BCA
Business Council of Australia

Supporting a practical and proportionate automated decision-making transparency framework

BCA's submission to the OAIC consultation on Guidance for Transparency in Automated Decision Making

June 2026

Contents

1. Executive summary 2
2. Key recommendations 3
3. Supporting a practical and proportionate ADM transparency framework 5
   - 3.1 Clarity and certainty regarding the scope of the obligation 5
   - 3.2 Recognising differing levels of automation and human involvement 6
   - 3.3 Maintaining a focus on decisions with material impacts on individuals 7
   - 3.4 Practical implementation in technology supply chains 8
   - 3.5 Protection of commercial confidentiality and intellectual property 8
   - 3.6 Ensuring transparency requirements are proportionate and meaningful 9
4. Promoting regulatory coherence 10
5. Conclusion 11

1. Executive summary

The Business Council of Australia (BCA) welcomes the opportunity to provide input into the Office of the Australian Information Commissioner's (OAIC) consultation on guidance for the new automated decision-making (ADM) transparency obligation.

The BCA represents more than 120 of Australia's largest businesses across sectors including finance, resources, health, care, retail, technology, infrastructure and energy. These organisations are increasingly investing in data-driven technologies, automation and artificial intelligence (AI) to improve productivity, enhance customer experiences, support better decision-making and deliver innovative products and services. The BCA recognises that trust is fundamental to the successful adoption of these technologies and supports transparency as an important mechanism for building and maintaining that trust.

The BCA supports the objective of the new ADM transparency obligation and recognises that individuals should have appropriate visibility into significant decisions that affect them, particularly where those decisions involve automated systems and the use of personal information. The key issue for this consultation is not whether transparency should exist, but how the obligation is interpreted and implemented in practice.

The OAIC's guidance will determine the practical operation of the regime. While the legislative objective is clear, a number of key concepts underpinning the obligation remain broad and open to interpretation, including "computer program", "decision", activities that are "substantially and directly related" to decision-making, and decisions that "significantly affect" an individual's rights or interests. Greater clarity regarding the scope and application of these concepts will be critical to ensuring organisations can implement the new requirements consistently and effectively.

The BCA supports an approach that is risk-based, proportionate and aligned with the intent of the legislation. In particular, the guidance should provide clear and practical boundaries regarding the technologies and decisions that fall within scope, maintain a focus on decisions that materially affect individuals, recognise the role of meaningful human oversight, establish practical expectations for organisations using third-party technologies, protect commercially sensitive information and ensure disclosure requirements remain proportionate and meaningful for consumers.

A key concern for business is that an overly broad interpretation of the obligation could extend its application well beyond the intended higher-risk uses of ADM. If too many lower-risk software tools, business processes and technologies are brought within scope, organisations may face significant compliance costs and uncertainty without a corresponding benefit for consumers. More importantly, this risks diluting the effectiveness of the regime by directing attention and compliance effort away from those decisions that have the greatest impact on individuals' rights and interests as per Parliament’s intention.

The BCA also notes that automated and AI-enabled technologies are increasingly used to improve productivity, support economic growth and deliver better outcomes for consumers. It is therefore important that the implementation of the transparency obligation supports responsible innovation and technology adoption while maintaining appropriate safeguards. Regulatory settings that are uncertain, overly broad or difficult to operationalise may discourage investment in new technologies and slow the deployment of innovations that can deliver significant economic and social benefits.

The BCA encourages the OAIC to develop guidance that supports transparency and accountability while remaining practical, proportionate and consistent with the legislative framework. The guidance should also align, where appropriate, with broader Commonwealth initiatives relating to AI governance, privacy reform and automated decision-making to support a coherent national approach to the responsible use of AI and emerging technologies.

2. Key recommendations

Recommendation 1. Provide greater clarity on scope and definitions

The OAIC should provide greater clarity on the scope of the ADM transparency obligation and the interpretation of key concepts within the legislation. In particular, the guidance should clarify the boundaries of the obligation, including what constitutes a "computer program", when a decision is "substantially and directly related" to its use, and the types of technologies and use cases intended to fall within scope.

Recommendation 2. Distinguish decision support from decision automation

The OAIC should provide guidance that recognises automated decision-making exists on a spectrum and distinguishes between systems that support human decision-makers and systems that autonomously make, materially determine or substantially replace a decision.

