AUSTRALIA'S 2026 ADM

TRANSPARENCY MANDATES

From privacy-policy disclosure to enterprise-wide operational evidence

What the law requires — and how Cetus AI supports defensible readiness across the enterprise

CUSTOMER WHITE PAPER • AUGUST 2026

DOCUMENT REFERENCE	CETUS-AI-WP-ADM-2026-01
DOCUMENT TYPE	Customer white paper
VERSION	1.0
PREPARED	25 August 2026
PREPARED BY	Cetus AI • Brookwater, Queensland

Cetus AI • ABN 16 695 570 819 • hello@cetusai.com.au • cetusai.com.au

Document control

This page defines the status, audience and limitations of the white paper before the reader enters the substantive content.

Document attribute	Value
Document title	Australia's 2026 ADM Transparency Mandates
Subtitle	From privacy-policy disclosure to enterprise-wide operational evidence
Document reference	CETUS-AI-WP-ADM-2026-01
Document version	1.0
Prepared date	25 August 2026
Prepared by	Cetus AI, Brookwater, Queensland
Business details	ABN 16 695 570 819 | hello@cetusai.com.au | cetusai.com.au
Classification	Customer White Paper — General Information
Intended audience	CIO, CISO, privacy, legal, risk, audit, architecture, AI platform and business decision owners
Coverage	Federal APP 1.7–1.9, WA IPP 10 and enterprise evidence architecture

IMPORTANT This white paper provides general information and is not legal, regulatory, security, audit or financial advice. Cetus AI supports technical controls and evidence; the organisation remains responsible for legal classification, disclosure wording, decision quality and compliance.

Version history

Version	Date	Status	Description
1.0	25 August 2026	Current	Initial customer white paper on the 2026 ADM transparency mandates and enterprise-wide Cetus AI evidence approach

Table of contents

The table of contents updates automatically when fields are refreshed in Microsoft Word.

Executive summary

Australia's automated decision-making transparency landscape changes materially in 2026. At the federal level, new **Australian Privacy Principle 1.7–1.9** provisions commence on **10 December 2026**. An APP entity must include additional information in its APP Privacy Policy when it has arranged for a computer program to use personal information to make, or materially assist, a decision that could reasonably be expected to significantly affect an individual's rights or interests.[1] [2]

The federal amendments are important but frequently overstated. They create a **privacy-policy transparency obligation**. They do not, by themselves, establish a universal right to transaction-level notification, a detailed algorithmic explanation, human intervention or appeal. Those rights may arise from other laws, contracts, policies or sector obligations, but they are not created generally by APP 1.7–1.9.[1] [3]

Western Australia introduced a separate public-sector regime on **1 July 2026**. Information Privacy Principle 10 requires affected WA public-sector entities and relevant contracted service providers to undertake stronger measures, including an ADM impact assessment, individual notification, understandable information on request and a process for human intervention.[4] The two regimes must not be treated as one national mandate.

The immediate legal output may be a privacy-policy disclosure, but the underlying enterprise task is much larger. An organisation cannot keep that disclosure accurate unless it knows where automated and materially assisted decisions occur, which personal information is used, which systems and vendors contribute, whether a human remains genuinely accountable, and how the decision changes when models, agents, rules, data sources or workflows are updated.

This is an operational evidence problem. The relevant facts may be distributed across Salesforce, SAP, MuleSoft, identity platforms, AI services, policy engines, human-review systems and security logs. Direct AI-platform evidence may also sit inside SAP AI, Salesforce Agentforce, Azure OpenAI, AWS Bedrock or Microsoft Copilot. Each platform can provide useful telemetry, but no single source necessarily explains the complete decision.

Cetus AI addresses this gap through **Songlines Control®**, an additive policy and evidence layer designed to work with the enterprise stack rather than replace it. Cetus AI can support discovery, ADM inventory, policy evaluation, direct AI-platform telemetry, cross-system correlation, human-review evidence, change monitoring and review-ready exports. It does not independently determine legal scope or certify compliance; those responsibilities remain with the organisation's privacy, legal, risk, technology and business owners.[5] [6]

The practical proposition is straightforward:

> EVIDENCE RESULT **Make the public disclosure accurate by making the underlying decision process visible, controlled and reviewable.**

1. Why 2026 changes the governance question

AI governance discussions often begin with principles, policies, committees and risk registers. Those foundations remain necessary, but the 2026 transparency requirements force organisations to answer a more operational set of questions. Which computer programs contribute to decisions? What personal information do they use? What is the program's actual role? Could the outcome significantly affect a person? Which categories must be disclosed publicly, and who keeps those disclosures current?

