

Evidence-First Generative AI for Regulated Enterprise Work

A technical and operating framework for responsible intelligence: grounding, refusal, permissions, audit and human accountability.

Published August 2026 | CORIIO Editorial

Executive summary

Enterprise teams do not need another place to copy information. They need a governed way to connect records that use different identifiers, owners, permissions, definitions and clocks. This paper examines responsible intelligence as an operating architecture. The central proposition is that grounding, refusal, permissions, audit and human accountability must remain explicit from source ingestion through retrieval, reasoning, decision and action. When those mechanics disappear, fluent answers become difficult to verify and risky to operationalize.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

The problem beneath the interface

A cross-system question usually begins as an information-reconstruction exercise. Teams locate records, reconcile names, determine which source is authoritative, check whether the information is current and ask who can see it. A conversational interface can hide this work but cannot eliminate it. For responsible intelligence, the product must perform the reconstruction in a way that remains visible to a reviewer. That requires a model of authority and evidence, not only a model of semantic similarity.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Architecture

CORIIO treats connected records as evidence. Each representation keeps a source identifier, version, observed time and lifecycle state. Entity resolution connects shared customers, products, capabilities, services, contracts, incidents and work while retaining ambiguous matches for review. Typed relationships carry provenance and time. The result is a graph that can support retrieval without asserting more certainty than its evidence allows. In responsible intelligence, the governing mechanics are grounding, refusal, permissions, audit and human accountability.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Permission and policy

Authentication identifies the actor; it does not determine every record the actor may use. CORIIO composes organization scope, source-system access, role, attributes, purpose and field policy before retrieving evidence. This pre-retrieval boundary matters because content that enters a model context has already crossed a disclosure boundary. Permission revocation, source deletion and policy change must therefore invalidate or suppress affected evidence throughout indexes, graph relationships, caches and generated artifacts.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Evaluation framework

Evaluation begins with held-out enterprise questions and known evidence sets. A case defines relevant evidence, inaccessible evidence, a known conflict, freshness expectations and an acceptable refusal. The system is scored on evidence coverage, citation validity, permission fidelity, conflict handling, refusal, determinism where required and decision usefulness. For responsible intelligence, the evaluation must also exercise grounding, refusal, permissions, audit and human accountability. A successful demonstration is not sufficient; lifecycle and degraded-state behavior are part of the capability.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Implementation sequence

Start with one bounded question and identify the accountable decision owner. Inventory authoritative sources, identifiers, access boundaries and freshness needs. Configure the minimum connector scopes. Measure initial and incremental completeness. Define entity-resolution and conflict rules. Establish the answer or workflow evidence contract. Exercise permission changes, deletion, throttling and dependency failure. Only then broaden the domain. This sequence keeps responsible intelligence tied to an observable operating outcome rather than an open-ended data program.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Governance and accountability

Governance should be implemented as product behavior. Quality signals change retrieval ranking. Policy can refuse a query or require review. Tool grants and workflow approvals constrain action. Audit records retain actors, organization, evidence and correlation history. Humans retain strategic and irreversible authority. The aim is not to remove judgment; it is to give judgment a clearer evidence base and a durable decision record.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Architecture review checklist

Reviewers should ask: Which system owns each claim? How is identity resolved? When does a record become stale? Can inaccessible evidence enter ranking? How is conflict represented? Which calculations are deterministic? What operation is reversible? Who approves the consequence? How does the system recover? What evidence proves the outcome? For responsible intelligence, answers should explicitly cover grounding, refusal, permissions, audit and human accountability.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?

Conclusion

World-class enterprise intelligence is not defined by how confidently a model writes. It is defined by whether the organization can understand why an answer exists, which evidence supports it, which boundaries constrained it and who remains accountable for what happens next. A durable approach to responsible intelligence makes grounding, refusal, permissions, audit and human accountability inspectable and testable. That is the standard CORIIO is designed to meet.

Practical questions

Authority	Which source and owner can approve this information?
Evidence	Can each material conclusion be traced to a current source fragment?
Boundary	Which organization, role, source ACL or policy restricts use?
Failure	What remains safe and explainable when a dependency is unavailable?