Recommendation 3. Focus on high-impact decisions

The OAIC should adopt a risk-based interpretation that focuses on decisions with direct and material consequences for individuals' rights or interests, such as significant employment decisions and decisions related to credit, insurance, healthcare, housing, access to public benefits and other essential services. This would be consistent with the examples in the Explanatory Memorandum and would reduce the risk of an overly broad interpretation of the obligation.

Recommendation 4. Support efficient compliance across technology value chains

The OAIC should support approaches that enable transparency information to be shared efficiently across technology value chains and recognise the different roles played by technology providers and deployers. Guidance should avoid creating duplicative assessment and documentation obligations where relevant information can be provided once and relied upon by multiple organisations.

Recommendation 5. Protect commercially sensitive information

The OAIC should clearly confirm that compliance with the ADM transparency obligation does not require disclosure of proprietary algorithms, source code, model architecture, trade secrets or other commercially sensitive information, nor of any of the foregoing relevant to providing, maintaining and improving the safety, integrity and security of a service.

Recommendation 6. Keep transparency proportionate and meaningful

The OAIC should ensure guidance remains closely aligned with the statutory requirements, clearly distinguish between mandatory obligations and examples of good practice, and encourage concise, consumer-focused disclosures that improve understanding without creating unnecessary complexity.

Recommendation 7. Promote regulatory consistency and interoperability

The OAIC should seek consistency with broader Commonwealth, state and territory AI, privacy and ADM initiatives, and interoperability with established international frameworks, without exceeding the requirements of the Australian framework or duplicating existing sector-specific regulation. This would:

- promote consistency in terminology, concepts and compliance expectations
- reduce duplication and fragmentation
- support a coherent national approach.

Recommendation 8. OAIC consult industry on draft guidance and implementation materials

The OAIC should consult closely with industry on draft guidance, rules and implementation materials before they are finalised, including through exposure drafts, worked examples and practical implementation scenarios. This would help ensure the guidance is workable across different sectors, reflects how ADM and AI-enabled systems are used in practice, and supports consistent compliance without creating unnecessary complexity or unintended consequences.

3. Supporting a practical and proportionate ADM transparency framework

The BCA supports the objective of improving transparency regarding the use of automated decision-making and recognises the important role transparency plays in building trust and confidence in the responsible use of personal information, AI and other data-driven technologies. The ADM transparency obligation has the potential to improve consumer understanding of significant decisions that affect them while supporting accountability in the use of emerging technologies.

As the OAIC develops guidance for the new obligation, it will be important to ensure the framework remains practical, proportionate and aligned with the intent of the legislation. In particular, the guidance should provide organisations with sufficient certainty regarding the technologies and decisions that fall within scope, maintain a focus on decisions that materially affect individuals, recognise the role of meaningful human oversight and support practical implementation across increasingly complex technology value and supply chains.

If applied too broadly, the obligation risks becoming less effective by capturing low-risk activities and diverting attention from the higher-impact decisions where transparency matters most. It may also create uncertainty that discourages the responsible adoption of new technologies and ways of working that benefit consumers, businesses and the broader economy.

The following sections outline the BCA's key implementation considerations and recommendations for the OAIC's guidance.

3.1 Clarity and certainty regarding the scope of the obligation

A central issue arising from the consultation is the breadth of concepts such as "computer program", "decision" and activities that are "substantially and directly related" to decision-making. The OAIC's discussion suggests the obligation may apply to a wide spectrum of technologies, including AI systems, analytics tools, rules-based systems and other software that supports business processes.

If the scope of the obligation is interpreted too broadly, businesses may face significant uncertainty in determining which systems are captured, leading to inconsistent implementation and diverting resources towards assessing low-risk technologies rather than focusing on higher-risk uses of ADM. The guidance should therefore provide clear, practical and objective indicators that help organisations identify with confidence when disclosure obligations are triggered. In this regard, clear guidance and worked examples that delineate the meaning of the terms "substantially" and "directly related" will help ensure the transparency and accountability objectives of the disclosure obligation are met and avoid unnecessary and potentially confusing disclosure statements where organisations are unclear whether a decision-making process is included.