These questions are not limited to generative AI. The federal provisions are technology-neutral. The OAIC's May 2026 Issues Paper states that the concept of a computer program can include rule-based systems, conventional software, machine learning, artificial intelligence and generative AI.[3] A legacy scoring engine, a recruitment-ranking tool, a refund workflow, an AI assistant and an autonomous agent can therefore raise the same threshold question if they use personal information in connection with a significant decision.

The practical governance unit is not simply "the model". It is the **decision chain**: the business purpose, affected individual, personal-information categories, application workflow, orchestration, model or rule execution, policy conditions, human involvement, final outcome and evidence retained.

This distinction matters because an organisation may have strong controls over individual components while lacking a reliable view of the end-to-end process. Identity teams may know who authenticated. MuleSoft may know which APIs were called. Salesforce may retain case history. SAP may hold authoritative account facts. An AI platform may record the model deployment or agent action. A privacy team may maintain the public policy. Unless those elements are connected, the organisation may struggle to demonstrate why its disclosure remains accurate.

2. The federal APP 1.7 trigger: three elements must be present

The federal obligation applies when all three APP 1.7 elements are satisfied.[1]

Statutory element	Enterprise interpretation	Evidence required to assess it
The entity has arranged for a computer program to make a decision, or do something substantially and directly related to making it.	The organisation has designed, procured, authorised, configured or integrated a program that decides, scores, ranks, recommends, triages, summarises or otherwise materially guides the decision.	System owner, vendor, approved purpose, workflow map, program role, human reliance, override practice and contractual responsibilities.
The decision could reasonably be expected to significantly affect an individual's rights or interests.	The potential effect is meaningful rather than trivial and may be adverse or beneficial.	Decision category, affected population, contractual or statutory rights, access to services, financial effect, employment or opportunity impact, duration and vulnerability factors.
The program uses personal information about the individual.	Personal information about the affected person is used in the program's operation.	Information categories, sources, purpose, data lineage, classification, retention and third-party transfers.

An organisation should assess these elements separately and document the reasoning. The presence of AI does not automatically trigger APP 1.7. Equally, the absence of advanced AI does not automatically remove a process from scope. The role of the program in the actual decision is what matters.

Human involvement does not automatically remove scope

APP 1.7 captures decisions made by a computer program and processes in which a program does something **substantially and directly related** to making the decision. The OAIC's consultation material describes "substantially" as concerning whether the program is a key factor facilitating the human decision and "directly" as requiring a direct connection to that decision.[3]

A human clicking "approve" after relying heavily on a recommendation may therefore remain within scope. Conversely, a calculation that is incidental to a broader judgement may not meet the materiality threshold. Organisations need evidence

of how the workflow operates in practice: whether reviewers see the underlying facts, whether they can override the output, how often they do so, and whether the automated result is advisory or effectively determinative.

Significant effect is broader than an adverse outcome

APP 1.9 confirms that the effect may be **adverse or beneficial**. Examples in the legislation and explanatory material include statutory benefits, contractual rights, access to healthcare or essential services, banking and credit, education, employment opportunities and materially differentiated pricing.[1] [3]

Context matters. The same type of decision may have different significance depending on the person affected, the duration of the effect, the sensitivity of the information, the financial or health consequences, and whether the person is experiencing vulnerability.[3]

3. What APP Privacy Policies must disclose

When APP 1.7 applies, APP 1.8 requires the entity's APP Privacy Policy to describe three categories.[1]

Required category	What the policy must describe	What the legislation does not require the policy to publish
Kinds of personal information used	The categories of personal information used in the relevant programs.	Every field, individual record, full prompt, source code or model weight.
Kinds of solely automated decisions	The categories of decisions made entirely by the computer program.	Every transaction-level outcome or detailed decision logic.
Kinds of materially assisted decisions	The categories of decisions where the program performs something substantially and directly related to making the decision.	Commercial-in-confidence system information or a universal technical explanation.

