to-Inference classification model

Availability does not authorise every analytical use of an Evidence Object. Evidence eligibility is task-bound and depends on identity, provenance, version, time, authority, jurisdiction, rights, relationships, and intended use. These conditions are made operational through an Evidence Context Envelope, or ECE, for each analytical run.

Before analysis, the system must resolve the object, determine the applicable version and temporal state, preserve its institutional role, and record its permitted use. Retrieval supplies candidate material; governed analysis determines whether and in what role it may be used. Every subsequent retrieval, calculation, classification, summary, or proposed relationship is a governed transformation that must preserve the objects used, selection conditions, and claim status.

8.2 THE EVIDENCE-TO-INFERENCE BOUNDARY

8.2.1 DEFINITION

The Evidence-to-Inference Boundary is implemented as a multi-axis claim model. It records how a Material Analytical Claim was produced, the condition of its evidentiary resolution, its review disposition and its temporal applicability. These dimensions are not quality rankings, probabilities of truth or substitutes for institutional judgement.

Each Material Analytical Claim must carry independent `claim_basis`, `resolution_state` and `validation_status` fields, together with the applicable publication, version, effective, knowledge-boundary and applicability states.

8.2.2 CLAIM BASIS

Code	Display label	Required treatment
DIRECTLY_SUPPORTED	Directly-Supported	Resolve to exact Evidence Object location, version, status, context and qualification expressly supporting the claim
DETERMINISTICALLY_DERIVED	Deterministically-Derived	Record inputs, rules, parameters, units, rounding, exclusions, method version and reproducibility information
MACHINE_INFERRED	Machine-Inferred	Record Evidence Object basis, model or method class and version, configuration, limitations and non-source-stated character
COMPOSITE	Composite	Use only where the claim cannot reasonably be decomposed; identify every component basis and preserve inferential or qualified components

Hypothesis is not a claim basis. It is represented as `resolution_state` = UNRESOLVED and, where useful, `intended_use` = HYPOTHESIS_FOR_INVESTIGATION.

8.2.3 RESOLUTION STATE

Code	Meaning
SUPPORTED	The recorded evidence and method support the claim within the declared scope; this does not establish truth
QUALIFIED	Support exists, but a material limitation, condition, exception or scope restriction remains

Code	Meaning
CONFLICTED	Applicable evidence or methods remain materially incompatible after context checks
UNRESOLVED	The available eligible basis is insufficient to resolve the claim

8.2.4 VALIDATION AND TEMPORAL STATE

validation_status records whether the claim is UNREVIEWED, under REVIEW_IN_PROGRESS, REVIEWED, ACCEPTED_FOR_DECLARED_USE, ACCEPTED_WITH_QUALIFICATION, DISPUTED or REJECTED. Review may authorise reliance only for the declared use; it does not change claim basis, assertion origin, source authority or truth.

Temporal state is not a single label. Each claim records the applicable source_publication_status, version_state, effective_state, knowledge_boundary_state and applicability_state described in Chapter 7.

8.2.5 NON-ELEVATION RULE

A deterministic transformation does not elevate the evidentiary condition of its inputs. Where a material input is Machine-Inferred, Qualified, Conflicted, Unresolved, Stale or Pending Revalidation, that condition must remain visible in the resulting claim record.

For a COMPOSITE claim, the output must not be presented as wholly Directly-Supported or wholly Deterministically-Derived if any material component is Machine-Inferred, Qualified, Conflicted or Unresolved. Decomposition is preferred where it makes the different bases inspectable.

8.3 CLAIM-LEVEL ANALYTICAL LINEAGE

8.3.1 ANALYTICAL LINEAGE AS A REQUIREMENT

Document-level citation hides material differences among direct support, deterministic derivation and machine inference. AIEDSM therefore requires claim-level analytical lineage linking each Material Analytical Claim to its inputs, transformations, relationship assertions, states, limitations and TIIR.

The lineage must answer: what exact claim was produced; which Evidence Objects and versions were selected, qualified, excluded or considered; what transformations occurred; what claim_basis and resolution_state apply; what validation and temporal states govern use; and what conflict, gap, rights or limitation conditions remain.

The resulting Analytical Basis Record contains inspectable records. It is not private model reasoning and does not require disclosure of hidden chain-of-thought.

8.3.2 MINIMUM CLAIM RECORD

A Material Analytical Claim record must include:

1. claim identifier, exact text, version and creation time;
2. declared purpose and the approved significance rule or escalation threshold applied;
3. claim_basis and resolution_state;
4. validation_status and all applicable temporal-state dimensions;
5. supporting, qualifying and contradicting Evidence Object identifiers, versions and exact locations;
6. relationship assertion identifiers and versions;
7. transformation events, methods, models, tools, configurations and parameters;
8. search, access and selection boundaries where absence or coverage is material;
9. conflict, gap, rights and limitation records;
10. reviewer disposition where applicable; and
11. TIIR identifier and version.

Structured lineage may remain separate from readable prose, provided that both resolve to one another.

8.3.3 TRANSFORMATION EVENTS

Each material transformation, including extraction, conversion, comparison, graph traversal, classification, translation, or generation, should be recorded as an event with inputs, outputs, software actor, time, and material configuration. Intermediate summaries and agent messages remain derived analytical objects, not primary evidence.

8.3.4 TRACEABILITY DOES NOT ESTABLISH CORRECTNESS

Analytical lineage makes identity, provenance, and support inspectable. It can still record a mistaken calculation, misclassification, incomplete retrieval set, or inappropriate interpretation. Traceability supports governance and review; it does not establish substantive correctness.

8.4 EVIDENCE DIVERGENCE SIGNAL

8.4.1 DEFINITION AND PURPOSE

An Evidence Divergence Signal is a machine-detectable indication that relevant objects, claims, values, classifications, or relationships may not align. It triggers examination and prevents material differences from being averaged, merged, or summarised away. It does not establish conflict or error because scope, timing, version, jurisdiction, measurement, or authority may explain the difference.

8.4.2 DIVERGENCE DIMENSIONS

An Evidence Divergence Signal may be generated across one or more dimensions:

1. **semantic divergence:** claims appear to express materially different meanings;
2. **numerical divergence:** values differ beyond a declared tolerance;
3. **classification divergence:** objects assign different types, categories, or statuses;
4. **temporal divergence:** claims apply to different dates, periods, or effective states;
5. **version divergence:** the selected objects belong to different revision states;
6. **authority divergence:** sources occupy different institutional roles;
7. **jurisdictional divergence:** the claims arise under different legal or geographic boundaries;
8. **relationship divergence:** objects assert different relationships to the same entity, requirement, or claim;
9. **provenance divergence:** apparently equivalent claims depend on different source chains; and
10. **rights divergence:** the objects permit different forms of access, reuse, or analytical processing.

A signal records its dimension, objects, detection rule, applicable threshold, and time without overwriting source objects or relationships.

8.4.3 SIGNAL HANDLING

The system preserves each object, tests contextual explanations and authorised relationships, applies the declared significance criteria, preserves that result and its basis, and either records

closure or routes the signal to an Evidence Conflict State and Institutional Review. A proposal and a final publication may diverge without constituting two competing current requirements.

8.5 EVIDENCE CONFLICT STATE

8.5.1 DEFINITION

An Evidence Conflict State exists when relevant and sufficiently applicable objects or claims remain materially incompatible after context checks. It is a governed analytical state, not an instruction for the system to choose a winner.

8.5.2 CONFLICT TYPES

AIEDSM recognises the following principal conflict types:

1. **Content Conflict:** two applicable claims cannot both be accepted as stated;
2. **Version Conflict:** the applicable version cannot be determined or multiple incompatible versions are presented as current;
3. **Temporal Conflict:** claims depend on incompatible effective periods or cut-off dates;
4. **Authority Conflict:** sources with different institutional roles make competing claims and no authorised priority rule resolves them;
5. **Jurisdiction Conflict:** different jurisdictional rules apply to the same analytical subject;
6. **Applicability Conflict:** the criteria determining which requirement or interpretation applies are themselves disputed or incomplete;
7. **Method Conflict:** different declared calculation or classification methods produce materially different results; and
8. **Provenance Conflict:** the apparent source history contains incompatible or unverifiable chains.

A conflict may have multiple types. The record preserves its rationale and objects.

8.5.3 NO FORCED RESOLUTION

AIEDSM prohibits silent forced resolution based on frequency, search rank, embedding similarity, confident wording, or retrieval order. Resolution requires an explicit supersession relationship, priority rule, effective date, approved mapping, reproducible method, or accountable decision. Otherwise the output preserves the competing positions, contexts, and analytical consequence.

8.5.4 CONFLICT STATE TRANSITIONS

An Evidence Conflict State may transition to:

1. **Explained:** the difference is attributable to scope, time, version, or another recorded distinction;
2. **Resolved by Authoritative Relationship:** a valid supersession, amendment, correction, or other authoritative relation determines the applicable object;
3. **Resolved by Declared Method:** an approved analytical rule determines the treatment without claiming that a substantive institutional disagreement has disappeared;
4. **Escalated for Institutional Review:** an accountable reviewer must determine treatment;
or
5. **Unresolved:** no sufficient basis exists for resolution.

Each transition records the actor, rule, time, and affected claims. A machine recommendation may support review but has no independent decision authority.

8.6 EVIDENCE COVERAGE GAP ANALYSIS

8.6.1 DEFINITION

Evidence Coverage Gap Analysis compares a governed Expected Evidence Profile with the Evidence Objects available and eligible within a declared search, access, rights, language and temporal boundary. It describes the evidence environment; it does not establish legal compliance, materiality, assurance sufficiency or the existence of an unlocated item.

8.6.2 EXPECTED EVIDENCE PROFILE GOVERNANCE

An Expected Evidence Profile, or EEP, is a versioned and governed specification of the evidence roles, object types, relationships, periods, authority conditions and coverage rules expected for a declared analytical task.

A governed EEP must record its identifier, version, lifecycle and change history; author or controller; authority basis and exact source; purpose, use, scope, subject, jurisdiction and period; expected elements; admissible object types, roles, authority, time and relationships; coverage and substitution rules; search, access, rights, language and exclusion boundaries; cut-off and Known-As-Of mode; approved significance criteria; assertion origin and validation status; approval, review and expiry; and known limitations.

An AIEDSM system must not create an EEP and then report an unqualified gap as though the expectation were source-declared or institutionally authorised. A MACHINE_PROPOSED or HUMAN_PROPOSED EEP cannot support a governed coverage-gap conclusion until accepted for the declared use by an authorised reviewer. A missing, expired, disputed or unaccepted EEP requires resolution_state = UNRESOLVED.

8.6.3 COVERAGE STATES

Code	Meaning
COVERED	An eligible Evidence Object satisfies the declared condition
PARTIALLY_COVERED	Some but not all declared conditions are satisfied
COVERED_WITH_QUALIFICATION	Evidence is available with a material limitation, conflict or lower institutional role
NOT_LOCATED	No qualifying object was located within the declared boundary; this does not mean does not exist
NOT_ACCESSIBLE	An object may exist but could not be accessed under applicable controls
NOT_APPLICABLE	The expectation does not apply within the declared context
CONFLICTED	Available objects cannot support one coverage determination
UNRESOLVED	Information or EEP governance is insufficient to classify coverage

8.6.4 GAP TYPES AND SUFFICIENCY BOUNDARY

Gap types include identity, provenance, version, temporal, authority, relationship, content, coverage-period, rights and required-review gaps. Every gap must cite the EEP element and boundary against which it was identified.

