# KURIOM

## Public Submission

### Automated Decision-Making Transparency Obligation (APP 1)

*Issues Paper · Office of the Australian Information Commissioner*

| Field | Detail |
|---|---|
| Submitted by | [redacted] Kuriom |
| Contact | [redacted] kuriom.ai |
| Submission date | June 2026 |
| Submitted to | privacyreformtranche1@oaic.gov.au |
| Consultation deadline | Monday 15 June 2026 |
| Issues Paper | Automated Decision-Making Transparency Obligation (APP 1) · 18 May 2026 |
| Architecture reference | SSRN:6782100 (submitted May 2026) |
| Patent applications | 64/008,411 · 64/043,572 · 64/049,750 · 64/059,101 |

*Sequence Knowledge Authorization · AI · Patent Pending*

# 1. Submitter Background

Kuriom is a pre-execution sequence authorization platform for AI agents. Kuriom evaluates the complete intended sequence of AI agent actions as an indivisible unit against an organisation's validated knowledge foundation before any action executes. The evaluation is deterministic and binary: authorized or withheld. The result is recorded in an immutable ledger before any consequence occurs.

Kuriom is AI-model agnostic and has no commercial agreements with any AI agent developer or model provider. The governance evidence Kuriom produces is generated by a system with no financial relationship with the AI agent it governs.

The Kuriom architecture is documented in a working paper submitted to SSRN (SSRN:6782100, submitted May 2026) and covered by four patent applications:

- 64/008,411 (March 17, 2026) — Sequence Authorization Unit, Governance Authorization Token, Knowledge Governance Immutability Framework
- 64/043,572 (April 19, 2026) — Spawning chain governance, agent identity verification
- 64/049,750 (April 26, 2026) — Byzantine fault-tolerant independent auditor architecture
- 64/059,101 (May 6, 2026) — Complete lifecycle governance including instruction source validation, context source anomaly detection, knowledge-grounded runtime rule generation, and execution evidence token

The production codebase comprises nine named components, eleven binaries, and sixteen validated scenarios running against live AWS DynamoDB and Neo4j AuraDB infrastructure. The knowledge_state_id — the cryptographic fingerprint identifying the exact knowledge foundation that grounded an AI agent's action — is proved injective in Lean 4 with zero sorry instances. This is a mathematical guarantee, not a heuristic.

Kuriom submits this response because the four definition questions in the Issues Paper — the meaning of computer program, substantially and directly related, significantly affect rights or interests, and making a decision — have precise architectural answers at the pre-execution authorization layer that the OAIC's guidance development will benefit from.

# 2. The Architectural Gap This Submission Addresses

The ADM obligation under APP 1.7 requires disclosure of the kinds of personal information used in, and the kinds of decisions made by, computer programs that significantly affect individual rights or interests. The obligation is correctly framed as a transparency requirement.

This submission draws the OAIC's attention to a structural gap beneath the transparency obligation: most current ADM governance frameworks evaluate what the computer program did. They do not evaluate what the computer program knew — and whether that knowledge was verified, current, and within authorized scope — at the moment it acted.

> The distinction is architecturally significant. An AI agent making a credit decision, a clinical recommendation, or an insurance determination may be acting on knowledge that is stale, from an unverified source, or outside the scope the organisation authorized it to use. The decision may appear to be within policy. The knowledge foundation that grounded it was not evaluated before the decision executed.

The APP 1.7 transparency obligation requires disclosure of what the computer program did. The harder governance question — and the one this submission addresses — is whether the knowledge the computer program relied on was authorized before it acted. These are different obligations at different layers. Both are necessary.

# 3. Question 1 — Meaning of Substantially and Directly Related

> **Question 1: What are the relevant factors to enable an entity to assess whether a computer program substantially facilitates and is directly connected to a human decision-maker's decision?**

The Issues Paper lists five candidate factors: degree of reliance, ability and likelihood of human override, nature of the output (advisory vs determinative), transparency and explainability, and integration into decision-making workflow.