The statutory wording focuses on "kinds" or categories. The OAIC's May 2026 Issues Paper nevertheless indicates that disclosure should be meaningful: clear plain language, logical grouping, sufficient specificity and a route to further information.[3] The objective is not to publish sensitive implementation detail. It is to give people a useful understanding of how personal information is used in significant automated or materially assisted decisions.

The policy must also remain current. This turns a drafting exercise into an operating obligation. If a new model is deployed, a workflow changes from advisory to effectively determinative, additional personal information is introduced, or a vendor adds an agent action, the underlying assessment and public wording may need reconsideration.

4. What the federal amendments do not create

The federal provisions should not be described as a complete automated-decision rights regime. APP 1.7–1.9 do not, by themselves, create a universal requirement for individual notification each time ADM is used, a case-specific explanation, mandatory human intervention, a new appeal right or publication of source code.[1] [3]

Other obligations may still apply. Employment, discrimination, consumer, administrative, credit, health, contractual, sector and internal governance requirements can impose additional duties. The correct approach is to treat APP 1.7–1.9 as one component of a wider governance environment rather than the only applicable rule.

The federal enforcement position is also material. Part 15 adds APP 1.7 to the section 13K civil-penalty category. Section 13K provides for a maximum penalty of 200 penalty units, with the applicable value and corporate treatment to be checked at the relevant time.[1] [7]

5. Western Australia IPP 10 is a separate and stronger regime

Western Australia's IPP 10 commenced on **1 July 2026** for WA public-sector entities and relevant contracted service providers where ADM uses qualifying personal information to make a significant decision.[4] Its structure is materially different from the federal APP Privacy Policy obligation.

Requirement	Federal APP 1.7–1.9	WA IPP 10
Commencement	10 December 2026	1 July 2026
Primary coverage	APP entities under the federal Privacy Act	WA public-sector entities and relevant contracted service providers
Primary transparency mechanism	Categories published in the APP Privacy Policy	Individual notification and information on request
Impact assessment	Not expressly prescribed by APP 1.7–1.9	Required, including harm, bias, discrimination, human intervention and statutory compliance
Human intervention	Not created generally by APP 1.7–1.9	A request process must be available
Ongoing review	Not expressly prescribed by APP 1.7–1.9	Periodic evaluation and assessment updates are required

An organisation may encounter both regimes. A private-sector technology provider may be subject to the federal Privacy Act for its own activities while also supporting a WA government customer's IPP 10 obligations through contract and system design. Architecture, procurement and evidence requirements should therefore identify which entity, decision, data and jurisdiction are involved.

6. Why enterprise implementation is difficult

The enterprise challenge is not usually the absence of logs. It is the fragmentation of meaning and accountability across systems.

Enterprise component	What it knows	What remains unresolved without correlation
Salesforce or another business application	Case context, consent, customer interaction, workflow status and human action.	Which source facts, AI service, policy and model output materially influenced the decision.
SAP or another system of record	Contract, account, payment, service, workforce or supplier facts.	How those facts were transformed, scored or used by a downstream AI-enabled workflow.
MuleSoft or another integration platform	API calls, transformations, routes, correlation identifiers and response status.	The full model or agent execution context when calls occur directly inside an AI platform.
SAP AI or Salesforce Agentforce	Model or agent identity, grounding, actions, control outcomes and handoff.	The complete decision context across other systems and the approved public disclosure category.
Azure OpenAI, AWS Bedrock or Microsoft Copilot	Platform-native model, deployment, guardrail, session, grounding, action or audit events.	The organisation-wide decision purpose, significant-effect assessment, final human outcome and policy mapping.
Identity, SIEM, DLP and GRC systems	Authentication, security events, data controls, risks and documented policies.	A single decision-centred record connecting authority, data, model, policy, action and review.

Third-party systems increase the difficulty. APP 1.7 uses the phrase **“has arranged for”**, which can capture vendor-operated systems where the organisation has procured, authorised or integrated the relevant decision process.[3] Responsibility therefore cannot be delegated simply by purchasing a service.

Legacy systems are also relevant. The application provision applies to decisions made after commencement even where the arrangement, personal information or acquisition predates the commencement date.[1] Organisations should not limit discovery to newly purchased AI tools.