Object count, citation count, edge density, frequency and centrality do not establish sufficiency. Density must remain separate from authority, relevance, independence, temporal validity, completeness and conflict. No coverage state is by itself a compliance, assurance, audit, certification or materiality conclusion.

8.7 INTEGRATED ANALYTICAL CONTROL SEQUENCE

An AIEDSM analytical run follows the dependencies below. Compatible controls may execute concurrently or iteratively, but the represented distinctions and immutable events must survive.

Stage	Mandatory dependency and output
1. Register Task and Declared Use	Question, purpose, scope, users, jurisdiction, cut-off, Known-As-Of mode, controller, rights and significance criteria
2. Load Governed EEP, where applicable	Resolve approved profile and version; without one, do not report a governed coverage-gap conclusion
3. Candidate Evidence Discovery	Retrieve candidates; discovery establishes neither eligibility nor support
4. Resolve Object and Representation	Resolve ISA identity, version, time, authority, rights and representation observation
5. Construct ECE	Bind task context and Field Assertion Records without altering the Evidence Object
6. Apply Eligibility and Selection	Record Included, Included with Qualification, Excluded or Unresolved and the controlling reason
7. Perform Bounded Analysis	Retrieve, compare, calculate, traverse, classify, translate, summarise or infer with transformation events
8. Construct Relationship Assertions	Apply AIEDSM-RV 1.0 and preserve origin, validation, time, version and lifecycle
9. Classify Claims and Build Lineage	Apply the multi-axis claim model to every Material Analytical Claim
10. Evaluate Divergence, Conflict and Coverage	Apply approved significance criteria; preserve unresolved states; permit governed iteration
11. Create Candidate TIIR	Package as DRAFT or PENDING_INSTITUTIONAL_REVIEW
12. Conduct Institutional Review	Review under an existing mandate and record rationale, qualification, restriction and disagreement
13. Release or Dispose	Only approved states permit reliance, within their declared boundary
14. Link Accountable Decision	Record acceptance, departure or rejection rationale in a separate Decision Record
15. Monitor and Revalidate	Apply triggers and block new consequential reliance on stale, expired or pending-revalidation records

Divergence, conflict, gap or unresolved-state analysis may require renewed discovery, object resolution, ECE revision, relationship testing or claim decomposition. The lifecycle is governed by dependencies and event preservation, not by an irreversible one-way pipeline.

8.8 CHAPTER CONCLUSION

AIEDSM does not prohibit inference. It keeps the boundary between evidence and inference visible so that a machine proposition cannot silently acquire the identity, authority, applicability, or

status of its sources. The Evidence-to-Inference Boundary, Evidence Divergence Signal, Evidence Conflict State, Evidence Coverage Gap Analysis, and claim-level analytical lineage convert a machine answer into an inspectable analytical record without claiming that inspectability guarantees correctness.

CHAPTER 9: AIEDSM REFERENCE ARCHITECTURE AND IMPLEMENTATION BOUNDARY




Figure 4. AIEDSM reference analytical lifecycle

9.1 PURPOSE OF THE REFERENCE ARCHITECTURE

The AIEDSM Reference Architecture separates evidence governance from execution technology. It does not prescribe a provider, database, model, embedding method, orchestration framework, agent protocol, or interface. Identifiers, provenance vocabularies, catalogues, signatures, graphs, indexes, retrieval systems, models, agents, and protocols may support an implementation; none independently establishes AIEDSM.

The architecture has seven functional layers:

1. Institutional Source Layer
2. Evidence Object and Resolution Layer
3. Evidence Context Envelope Layer
4. Provenance and Relationship Layer
5. Retrieval and Machine Analysis Layer

6. Analytical Output and Control Layer

7. Institutional Review and Decision Layer

Each layer has a distinct governance purpose so technical capability is not mistaken for institutional authority.

9.2 LAYER 1: INSTITUTIONAL SOURCE LAYER

Within the EMJ.LIFE MWP01-MWP04 architecture, the source layer begins with eligible institutional or participation-derived activity entering MWP01-PADV. PADV forms the governed activity/data record that may proceed to MWP03-ISA for Evidence Object admission. Where separate NTCC conditions are met, MWP02-NTCC may be produced as a non-tradable recognition output of PADV; it is not a stage between PADV and ISA and it does not create an alternative route.

The layer preserves the originating actor, event, verification conditions, time, rights and PADV lineage. Revision, restriction, withdrawal or deletion is represented without rewriting history. Historical content remains available only where lawful, permitted, proportionate and consistent with retention obligations. Where content must be deleted, a non-content tombstone or immutable event record should be retained where lawful and permitted.

9.3 LAYER 2: EVIDENCE OBJECT AND RESOLUTION LAYER

AIEDSM inherits only Evidence Objects admitted and governed through MWP03-ISA. It does not redefine the object or treat every record, document, chunk, token, relationship edge, model output or TIIR as one.

The layer provides the controlled entry point for machine analysis through persistent identity or a controlled registry identifier; canonical and representation resolution; object type; version, publication and effective states; provenance and authority; rights and permitted use; and versioned AIEDSM-RV 1.0 relationships.

Different Word, PDF, DOI, HTML, metadata or API representations may resolve to the same institutional object, but only through a recorded representation mapping and observation. Similar copies must not be collapsed without resolvable identity or an authorised relationship.

9.4 LAYER 3: EVIDENCE CONTEXT ENVELOPE LAYER

The ECE is the task-bound control object connecting an eligible Evidence Object and exact representation to a specified analytical run. It is not a duplicate source record.

It binds the object, representation observation, purpose, declared and prohibited uses, cut-off and Known-As-Of mode; jurisdiction and temporal states; authority and source role; rights and access; Field Assertion Records; versioned relationship assertions; EEP where applicable; selection rationale; processing events; and revalidation rule.

An ECE has its own identity, schema version and time. It remains immutable for a completed run. A material change to scope, source version, field assertion, right, cut-off or declared use creates a new ECE version. Every fragment or agent hand-off remains linked to the applicable ECE and Evidence Object.

9.5 LAYER 4: PROVENANCE AND RELATIONSHIP LAYER

The Provenance and Relationship Layer connects Evidence Objects, Field Assertion Records, version events, claims, actors, methods, reviews, outputs and decisions through inspectable records.

AIEDSM separates `relationship_type`, `assertion_origin`, `validation_status` and `relationship_lifecycle_status`. A Source-Declared supersession differs from a Machine-Proposed similarity edge even if both connect the same nodes. Review may alter `validation_status`, but it must not rewrite origin.

Every material graph path is a Derived Machine Representation with its edge sequence, versions, assertions, path rule, time and limitations. Graph presence, centrality and proximity do not establish evidentiary weight, institutional equivalence or authority.

9.6 LAYER 5: RETRIEVAL AND MACHINE ANALYSIS LAYER

9.6.1 FUNCTIONAL SCOPE

The Retrieval and Machine Analysis Layer may perform:

1. identifier-based resolution;
2. keyword or semantic retrieval;
3. metadata and graph filtering;
4. version and temporal selection;
5. relationship traversal;
6. comparison and classification;
7. deterministic calculation;

8. divergence detection;
9. conflict-state analysis;
10. evidence coverage analysis;
11. summarisation; and
12. bounded machine inference.

The layer may use rules, statistical methods, language models, vector search, graph algorithms, or combinations of these methods. AIEDSM does not prescribe one model architecture.

9.6.2 RETRIEVAL-AUGMENTED GENERATION

Retrieval-Augmented Generation, or RAG, is an implementation technique, not an evidence-governance method. Retrieval should consider type, version, status, time, jurisdiction, authority, rights, and relationship state, while recording queries, filters, retriever version, objects, fragments, material rankings, and exclusions.

Chunking must not sever source identity or ECE context. Generated citations require validation against retrieval records. Because retrieval neither establishes factuality nor completeness, the Evidence-to-Inference Boundary and claim-level lineage apply after generation.

9.6.3 KNOWLEDGE GRAPH AND GRAPH RAG

Knowledge graphs can support typed traversal, version reconstruction, provenance queries, entity resolution, conflict exposure, and coverage analysis. GraphRAG can extend analysis through graph indexes and generated community summaries. Model-extracted entities, claims, relationships, and summaries remain intermediate machine-derived artefacts. Their source units, extraction events, model or configuration, and assertion origins must remain traceable to the relevant Evidence Objects.

9.6.4 DETERMINISTIC AND PROBABILISTIC COMPONENTS

AIEDSM permits deterministic and probabilistic components within the same workflow, but their outputs must remain distinguishable.

A deterministic applicability filter and a model-proposed thematic relationship have different output states even when combined in one narrative.

9.7 AI AGENT ARCHITECTURE

9.7.1 AGENTS AS SOFTWARE ACTORS

An AI agent may retrieve, compare, calculate, classify, summarise, propose relationships, detect gaps, and coordinate tools. It remains a software actor without independent institutional authority.

Each participating agent should have a resolvable identity and execution record including:

1. agent identifier and version;
2. responsible operator or controller;
3. assigned role and permitted actions;
4. authorised data and tool scope;
5. run identifier and time;
6. received inputs and ECE references;
7. generated outputs and hand-offs;
8. exceptions, failed calls, or inaccessible evidence; and
9. escalation requirements.

Tool access does not expand authority. Retrieval permission does not confer authoritative interpretation, approval, or publication power.

9.7.2 MULTI-AGENT HANDOFFS

Each multi-agent hand-off is an analytical event. The receiving agent should distinguish:

1. a source Evidence Object;
2. an Evidence Context Envelope;
3. a deterministic result;
4. a machine inference;
5. an intermediate summary;
6. a proposed relationship; and
7. an instruction or decision issued by an accountable authority.

An agent message is not primary evidence. Unresolvable lineage requires qualification or exclusion from evidentiary support.

9.8 MCP AND AGENT-TO-AGENT COMMUNICATION

9.8.1 PROTOCOL-NEUTRAL POSITION

Model Context Protocol, or MCP, Agent-to-Agent communication, or A2A, and comparable protocols may transport context, expose tools, coordinate tasks, or exchange messages. In this chapter, MCP means Model Context Protocol and not the Minimum Cooling Period used in MWP03-ISA. AIEDSM remains implementation-neutral.

MCP or A2A does not establish:

1. Evidence Object identity;
2. source authenticity;
3. version applicability;
4. institutional authority;
5. evidentiary sufficiency;
6. correctness of an inference;
7. permission to use protected material; or
8. authority to make an institutional decision.

These conditions arise, if at all, through the evidence and governance layers.

9.8.2 EVIDENCE-PRESERVING TRANSPORT

Where MCP, A2A, an API, message queue or another transport mechanism is used, every transmitted analytical object must carry or resolve to:

1. object and envelope identifiers;
2. representation and location references;
3. version and status;
4. relevant provenance;
5. epistemic state;
6. relationship type, assertion origin, validation status and lifecycle status;
7. rights or use restrictions;
8. transformation history; and
9. the identity and role of the sending software actor.

A payload may use governed references instead of duplicating every field, but the receiver must not discard material context.

9.8.3 MCP AND A2A AS REFERENCE IMPLEMENTATIONS

MCP and A2A may appear in implementation profiles or reference implementations, but are not foundational definitions. Protocol reference does not imply endorsement, certification, compatibility, or conformance.