Kuriom's pre-execution architecture provides a sixth factor that the OAIC should include in its guidance:

### Factor 6 — Knowledge Foundation Reliance

A computer program is substantially related to making a decision when the knowledge it relies on is the primary input to that decision — regardless of whether a human makes the final determination. The degree to which the human decision-maker relies on knowledge generated, retrieved, or processed by the computer program is a direct measure of substantiality.

This factor is architecturally distinct from the five listed in the Issues Paper because it operates at the knowledge layer rather than the action layer. A computer program may generate an output that a human formally approves — satisfying human oversight requirements — while the human's actual decision is entirely dependent on knowledge the computer program assembled. The substantiality is in the knowledge dependency, not the formal approval step.

The NACIA edge case in the Issues Paper illustrates this precisely. GPTea summarises candidate profiles and recommends eligibility. NACIA staff always make the final decision. But if the human decision-maker would not form a view without GPTea's recommendation — and the recommendation is grounded in patient or applicant knowledge that GPTea assembled — the computer program is substantially related to the decision through its knowledge contribution, not only through its output.

**Ranking of factors — Kuriom's recommended order:**

| Rank | Factor | Rationale |
|---|---|---|
| 1 | Degree of reliance on ADM output | Primary measure of substantiality — if the human would not form a view without the output, the program is substantially related |
| 2 | Knowledge foundation reliance (new) | The knowledge the program assembled is the substance of the decision — even where a human formally approves |
| 3 | Integration of ADM into decision-making workflow | Structural embeddedness determines whether bypass is operationally possible |
| 4 | Nature of output — advisory vs determinative | Determinative outputs have higher substantiality but advisory outputs with high reliance scores are equally substantial in practice |
| 5 | Transparency and explainability of outputs | Affects accountability but does not determine substantiality |
| 6 | Ability and likelihood of human override | Override capability is necessary but not sufficient — overrides that never occur in practice do not reduce substantiality |

## 4. Questions 2 and 3 — The NACIA Edge Case

> **Question 2: Would you consider NACIA's described uses of GPTea to be substantially and directly related to making decisions? Question 3: Can you provide your reasons?**

Answer to Question 2: Yes.

GPTea is substantially and directly related to NACIA's eligibility decisions for the following reasons.

### Reason 1 — Knowledge assembly is the decision substance

GPTea summarises candidate profiles. The summary is not a neutral aggregation — it selects, weights, and presents information from the applicant's record in a form that shapes what the human decision-maker attends to. The human decision-maker's view of the applicant is mediated entirely by GPTea's summary. The substance of the decision — what information was considered and in what form — is determined by GPTea before any human review begins.

### Reason 2 — The Composition Gap applies

NACIA staff are authorized to make individual eligibility determinations. GPTea is authorized to summarise profiles. But the combination — a summary that shapes a determination affecting an individual's access to aged-care support — was never evaluated as a unit before it executed. Individual authorizations do not constitute authorization of the sequence. The combination must be evaluated as an indivisible unit.

### Reason 3 — Pre-execution authorization is the missing governance layer

NACIA's governance framework, as described, ensures human oversight of the final determination. It does not verify, before GPTea generates a summary, whether the knowledge GPTea is drawing on is current, complete, and from a verified source. A summary grounded in stale or inaccurate data shapes a human determination that may itself be accurate — but the governance gap is upstream of the human review, not at the review itself.

> For the OAIC's guidance: the ADM transparency obligation should address not only what the computer program did — the output — but what knowledge the computer program relied on to produce that output. Disclosure of the kinds of personal information used in the operation of such computer programs (APP 1.8(a)) should extend to the provenance, currency, and verification status of that information at the moment the program acted on it.

## 5. Questions 4 and 6 — Significance Threshold and Architectural Definition

> **Question 4: What factors increase the likelihood that a decision could affect individual rights or interests?**
> **Question 6: What constitutes a significant service or support?**

Kuriom's Constitutional Prohibition Framework (CPF) provides an architectural definition of the significance threshold that the OAIC's guidance may find useful.