Finally, direct AI-platform execution creates a visibility gap. An AI action may originate inside Agentforce, SAP AI, Microsoft Copilot, Azure OpenAI or AWS Bedrock without every request crossing MuleSoft. Treating the integration layer as the only evidence source will therefore produce an incomplete view of the enterprise AI estate.

7. The enterprise ADM evidence model

A defensible readiness program should connect the public disclosure to an accountable internal evidence model. The federal law does not prescribe a particular register or technology platform. The following artefacts are practical controls that support accurate scope assessment, policy maintenance and review.[2] [3]

Evidence artefact	Purpose
ADM register	Records systems, owners, vendors, decision categories, affected people, personal-information categories, automation role, human involvement and significant-effect rationale.
Personal-information lineage	Shows which data categories were used, where they originated, how they were transformed and which systems or AI platforms received them.
Decision-role assessment	Distinguishes solely automated outcomes from ranking, recommendation, scoring, triage, summarisation or other material assistance.
Policy and version map	Links the workflow to approved conditions, rule or model versions, deployment dates, exceptions and change approvals.
Decision evidence record	Correlates identity, application, source data, orchestration, AI execution, policy, human review and final outcome.
Disclosure mapping	Connects each in-scope workflow to the relevant public privacy-policy category and publication version.
Change and exception record	Identifies what changed, whether reassessment occurred, who approved it and what evidence remains outstanding.

The evidence model should be proportionate. Enterprise-wide collection should default to metadata, identifiers, timestamps, hashes, policy results and source references. Full prompts, outputs or personal-information payloads should be retained only where there is an approved purpose, appropriate authority, access control and retention rule.

8. Architecture: six layers and one control spine

The recommended Cetus AI architecture preserves existing enterprise responsibilities while making direct AI-platform evidence visible. It should be represented as six horizontal layers with a cross-cutting identity, purpose, classification, policy and human-authority spine.