9.9 LAYER 6: ANALYTICAL OUTPUT AND CONTROL LAYER

The Analytical Output and Control Layer applies the multi-axis claim model, builds claim-level lineage, preserves divergence and conflict, performs governed coverage analysis where a current EEP exists, exposes limitations and constructs a Candidate TIIR.

AIEDSM is fail-closed at the boundary of unauthorised processing, release and consequential reliance. An unresolved mandatory condition must create a Restricted or Unresolved state; it must not be converted into an assumed value or omitted to permit completion.

Mandatory gate	Required result when unsatisfied
Rights and permitted use	Affected processing or release must not proceed; record RESTRICTED or the applicable failure state
Identity and provenance	TIIR must not become APPROVED_FOR_DECLARED_USE
Temporal and applicability	Object or claim must not be presented as current or applicable; qualify, restrict or escalate
Claim classification and lineage	Release for declared use must not proceed
Relationship authority	Machine/Human-Proposed relation cannot become source-declared, institutionally mapped, equivalent or substitutable
Expected Evidence Profile	Gap conclusion must not be released without current accepted EEP and search boundary
Material conflict	Conflict remains visible; conclusive dependent claim requires authorised disposition or treatment rule
Institutional Review	Required but incomplete, expired or unauthorised review blocks approved states

Mandatory gate	Required result when unsatisfied
Revalidation	STALE, PENDING_REVALIDATION, EXPIRED, SUPERSEDED or WITHDRAWN TIIR cannot support a new consequential decision

Where uncertainty can be safely preserved without prohibited processing, the record remains RESTRICTED, UNRESOLVED or PENDING_INSTITUTIONAL_REVIEW until governed disposition. A gate outcome is not a compliance or truth conclusion.

9.10 LAYER 7: INSTITUTIONAL REVIEW AND DECISION LAYER

The accountable institution evaluates the output, objects, conflicts, gaps, and method. It may approve, qualify, reject, return, supersede, restrict, or seek specialist judgement. No model, agent, retriever, graph, protocol, or score removes accountable authority from an institutionally consequential decision.

9.11 CONCEPTUAL DEMONSTRATION 1: CONTROLLED PUBLICATION RETRIEVAL

9.11.1 SCENARIO

An AI system is asked to state the MWP03-ISA Evidence Object dependency used by AIEDSM. Discoverable representations include a controlled Word master, the Version of Record PDF, the DOI record, the canonical landing page, search snippets, an earlier draft and third-party summaries.

9.11.2 POPULATED CONTROL RECORD

Control	Illustrative populated value
Task and declared use	Verify the canonical MWP03 dependency for public methodology citation; cut-off 10 September 2026
Candidate discovery	Word, PDF, DOI record, landing page, snippets, earlier drafts and third-party summaries
Resolved Evidence Object	MWP03-ISA V3.0; DOI 10.64969/padv.isa.2026.v3 ; published 10 September 2026
Selected representation	Version of Record PDF; SHA-256 ea32924428d8c9339027e9f98767f24950703d1773a41e9ff5dc02fb8ae2848

Control	Illustrative populated value
ECE conditions	Publisher EMJ LIFE HOLDINGS PTE. LTD.; Singapore; CC BY-ND 4.0; canonical landing page publication-isa
Relationship assertion	MWP04 DEPENDS_ON MWP03; assertion_origin = INSTITUTIONALLY_MAPPED; validation_status = ACCEPTED_FOR_DECLARED_USE
Claim state	claim_basis = DIRECTLY_SUPPORTED; resolution_state = SUPPORTED; temporal state records V3.0 as CURRENT at cut-off
Excluded support	Search snippets and third-party summaries support discovery only; earlier drafts remain HISTORICAL
Candidate TIIR	DRAFT until Institutional Review; limitations preserve any unavailable source or resolver state
Review and decision	Separate Review Event verifies citation and rights; any publication decision remains institutionally accountable

9.11.3 DEMONSTRATION BOUNDARY

This populated example demonstrates record structure and control application. It does not constitute empirical validation, assurance, product testing or external adoption.

9.12 CONCEPTUAL DEMONSTRATION 2: CONSULTATION RELATIONSHIP RECONSTRUCTION

9.12.1 SCENARIO

An AI system is asked to reconstruct the relationship among:

1. a consultation notice;
2. an exposure draft;
3. supporting materials;
4. an institutional response;
5. a feedback statement;
6. a final publication; and
7. subsequent corrections or guidance.

9.12.2 AIEDSM TREATMENT

Each publication remains a separate object with its issuer, date, version, status and role. PROPOSES, RESPONDS_TO, FINALISES, SUPERSEDES, CORRECTS and INTERPRETS require their AIEDSM-RV 1.0 controls, assertion origin and provenance. Semantic similarity may locate passages but cannot establish acceptance, adoption, substitution or equivalence. Divergence triggers context comparison; authoritative relationships may establish current treatment, while unclear treatment remains UNRESOLVED. The record classifies claims on the independent basis, resolution, validation and temporal axes.

9.12.3 DEMONSTRATION BOUNDARY

This is a conceptual demonstration of the methodology. It does not represent regulatory interpretation, evidence of external adoption, or empirical proof of system performance.

9.13 CHAPTER CONCLUSION

The architecture positions evidence governance above implementation technology and below accountable institutional decision. RAG, graphs, agents, MCP, and A2A can extend capacity but do not create authority. They operate within a common structure of resolvable Evidence Objects, ECEs, typed and provenance-aware relationships, bounded inference, claim-level lineage, reviewable records, and accountable decisions.

PART IV: INSTITUTIONAL OUTPUT AND GOVERNANCE

CHAPTER 10: TRACEABLE INSTITUTIONAL INTELLIGENCE AND INSTITUTIONAL ACCOUNTABILITY




Figure 5. Traceable Institutional Intelligence Record and accountability records

10.1 FROM MACHINE OUTPUT TO INSTITUTIONAL INTELLIGENCE

Fluency, timeliness, and citations do not make an AI answer institutional intelligence. AIEDSM requires an inspectable analytical basis, explicit epistemic states, preserved context, and accountable review. Traceable Institutional Intelligence describes a reviewable output, not a machine possessing institutional intelligence.

10.2 DEFINITION OF TRACEABLE INSTITUTIONAL INTELLIGENCE

Traceable Institutional Intelligence is a bounded analytical output in which material claims remain resolvably connected to applicable Evidence Objects, transformations, temporal and jurisdictional context, institutional relationships, epistemic states, limitations, and review status.

Traceable Institutional Intelligence should exhibit the following properties:

1. **resolvable basis:** supporting Evidence Objects and representations can be identified;
2. **context preservation:** source authority, version, status, time, jurisdiction, relationships, and rights remain inspectable;

3. **claim-level differentiation:** direct support, deterministic derivation, machine inference, and unresolved propositions are separated;
4. **transformation visibility:** material analytical operations are recorded;
5. **conflict visibility:** relevant conflicts are not silently removed;
6. **coverage transparency:** evidence gaps are reported against a declared profile;
7. **limitation disclosure:** system and evidence limitations are visible;
8. **reviewability:** an accountable reviewer can inspect the basis and disposition; and
9. **decision separation:** the output does not present itself as the final institutional decision.

These properties support governance but do not guarantee correctness, completeness, legal validity, assurance, or freedom from model error.

10.3 THE TRACEABLE INSTITUTIONAL INTELLIGENCE RECORD

10.3.1 DEFINITION

The Traceable Institutional Intelligence Record, or TIIR, packages the output and claim-level lineage for Institutional Review. It is not automatically an Evidence Object. That status requires an applicable evidence-generation, approval, publication, or acceptance process outside AIEDSM.

10.3.2 MINIMUM TIIR COMPONENTS

The minimum TIIR contains eleven separately resolvable components:

Component	Minimum content
1. Record Identity and Lifecycle	Identifier, version, record_status, creation time, change history, controller and revalidation rule
2. Task and Declared Use	Purpose, questions, scope, users, jurisdiction, cut-off, Known-As-Of mode, permissions, prohibitions and significance criteria
3. Method and Transformation Record	Methods, models, tools, rulesets, configurations, material parameters and events
4. Evidence Basis	Evidence Object, representation observation and ECE identifiers; inclusion, qualification, exclusion, search, access and rights boundaries
5. Claim Register	Material claims with claim_basis, resolution_state, validation_status, temporal states, exact support, transformations and limitations

Component	Minimum content
6. Relationship Record	AIEDSM-RV 1.0 relationships, assertion origins, validation and lifecycle states, versions and material paths
7. Divergence and Conflict Record	Signals, conflict states, treatments, dispositions and unresolved consequences
8. Coverage Record	Governed EEP and version, coverage states, gap types, search boundary and limitations, where applicable
9. Limitations and Use Conditions	Known limitations, exclusions, restrictions, unresolved states and reliance conditions
10. Institutional Review Record	Separate Review Event with mandate, competence, scope, method, findings, disposition and rationale
11. Decision Linkage	Reference to a separate Decision Record, including acceptance, departure or rejection rationale

A Candidate TIIR may combine readable content, metadata and graph records. No single serialisation is prescribed. Review and decision records remain distinct from the analytical output they govern.

10.3.3 TIIR RECORD LIFECYCLE

record_status is independent from claim basis, claim resolution, validation status and temporal applicability. Controlled values are:

Code	Release consequence
DRAFT	Work in progress; not authorised for declared use
PENDING_INSTITUTIONAL_REVIEW	Awaiting required review; not authorised for declared use
RETURNED_FOR_REVISION	Must be revised before renewed review
APPROVED_FOR_DECLARED_USE	Authorised only for the recorded purpose, scope, period and conditions
APPROVED_WITH_QUALIFICATION	Authorised only with visible recorded qualifications
RESTRICTED	Processing, access, reliance or release is limited by a recorded control
STALE	A trigger indicates the analysis may no longer represent the current evidence environment; no new consequential reliance

Code	Release consequence
PENDING_REVALIDATION	Impact assessment or renewed review is under way; no new consequential reliance
EXPIRED	The authorised use period has ended
REJECTED	Not authorised for the declared use
SUPERSEDED	Replaced by an identified later TIIR for a declared scope
WITHDRAWN	Withdrawn by the responsible institutional authority

Status changes are immutable events. Approval is use-specific and time-bound. Internal research approval does not authorise public, legal, regulatory, assurance, investment, procurement or other use unless expressly stated.

10.4 INSTITUTIONAL REVIEW

10.4.1 PURPOSE

Institutional Review is accountable examination of a TIIR before declared use. Depth reflects consequence, with a different review expected for low-impact research than for procurement, reporting, capital allocation, public policy, or assurance-related use.

10.4.2 REVIEW QUESTIONS

Institutional Review should consider:

1. Were the correct Evidence Objects and versions selected?
2. Was the analytical cut-off date appropriate?
3. Were authority, jurisdiction, status, and rights correctly represented?
4. Do cited locations materially support the claims attributed to them?
5. Are deterministic transformations reproducible?
6. Are machine inferences clearly labelled and appropriately bounded?
7. Were divergence signals and conflicts handled without silent resolution?
8. Was Evidence Coverage Gap Analysis based on a valid Expected Evidence Profile?
9. Are material omissions or exclusions disclosed?
10. Is specialist judgement required?

11. Is the proposed use within the approved boundary?

Review may use validation, sampling, domain expertise, legal review, model evaluation, or other controls. No universal procedure is prescribed.