### The CPF significance threshold

A decision significantly affects individual rights or interests when:

- It crosses a Mode 1 hard block — an action that is constitutionally prohibited regardless of context. Examples: processing a transaction against a sanctioned counterparty, issuing a clinical recommendation without required expert review, making a credit determination without current regulatory compliance data.
- It triggers a Mode 2 composite conditional block — individually authorized actions combining to produce an outcome the organization never authorized. The combination crosses the significance threshold even when each individual action does not.

For the OAIC's guidance, this architectural framing has a practical application: the significance threshold should be assessed at the sequence level, not the individual action level. An AI agent making three individually authorized determinations — query record, assess eligibility, update status — may produce a combination that significantly affects an individual's access to a service, even though no single step crosses the threshold alone.

### Factors that increase significance — ranked

| Rank | Factor | Architectural basis |
|---|---|---|
| 1 | Financial outcomes | Direct impact on rights under |

| | | |
|---|---|---|
| | | contract (APP 1.9(d)(ii)) — credit, insurance, pricing |
| 2 | Access to significant services | Healthcare, aged care, housing, employment — access determination is the primary significance trigger |
| 3 | Sensitive information | Disability, health, financial hardship — higher significance where vulnerable attributes are processed |
| 4 | Sequential action combinations | Individually authorized steps combining to produce unauthorized outcomes — the Composition Gap |
| 5 | Vulnerable persons | Children, persons experiencing vulnerability — significance threshold lower as per EU GDPR guidance cited in Issues Paper |
| 6 | Intrusive practices | Behavioral profiling, differential pricing, targeted exclusion |

### Significant services — recommended scope

Kuriom recommends the OAIC define significant services to include, at minimum:

- Healthcare and aged care — clinical decisions, care eligibility, medication management
- Financial services — credit assessment, insurance underwriting, access to banking
- Employment — hiring, promotion, performance assessment, termination
- Government benefits — housing, social support, immigration, licensing
- Critical utilities — access to telecommunications, essential utilities, digital identity

## 6. Question 11 — Meaning of Arranged For

> **Question 11: Are there any scenarios relevant to the meaning of 'arranged for' that you consider require guidance?**

Yes. The following scenario requires guidance and is directly addressable from Kuriom's architecture.

### Scenario — Third-party AI knowledge foundation

An entity procures a third-party AI system to assist with eligibility assessments. The entity configures the system's permitted actions and approves its deployment. The third-party vendor maintains and updates the knowledge the AI system draws on — regulatory data, eligibility criteria, clinical protocols. The entity has not verified whether that knowledge is current, complete, or from an authorized source at the time each determination is made.

The entity has arranged for the computer program to make decisions. But the knowledge foundation the program relies on is controlled by the vendor, not the entity. The entity cannot currently answer the question: was the knowledge my AI agent used authorized and current at the moment it made this determination?

This scenario requires guidance on three points:

### Point 1 — Knowledge source authorization is part of 'arranged for'

An entity that arranges for a computer program to make decisions must also arrange for the knowledge that computer program relies on to be verified and within authorized scope. Arranging for the program without arranging for its knowledge foundation does not constitute adequate governance of the ADM.

### Point 2 — The pre-execution record is the evidence standard

The APP 1.7 disclosure obligation requires entities to describe the kinds of personal information used in the operation of such computer programs. This is a description of categories. It does not require — but the OAIC's guidance should encourage — a per-instance record of what specific information was used, from what source, at what moment, with what verification status. The per-instance record is the evidence standard that makes the category disclosure meaningful.

Kuriom's Governance Authorization Token (GAT) is that per-instance record. It is cryptographically signed, written to an immutable ledger before any action executes, and identifies the exact knowledge_state_id — the fingerprint of every knowledge node active and validated at the moment of authorization. This fingerprint is proved injective in Lean 4: no two different knowledge states can produce the same fingerprint. The record is mathematically tamper-resistant, not just procedurally documented.