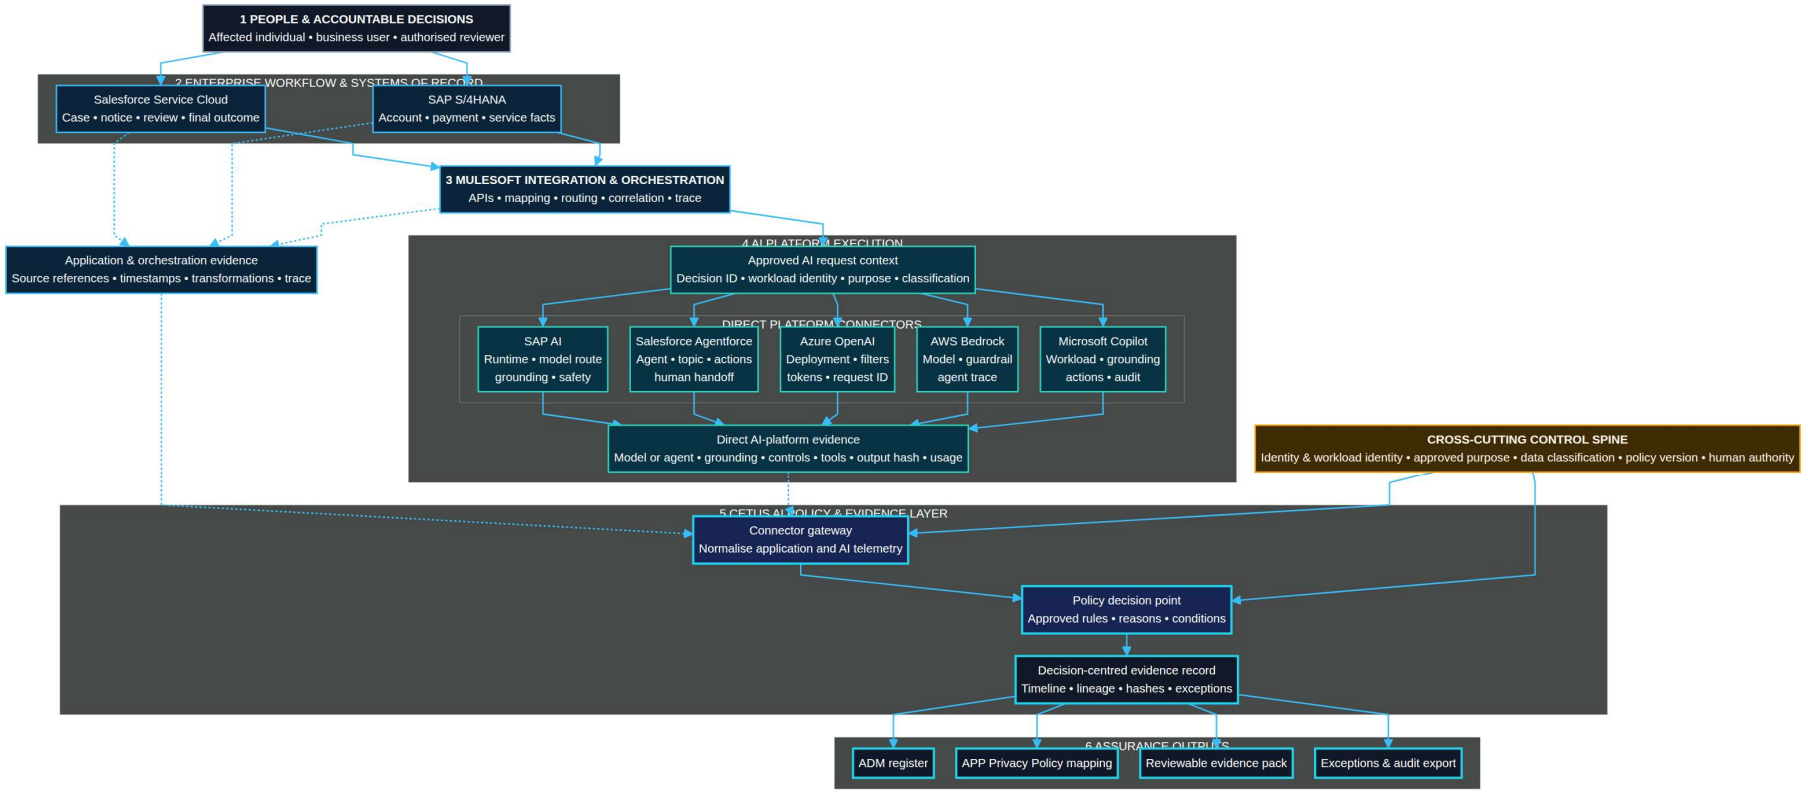

Figure 1. Illustrative enterprise ADM evidence architecture across business applications, orchestration, direct AI-platform execution, Cetus AI policy and evidence, and assurance outputs.

Figure 1. Illustrative enterprise evidence architecture. Solid lines represent operational flow; dotted lines represent asynchronous evidence events. The diagram is a logical pattern, not a mandatory Privacy Act architecture or a claim of live deployment across every named platform.

Layer	Representative components	Governance purpose
1. People and accountable decisions	Affected individual, business user, authorised reviewer and accountable owner	Establish whose rights or interests may be affected and who retains authority.

Layer	Representative components	Governance purpose
2. Enterprise workflows and systems of record	Salesforce, SAP and other business applications	Preserve authoritative case context, source facts and final workflow status.
3. Integration and orchestration	MuleSoft APIs, transformations, routing and correlation	Carry decision, correlation and trace identifiers between systems.
4. AI-platform execution	SAP AI, Agentforce, Azure OpenAI, AWS Bedrock and Microsoft Copilot	Capture model, agent, grounding, guardrail, tool-use, action and execution evidence directly from each platform.
5. Cetus AI policy and evidence layer	Connector gateway, policy decision point, event normalisation, correlation, integrity, retention and evidence export	Produce a decision-centred governance record without replacing the enterprise platforms.
6. Assurance outputs	ADM register, disclosure mapping, evidence pack, exceptions, audit export and board reporting	Make the record usable by privacy, legal, risk, audit, technology and business owners.

The cross-cutting spine should carry enterprise identity, workload identity, approved purpose, data classification, policy version, change approval and human-review authority. These attributes allow events from different platforms to be understood as part of the same accountable decision.

Direct AI-platform connector evidence

Direct connectors should record only the evidence available and authorised within the customer's licensed platform configuration. The following field groups are **illustrative normalised connector contracts**, not complete or vendor-certified native schemas.

