RJL SOLUTIONS

CONTEXTUAL GOVERNANCE

Context Before Query

Contextual governance for enterprise AI: business meaning before the call, enforceable policy at the data.

CLEAR OWNERSHIP · REPEATABLE CONTROLS · OPERATING EVIDENCE

Richard White
RJL Solutions · rjlsol.com

ARCHITECTURE UPDATE
29 OCTOBER 2026

01 / THE EXECUTIVE CASE

An authorized answer can still be wrong.

Access control won't tell an agent which date, unit, definition or version serves a business question.

A person who's unsure about a field can ask its owner. An agent can ask for clarification or escalate too. Nobody should rely on it to infer missing business meaning, though. A query that looks sensible can mix preliminary and final values, join incompatible identifiers or use the wrong currency.

QUESTION	CONTEXT NEEDED	FAILURE WITHOUT IT
What does it mean?	Definitions, grain, dates, units and valid relationships.	A plausible calculation answers a different business question.
May it be used here?	Approved purpose, rights and policy that the runtime can enforce.	Authorized data is used outside its permitted task or scope.
Is it fit to use now?	Product version, completeness, freshness and finality.	A current-looking answer uses stale or provisional evidence.

The design is straightforward. Get business context before the data call. Require that context in the runtime. Enforce policy at the serving service, and keep evidence linking the answer to what was used. Catalog metadata improves interpretation. It doesn't grant data access.

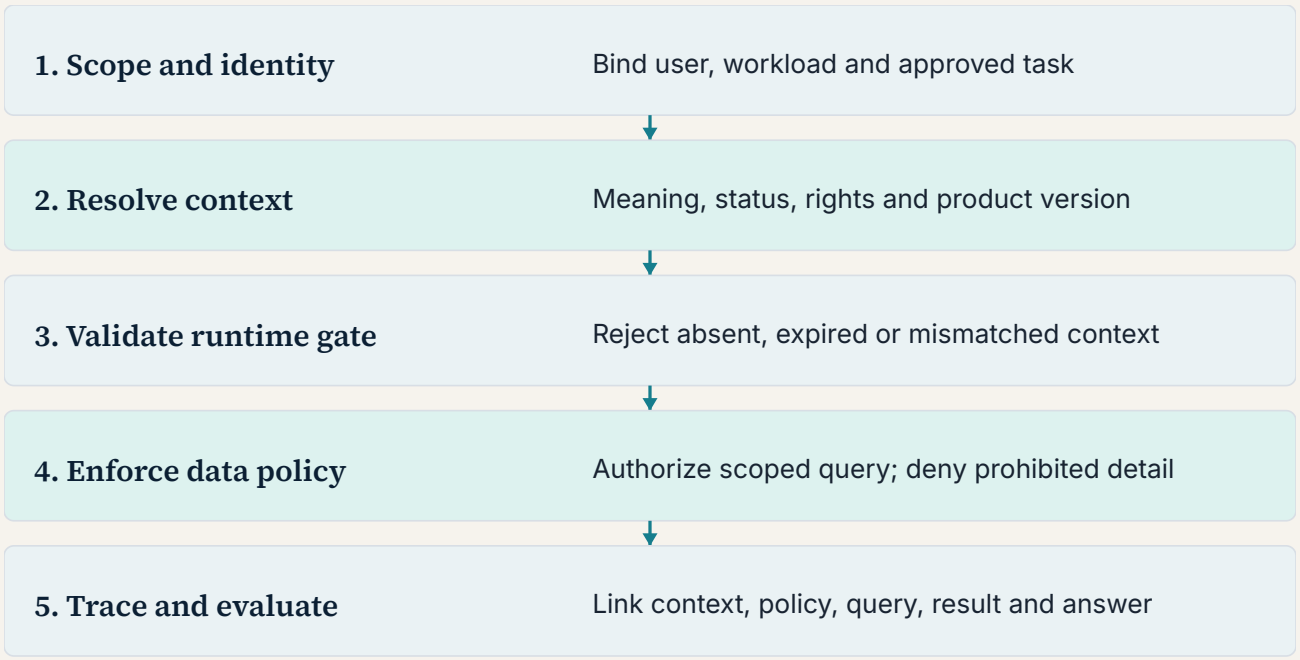
THE CONTRIBUTION Connect catalog context, scoped tools, data policy and operational evidence into one enforceable path. Protocol connectivity on its own doesn't give you this operating design.