Parliament deliberately chose the qualifier “significantly” to exclude trivial or routine effects on individuals. The Explanatory Memorandum's examples of decisions that meet this threshold (granting admission to a country, entitlement to housing benefits, contracts for life insurance, access to healthcare) all share common characteristics: they involve discrete, binary outcomes with direct material consequences for an individual’s rights or interests such as livelihood, liberty, health, or shelter.

“Significantly affect” should, accordingly, be tied to a discrete “decision” that has a direct, material, impact on an individual’s rights or interests that produces a definitive outcome (for example, to approve or deny, grant or refuse, include or exclude), and not to the accumulation of micro-interactions, or the theoretical possibility of harm, or individualised and context-sensitive recommendations. This is consistent with the examples in the Explanatory Memorandum (government benefits, insurance contracts, healthcare access), which all involve discrete, binary outcomes with direct material consequences on an individual’s rights or interests.

It is also consistent with how the term “decision” is interpreted in the Privacy Act, where it carries its ordinary meaning, informed by its use across administrative law, corporations law, and the Administrative Review Tribunal

Act 2024. In each of these contexts, a "decision" is a discrete, operative determination directed at an individual that produces a definitive outcome (granting/refusing a benefit, imposing a condition, issuing or revoking a licence).

To that end, the guidance should make clear that ordinary common, or operational uses of AI, automated tools or other software – including productivity, workflow, cybersecurity, analytics and CRM tools, and general-purpose or commonly available AI features such as drafting, summarisation, personalisation, translation and coding – does not, of itself, fall within scope, and should align the concept of "making a decision" with the definition in the Administrative Review Tribunal Act 2024. Activities such as summarising, organising or filtering information for human review, or generating preliminary outputs (such as in a chatbot), lack a direct causal link and should not, of themselves, trigger the core requirements.

As a starting point, the guidance could treat a system as "substantially" related only where it replaces, or substantially replaces, human decision-making, and "directly related" as requiring a direct causal link to the decision rather than preparatory, ancillary or supporting steps. The guidance should also exclude indirect or downstream impacts, and the inclusion of a non-exhaustive list of technologies and activities that are out of scope could provide greater certainty. For example, without a clear distinction, governance frameworks could be expected to capture everyday tools such as document summarisers, personalised ads and content, and email prioritisation, placing them under the same governance as higher-risk uses such as AI-assisted recruitment – an outcome that is unworkable and would discourage proper governance where it matters most.

3.2 Recognising differing levels of automation and human involvement

The consultation paper appropriately explores situations where automated systems assist human decision-makers rather than making decisions autonomously. For many organisations, AI and analytics tools are deployed to improve the quality, consistency and efficiency of decision-making, while accountability remains with a human decision-maker. The extent of human oversight is relevant because it influences both the risk profile of a system and the level of control exercised over outcomes.

Guidance that does not adequately distinguish between fully automated decisions and decisions subject to human review, input or judgement may unintentionally discourage the adoption of governance practices that improve accountability. The OAIC should recognise differing levels of automation and ensure the guidance reflects the important role that genuine human-centred decisions and oversight can play in mitigating risk.

The guidance should not classify entire categories of technologies (such as generative AI, large language models, or machine learning systems) as inherently "substantially and directly related" to decision-making. Whether a particular tool meets this threshold depends on its actual role in a specific decision workflow, including whether the human decision-maker has access to underlying source and other relevant materials and retains genuine discretion to depart from the system's output.

A practical example of decision support is a clinical decision support system, which is often built into electronic medical records and supports clinicians to make decisions; for example, by flagging a drug interaction, but does not make decisions autonomously.

The guidance should distinguish between systems that merely inform, advise, organise or support human judgement and systems that materially determine, or substantially replace, the decision-making process. The key question should be whether the human decision-maker retains genuine access to relevant information and authority to depart from the system's output. For greater clarity and ease of implementation, some BCA members consider that the obligation should be confined to circumstances where AI and ADM processes make a decision end-to-end, without human review or input, so that the mere use of AI or software tools within a human-decision making process would not, of itself, trigger the obligation and "substantially and directly related to" would be interpreted narrowly.