Platform	Representative evidence categories
SAP AI	Service and tenant token, resource group, deployment or configuration, model route, approved grounding references, prompt-template hash, safety result, token or compute usage, output hash and decision contribution.
Salesforce Agentforce	Org token, agent and version, session, topic, approved grounding, user or workload authority, action invoked, target system, control outcome, output hash and human handoff.
Azure OpenAI	Azure resource and deployment, model version, region, request ID, workload identity, content-filter result, prompt/output hashes and token usage.
AWS Bedrock	Account and region token, model or inference profile, request ID, guardrail identifier/version/action, agent or knowledge-base references, action trace and usage.
Microsoft Copilot	Tenant and workload, Copilot or agent identity, session, Graph or connector grounding, action invocation, DLP or policy result, audit event and human handoff.

Direct evidence feeds can be asynchronous. Cetus AI does not need to become a mandatory transit point for every AI interaction. MuleSoft remains the integration and orchestration layer where it is already used; direct connectors close the telemetry gap for platform-native AI execution.

9. Worked enterprise scenario: one decision across the stack

Consider a clearly illustrative energy-hardship decision. A customer applies through an authenticated portal for a payment arrangement and temporary protection from disconnection. Salesforce records the application, consent, vulnerability context and human review. SAP provides the contract account, payment history, arrears and service status. Agentforce summarises approved case context, invokes a MuleSoft assessment action and creates the human-review handoff. MuleSoft retrieves the required facts and carries shared decision, correlation and trace identifiers. SAP AI produces an affordability-band contribution grounded in approved source facts. A policy service produces a conditional recommendation. An authorised officer reviews the record, adjusts the proposed monthly payment and makes the final decision.

The example illustrates a **materially assisted human decision**, not a solely automated final outcome. Its potential significance arises from contractual payment terms and continuing access to an essential service. The evidence record can connect:

Evidence question	Illustrative answer
Who and what decision?	A tokenised affected individual; hardship payment arrangement and disconnection hold.
What personal information?	Application, consent, account, payment, service and vulnerability categories.
Which systems contributed?	Salesforce, SAP, MuleSoft, Agentforce, SAP AI, policy service and human-review workflow.
What did AI do?	Agentforce prepared the approved case context and action; SAP AI generated an affordability-band input.
What policy applied?	The approved hardship rule version and human-review condition.
Was a human genuinely involved?	The authorised officer reviewed the source evidence, changed the recommended payment and recorded the reason.
What evidence remains?	Shared identifiers, timestamps, source references, policy result, model/agent metadata, handoff, final outcome, exception and integrity hashes.

IMPORTANT **Illustrative-data boundary:** This scenario contains synthetic records only. It is not a customer deployment, testimonial, legal opinion, regulatory finding or certification of compliance.

The value of the example is not the particular decision. It is the structure: the organisation can reconstruct how business facts, orchestration, AI execution, policy and human judgement combined into one outcome.

10. How Cetus AI supports enterprise-wide readiness

Cetus AI supports the operating model required to keep ADM transparency accurate and reviewable across a changing enterprise estate. Songlines Control® provides six connected capabilities.[5] [6] [8]

Capability	Enterprise function	ADM readiness contribution
Discover and inventory	Identify workflows, vendors, models, agents, direct platform services and accountable owners.	Establishes the population that must be assessed and connected to public disclosure.
Assess the decision role	Record automation type, significant-effect reasoning, personal-information categories and human authority.	Supports the APP 1.7 threshold analysis without pretending to replace legal judgement.
Evaluate policy	Apply approved identity, purpose, data-classification, model, budget and human-review conditions to selected workloads.	Connects documented policy to operational decision conditions.
Collect direct telemetry	Ingest approved events from MuleSoft, Salesforce, SAP, SAP AI, Agentforce and other AI platforms.	Captures evidence that would otherwise remain distributed inside application and platform logs.
Correlate and protect evidence	Normalise identifiers, order events, map lineage, retain hashes, record exceptions and apply access controls.	Produces a reviewable decision-centred record with defined integrity and minimisation controls.
Monitor change and export assurance	Track model, agent, policy, workflow and disclosure changes; produce registers, evidence packs and exception reports.	Helps keep public wording and internal controls aligned as the estate evolves.