NIST AI risk guidance and OWASP security guidance inform the surrounding evaluation and controls [1, 2, 4]. The architecture and readiness tests in this paper are my implementation recommendations.

02 / THE ACCESS PATH

Make context a runtime prerequisite.

The application and tools have to enforce the sequence. Leaving it to the model's willingness to follow an instruction won't do.



Illustrative control path. Catalog context informs the request; the data service remains the enforcement point for that request.

The runtime records an approved task, user scope and workload identity. It resolves current catalog context and binds it to the product version, task and authorization scope. Before executing the call, a data tool either accepts a validated reference to that context or fetches the context itself. Nothing else gets through.

Trusted runtime state or a protected token can back the context reference. Validate it server-side. Check its expiry, and check that the product and version match. Never accept free-form metadata, or a purpose string from the model, as evidence of authorization.

NEGATIVE TEST — The tool denies a direct call whose context is absent, expired, altered or mismatched. If the policy or required context can't be evaluated, stop that operation and offer an approved fallback or escalation.

03 / CONTEXT CONTRACT

Publish meaning with the product.

Owners and stewards maintain business context; engineering makes its version and status available at the point of use.

CONTEXT FIELD	ILLUSTRATIVE RESERVE-PRODUCT CONTENT
Meaning and scope	Case reserves for specified disability claims; state whether IBNR is excluded.
Grain, dates and units	Product-line aggregate; valuation date defines the quarter; valid time and recorded time are kept apart; currency is specified.
Valid relationships	Approved claim/policy keys; joins and cardinality with known limitations.
Classification	Sensitivity of diagnosis, claimant identifiers and other fields.
Terms of use	Approved task and permitted purpose; consent, licensing and retention constraints.
Owner and contact	Accountable business owner, steward and incident contact.
Quality and freshness	Source completeness, valuation version, final/preliminary status and stale-data rule.
Guidance	Use final valuations for this comparison; return product-line aggregates only.

Treat context changes as product changes. A revised definition or permitted purpose can change an agent's behavior even when the schema stays the same. Apply versioning, accountable review and consumer-impact analysis.

Generate freshness and quality status from monitored evidence where you can. Keep authored guidance separate from measured status. Owners have to review ambiguous business language; a classification scan alone can't establish the intended use.

ADVISORY VERSUS MANDATORY	The context describes allowed behavior. Where a condition is mandatory, identify the runtime or serving control behind it, and test that control. A sentence in metadata is not a policy implementation.

04 / WORKED REQUEST

Compare the right reserve valuations.

Illustrative task: an actuarial analyst asks how disability case reserves changed during the third quarter.

STEP	BEHAVIOR AND EVIDENCE
Define the task	Confirm scope, product line, currency and whether the question concerns case reserves or a broader liability measure.
Resolve context	Find the certified product; obtain final June and September valuation versions, approved meaning and use restrictions.
Validate the call	Runtime checks user/workload scope, context version and the approved aggregate tool.
Enforce	Serving service allows only approved aggregates and denies claimant-level detail for this task.
Answer	Compare the two final valuations; cite product version, valuation dates, scope and limitations.
Record	Link context, policy decision, query, result version and answer under one request ID.

Suppose September is preliminary. The assistant must not silently mix it with a final June value and call the result a final quarter movement. It should explain the status and follow the product's rule for permitted provisional use or escalation.

Name the clock as well. "As of 30 September" can mean what was true on that date (valid time) or what the platform had recorded by then (recorded time). A late-reported claim changes the first after the fact. An audit or a restated report may need the second. The context carries both, and the answer states which one it used.

BANKING VARIATION	For commercial fee income, context defines segment scope, approved joins and local versus reporting currency. The aggregate interface blocks client-level output. Normalized amounts and product versions support a reproducible comparison.

These examples illustrate a control design. They describe no client deployment, achieved accuracy or regulatory approval.

05 / DATA ENFORCEMENT

Enforce the restriction you actually state.

Masking fields and forbidding detail rows are different requirements.

REQUIREMENT	EXAMPLE ENFORCEABLE CONTROL
Only authorized identity and scope	Server-side user/workload authorization; least-privilege credentials and tenant isolation.
No claimant-level rows for this task	Aggregate-only view or narrowly scoped query tool; prohibited detail calls denied.
Sensitive fields protected	Appropriate field controls such as masking or projection restrictions.
Use limited to an approved purpose	Trusted workflow/task authorization linked to the serving decision; permitted-use constraints.
Final and current version required	Version/status checks at call time; stale or provisional data handled under explicit policy.