Recommendation 1. Provide greater clarity on scope and definitions

The OAIC should provide greater clarity on the scope of the ADM transparency obligation and the interpretation of key concepts within the legislation. In particular, the guidance should clarify the boundaries of the obligation, including what constitutes a "computer program", when a decision is "substantially and directly related" to its use, and the types of technologies and use cases intended to fall within scope.

Recommendation 2. Distinguish decision support from decision automation

The OAIC should provide guidance that recognises automated decision-making exists on a spectrum and distinguishes between systems that support human decision-makers and systems that autonomously make, materially determine or substantially replace a decision.

3.3 Maintaining a focus on decisions with material impacts on individuals

The legislation is directed at decisions that could reasonably be expected to significantly affect an individual's rights or interests. The consultation paper raises important questions about how broadly this concept should be interpreted, including in relation to matters such as differential pricing, targeted content and other commercial decisions. A broad interpretation risks capturing a large range of routine commercial activities that may have none or only a limited or temporary impact on consumers. This could dilute the purpose of the regime and reduce its effectiveness by treating fundamentally different categories of decisions in the same way. The guidance should maintain a tight focus on decisions that materially affect access to employment, financial products, insurance, healthcare, education, essential services or other matters with significant consequences for individuals' rights or interests, whereby the "decision" is a discrete decision made about the individual.

The test should be based on objective and predictable criteria: would a reasonable person expect this type of decision to have a material impact on an individual's rights or interests?

For example, the final guidance should draw a distinction between clinical decision-making and healthcare administration: while healthcare is inherently a high-impact sector, there are many instances where automated decision-making may be used for low-risk administrative activities (for example, bed management, rostering or theatre scheduling) that should not be treated as having a significant effect on an individual. Care should also be taken not to expand the concept of an individual's "interests" too broadly – for example, treating access to the lowest available price as an "interest" would read the obligation extremely widely – and the guidance should remain workable for businesses, which may not always know whether a particular customer is vulnerable.

Recommendation 3. Focus on high-impact decisions

The OAIC should adopt a risk-based interpretation that focuses on decisions with direct and material consequences for individuals' rights or interests, such as significant employment decisions and decisions related to credit, insurance, healthcare, housing, access to public benefits and other essential services. This would be consistent with the examples in the Explanatory Memorandum and would reduce the risk of an overly broad interpretation of the obligation.

3.4 Practical implementation in technology supply chains

The consultation explores how the APP 1 obligation applies where organisations use third-party technologies as part of their decision-making processes. In practice, many organisations rely on externally supplied software, cloud services, analytics platforms and AI-enabled products to support business operations and decision-making. As technology ecosystems become increasingly interconnected and more complex, transparency obligations should be implemented in a way that reflects how these systems are developed, supplied and deployed in practice.

The BCA supports an approach that maintains accountability for decisions affecting individuals. Organisations that deploy AI or ADM tools should understand the purpose for which those tools are used, how they are incorporated into their decision-making processes, what personal information is involved, what governance arrangements apply, and the role of human oversight. This is consistent with good governance and is essential to maintaining trust in the responsible use of technology. Because the behaviour of AI systems depends not only on their code but on how they are trained, tuned and updated over time, a deploying organisation often cannot, on its own, fully describe every component that shapes a decision affecting an individual, particularly in multi-vendor environments.

The guidance should also recognise the different roles played by technology providers and deployers. The obligation, and the term "arranged for", is most appropriately confined to the entity that determines the purposes and means of the processing (the controller or deployer); processors supplying back-end software on a customer's instructions, and, for AI, the developer or provider of a model as distinct from its deployer, are not the accountable party. This approach is consistent with the EU Artificial Intelligence Act[1] and jurisdictional AI laws in the United States such as the Colorado AI Act.[2] Standardising terminology would also help to avoid confusion where two businesses rely on the same hyperscaled service but describe it differently in their respective policies.

In multi-party arrangements, the disclosure obligation should attach to the party that has "arranged for" the ADM. In the BCA's view, that is ordinarily the organisation that specifies and controls the decision, which remains accountable for transparency however the function is delivered. The principal question for the guidance to resolve is the outsourcing case: where a supplier independently chooses to use ADM to perform a function it has been engaged to deliver, it should be the supplier that has arranged for the program and carries the obligation as an APP entity in its own right. The OAIC's guidance should settle when this is so to locate responsibility effectively in practice. Similarly, clearer guidance would also assist in determining that an APP entity has "arranged for" ADM where it decides to procure and utilise an off-the-shelf software-as-a-service tool that uses automated decision-making.