10.4.3 REVIEWER ACCOUNTABILITY

Institutional Review documents and operationalises accountability under an existing institutional mandate for the declared use. It does not create the mandate, transfer source authority, establish truth or convert a machine output into primary evidence.

A Review Event must record reviewer identity, role and institution; mandate basis; competence and required specialist role; independence and conflict-of-interest declaration; TIIR, claim, Evidence Object and use scope; review method, depth, sampling and exclusions; findings, qualifications, unresolved issues and disagreements; disposition and rationale; revisions, restrictions, conditions and escalation; review date, effective period, expiry and revalidation triggers; and reviewer attestation or immutable event reference.

10.5 INSTITUTIONALLY ACCOUNTABLE DECISION

10.5.1 DEFINITION

An **Institutionally Accountable Decision** is attributable to an authorised human or organisational authority that has considered the evidence, TIIR, unresolved conditions, and governance requirements. A TIIR informs but does not make the decision.

10.5.2 DECISION SEPARATION

Analysis and decision remain separate. A system may recommend, rank, identify risk or detect a gap; the accountable authority determines its influence under an existing mandate.

A Decision Record must identify:

1. accountable decision-maker or institutional function;
2. authority basis and declared decision scope;
3. TIIR and Evidence Object identifiers and versions considered;
4. material conflicts, gaps, qualifications, restrictions and specialist advice;
5. decision, rationale, effective date, conditions, review date and expiry;

6. explicit rationale for accepting, departing from or rejecting the TIIR's analysis or recommendation; and

7. appeal, override, amendment, supersession and revalidation mechanisms.

This record prevents consequential choices from being attributed to a model when responsibility belongs to the institution.

10.6 GOVERNANCE IMPLICATIONS

10.6.1 INSTITUTIONAL CAPABILITY

AIEDSM requires governance capabilities for identity, metadata, provenance, versions, rights, relationships, review, and decision accountability, not merely an AI model.

An institution should be able to determine:

1. which sources are authorised for which analytical purposes;
2. how current and historical objects are resolved;
3. how machine-proposed relationships are separated from authoritative ones;
4. how analytical outputs are reviewed and released;
5. how model and tool changes are recorded; and
6. who remains accountable for consequential decisions.

10.6.2 EVIDENCE SOVEREIGNTY

Evidence sovereignty means that the institutional identity, provenance, status, authority, rights, and relationships of evidence are not surrendered to the model or platform that processes it.

Derived records may not silently alter source objects, erase issuers, replace canonical identity, or present machine content as source-issued. Machine readability does not remove rights, permissions, or use restrictions.

10.6.3 INDEPENDENT DECISION AUTHORITY

Analytical capacity and decision authority remain distinct. Greater volume, speed, or relational depth does not create absent authority.

10.6.4 REVALIDATION AND CHANGE IMPACT

The following events require a recorded impact assessment and, where a material dependency exists, transition of an approved TIIR to STALE or PENDING_REVALIDATION:

Research Program: Evidence Infrastructure Research Series

Trigger class	Illustrative trigger
Object or representation	Revision, amendment, correction, supersession, consolidation, withdrawal or material stable-locator change
Integrity	Digest mismatch or changed validation result
Institutional state	Publication, effectivity, applicability, jurisdiction, authority or version change
Rights and access	Licence, confidentiality, retention, access or permitted-use change
Relationship	Material assertion disputed, rejected, superseded, retired, withdrawn or reclassified
Evidence environment	New evidence creates material divergence, conflict, qualification or coverage change
Expected Evidence Profile	EEP change, expiry, dispute or withdrawal
Declared task	Change of question, use, audience, scope, jurisdiction, cut-off or Known-As-Of mode
Method or system	Material model, tool, rule, prompt template, configuration, threshold or parameter change
Review governance	Mandate, competence, independence or approval-boundary change or expiry
Anomaly	Provenance, security, integrity or processing anomaly discovered

A trigger must not silently overwrite the original TIIR. The system preserves the state used by the earlier analysis, identifies affected fields, relationships, claims, reviews and decisions, creates a change-impact event and blocks new consequential reliance until governed disposition.

10.7 INTEGRATION WITH PADV, NTCC, AND ISA

AIEDSM is the fourth methodology within the EMJ.LIFE Evidence Infrastructure Research Series. The four publications remain distinct in purpose and governance.

10.7.1 MWP01-PADV - EVIDENCE FORMATION ROUTE

Within this methodology architecture, PADV is the sole route through which eligible participation-derived activity is structured for admission as an ISA-governed Evidence Object. PADV formation does not guarantee ISA admission.

10.7.2 MWP02-NTCC - CONDITIONAL PADV OUTPUT

NTCC is a governed, non-tradable recognition output that may arise from an eligible PADV record when the separate NTCC conditions are met. NTCC is not an analytical stage, not a mandatory step between PADV and ISA and not an alternative Evidence Object formation route.

10.7.3 MWP03-ISA - EVIDENCE OBJECT GOVERNANCE AND STRUCTURAL INTEROPERABILITY

ISA governs admission, classification, identity, relationships, routing, structural interoperability and continuing status of the resultant Evidence Object without rewriting external standards.

10.7.4 MWP04-AIEDSM - MACHINE EVIDENCE INTELLIGENCE

AIEDSM governs machine analysis of ISA-defined Evidence Objects while preserving institutional context, authority boundaries, temporal validity, rights, relationship provenance and the separation between evidence support and machine inference.

The canonical architecture is PADV to ISA-governed Evidence Object to AIEDSM. NTCC appears only as a conditional non-tradable recognition output from PADV and remains linked to its originating PADV record. AIEDSM does not replace, bypass or create an alternative to PADV formation or ISA governance.

10.8 METHODOLOGICAL CONTRIBUTION AND EXTERNAL ALIGNMENT

AIEDSM integrates established capabilities rather than claiming invention of their underlying technologies. The crosswalk below identifies what the cited prior art already supplies and the narrower institutional control problem addressed by AIEDSM.

Prior art	Existing contribution	AIEDSM integration boundary
FAIR	Machine-oriented findability, accessibility, interoperability, reuse, identifiers, metadata, licences, provenance and qualified references	Governs which institutional conditions survive machine analysis of ISA Evidence Objects
RDF, SKOS and SHACL	Graph statements, concept schemes, mappings and graph constraints	Graph assertion or validation does not establish truth, evidentiary weight or equivalence
PROV-O and DCAT 3	Provenance, cataloguing, versions, dataset series, relations and digests	Metadata does not decide task applicability, source authority or claim support
OWL-Time	Instants, intervals, durations and temporal topology	Temporal encoding does not determine legal effect or task applicability
ODRL	Permissions, prohibitions, duties and policy constraints	Policy expression does not prove rights clearance or enforce every use

Prior art	Existing contribution	AIEDSM integration boundary
C2PA 2.4	Signed claims, manifests, assertions, content bindings and validation states	Integrity and authenticity do not establish truth, authority or sufficiency
NIST AI RMF and GAI Profile	Cross-sector AI risk governance, transparency, evaluation and lifecycle actions	Do not define ISA Evidence Objects, ECE, claim model or TIIR
RAG and GraphRAG	Retrieval-conditioned generation, graph indexing and corpus summarisation	Retrieval and graph proximity do not convert output into source-supported evidence
VeriTrail	Detection and tracing of unsupported content across generation steps	AIEDSM additionally preserves identity, authority, time, rights, relationships, review and use
MCP and A2A	Context/tool connections and inter-agent discovery, messages, artefacts, tasks and hand-offs	Transport does not create authority, alter claim basis or transfer accountability

These boundaries are cumulative. Conformance with an external technical specification, successful validation, cryptographic integrity, retrieval from a cited source or human review can strengthen a record, but none independently converts machine analysis into institutionally authoritative evidence.

10.9 CLOSING STATEMENT

AI does not make institutional evidence meaningful. Identity, provenance, time, authority, relationships, rights, and governed use do. AIEDSM preserves those conditions across retrieval, comparison, calculation, graph analysis, pattern identification, gap detection, conflict exposure, summarisation, and bounded inference.

The method concludes with a clear boundary:

AI may expand analytical capacity, but it must not erase the institutional conditions that make evidence meaningful.

AIEDSM V1.0 defines a controlled methodology for analysing ISA-governed Evidence Objects. It does not certify evidence, determine legal compliance or materiality, issue assurance, create institutional equivalence or transfer decision authority to a machine. Traceable Institutional Intelligence supports judgement, Institutional Review governs declared reliance and an Institutionally Accountable Decision remains attributable to the authorised human or institutional function.

APPENDICES

APPENDIX A: CONTROLLED TERMINOLOGY AND FOUNDATIONAL PROPOSITIONS

A.1 CONTROLLED TERMINOLOGY

AIEDSM V1.0 uses one controlled methodological hierarchy: one Core Doctrine, five Operational Pillars, twelve Foundational Propositions and thirteen Control Principles. The terms below are governed by AIEDSM Controlled Taxonomy v1.0.

Controlled term	Locked meaning
AIEDSM	AI Evidence Data Science Methodology. The EMJ.LIFE methodology governing context-preserving machine analysis of ISA-governed Evidence Objects and the production, review and revalidation of Traceable Institutional Intelligence.
AIEDSM Controlled Taxonomy v1.0	The versioned set of node classes, relationship types, assertion origins, validation states, temporal-state dimensions, claim-state dimensions, coverage states and record lifecycle states defined in this publication. Namespaced extensions must not redefine a v1.0 term.
Analytical Basis Record	The inspectable support, transformation, method, state, limitation and review record for a Material Analytical Claim. It is not private model reasoning and does not require disclosure of hidden chain-of-thought.
Assertion Origin	The immutable recorded basis through which a field value or relationship entered the analytical environment: SOURCE_DECLARED, INSTITUTIONALLY_MAPPED, DETERMINISTICALLY_DERIVED, MACHINE_PROPOSED or HUMAN_PROPOSED.
Authority Non-Transfer Principle	The principle that machine processing must not create, increase, confer or transfer institutional authority. An output may preserve, cite or characterise source authority; it does not inherit that authority.
Authorised Substitutability	A directional determination by an authority competent for the declared use that one object may be used in place of another within a recorded purpose, scope, jurisdiction, period, conditions and revocation rule.
Candidate TIIR	A Traceable Institutional Intelligence Record in DRAFT or PENDING_INSTITUTIONAL_REVIEW status. It is not authorised for declared use until a permitted approval state is recorded.
Claim Basis	The independent classification of how a Material Analytical Claim was produced: DIRECTLY_SUPPORTED, DETERMINISTICALLY_DERIVED, MACHINE_INFERRED or, only where reasonable decomposition is unavailable, COMPOSITE.