Additive rather than disruptive

The architecture is deliberately additive. Salesforce and SAP remain authoritative business systems. MuleSoft continues to orchestrate. Azure, AWS, Microsoft, SAP and Salesforce continue to execute AI services. Existing identity, SIEM, DLP and GRC platforms continue to perform their established roles. Cetus AI adds the decision-centred policy and evidence layer that connects them.[5]

Visibility before enforcement

Customers do not need to begin with a new enterprise-wide runtime control plane. A practical implementation can start with read-only evidence and an ADM inventory. Selected workflows can later use a policy decision point or sidecar pattern where runtime conditions are justified. Broader orchestration should be introduced only where it aligns with the workload's trust boundary and business architecture.[5] [6]

Compliance support—not compliance by assertion

Songlines Control® can provide technical evidence, policy records, version history, human-review events and disclosure mappings. It cannot independently determine whether an entity is legally covered, whether a decision has a significant effect, whether human involvement is substantial, or what wording counsel should approve. Those remain organisational decisions.

11. Privacy, security and evidence minimisation by design

An ADM evidence program should not create a new uncontrolled repository of personal information. The collection design should begin with purpose, necessity and least privilege.

Evidence mode	Information retained	Recommended use
Metadata only	Identifiers, model or agent version, timestamps, control results, token counts, hashes and source references.	Preferred default for broad enterprise collection.
Hashed content	Metadata plus prompt, grounding and output hashes.	Use when content comparison is needed without copying the content.
Approved excerpt	Selected encrypted prompt, grounding or output excerpt.	Use only for defined higher-risk workflows and authorised reviewers.
Full payload	Complete input and output content.	Exceptional; requires explicit purpose, legal and privacy approval, restricted access and controlled retention.

Source identifiers should be tokenised where practical. Connectors should use approved workload identities and least-privilege scopes. Evidence stores should use customer-defined retention, protected timestamps, access logging, role separation and appropriate integrity controls. Reviewers should see only the information necessary for their role.

Cetus AI's public security position describes tenant isolation, role-based access, PII controls, session hardening, cryptographically signed evidence and flexible deployment patterns.[9] Security and assurance claims should be validated against the customer's contracted scope and the deployment being considered.

12. Accountability must span privacy, architecture and the business

ADM transparency cannot be owned by one team. The privacy function may own the public policy, but architecture knows where systems connect, procurement knows what vendors can change, security knows which identities and data controls apply, and the business owns the decision and its consequences.

Accountability	Core responsibility
Business decision owner	Define the purpose, affected population, final authority, acceptable outcomes and ongoing effectiveness measures.
Privacy and legal	Confirm entity coverage, statutory scope, significant-effect analysis, disclosure wording and interaction with other laws.
Enterprise architecture	Maintain the end-to-end workflow, system boundaries, correlation model, integration pattern and resilience design.
AI platform owner	Approve model or agent use, connector access, telemetry availability, version control, guardrails and change notification.
Security and identity	Define workload authority, least privilege, data classification, DLP, audit access and evidence protection.
Risk, compliance and audit	Challenge the control design, review exceptions, test evidence retrieval and report assurance results.
Cetus AI / delivery partner	Configure approved connectors, policy and evidence mappings; support correlation, reporting and operational handover.

This shared operating model prevents a common failure: publishing a technically correct policy statement that becomes inaccurate because system changes were never connected to privacy review.

13. A practical readiness pathway

The recommended approach is to start with one significant decision and prove the evidence model before scaling.

Stage	Key activity	Practical output
1. Discover	Identify programs, vendors, AI platforms and workflows that use personal information in decisions.	Accountable ADM inventory and initial scope population.
2. Assess	Test arrangement, significant effect, personal-information use, automation role and human authority.	Documented classification, decision boundary and evidence gaps.
3. Disclose	Group the relevant personal-information and decision categories in clear, legally reviewed wording.	Updated APP Privacy Policy and mapping from each in-scope process to the published category.
4. Operate	Connect evidence sources, monitor change, test retrieval, manage exceptions and reassess material updates.	Continuous decision evidence, change history, exception reporting and assurance review.

Selecting the first pilot

The first workflow should be material enough to demonstrate value but bounded enough to analyse. It should have a clear business owner, identifiable personal-information categories, a known integration path, accessible evidence sources and a meaningful human-review question. The pilot should measure evidence completeness, correlation accuracy, retrieval time, minimisation, reviewer usability and change-control readiness—not an invented compliance score.