Recommendation 4. Support efficient compliance across technology value chains

The OAIC should support approaches that enable transparency information to be shared efficiently across technology value chains and recognise the different roles played by technology providers and deployers. Guidance should avoid creating duplicative assessment and documentation obligations where relevant information can be provided once and relied upon by multiple organisations.

3.5 Protection of commercial confidentiality and intellectual property

The BCA supports transparency of ADM use and the categories of personal information involved in decision-making processes. However, transparency obligations should not create pressure for businesses to disclose commercially sensitive information, proprietary methodologies or intellectual property. Many organisations invest substantially in the development of algorithms, predictive models, risk frameworks and technology-enabled

[1] https://artificialintelligenceact.eu
[2] https://coag.gov/ai/

business processes that underpin innovation and competition. Guidance that blurs the distinction between transparency and disclosure of proprietary information may discourage investment and create uncertainty regarding the protection of valuable intellectual assets. The OAIC should clearly reinforce that meaningful transparency can be achieved without requiring disclosure of commercially sensitive system design or operation. In particular, APP 1.7 should not require disclosure of source code, model weights, proprietary algorithms or training datasets, nor of safety- and security-protective information, the disclosure of which could negatively impact on the safety and integrity of services. Transparency can be successfully achieved without requiring disclosure of proprietary system design.

Recommendation 5. Protect commercially sensitive information

The OAIC should clearly confirm that compliance with the ADM transparency obligation does not require disclosure of proprietary algorithms, source code, model architecture, trade secrets or other commercially sensitive information, nor of any of the foregoing relevant to providing, maintaining and improving the safety, integrity and security of a service.

3.6 Ensuring transparency requirements are proportionate and meaningful

The ADM transparency obligation is intended to improve individuals' understanding of how automated decision-making affects them. The Privacy Act establishes a clear set of disclosure requirements, focused on the categories of personal information used and the types of decisions that are made or supported by ADM. The consultation also explores a range of broader transparency concepts drawn from international frameworks and other governance regimes. While these concepts may provide useful context, it is important that guidance remains closely aligned with the requirements enacted by Parliament and does not create an expectation that organisations provide significantly more information than the legislation requires.

Some international frameworks confer individual rights to explanation and contestation that do not form part of this obligation. For example, the UK Information Commissioner's Office's six-element framework for explaining AI decisions was developed for a regime that imposes such rights (under GDPR Articles 13 to 15 and 22).[3] These rights do not exist under APPs 1.7 to 1.9, and importing the substance of that framework into Australian guidance would risk gold-plating the statutory text beyond Parliament's intent.

Excessively detailed or technical disclosures risk increasing compliance costs while reducing consumer understanding, as individuals may be presented with large volumes of information that are difficult to interpret or act on. The guidance should therefore encourage disclosures that are proportionate, practical and focused on helping individuals understand when ADM is being used, the nature of its role in decision-making, and the potential impact on them.

Recommendation 6. Keep transparency proportionate and meaningful

The OAIC should ensure guidance remains closely aligned with the statutory requirements, clearly distinguish between mandatory obligations and examples of good practice, and encourage concise, consumer-focused disclosures that improve understanding without creating unnecessary complexity.

[3] https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence/

4. Promoting regulatory coherence

The ADM transparency obligation is being introduced at a time when governments in Australia and overseas are developing overlapping frameworks for privacy, AI governance and automated decision-making. Australian businesses operating across domestic and international jurisdictions are increasingly required to comply with multiple regimes that address similar risks through different terminology, thresholds and compliance expectations. This creates a strong need for the OAIC's guidance to support coherence, both within Australia and internationally.

Domestically, the guidance should align where appropriate with Commonwealth, state and territory initiatives relating to AI, privacy and automated decision-making. Consistency in terminology and concepts will help reduce duplication, support clearer compliance pathways and avoid creating parallel obligations that apply to the same systems or decisions. This is particularly important in sectors that are already subject to specific regulatory frameworks. For example, clinical algorithms and diagnostic AI may already be regulated under the Therapeutic Goods Administration's Software as a Medical Device framework. The OAIC's guidance should avoid creating unnecessary dual regulation where existing sector-specific regimes already address relevant risks.