Controlled term	Locked meaning
Claim-Level Analytical Lineage	The resolvable connection between a Material Analytical Claim and the exact Evidence Objects, source locations, transformations, relationship assertions, states, methods and limitations on which it depends.
Decision Record	A record attributable to an authorised decision-maker or institutional function that states its mandate, the TIIR and evidence considered, the decision and its rationale, including any acceptance of, departure from or rejection of the TIIR.
Deterministically-Derived Result	A result produced from identified inputs through an explicit, versioned and reproducible calculation, rule, aggregation or transformation. Reproducibility does not establish truth or elevate the state of any input.
Directly-Supported Claim	A claim whose material content is expressly supported by one or more identified Evidence Objects at exact locations in the applicable version and context.
Evidence Conflict State	A governed analytical state in which sufficiently applicable Evidence Objects, claims or methods remain materially incompatible after version, temporal, jurisdictional, scope, authority and relationship checks.
Evidence Context Envelope	A versioned, task-bound control object binding an ISA-governed Evidence Object and selected representation to the analytical purpose, declared use, cut-off, jurisdiction, authority, temporal states, rights, relationships, field attribution, selection rationale, processing events and revalidation rule.
Evidence Coverage Gap Analysis	The governed comparison of a current and accepted Expected Evidence Profile with the Evidence Objects available and eligible within a declared search, access, rights, language and temporal boundary.
Evidence Divergence Signal	A machine-detectable indication that relevant Evidence Objects, claims, values, classifications or relationships may not align. A signal requires contextual examination and is not itself a conflict or error conclusion.
Evidence Object	An institutional evidence unit admitted and governed through MWP03-ISA after the MWP01-PADV formation route. AIEDSM does not redefine a document, chunk, token, graph edge, model output or TIIR as an Evidence Object.
Evidence Sovereignty	The governance condition under which evidence identity, provenance, version, status, authority, rights and relationships remain attributable to their legitimate source and are not surrendered to the model or platform processing them.

Controlled term	Locked meaning
Evidence-to-Inference Boundary	The methodological boundary implemented through the Multi-Axis Claim Model, which keeps claim basis, resolution state, validation status and temporal state independently inspectable.
Expected Evidence Profile	A versioned and governed specification of the evidence roles, object types, relationships, periods, authority conditions, search boundaries and coverage rules expected for a declared analytical task.
Field Assertion Record	The field-level record of value or field state, assertion origin, source reference, asserting actor, method or model version, assertion time, valid time, validation status and any superseding assertion.
Field State	The explicit condition of an ECE field: KNOWN, UNKNOWN, NOT_APPLICABLE, UNAVAILABLE, RESTRICTED, CONFLICTING or INVALID. A blank value is not a field state.
Institutional Correspondence	A bounded, normally directional mapping that objects or concepts perform related functions for a declared analytical purpose. It does not represent substitutability or equivalence.
Institutional Equivalence	A reciprocal determination by an authority competent for the declared equivalence that two objects or concepts may be treated as equivalent for a defined purpose, scope, jurisdiction and period.
Institutional Review	Accountable examination of a Candidate TIIR under an existing institutional mandate before reliance for a declared use. Review documents and operationalises accountability; it does not create the mandate or transfer source authority.
Institutionally Accountable Decision	A decision attributable to an authorised human or institutional function acting under an existing mandate after considering the relevant evidence, TIIR, unresolved conditions and governance requirements.
Machine-Inferred Proposition	A proposition produced through probabilistic, semantic, relational or model-based analysis that is not directly stated by the supporting Evidence Objects. Acceptance for use does not change its claim basis.
Material Analytical Claim	An output proposition whose error, omission, ambiguity, qualification or use could reasonably affect the declared institutional purpose, an Institutional Review or an Institutionally Accountable Decision.
Multi-Axis Claim Model	The independent classification of each Material Analytical Claim by claim_basis, resolution_state, validation_status and applicable temporal-state dimensions. The axes are not probabilities of truth or quality rankings.

Controlled term	Locked meaning
Relationship Assertion	A directional, typed, scoped and versioned proposition connecting identified nodes, with its assertion origin, provenance, evidence basis, valid and transaction time, authority boundary, validation status, lifecycle status and limitations.
Resolution State	The condition of a claim's evidentiary resolution: SUPPORTED, QUALIFIED, CONFLICTED or UNRESOLVED. It remains separate from claim basis and validation status.
Review Event	An immutable record of reviewer identity, mandate, competence, independence, scope, method, findings, disagreements, disposition, rationale, conditions, time and expiry.
Traceable Institutional Intelligence	A bounded analytical output in which Material Analytical Claims remain resolvably connected to applicable Evidence Objects, transformations, temporal and jurisdictional context, institutional relationships, claim states, limitations and review disposition.
Traceable Institutional Intelligence Record	The controlled record that packages Traceable Institutional Intelligence, claim-level lineage, analytical conditions, conflict, coverage, limitations, use conditions, lifecycle status and links to separate review and decision records.

A.2 FOUNDATIONAL PROPOSITIONS

The following twelve propositions are the authoritative consolidated statement of the AIEDSM Foundational Propositions. Lists elsewhere in this publication summarise or operationalise this set and do not create an alternative hierarchy.

PROPOSITION 1: EVIDENCE ELIGIBILITY IS TASK-BOUND

Eligibility depends on declared purpose, scope, time, jurisdiction, authority, rights, relationships and required evidence role.

PROPOSITION 2: DATA IS NOT AUTOMATICALLY EVIDENCE

Data becomes eligible for evidentiary use only when the conditions required for the declared task are sufficiently established.

PROPOSITION 3: SEMANTIC SIMILARITY DOES NOT ESTABLISH INSTITUTIONAL EQUIVALENCE

Similarity supports discovery and comparison; equivalence requires an authorised, scoped and time-bound determination.

PROPOSITION 4: AUTHORITY DOES NOT TRANSFER THROUGH MACHINE PROCESSING

Quotation, transformation, calculation, linkage, review or fluent expression does not transfer source authority to an Analytical Output.

PROPOSITION 5: TRACEABILITY DOES NOT ESTABLISH CORRECTNESS

Inspectable identity, provenance, support and transformation records do not establish substantive correctness.

PROPOSITION 6: INTEGRITY DOES NOT ESTABLISH TRUTH

A technical integrity result may show that specified content has not changed; it does not establish truth, authority or applicability.

PROPOSITION 7: MACHINE INFERENCE MUST NOT BE EPISTEMICALLY ELEVATED

A Machine-Inferred Proposition remains machine-inferred after review or acceptance. Review may authorise a declared use; it does not change claim basis, assertion origin or source authority.

PROPOSITION 8: ABSENCE REQUIRES A GOVERNED EXPECTATION AND DECLARED SEARCH BOUNDARY

A coverage gap may be reported only against a governed Expected Evidence Profile and an inspectable search, access, rights, language and temporal boundary.

PROPOSITION 9: CONFLICT REMAINS VISIBLE UNTIL GOVERNED DISPOSITION

Competing evidence remains visible until an authorised relationship, declared method or accountable decision determines its treatment.

PROPOSITION 10: EVIDENCE DENSITY DOES NOT ESTABLISH EVIDENCE SUFFICIENCY

Volume, frequency, centrality or retrieval rank cannot substitute for authority, relevance, independence, temporal validity and coverage.

PROPOSITION 11: TECHNICAL CAPABILITY DOES NOT CREATE INSTITUTIONAL AUTHORITY

Models, protocols, agents, graphs, signatures and retrieval systems do not establish institutional identity, truth, applicability or decision authority.

PROPOSITION 12: CONSEQUENTIAL DECISIONS REMAIN INSTITUTIONALLY ACCOUNTABLE

AI may inform action, but consequential use remains attributable to an authorised human or institutional function acting under an existing mandate.

APPENDIX B: INSTITUTIONAL BOUNDARY AND EVIDENCE SOVEREIGNTY STATEMENT

B.1 INSTITUTIONAL BOUNDARY

AIEDSM is an EMJ.LIFE methodology for governing machine analysis of institutional evidence. It is not an international standard, regulatory rule, legal interpretation, assurance standard, audit procedure, certification scheme, conformity assessment method, or professional opinion.

AIEDSM does not:

1. determine legal or regulatory compliance;
2. issue assurance, audit, verification, or certification conclusions;
3. determine materiality;
4. establish the truth of an Evidence Object or analytical output;
5. create authority not present in the source or governance structure;
6. replace legal, accounting, assurance, scientific, regulatory, editorial, investment, or other professional judgement;
7. convert correlation into causation;
8. treat editorial analysis as primary evidence;
9. treat machine inference as source-issued evidence;
10. guarantee completeness, factual correctness, or freedom from model error;
11. guarantee that hallucination, bias, omission, or retrieval failure will not occur;
12. mandate a particular AI model, vendor, protocol, database, or software architecture;
or
13. claim external adoption, endorsement, or conformance.

Implementation of AIEDSM controls may reduce certain forms of evidence ambiguity and make analytical limitations more inspectable. The effect depends on source quality, system design, implementation, testing, permissions, review, and institutional governance.

B.2 EVIDENCE SOVEREIGNTY STATEMENT

Evidence sovereignty under AIEDSM is governed by the following conditions:

1. **Source identity remains attributable.** The original issuer, publisher, custodian, and source role must not be replaced by the identity of the processing system.
2. **Canonical identity remains resolvable.** Copies, fragments, embeddings, and summaries must remain linked to the applicable Evidence Object where permitted.
3. **Version history remains visible.** Revision, replacement, withdrawal, and supersession must not be erased by presenting only a current fragment.
4. **Authority remains non-transferable.** Machine outputs may cite authoritative sources but do not become authoritative source statements.
5. **Rights remain operative.** Machine readability does not remove licence, confidentiality, data protection, access, attribution, or permitted-use conditions.
6. **Source objects remain unaltered by analysis.** Analytical corrections, annotations, and proposed relationships must be stored as separate governed records.
7. Institutional relationships remain typed. Source-declared, institutionally mapped, derived, machine-proposed, and human-proposed relationships and their independent validation statuses must remain distinguishable.
8. **Uncertainty remains representable.** Conflict, incompleteness, inaccessibility, and unresolved applicability must not be removed to produce a more fluent answer.
9. **Release remains governed.** Analytical outputs may be restricted, reviewed, qualified, rejected, or approved only for a declared use.
10. **Decision authority remains accountable.** The institution retains responsibility for consequential use of the analytical output.

B.3 EXTERNAL STANDARDS AND TECHNOLOGY BOUNDARY

AIEDSM may draw upon or be implemented with FAIR principles, W3C PROV-O, DCAT, RDF, SHACL, SKOS, C2PA, retrieval-augmented generation, knowledge graphs, GraphRAG, AI agents, MCP, A2A, NIST AI RMF, and other standards or technologies.

Reference to these resources does not imply that AIEDSM replaces them or conforms to them. Formal conformance requires requirements and validation mechanisms defined by the relevant specification or a separately published implementation profile.

APPENDIX C: EVIDENCE CONTEXT ENVELOPE MINIMUM FIELD FRAMEWORK

C.1 PURPOSE, REQUIREMENT LEVELS AND FIELD STATES

This minimum ECE field framework is the AIEDSM-ECE 1.0 conceptual profile. It is implementation-neutral and does not prescribe one serialisation. An implementation may use different names or storage arrangements only if each controlled distinction remains machine-resolvable and the mapping to this profile is documented.

Required means present in every ECE. Conditional means required whenever the field is applicable or material to the declared use. Optional means additional governance detail. Structural completeness does not establish value accuracy, truth, authority or permission.

Field state	Meaning
KNOWN	A usable value is recorded with its attribution basis
UNKNOWN	The value has not been established
NOT_APPLICABLE	The field does not apply under a recorded rule and context
UNAVAILABLE	The value may exist but could not be obtained within the declared search or access boundary
RESTRICTED	The value exists or may exist but cannot be exposed or used under applicable controls
CONFLICTING	Materially incompatible candidate values remain unresolved
INVALID	The supplied value fails the declared structural or domain rule

Every Required or applicable Conditional field must carry a field_state. A blank value is not a substitute for a field state.