### Point 3 — Non-human actor identity requires separate treatment

Current IAM frameworks are designed for human actors. AI agents operating under human-delegated authority are not humans. The entity that arranged for the AI program must be able to demonstrate that the agent's identity, pathway, and scope were registered and authorized before each sequence executed — not just that a human approved the deployment at the outset.

The Issues Paper notes that 'arranged for' recognizes that one entity may operate the program while another is responsible for arranging for it to make or assist a decision. The guidance should extend this to the knowledge layer: the entity responsible for the arrangement bears responsibility for the currency and authorization of the knowledge foundation the program relies on — not only for the program's output.

## 7. The Mathematical Proof Basis for Pre-Execution Evidence

Kuriom's submission draws on a formal mathematical proof that has direct relevance to the evidentiary standard the OAIC should consider for ADM governance.

### The knowledge_state_id injectivity proof

The knowledge_state_id is a cryptographic fingerprint computed as the sum of 2^node_id for every knowledge node active and validated at the moment of pre-execution authorization. This fingerprint identifies exactly which knowledge nodes grounded the AI agent's action.

The injectivity of knowledge_state_id is proved in Lean 4 with zero sorry instances. Injectivity means: no two different knowledge states can produce the same fingerprint. Formally:

> $\forall\ S_1\ S_2$ : KnowledgeState, knowledge_state_id($S_1$) = knowledge_state_id($S_2$) → $S_1 = S_2$ This is a mathematical guarantee. It is not probabilistic. It cannot be forged by constructing a different knowledge state that produces the same fingerprint.

For the OAIC's guidance development, this proof has three practical implications:

- The pre-execution record is independently verifiable. Any party — including the individual affected by the decision, a regulator, or a court — can verify the fingerprint against the knowledge state without contacting Kuriom. The evidence does not rely on the issuer's integrity.
- The record predates the decision. The GAT is written to the KGIF immutable ledger before any action executes. This is not reconstruction after the fact. It is evidence that existed before the decision was made. This is the standard the OAIC's guidance should set for pre-execution ADM governance.
- The record is tamper-resistant without relying on trust. A rolling SHA-256 chain hash across all KGIF events means any modification to any record in the chain is detectable by any verifier. The integrity of the evidence does not depend on the custodian.

The OAIC's guidance should distinguish between two categories of ADM evidence:

| Category | Description | Adequacy for APP 1.7 |
|---|---|---|
| Post-execution logs | Records assembled from logs after the decision was made. Describes what happened. Cannot prove the knowledge foundation was authorized before it happened. | Necessary but not sufficient |
| Pre-execution authorization records | Records written before any action executes. Proves the knowledge foundation was evaluated and authorized. Independently verifiable without relying on the issuer. Mathematically tamper-resistant. | Sufficient for the evidentiary standard APP 1.7 requires |

## 8. Regulatory Alignment

Kuriom's architecture directly addresses the Australian regulatory obligations converging on ADM governance:

| Obligation | Requirement | Kuriom Architecture Response |
|---|---|---|
| Privacy Act APP 1.7–1.9 (December 10, 2026) | Disclose kinds of personal information used and kinds of decisions made by computer programs significantly affecting individual rights | GAT per-instance record: exact knowledge nodes active at authorization · knowledge_state_id fingerprint · Lean 4 injectivity proof |
| APRA CPS 230 (In force 2025) | Operational resilience — material service provider governance — pre-execution AI governance evidence | Five-condition gate evaluates knowledge completeness, currency, integrity, boundary, and prohibition before any AI action in a critical operation |
| APRA Enforcement Intent (April 30, 2026) | Adequate identification, management, and control of AI risks | Pre-execution GAT is the evidence that governance controls were operative before the AI acted — not documentation assembled after supervisory review begins |
| FAR — Financial Accountability Regime (March 2024 ADIs) | Individual accountability — accountable persons must take reasonable steps to prevent governance failures | KAR (Knowledge Attribution Record) names who authorized what knowledge foundation — individual accountability chain per decision |
| ASIC Licensee Obligations (May 8, 2026) | Decision-making and escalation processes fit for purpose — architectural not procedural governance | CPF Mode 1 hard blocks and Mode 2 composite conditions enforce prohibition architecture before any sequence executes |
| Western Australia POLA Act (July 1, 2026) | ADM using personal information — significant decisions — notification, human intervention on request | HITL escalation path: GAI audit package routes to expert panel with full knowledge foundation evidence for any non-authorized sequence |