International interoperability is also important. Many Australian businesses already operate in overseas markets or provide services to customers in jurisdictions with established privacy, AI and automated decision-making requirements. Guidance that is materially inconsistent with leading frameworks such as the GDPR, the EU AI Act and US state privacy and AI laws would require businesses to design bespoke Australian compliance processes, significantly increasing cost and complexity without delivering proportionate benefits for Australian consumers undermining the competitiveness of Australian businesses. Where appropriate, the OAIC should draw on international concepts that support clear thresholds, role-based accountability and practical compliance.

Interoperability does not mean importing overseas regimes wholesale. Australia's framework should remain aligned with the text and intent of the Privacy Act. However, where international approaches provide useful concepts, such as focusing on decisions with legal or similarly significant effects, distinguishing between providers and deployers, or enabling standardised technical documentation to support downstream compliance, the OAIC should seek to avoid unnecessary divergence. This would support consumer transparency while reducing fragmentation for businesses operating across multiple markets.

Recommendation 7. Promote regulatory consistency and interoperability

The OAIC should seek consistency with broader Commonwealth, state and territory AI, privacy and ADM initiatives, and interoperability with established international frameworks, without exceeding the requirements of the Australian framework or duplicating existing sector-specific regulation. This would:

- promote consistency in terminology, concepts and compliance expectations
- reduce duplication and fragmentation
- support a coherent national approach.

5. Conclusion

The BCA supports the objective of improving transparency regarding the use of automated decision-making and recognises the important role transparency plays in maintaining trust and confidence in the responsible use of personal information, data-driven technologies and AI-enabled systems. The new APP 1 obligation provides an important opportunity to strengthen transparency for individuals while supporting the continued adoption of technologies that deliver benefits for consumers, businesses and the broader economy.

As the OAIC develops guidance for the new transparency obligation, it will be important to ensure the framework remains practical, proportionate and aligned with the intent of the legislation. In particular, guidance should provide clear boundaries regarding the scope of the obligation, focus on decisions that materially affect individuals' rights or interests, recognise the role of meaningful human oversight, establish practical expectations for the use of third-party technologies, protect commercially sensitive information and the safety, integrity and security of services and support transparency requirements that are both proportionate and meaningful for consumers.

The BCA is concerned that an overly broad interpretation of key concepts within the legislation could result in a wide range of low-risk technologies and routine business processes being brought within scope. Such an outcome risks diluting the policy intent of the regime by directing attention and compliance effort away from higher-risk uses of automated decision-making where transparency is most important. It would also create uncertainty for organisations seeking to deploy new technologies and discourage innovation, investment and the responsible adoption of AI-enabled tools that can improve services, productivity and customer outcomes.

Recommendation 8. OAIC consult industry on draft guidance and implementation materials

The OAIC should consult closely with industry on draft guidance, rules and implementation materials before they are finalised, including through exposure drafts, worked examples and practical implementation scenarios. This would help ensure the guidance is workable across different sectors, reflects how ADM and AI-enabled systems are used in practice, and supports consistent compliance without creating unnecessary complexity or unintended consequences.

BCA

Business Council of Australia

BUSINESS COUNCIL OF AUSTRALIA
GPO Box 1472
Melbourne 3001
T 03 8664 2664
www.bca.com.au

© Copyright June 2026 Business Council of Australia ABN 75 008 483 216

All rights reserved. No part of this publication may be reproduced or used in any way without acknowledgement to the Business Council of Australia.

The Business Council of Australia has taken reasonable care in publishing the information contained in this publication but does not guarantee that the information is complete, accurate or current. In particular, the BCA is not responsible for the accuracy of information that has been provided by other parties. The information in this publication is not intended to be used as the basis for making any investment decision and must not be relied upon as investment advice. To the maximum extent permitted by law, the BCA disclaims all liability (including liability in negligence) to any person arising out of use or reliance on the information contained in this publication including for loss or damage which you or anyone else might suffer as a result of that use or reliance.