C.2 ENVELOPE CONTROL FIELDS

Field	Level	Purpose
ece_id	Required	Persistent identifier for the Evidence Context Envelope
ece_schema_version	Required	Version of the AIEDSM-ECE profile or documented mapping used
ece_created_at	Required	Date and time the envelope was created, with time zone
record_status	Required	Controlled lifecycle status of the ECE record

Field	Level	Purpose
analytical_run_id	Required	Identifier of the analytical run using the envelope
task_id	Required	Identifier of the registered analytical task
task_purpose	Required	Declared purpose for which the object is being used
analytical_cutoff	Required	Temporal boundary governing version, discovery and Known-As-Of treatment
known_as_of_mode	Required	Declared historical or current knowledge-boundary mode
responsible_controller	Required	Human or institutional function responsible for the run
declared_use	Required	Specific use for which analysis may be produced or reviewed
prohibited_uses	Required	Uses that remain outside the authorised boundary
expected_evidence_profile_id	Conditional	Applicable governed EEP identifier and version
revalidation_rule_id	Conditional	Rule defining change-impact and renewed review triggers
next_review_at	Conditional	Scheduled review or expiry point, with time zone

C.3 EVIDENCE OBJECT AND REPRESENTATION IDENTITY FIELDS

Field	Level	Purpose
evidence_object_id	Required	ISA-governed identifier of the Evidence Object
object_binding_method	Required	Method and event through which the ECE resolved to the Evidence Object
canonical_reference	Conditional	Resolvable canonical reference where one lawfully exists
controlled_registry_id	Conditional	Governed internal identifier when no public canonical reference exists
canonical_resolution_state	Required	Resolved, Qualified, Unresolved, Restricted or Failed
object_type	Required	AIEDSM Controlled Taxonomy node class and applicable ISA classification
object_title	Required	Human-readable title from the governed source or registry
issuer_or_origin	Required	Institution, entity, person or governed process responsible for origin
authority_role	Required	Institutional role of the source for the declared task

Field	Level	Purpose
primary_source_reference	Conditional	Resolvable primary-source reference for a derived or secondary object
selected_representation_id	Required	Identifier of the exact file, page, response or representation analysed
representation_observation_id	Required	Identifier of the observation record for the selected representation

C.4 VERSION, PUBLICATION, EFFECTIVE AND KNOWLEDGE-BOUNDARY FIELDS

Field	Level	Purpose
object_version	Required	Version identifier used in the run
native_status_label	Required	Unmodified status label and source-declared basis
source_publication_status	Required	DRAFT, PROPOSED, FINAL, WITHDRAWN, OTHER or UNKNOWN
version_state	Required	CURRENT, SUPERSEDED, HISTORICAL or UNKNOWN
effective_state	Required	NOT_YET_EFFECTIVE, EFFECTIVE, EXPIRED, NOT_APPLICABLE or UNRESOLVED
knowledge_boundary_state	Required	KNOWN_BY_CUTOFF, INSTITUTIONALLY_AVAILABLE_BUT_NOT_SYSTEM_AVAILABLE, INGESTED_AFTER_CUTOFF, NOT_YET_PUBLISHED_AT_CUTOFF or UNKNOWN
institutionally_available_at	Conditional	When the representation became available under the applicable access conditions
system_accessed_at	Required	When the analytical environment first obtained access
system_ingested_at	Required	When the representation entered the analytical store or index
effective_from	Conditional	Beginning of effect under the declared source or authority basis
effective_until	Conditional	End or expiry of effect
observation_or_coverage_period	Conditional	Period described by the Evidence Object
applicability_state	Required	APPLICABLE, CONDITIONALLY_APPLICABLE, NOT_APPLICABLE, DISPUTED or UNRESOLVED
time_precision_and_zone	Required	Precision, time standard, zone or UTC offset and endpoint inclusivity

Field	Level	Purpose
time_uncertainty	Conditional	Estimated, inferred or disputed time and its attribution basis

A Known-As-Of analysis must distinguish institutional availability from system acquisition. Material first ingested after the cut-off must not be silently included in a reconstruction of the system's earlier knowledge, even when the source bears an earlier publication date.

C.5 JURISDICTION, APPLICABILITY AND AUTHORITY FIELDS

Field	Level	Purpose
jurisdiction	Conditional	Legal, regulatory, geographic or organisational boundary
applicability_scope	Required	Entities, activities, periods, claims or subject matter to which the object applies
applicability_basis	Required	Source declaration, rule, institutional determination or proposal supporting treatment
applicability_conditions	Conditional	Thresholds, exceptions, transition rules and exclusions
normative_role	Required	Requirement, guidance, interpretation, response, entity assertion, editorial analysis or another declared role
authority_type	Required	External regulatory, framework-owner, contractual, internal-governance or other authority class
authority_basis	Required	Mandate, source or governance instrument from which authority derives
authority_scope	Required	Purpose, subject, jurisdiction and period within which authority operates
authority_limitations	Conditional	Conditions and exclusions limiting authority or use

C.6 PROVENANCE, INTEGRITY AND FIELD ATTRIBUTION

Field	Level	Purpose
provenance_record_reference	Required	Resolvable source and processing lineage record
retrieved_at	Required	Time the selected representation was retrieved
retrieval_source	Required	Resolver, repository, API or controlled store used

Field	Level	Purpose
representation_reference	Required	Reference to the exact representation analysed
representation_media_type	Required	File or response format
source_last_modified_at	Conditional	Source-declared modification value where available
content_digest	Conditional	Digest and algorithm for the observed representation
snapshot_or_preservation_reference	Conditional	Retained representation where rights permit
integrity_validation_state	Conditional	Technical validation result without truth implication
field_assertion_records	Required	Field-level attribution records for all material assignments

C.6.1 MINIMUM FIELD ASSERTION RECORD

Each material ECE field assignment must resolve to a Field Assertion Record. Global object-level provenance does not replace field-level attribution where a value is assigned, normalised, mapped, derived or proposed.

Field	Minimum requirement
field_assertion_id	Persistent identity for the field assertion
ece_id and ece_version	Envelope identity and applicable version
field_path	Exact machine-readable field name or path
value	Recorded value, omitted only where field_state precludes a value
field_state	One controlled state from C.1
assertion_origin	SOURCE_DECLARED, INSTITUTIONALLY_MAPPED, DETERMINISTICALLY_DERIVED, MACHINE_PROPOSED or HUMAN_PROPOSED
source_reference	Exact Evidence Object, source location, rule, event or attribution basis
asserting_actor_id	Identified source, institution, person, system, method or model
method_or_model_id_and_version	Required for mapped, derived or machine-proposed values
asserted_at	Transaction timestamp with time zone
valid_time	Applicable effective interval or point
validation_status	Current review disposition, separate from assertion origin

Field	Minimum requirement
review_event_id	Required when validation_status is not UNREVIEWED
confidence_or_score	Optional technical signal; never a substitute for origin, validation, authority or field state
supersedes_field_assertion_id	Required where this assertion replaces an earlier assertion

C.7 RELATIONSHIP FIELDS

Field	Level	Purpose
relationship_assertion_id	Conditional	Persistent identity for a task-relevant assertion
assertion_version	Conditional	Version or immutable assertion event used
source_node_id_and_version	Conditional	Originating node and applicable version
target_node_id_and_version	Conditional	Receiving node and applicable version
relationship_type	Conditional	AIEDSM-RV 1.0 predicate or governed namespaced extension
relationship_scope	Conditional	Claim, section, field, entity, period, jurisdiction, purpose or other bounded scope
assertion_origin	Conditional	Immutable origin from the controlled assertion-origin taxonomy
assertion_provenance	Conditional	Asserting source, person, institution, method, model or workflow
evidence_basis	Conditional	Evidence Object locations or reproducible method supporting the assertion
valid_time_and_transaction_time	Conditional	Effective conditions and assertion-record time
authority_basis_and_limit	Conditional	Authority of the asserting actor and its boundary
validation_status	Conditional	Independent review disposition
relationship_lifecycle_status	Conditional	ACTIVE, CONTESTED, SUPERSEDED, RETIRED or WITHDRAWN
conditions_and_exclusions	Conditional	Purpose, direction, version and cases where the assertion does not apply
superseding_assertion_reference	Conditional	Later assertion or event that modifies or replaces this assertion

C.8 RIGHTS, ACCESS AND USE FIELDS

Field	Level	Purpose
access_state	Required	PUBLIC, AUTHENTICATED, RESTRICTED, CONFIDENTIAL or INACCESSIBLE
licence_or_rights_reference	Conditional	Resolvable licence, policy, contract or rights statement
permitted_machine_use	Required	Declared permission for the requested processing
prohibited_machine_use	Required	Expressly excluded processing or release
attribution_requirement	Conditional	Required attribution form and reference
confidentiality_class	Conditional	Applicable confidentiality or sensitivity state
retention_and_deletion_rule	Conditional	Retention, deletion, tombstone and downstream-impact requirements
rights_resolution_state	Required	RESOLVED, RESTRICTED, PROHIBITED or UNRESOLVED

C.9 ANALYTICAL SELECTION, PROCESSING AND REVALIDATION FIELDS

Field	Level	Purpose
selection_state	Required	INCLUDED, INCLUDED_WITH_QUALIFICATION, EXCLUDED or UNRESOLVED
selection_rationale	Required	Rule and reason supporting the selection treatment
exclusion_reason	Conditional	Reason a candidate was excluded
retrieved_fragment_locations	Conditional	Exact locations supplied to an analytical component
processing_event_ids	Conditional	Immutable extraction, conversion, comparison, traversal, classification, translation or generation events
method_model_or_tool_id	Conditional	Identifier of the analytical component
method_model_or_tool_version	Conditional	Applicable component version
material_configuration_reference	Conditional	Configuration, rule, threshold, parameter or prompt-template reference
material_claim_ids	Conditional	Claims that use the ECE
known_limitations	Conditional	Limitations affecting interpretation or use

Field	Level	Purpose
institutional_review_required	Required	Whether output requires Institutional Review before declared use
revalidation_rule_id	Conditional	Applicable change-trigger and impact-propagation rule
change_impact_event_ids	Conditional	Recorded events affecting the ECE or dependent outputs

C.10 ENVELOPE VALIDATION PRINCIPLES

An implementation must test that:

1. the Evidence Object, exact representation and observation record resolve;
2. the selected version and publication state are consistent with the analytical cut-off;
3. institutional availability, system access and system ingestion remain distinguishable;
4. effective state, applicability, time precision and uncertainty are explicit;
5. every material field carries a value or explicit field_state and resolves to a Field Assertion Record;
6. authority and normative role are supported by an attribution basis rather than inferred from semantic content alone;
7. rights and permitted-use conditions authorise the declared processing and release;
8. relationship type, assertion origin, validation status and lifecycle status remain distinct;
9. Machine-Proposed and Human-Proposed values or relationships remain visibly identified after review;
10. transformation events preserve the link to the ECE, Evidence Object, selected representation and source location;
11. unresolved mandatory conditions produce Restricted or Unresolved treatment rather than assumed values; and
12. revalidation triggers propagate to affected claims, TIIRs, reviews and decisions.

Passing these structural tests indicates only that specified records are present and inspectable. It does not establish value accuracy, substantive correctness, authority, assurance, legal compliance or external conformance.