An aggregation policy with minimum group sizes can reduce disclosure of individual records. When one person has several records, you may need entity-aware controls. These controls have limits. Repeated queries and other information can still create inference risks. Evaluate the actual platform, and monitor for misuse [3].

The platform should deny a forbidden detail request. Returning masked detail and claiming to have enforced aggregation-only use is something else entirely. Another authorized task may still need masking, but masking addresses a different policy.

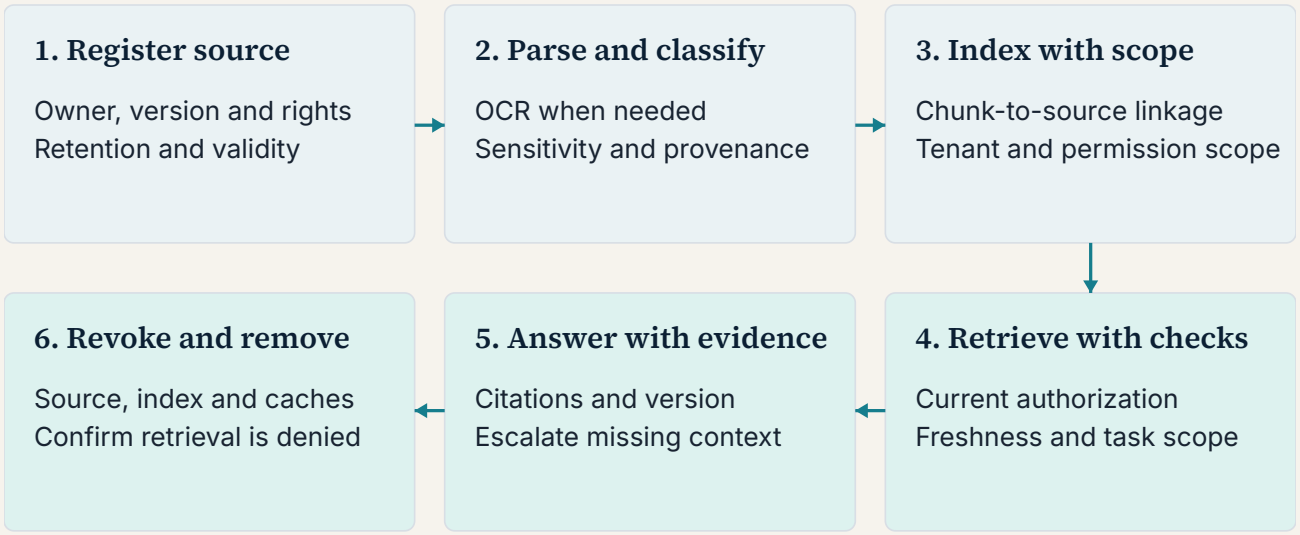
Use one policy authority with consistent enforcement across warehouse, API, source-query and retrieval paths. Test the alternate routes. A catalog entry or a successful login doesn't authorize every use or every field.

BOUND THE SAFETY CLAIM — Controls can block defined prohibited operations. They can't guarantee factual correctness, remove every inference risk or prevent all misuse of permitted output.

06 / KNOWLEDGE CONTENT

Keep document rights alive through retrieval.

A procedure, transcript or scanned document needs a versioned source and current authorization. An embedding alone isn't enough.



Illustrative knowledge lifecycle. Revocation affects future retrieval and cached serving paths.

Register ownership, provenance, validity dates and use restrictions. Parse the material, using OCR or transcription where appropriate, and validate important extracted content. Keep each indexed chunk linked to its source version and sensitivity.

At retrieval time, check current user and task authorization. Filter by tenant, document scope, validity and source status before returning anything. A right established at ingestion can be revoked later [4].

Carry source retirement and permission changes through to indexes, retrieval caches and answer-serving caches. Define a measurable revocation commitment, and test it. Revocation can't take back information already disclosed, so limit unnecessary retention of, and access to, conversation history.

KNOWLEDGE IS NOT AN INSTRUCTION CHANNEL

Treat retrieved material as untrusted input. Restrict tool authority, separate trusted configuration from content and test prompt-injection scenarios. No document should ever be able to grant a new privilege.

07 / FAILURE BEHAVIOR

Decide how the system stops or degrades.

Failure handling is part of the product promise. It isn't an exception to it.