## 9. Recommendations for OAIC Guidance

Based on the architecture described above, Kuriom makes the following recommendations for the OAIC's ADM guidance:

### Recommendation 1 — Add knowledge foundation reliance as a factor in the substantially and directly related assessment

The OAIC's guidance on Question 1 should include, alongside the five factors listed in the Issues Paper, the degree to which the human decision-maker relies on knowledge assembled by the computer program. A program that assembles the knowledge a human relies on is substantially related to the decision — even where a human formally approves the outcome.

### Recommendation 2 — Address the Composition Gap in guidance on sequence-level ADM

The OAIC's guidance should address the condition where individually authorized computer program actions combine to produce a decision that significantly affects an individual's rights or interests — even though no individual action crosses the significance threshold alone. This is the Composition Gap. Guidance should clarify that APP 1.7 captures the sequence as a whole, not only individual steps within it.

### Recommendation 3 — Distinguish post-execution logs from pre-execution records in the evidence standard

The OAIC's guidance should distinguish between post-execution audit trails — records assembled from logs after a decision was made — and pre-execution authorization records — records that existed before any action executed. For decisions that significantly affect individual rights or interests, the OAIC should encourage entities to implement pre-execution records as the higher standard of governance evidence.

### Recommendation 4 — Extend 'arranged for' to cover knowledge foundation authorization

The OAIC's guidance on 'arranged for' should make clear that an entity arranging for a computer program to make decisions must also arrange for the verification and currency of the knowledge that program relies on. Arranging for the program without governing its knowledge foundation does not constitute adequate governance of the ADM for the purposes of APP 1.7.

### Recommendation 5 — Encourage per-instance evidence, not only category disclosure

APP 1.8 requires disclosure of kinds of personal information used and kinds of decisions made. The OAIC's guidance should encourage entities to implement per-instance records — the specific information used, from what source, with what verification status, at what moment — as best practice beyond the minimum category disclosure. This is the evidence standard that will be required in dispute resolution, regulatory examination, and judicial review.

## 10. Conclusion

The ADM transparency obligation commences December 10, 2026. The OAIC's guidance will define what adequate transparency looks like in practice. This submission argues that adequate transparency requires more than disclosure of what the computer program did. It requires evidence of what the computer program knew — and whether that knowledge was authorized before it acted.

The pre-execution authorization architecture described in this submission — validated in a production codebase, covered by four patent applications, and grounded in a mathematical proof — provides the OAIC with a concrete model for what pre-execution ADM governance looks like in practice. The knowledge_state_id injectivity proof demonstrates that pre-execution evidence can be mathematically tamper-resistant, independently verifiable, and chronologically prior to any decision consequence.

Kuriom is available to brief the OAIC on the technical architecture and its regulatory alignment at any stage of the guidance development process.

| Contact | Detail |
|---|---|
| Submitter | [redacted] Kuriom |
| Email | [redacted] |
| Website | kuriom.ai |
| Architecture reference | SSRN:6782100 (submitted May 2026) |
| Patent applications | 64/008,411 · 64/043,572 · 64/049,750 · 64/059,101 |
| Submission email | privacyreformtranche1@oaic.gov.au |
| Consultation deadline | Monday 15 June 2026 |

**Kuriom · Sequence Knowledge Authorization · AI** Patent Applications: 64/008,411 · 64/043,572 · 64/049,750 · 64/059,101
SSRN:6782100 · [redacted] · kuriom.ai Public Submission · June 2026