APPENDIX D: TAXONOMY, STATE, LIFECYCLE AND TESTING FRAMEWORK

D.1 AIEDSM RELATIONSHIP VOCABULARY PROFILE 1.0

AIEDSM-RV 1.0 is the canonical predicate set for this publication. Every material assertion records its identifier and version, source and target nodes and versions, controlled type, direction, scope, immutable assertion origin, provenance, evidence basis, valid and transaction time, authority basis and limit, validation status, lifecycle status, conditions and any superseding assertion.

D.1.1 SUPPORT AND PROVENANCE RELATIONSHIPS

Type	Locked meaning	Mandatory control
CITES	Source explicitly references target	Preserve citation locus and version; citation is not support
HAS_PRIMARY_SOURCE	Secondary object identifies a primary source	Preserve characterisation and independently resolve the target where possible
SUPPORTS	Source materially supports an identified claim	Record exact claim, locus, context and limitations; not transitive
QUALIFIES	Source limits, narrows or conditions a claim or relationship	Preserve the qualified proposition and effect
CONTRADICTS	Source is materially incompatible after context checks	Preserve both positions, conflict type and disposition
DERIVED_FROM	Object or result was materially transformed from target	Preserve transformation, inputs, method, version, parameters and time

D.1.2 PUBLICATION AND VERSION RELATIONSHIPS

Type	Locked meaning	Mandatory control
PROPOSES	Draft or consultation introduces content for consideration	Do not treat as final, adopted or effective
RESPONDS_TO	Response addresses a proposal, question or request	Preserve responder, date, scope and status
FINALISES	Final publication completes a declared proposal process	Preserve authority and identify what was and was not finalised

Type	Locked meaning	Mandatory control
REVISION_OF	Newer object revises earlier object	Identify changed scope and continuing historical relevance
PRECEDES	Source is earlier in a declared sequence	Sequence does not imply revision, approval or causation
SUPERSEDES	Source authoritatively replaces target	Require authority basis, affected scope and effective time
AMENDS	Source authoritatively changes identified parts	Identify affected and continuing unaffected provisions
CORRECTS	Source authoritatively corrects target content	Preserve both versions, correction notice and downstream impact
CONSOLIDATES	Source combines identified earlier material	Do not erase component authority or history without a separate basis
WITHDRAWS	Source records cessation of target	Preserve issuer, basis, date, scope and residual or historical use

D.1.3 INSTITUTIONAL INTERPRETATION AND APPLICABILITY RELATIONSHIPS

Type	Locked meaning	Mandatory control
INTERPRETS	Source provides an interpretation of target	Preserve interpreting authority; do not transfer target authority
SUMMARISES	Source abridges target	Preserve omissions, version, actor or method and non-authoritative status
IMPLEMENTS	Source operationalises a requirement or policy	Preserve scope, jurisdiction, institution and effective conditions
APPLIES_TO	Connects object to entity, period, jurisdiction, purpose or use	Preserve basis, conditions and applicability state
REQUIRES	Source declares an evidence element, action or condition	Preserve authority, scope, period and exceptions
DEPENDS_ON	Claim, result or decision materially depends on target	Preserve direction and dependency conditions; not automatically transitive

D.1.4 REPRESENTATION, DISCOVERY AND EQUIVALENCE RELATIONSHIPS

Type	Locked meaning	Mandatory control
TRANSLATION_OF	Source is presented as a translation of target	Preserve language, translator, status and authoritative-language rule
ALTERNATE_REPRESENTATION_OF	Nodes represent one governed institutional object	Require shared identity or authorised mapping; no equal legal effect implied
SEMANTICALLY_SIMILAR_TO	Declared method detects semantic similarity	Discovery only; preserve method, version, score and threshold
RELATES_TO	Weak association supports discovery	Cannot alone establish support, applicability, substitution or equivalence
CORRESPONDS_TO	Directional mapping for a declared analytical purpose	Preserve authority, method, scope, period and exclusions
AUTHORISED_SUBSTITUTE_FOR_OR	Authority permits source to replace target for a defined use	Directional; preserve authority, jurisdiction, period, conditions and revocation
INSTITUTIONALLY_EQUIVALENT_TO	Authority establishes reciprocal equivalence within conditions	Require competent authority; never infer from similarity or correspondence

No relationship is presumed transitive unless a versioned profile expressly permits transitivity for defined node classes and scope. No inverse may be materialised without an explicit inverse rule and provenance. Symmetry must not be inferred from wording. INSTITUTIONALLY_EQUIVALENT_TO is reciprocal only within its authorised conditions; AUTHORISED_SUBSTITUTE_FOR and CORRESPONDS_TO remain directional.

A multi-edge path is a Derived Machine Representation and must preserve every edge, assertion origin, validation status, version, temporal condition and path rule. Earlier suffixed support and contradiction predicates and any unqualified equivalence predicate are retired aliases and must not appear in new records.

D.2 ASSERTION ORIGIN, VALIDATION AND RELATIONSHIP LIFECYCLE

Relationship type states what is asserted. `assertion_origin` states how it entered the analytical environment. `validation_status` states the review disposition. `relationship_lifecycle_status` states whether it remains operative. These dimensions must not be collapsed.

D.2.1 ASSERTION ORIGIN

Code	Display label	Meaning
SOURCE_DECLARED	Source-Declared	Expressly stated by an identified eligible source
INSTITUTIONALLY_MAPPED	Institutionally-Mapped	Established through an identified institution's authorised mapping process for a defined use
DETERMINISTICALLY_DERIVED	Deterministically-Derived	Follows from an explicit, versioned and reproducible rule applied to identified inputs
MACHINE_PROPOSED	Machine-Proposed	Proposed by a probabilistic, semantic, relational or model-mediated process
HUMAN_PROPOSED	Human-Proposed	Proposed by an identified person outside a source declaration or authorised institutional mapping

Assertion origin is immutable. Review, acceptance, repetition or publication in a TIIR must not rewrite MACHINE_PROPOSED or HUMAN_PROPOSED as SOURCE_DECLARED or INSTITUTIONALLY_MAPPED.

D.2.2 VALIDATION STATUS

Code	Meaning
UNREVIEWED	No completed accountable review is recorded
REVIEW_IN_PROGRESS	Review has begun but no completed disposition exists
REVIEWED	Review is complete; this state alone does not authorise use
ACCEPTED_FOR_DECLARED_USE	An authorised reviewer accepted use for the recorded purpose, scope, period and conditions
ACCEPTED_WITH_QUALIFICATION	Use is accepted only with recorded limitations or conditions
DISPUTED	A material challenge remains recorded
REJECTED	The assertion or claim is not accepted for the declared use

D.2.3 RELATIONSHIP LIFECYCLE STATUS

Code	Meaning
ACTIVE	Operative for its declared scope, subject to validation and temporal conditions

Code	Meaning
CONTESTED	A material challenge or competing assertion remains active
SUPERSEDED	Replaced by an identified later assertion for a declared scope and time
RETIRED	Not used for current analysis but preserved for reconstruction
WITHDRAWN	Withdrawn by the asserting authority or accountable controller under recorded conditions

D.3 MULTI-AXIS CLAIM AND TIIR LIFECYCLE FRAMEWORK

The Evidence-to-Inference Boundary is implemented through independent state dimensions. The dimensions are not quality rankings, probabilities of truth or substitutes for institutional judgement.

D.3.1 CLAIM BASIS

Code	Required treatment
DIRECTLY_SUPPORTED	Resolve to exact Evidence Object location, version, status, context and qualification expressly supporting the claim
DETERMINISTICALLY_DERIVED	Record inputs, rules, parameters, units, rounding, exclusions, method version and reproducibility
MACHINE_INFERRED	Record evidence basis, model or method, version, material configuration, limitations and non-source-stated character
COMPOSITE	Use only where reasonable decomposition is unavailable; identify every component basis and preserve inferential or qualified components

Hypothesis is not a claim basis. It is represented through resolution_state = UNRESOLVED and, where useful, intended_use = HYPOTHESIS_FOR_INVESTIGATION.

D.3.2 RESOLUTION STATE

Code	Meaning
SUPPORTED	Evidence and method support the claim within the declared scope; this does not establish truth
QUALIFIED	Support exists with a material limitation, condition, exception or scope restriction

Code	Meaning
CONFLICTED	Applicable evidence or methods remain materially incompatible after context checks
UNRESOLVED	The available eligible basis is insufficient to resolve the claim

D.3.3 VALIDATION AND TEMPORAL DIMENSIONS

Each Material Analytical Claim also records `validation_status` and the applicable `source_publication_status`, `version_state`, `effective_state`, `knowledge_boundary_state` and `applicability_state`. A score must not overwrite any controlled state.

D.3.4 NON-ELEVATION RULE

A deterministic transformation does not elevate the evidentiary condition of its inputs. Where a material input is Machine-Inferred, Qualified, Conflicted, Unresolved, Stale or Pending Revalidation, that condition must remain visible in the resulting claim record. Review may authorise a declared use; it must not rewrite assertion origin, claim basis, source authority or historical state.

D.3.5 TIIR RECORD LIFECYCLE

record_status	Release consequence
DRAFT	Work in progress; not authorised for declared use
PENDING_INSTITUTIONAL_REVIEW	Awaiting required review; not authorised for declared use
RETURNED_FOR_REVISION	Must be revised before renewed review
APPROVED_FOR_DECLARED_USE	Authorised only for the recorded purpose, scope, period and conditions
APPROVED_WITH_QUALIFICATION	Authorised only with visible recorded qualifications
RESTRICTED	Processing, access, reliance or release is limited by a recorded control
STALE	A trigger indicates the analysis may no longer represent the current evidence environment; no new consequential reliance
PENDING_REVALIDATION	Impact assessment or renewed review is under way; no new consequential reliance
EXPIRED	The authorised use period has ended
REJECTED	Not authorised for the declared use

record_status	Release consequence
SUPERSEDED	Replaced by an identified later TIIR for a declared scope
WITHDRAWN	Withdrawn by the responsible institutional authority

Status changes are immutable events. Internal research approval does not authorise public, legal, regulatory, assurance, investment, procurement or other use unless expressly stated.

D.3.6 FAIL-CLOSED GATES

Mandatory gate	Required result when unsatisfied
Rights and permitted use	Affected processing or release must not proceed; record RESTRICTED or the applicable failure state
Identity and provenance	The TIIR must not become APPROVED_FOR_DECLARED_USE
Temporal and applicability	Object or claim must not be presented as current or applicable; qualify, restrict or escalate
Claim classification and lineage	Release for declared use must not proceed
Relationship authority	A proposed relation cannot become source-declared, institutionally mapped, equivalent or substitutable without the required authority and validation record
Expected Evidence Profile	A governed gap conclusion must not be released without a current accepted EEP and declared search boundary
Material conflict	Conflict remains visible; a conclusive dependent claim requires authorised disposition or treatment rule
Institutional Review	Required but incomplete, expired or unauthorised review blocks an approved state
Revalidation	STALE, PENDING_REVALIDATION, EXPIRED, SUPERSEDED or WITHDRAWN TIIR cannot support a new consequential decision

D.3.7 REVALIDATION TRIGGERS

Trigger class	Illustrative event
Object or representation	Revision, amendment, correction, supersession, consolidation, withdrawal or material stable-locator change
Integrity	Digest mismatch or changed validation result
Institutional state	Publication, effectivity, applicability, jurisdiction, authority or version change

Trigger class	Illustrative event
Rights and access	Licence, confidentiality, retention, access or permitted-use change
Relationship	Material assertion disputed, rejected, superseded, retired, withdrawn or reclassified
Evidence environment	New evidence creates material divergence, conflict, qualification or coverage change
Expected Evidence Profile	EEP change, expiry, dispute or withdrawal
Declared task	Question, use, audience, scope, jurisdiction, cut-off or Known-As-Of mode changes
Method or system	Material model, tool, rule, prompt template, configuration, threshold or parameter change
Review governance	Mandate, competence, independence or approval boundary changes or expires
Anomaly	Provenance, security, integrity or processing anomaly is discovered

A trigger must not overwrite the original TIIR. The system preserves the earlier state, records a change-impact event, identifies affected fields, relationships, claims, reviews and decisions and blocks new consequential reliance until governed disposition.