FAILURE	REQUIRED BEHAVIOR
Context missing or expired	Deny the affected data call; refresh trusted context or escalate.
Policy or identity unavailable	Do not proceed without the required decision; use only a separately approved fallback.
Unclassified or unreviewed content	Restrict the affected asset until its required status and rights are established.
Source or product stale	Warn, restrict or stop according to the use-specific stale-data policy.
Definition or schema changes	Invalidate affected context; run impact review and release tests before supported use.
Prompt injection in content	Prevent privilege expansion; constrain tool calls and evaluate the attempted behavior.
Permission revoked	Deny future access; invalidate affected caches and test all supported routes.

"Fail closed" refers to a particular operation and a required control. It doesn't mean the whole enterprise goes dark. An application can offer a safe alternative, such as a manual review path, without bypassing authorization.

If cached context is permitted, bind it to identity, task scope and product version. Define its expiry, and invalidate it after relevant changes. A cache is an optimization. It's no evidence that authorization still holds.

TEST DEGRADED CONDITIONS	Run the failure cases before release and after material changes. Capture the denial, alert, user message and operator response under the same evidence model as a successful request.

08 / OWNERSHIP AND EVIDENCE

Give every control an operator.

Catalog, platform, application and business responsibilities must meet at the workflow boundary.

RESPONSIBILITY	ACCOUNTABLE DELIVERY ROLE
Meaning and intended use	Business data owner approves definitions and permitted business scope.
Metadata and guidance	Steward maintains reviewed context and issue routing.
Data and retrieval enforcement	Platform/serving owner implements and tests policy.
Runtime and tools	Application owner enforces task scope, context prerequisites and approvals.
Security monitoring	Security operations reviews relevant alerts and misuse evidence.
Release and risk review	Named authority approves risk-appropriate evidence; independent challenge where required.

Trace the task, product and version, context version, authorization decision, tool inputs, query or retrieval, result references and answer. Record enough to reconstruct a material outcome. Apply data minimization and retention rules. Accountability doesn't require logging every sensitive value or credential.

Measure the operating behavior

- Coverage of critical products with reviewed context and current quality status.
- Denied bypass, out-of-scope and revoked-access requests in controlled tests.
- Revocation propagation time and stale-context incidents.
- Incorrect interpretation, unsupported assertions and escalation on representative tasks.
- Time and effort to reconstruct an answer after a material issue.

A high context-lookup rate is useful evidence that people follow the process. It doesn't prove correct interpretation. Review outcomes and control tests together.

09 / READINESS

Require evidence before the first agent release.

The first scope can be small. Its controls still have to work.

READINESS QUESTION	EVIDENCE
Is the use bounded?	Named users, purpose, output, action limits and responsible business owner.
Is context usable?	Reviewed definitions, valid joins, product version and current status.
Can bypass occur?	Absent, altered, expired and mismatched context calls are denied.
Do rights hold?	Positive and negative authorization tests across tenant and serving paths.
Can the answer be evaluated?	Representative cases, expected evidence, ambiguity and failure tests.
Can the service be operated?	Monitoring, incident routing, restricted/stop mode and manual fallback.

Put period-end changes, missing sources, unusual joins and revoked permissions in the evaluation set. For an aggregate-only task, actively try to get detail rows. For a document task, try retired content and prompt-injection payloads.

Reassess any change to tools, models, prompts, context contracts or access rules. An earlier test result applies to the system and scope you tested. It isn’t permanent assurance for later versions.

THE RELEASE DECISION	Approve the actual workflow and its operating envelope. A generic AI policy, a protocol connection or a catalog completeness score can’t stand in for use-specific evidence.

10 / FIRST IMPLEMENTATION

Start with one question worth answering.

The first increment should prove context, enforcement and operation together.

FOCUS	WORK TO COMPLETE
Frame	Choose one decision or workflow; name the owner, users and outcome. Identify the products and document sources.
Prove	Write the context contract, narrow the tools and test representative questions plus prohibited requests.
Build	Implement the runtime gate, serving controls, version checks, traces and operational fallback.
Run and transfer	Monitor outcomes and failures, resolve metadata gaps and transfer accountable operation with a runbook.

Use existing catalog, identity and platform services where they meet the requirement. You don’t need a new vendor first. MCP can expose metadata and tools [5], but the access path still needs trusted authorization, context validation and release evidence.

Agree timing only after checking access, data condition, available capacity and risk review. A first-quarter plan can help sequence work. It shouldn’t imply every enterprise can release any agent in ninety days.