The organisation should also document what is already solved. Existing Salesforce, SAP, MuleSoft, identity, SIEM, DLP and GRC controls should be retained where effective. The Cetus AI design should target the genuine evidence and policy gap rather than duplicate mature controls.

14. Why Cetus AI

Cetus AI's value is not a claim that another dashboard will create compliance. The enterprise problem is the missing connection between documented policy, distributed technical execution and reviewable evidence.

Cetus AI is Australian by design. The company is based in Queensland and focuses on the governance, privacy, sovereignty and assurance expectations faced by Australian enterprises and government organisations.[10]

The architecture is additive. Cetus AI works with the platforms the enterprise already operates. It does not require Salesforce, SAP, MuleSoft, Azure, AWS or Microsoft to be displaced, and it does not require every AI interaction to traverse a new external gateway.[5]

The platform connects control with proof. Songlines Control® can combine identity, policy, AI-platform telemetry, orchestration, human review, exceptions and evidence integrity in one decision-centred record.[6] [8]

Adoption is staged. Organisations can begin with read-only discovery and evidence, validate one workflow, introduce runtime policy selectively and scale according to the trust boundary and operating model.[5] [11]

The legal boundary remains explicit. Cetus AI provides technical controls and evidence that support organisational compliance and assurance. Privacy, legal, governance and business owners remain responsible for legal interpretation, approval, fairness, decision quality and public disclosure.

15. Conclusion

Australia's 2026 ADM transparency mandates make automated and materially assisted decision-making a public accountability issue. The federal requirement is narrower than some commentary suggests, but implementing it well requires more than editing a privacy-policy page. Organisations must know where relevant decisions occur, which personal information is used, how systems and AI platforms contribute, whether a human retains meaningful authority and how the process changes over time.

The evidence is already present across the enterprise, but it is fragmented. Salesforce can explain the case. SAP can explain the source facts. MuleSoft can explain the orchestration. SAP AI, Agentforce, Azure OpenAI, AWS Bedrock and Microsoft Copilot can explain platform execution. Identity, policy, security and human-review systems explain authority and control. The governance problem is connecting those facts into one accountable record.

Cetus AI provides the additive policy and evidence layer required to make that connection. Songlines Control® supports discovery, classification, policy evaluation, direct telemetry, cross-stack correlation, evidence protection, disclosure mapping and ongoing change monitoring. It does not provide compliance by assertion. It enables the organisation to maintain and demonstrate the evidence on which responsible legal and governance decisions can be based.

> **EVIDENCE RESULT** **Start with one significant decision. Connect the systems that already know part of the story. Produce one reviewable record. Then scale the operating model across the enterprise.**

Discuss your ADM architecture

Cetus AI offers architecture reviews, ADM readiness assessments, evidence-pack pilots and Policy Decision Point integration. The first engagement is designed to identify what the organisation has already solved, where evidence remains fragmented and which decision provides the most useful pilot.

Email: hello@cetusai.com.au **Website:** www.cetusai.com.au **Company:** Cetus AI Pty Ltd | ABN 16 695 570 819

Important notice

This white paper provides general information and does not constitute legal, regulatory, security, audit or financial advice. The federal and Western Australian regimes have different scopes and duties. Organisations should obtain appropriate advice for their circumstances and verify current legislation and regulator guidance. Cetus AI and Songlines Control® support technical controls, evidence production and assurance activities; they do not independently determine legal compliance or replace organisational accountability. Platform capabilities and connector designs must be validated against the customer's licensed services, deployment configuration and contracted scope.

References

[1] Privacy and Other Legislation Amendment Act 2024 — Federal Register of Legislation

[2] OAIC — APP 1 Open and transparent management of personal information

[3] OAIC — Automated Decision-Making Transparency Obligation Issues Paper, May 2026

[4] WA Office of the Information Commissioner — Privacy and accountability in automated decision making

[5] Cetus AI — Enterprise Architecture

[6] Cetus AI — Songlines Control Platform Overview

[7] Privacy Act 1988, section 13K

[8] Cetus AI — Regulatory Evidence

[9] Cetus AI — Trust & Security

[10] Cetus AI — Enterprise AI Governance

[11] Cetus AI — Professional Services

[12] OAIC — The Privacy Act