D.4 EVIDENCE DIVERGENCE AND CONFLICT TESTING

Test divergence across semantic, numerical, classification, temporal, version, authority, jurisdictional, relationship, provenance and rights dimensions. Record the signal, affected objects and claims, detection rule or threshold, time, task treatment, disposition actor and closure or conflict reference.

A system applies only institutionally approved significance criteria or escalation thresholds. It does not independently determine legal, accounting, assurance, financial, sustainability or regulatory materiality.

Conflict control	Required test
Context check	Were version, period, jurisdiction, scope, authority, role and applicability compared before conflict was declared?
Source preservation	Do all competing Evidence Objects, exact claims and contexts remain visible?
No forced resolution	Does the system reject frequency, search rank, graph centrality, retrieval order, fluency or confidence as a resolution basis?

Conflict control	Required test
Governed disposition	Does Explained, Authoritative Relationship, Declared Method, Institutional Review or Unresolved treatment record its actor, rule, time and affected claims?
Downstream effect	Do dependent claims and TIIRs preserve the conflict and any changed disposition?

D.5 EXPECTED EVIDENCE PROFILE AND COVERAGE-GAP TESTING

Coverage testing requires a current and accepted Expected Evidence Profile with identity, version, author or controller, authority basis and source, purpose, scope, expected elements, admissible evidence roles, coverage rules, search and access boundary, cut-off, significance criteria, validation status, review, expiry and limitations.

A Machine-Proposed or Human-Proposed EEP cannot support a governed gap conclusion until accepted for the declared use. A missing, expired, disputed or unaccepted EEP requires UNRESOLVED treatment, not a statement of absence or non-compliance.

Coverage state	Meaning
COVERED	An eligible Evidence Object satisfies the declared evidence condition
PARTIALLY_COVERED	Some but not all declared conditions are satisfied
COVERED_WITH_QUALIFICATION	Evidence is available with a material limitation, temporal restriction, conflict or lower institutional role
NOT_LOCATED	No qualifying object was located within the declared boundary; this does not mean that it does not exist
NOT_ACCESSIBLE	An object may exist but could not be accessed under applicable controls
NOT_APPLICABLE	The expectation does not apply within the declared context
CONFLICTED	Available objects cannot support one coverage determination
UNRESOLVED	Evidence or EEP governance is insufficient to classify coverage

Gap types are identity, provenance, version, temporal, authority, relationship, content, coverage-period, rights and required-review gaps. Every gap must cite the EEP element and the search, access, rights, language and time boundary against which it was identified. Density and count do not establish sufficiency.

D.6 AIEDSM METHOD CONTROL TESTS

These reference checks are not certification and do not guarantee correctness.

Test	Operational question
1. Identity and representation resolution	Can every Evidence Object, selected representation and observation be resolved through its recorded identifier or controlled registry?
2. Version and temporal applicability	Do publication, version, effective, knowledge-boundary and applicability states match the cut-off and declared use?
3. Context survival	After retrieval, transformation, agent hand-off and generation, can every material output resolve to its ECE and Evidence Object?
4. Rights and permitted use	Were access, licence, confidentiality, retention, attribution and machine-use conditions resolved before processing and release?
5. Field attribution	Does every material ECE field carry a value or field state and resolve to its Field Assertion Record?
6. Authority Non-Transfer	Does the output preserve source authority without inheriting or expanding it?
7. Relationship provenance	Does every material assertion include type, direction, scope, origin, provenance, version, time, validation and lifecycle state?
8. Graph-path control	Does the engine block unsupported inverse and transitive propagation and reconstruct every material path?
9. Multi-axis claim classification	Does every Material Analytical Claim have claim_basis, resolution_state, validation_status and applicable temporal states?
10. Non-elevation	Do derived and composite claims preserve inferential, qualified, conflicted, unresolved and stale input conditions?
11. Conflict preservation	Are competing objects and consequences preserved without silent forced resolution?
12. Coverage-gap basis	Is every reported gap connected to a current accepted EEP and declared search, access, rights, language and time boundary?
13. Candidate TIIR and review gate	Is pre-review output labelled Candidate TIIR or TIIR Draft, and does required review block approval until complete?
14. Review and decision separation	Are Review Event and Decision Record distinct, attributable and explicit about rationale and mandate?

Test	Operational question
15. Revalidation and impact propagation	Do material changes make affected TIIRs stale or pending revalidation and prevent new consequential reliance?
16. Protocol neutrality	Do evidence identity, context and lineage survive a change of model, retriever, agent, MCP, A2A or transport mechanism?

D.7 TEST INTERPRETATION BOUNDARY

Passing indicates only that the specified control records and behaviours were present and inspectable within the declared scope. It does not establish:

1. factual correctness of every source, field, relationship, claim or output;
2. legal or regulatory compliance;
3. assurance, audit, verification or certification validity;
4. completeness beyond the declared search, access, rights, language and temporal boundary;
5. absence of hallucination, bias, omission, security failure or implementation defect;
6. institutional equivalence, authority or substitutability not separately established;
7. formal conformance with an external standard or technical specification; or
8. external adoption, endorsement or approval of AIEDSM.

An unresolved mandatory control must remain Restricted, Unresolved or Pending Institutional Review. It must not be converted into an assumed value or omitted to permit processing, release or consequential reliance.

APPENDIX E: REFERENCES

E.1 EMJ.LIFE METHODOLOGY FOUNDATIONS

EMJ LIFE HOLDINGS PTE. LTD. (2026). *PADV: Participation-Action-Data-Value ESG Behavioral Data Verification Methodology* (MWP01, Version 3.0).

EMJ LIFE HOLDINGS PTE. LTD. (2026). *NTCC: Non-Tradable Commitment Credit Methodology* (MWP02, Version 3.0).

Yu, Anderson. (2026). *ISA: Institutional Standards Architecture: A Structural Interoperability Methodology for Reusable Evidence Ecosystems* (MWP03, Version 3.0). EMJ LIFE HOLDINGS PTE. LTD. <https://doi.org/10.64969/padv.isa.2026.v3>

MWP03 canonical publication record: Document code MWP03; Version 3.0; Version of Record; published 10 September 2026; 271 pages; publisher EMJ LIFE HOLDINGS PTE. LTD.; Singapore.

MWP03 series: EMJ.LIFE Methodology White Paper Series; Evidence Infrastructure Research Series.

MWP03 author identifier: <https://orcid.org/0009-0002-2161-5808>

MWP03 DOI: [10.64969/padv.isa.2026.v3](https://doi.org/10.64969/padv.isa.2026.v3)

MWP03 canonical landing page: <https://www.emj.life/publication-isa>

MWP03 licence: Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0). <https://creativecommons.org/licenses/by-nd/4.0/>

MWP03 Version-of-Record PDF SHA-256:

ea32924428d8c9339027e9f98767f24950703d1773a41e9eff5dc02fb8ae2848

Yu, A. (2026). *Evidence Infrastructure Analysis 011: When the Reader Is No Longer Human*. Sustainability News Network. <https://www.sustainabilitynewsnetwork.net/article/when-the-reader-is-no-longer-human>

E.2 MACHINE-ACTIONABLE DATA AND KNOWLEDGE REPRESENTATION

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., et al. (2016). *The FAIR Guiding Principles for scientific data management and stewardship*. Scientific Data, 3, Article 160018. <https://doi.org/10.1038/sdata.2016.18>

World Wide Web Consortium. (2014). *RDF 1.1 Concepts and Abstract Syntax*. W3C Recommendation, 25 February 2014. <https://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/>

World Wide Web Consortium. (2009). *SKOS Simple Knowledge Organization System Reference*. W3C Recommendation, 18 August 2009. <https://www.w3.org/TR/2009/REC-skos-reference-20090818/>

World Wide Web Consortium. (2017a). *Shapes Constraint Language (SHACL)*. W3C Recommendation, 20 July 2017. <https://www.w3.org/TR/2017/REC-shacl-20170720/>

E.3 PROVENANCE, CATALOGUING, VERSIONING, TIME, RIGHTS, AND CONTENT CREDENTIALS

World Wide Web Consortium. (2013). *PROV-O: The PROV Ontology*. W3C Recommendation, 30 April 2013. <https://www.w3.org/TR/2013/REC-prov-o-20130430/>

World Wide Web Consortium. (2024). *Data Catalog Vocabulary (DCAT) - Version 3*. W3C Recommendation, 22 August 2024. <https://www.w3.org/TR/2024/REC-vocab-dcat-3-20240822/>

World Wide Web Consortium. (2017b). *Time Ontology in OWL*. W3C Recommendation, 19 October 2017. <https://www.w3.org/TR/2017/REC-owl-time-20171019/>

World Wide Web Consortium. (2018). *ODRL Information Model 2.2*. W3C Recommendation, 15 February 2018. <https://www.w3.org/TR/2018/REC-odrl-model-20180215/>

Coalition for Content Provenance and Authenticity. (2026). *Content Credentials: C2PA Technical Specification* (Version 2.4, April 2026). https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html

E.4 AI RISK GOVERNANCE AND GENERATIVE AI

Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>

Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.600-1>

E.5 RETRIEVAL, GRAPH ANALYSIS, AND TRACEABILITY

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. In Advances in Neural Information Processing Systems 33. <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., & Larson, J. (2024). *From Local to Global: A Graph RAG Approach to Query-Focused Summarization* (arXiv:2404.16130, Version 2 revised 19 February 2025). <https://doi.org/10.48550/arXiv.2404.16130>

Microsoft. (2026). *GraphRAG* (Version 3.1.2) [Computer software]. Released 21 August 2026. <https://github.com/microsoft/graphrag/releases/tag/v3.1.2>

Metropolitansky, D., & Larson, J. (2026). *VeriTrail: Closed-Domain Hallucination Detection with Traceability*. In International Conference on Learning Representations 2026. https://proceedings.iclr.cc/paper_files/paper/2026/hash/42ec51d1031d25eca238ccc67559fe84-Abstract-Conference.html

E.6 AGENT AND TOOL INTEROPERABILITY PROTOCOLS

Model Context Protocol. (2026). *Model Context Protocol Specification* (Protocol revision 2026-07-28). <https://modelcontextprotocol.io/specification/2026-07-28>

A2A Protocol Working Group. (2026). *Agent2Agent (A2A) Protocol Specification* (Protocol version 1.0; repository release v1.0.1). Linux Foundation. <https://github.com/a2aproject/A2A/tree/v1.0.1/specification>

E.7 REFERENCE USE BOUNDARY

The references above identify relevant foundations, standards, specifications, implementation technologies and research. Their inclusion does not indicate endorsement, adoption, certification, compatibility or conformance by the cited authors or institutions. AIEDSM does not claim to replace or normatively extend these resources. It uses them, where appropriate and subject to their licences and governance, as possible building blocks within an implementation-neutral methodology for machine analysis of ISA-defined Evidence Objects.