Extend when the evidence supports it

- A consumer understands the output and its limitations.
- The required restrictions hold under deliberate negative tests.
- Operators can detect, investigate, restrict and recover the service.
- Another workflow can reuse the pattern without bypassing its controls.

A COHERENT ENTERPRISE PATTERN

Business meaning improves interpretation. Runtime gates enforce the workflow. Data services enforce access and use restrictions. And operational evidence shows how the system really behaves.

ARCHITECTURE UPDATE / CONTEXT ASSEMBLY

Where the context record comes from.

A context record is a bounded runtime record assembled from trusted facts, declared intent and versioned product meaning.

FIELD / PROVENANCE	VALIDATION AND USE
Subject + workload / identity service	Verify authentication, tenant and acting service. Resolve an employee ID only where necessary; use an internal reference instead of exposing it to the model.
Function + legal entity / HR or directory	Use authoritative attributes with timestamps. Department informs interpretation; it does not automatically grant access.
Task + target / user and workflow	Confirm the invoice, claim or reporting scope. Distinguish self-service from acting for another person. User assertions are not entitlement evidence.
Meaning / product and semantic services	Bind product, metric, ontology mapping and source versions; record grain, dates and freshness requirements.
Authority / policy service	Evaluate current entitlements, purpose and delegation. Record the decision reference, scope and expiry. Recheck at sensitive tool calls.
Conversation / user-provided or inferred	Treat tone and preferences as interaction hints. They cannot establish identity, eligibility, urgency or permission.

The runtime creates the record after sign-in and task clarification, resolves the necessary attributes, then validates it server-side. Store the sensitive record in controlled runtime state; pass tools an opaque reference or validated token with audience, integrity and expiry checks. Refresh on expiry, role change, target change or a new material action.

If identity or a mandatory authorization fact is missing, stop the affected operation. Ask a clarifying question for ambiguous intent; use an approved fallback for unavailable sources. Never let the model fill a missing entitlement with a guess.

Design references checked 29 October 2026: Data mesh principles · Document-level access control · RAG security. These additions are RJL implementation recommendations, not a vendor certification or a compliance determination.

Sources and scope.

Public frameworks inform the design. The implementation choices, checklists and examples are RJL recommendations, drawn from Richard White's practice.

[1] *AI Risk Management Framework 1.0*. NIST AI 100-1, January 2023. Voluntary framework for use-specific AI governance and risk management. nist.gov / AI RMF 1.0

[2] *Generative AI Profile*. NIST AI 600-1, July 2024. Referenced for information integrity, privacy and security risk evaluation. nist.gov / Generative AI Profile

[3] *Aggregation policies*. Snowflake documentation. Illustrates controls separate from masking, with minimum-group and inference limitations. No vendor endorsement is implied. docs.snowflake.com / aggregation policies

[4] *RAG Security Cheat Sheet*. OWASP. Retrieval-time authorization, content isolation and cache/permission-change handling inform the design. cheatsheetseries.owasp.org / RAG security

[5] *Model Context Protocol*. Official specification. Connectivity is distinct from the application and serving controls described here. modelcontextprotocol.io / specification

Sources reviewed 10 October 2026; each is linked. Examples are illustrative; proposed measures and service commitments need an agreed baseline and risk assessment. This paper does not establish regulatory compliance or reproduce a licensed assessment standard. Companion papers at rjlsol.com/whitepapers.

Read the series

PAPER	WHAT IT COVERS
The Data and AI Operating Model	Ownership, teams, decision rights and funding.
Data Products in Practice	Boundaries, contracts, quality and certification.
The Federated Data Platform	Shared services, access paths, resilience and agents.
Context Before Query (this paper)	Context, authorization and enforceable policy for AI.
AI Governance Evidence	For regulated firms: what bank and insurance reviewers need to inspect.

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, and Run & Transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

Bring one proposed agent workflow and the products it would use. RJL Solutions helps define the context contract, enforcement path, failure behavior and evidence required before release.

OPERATING MODEL **The Data and AI Operating Model**
Ownership, teams, decision rights and funding.

PRODUCTS **Data Products in Practice**
Boundaries, contracts, quality and certification.

PLATFORM **The Federated Data Platform**
Shared services, access paths, resilience and agents.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

Examples are illustrative and describe no client engagement or client result. Measures, timings and service commitments are starting points that need an agreed baseline. Apply the design to your business purpose, existing controls and any obligations that apply to you.