SAP Australia Submission

June 2026

Consultation on Guidance for Transparency in Automated Decision Making

1. SAP Australia welcomes the opportunity to comment on the Office of the Australian Information Commissioner's (OAIC) consultation on guidance for the automated decision-making transparency obligation under Australian Privacy Principle (APP) 1.

2. This submission is provided from the perspective of a long-standing business-to-business enterprise technology provider whose software, cloud platforms and services form part of the operational backbone of the Australian economy. SAP has a 35-year legacy in Australia and supports organisations across all levels of government, large enterprise and small and medium-sized business, including many that deliver essential finance, procurement, workforce, asset management, logistics and public services. SAP provides enterprise technology to organisations, rather than consumer-facing services directly to individuals.

3. SAP's key points are that the guidance should:

    (a) remain focused on consequential uses of automated decision-making on individuals only,

    (b) align accountability with the entity that deploys and uses such systems,

    (c) promote meaningful, non-technical transparency, and

    (d) be workable in multi-party enterprise technology environments.

4. SAP supports the consultation's overarching direction and agrees that the new transparency obligation should strengthen trust, accountability and protection of individual rights where personal information is used in automated decision-making that may significantly affect an individual's rights or interests. From 10 December 2026, APP 1.7–1.9 will require APP entities to describe in their privacy policies the kinds of personal information used in relevant automated decision-making, and the kinds of decisions that are made, or substantially supported, by computer programs in ways that could significantly affect individuals' rights or interests.

5. SAP also considers that this consultation is an opportunity to support Australia's broader goals for safe and responsible AI adoption. SAP has embedded responsible AI practices over a number of years and has publicly positioned its approach around principles including safety, proportionality, fairness, privacy, human oversight, transparency, responsibility and multi-stakeholder governance, consistent with international frameworks such as the UNESCO Recommendation on the Ethics of Artificial Intelligence. SAP has also established

governance structures to support this work and emphasises transparency, explainability, privacy, data protection and cybersecurity in its approach to AI.

General observations

6. At a high level, SAP supports the objective of enhancing transparency in the use of automated decision-making systems. The new obligation sits within APP 1, which is concerned with open and transparent management of personal information, and is directed to what entities must say in their privacy policies about those automated decision-making practices.

7. The issues paper focuses on key threshold concepts relevant to that obligation, including "computer program", "arranged for", "substantially and directly related to making a decision", and decisions that may "significantly affect" an individual's rights or interests. SAP agrees that greater clarity on these concepts will be critical to ensuring that the obligation operates in a practical and proportionate way.

8. The consultation also usefully recognises that the obligation is about transparency in privacy policies, rather than creating a broader right to detailed technical explanation of automated systems. SAP supports that framing and considers that the guidance should remain focused on disclosures that are understandable, useful and proportionate.

9. SAP also considers that the guidance should make clear that the mere use of software, automation, analytics or AI-enabled functionality in ordinary business operations should not, without more, trigger the obligation. The focus should remain on consequential decision-making, not routine operational, administrative or productivity-enhancing functions.

Scope and accountability

10. In SAP's view, one of the most important issues for the guidance is how responsibility is allocated in enterprise technology environments where software providers, customers and other service partners may each play a role in the design, deployment and operation of a system.

11. For that reason, SAP considers that the guidance should clarify that the primary transparency obligation under APP 1.7 rests with the APP entity that affirmatively deploys and uses automated decision-making for a particular decision-making purpose affecting identifiable individuals. This would better align accountability with practical control over the relevant use case and with the entity best placed to understand how personal information is

used in the decision process. This approach would also be consistent with OAIC's broader practice of expecting the entity with the most direct relationship with affected individuals to carry primary responsibility for privacy notifications and related communications.

12. This is particularly important in the case of enterprise technology providers operating on a business-to-business model, where the provider makes capabilities available to organisational customers but does not itself decide the particular downstream purpose, workflow or affected population for which those capabilities are deployed.

13. A broader interpretation risks creating uncertainty for software providers whose technology may be capable of supporting automated decision-making, but who do not control the downstream context in which it is used. It may also create confusion for customers about where responsibility for privacy policy disclosures should sit.

14. SAP therefore encourages the OAIC to recognise a shared responsibility model. In this model, providers support transparency through appropriate documentation, governance artefacts and design choices, while the primary obligation to disclose relevant uses of automated decision-making remains with the APP entity using personal information in the relevant business or government context.

15. This would provide a more workable basis for compliance in modern cloud and AI ecosystems. It would also better reflect the policy intent of the obligation by linking disclosure responsibility to control over the purpose, context and likely effect of the relevant decision.

16. In this context, it may also be helpful for the guidance to outline the types of factors that may be relevant in assessing whether an entity has "arranged for" the use of a computer program. These could include which entity determines the purpose of the decision, configures the system in a particular use case, and controls the use of personal information in that context. Identifying such factors would support a more consistent and practical application of the obligation across multi-party environments.

Meaningful transparency

17. SAP strongly supports the objective of meaningful transparency for individuals. In SAP's view, this means that privacy policy disclosures should focus on the kinds of personal information used and the kinds of decisions made, or substantially supported, by automated means.

18. The guidance should draw a clear line between disclosing the function and effect of an automated decision-making system, and disclosing internal technical implementation. APP 1.7 should not be interpreted as requiring disclosure of source code, model weights, proprietary algorithms, training datasets, detailed workflow configurations, routing rules or

similarly technical and commercially sensitive information, and the explanatory materials recognise that commercial-in-confidence information about automated decision-making systems is not expected to be disclosed.

19. SAP also considers that the guidance should reinforce that transparency is not the same as technical exhaustiveness. The most effective disclosures are likely to be those that explain relevant uses of automated decision-making in clear and accessible language, rather than in highly technical terms.

20. Practical examples in the guidance would be especially useful in helping organisations distinguish between lower-risk automated functions and use cases that are sufficiently connected to a decision, and sufficiently significant in their effect, to warrant disclosure under APP 1.7. This would reduce the risk of both over-disclosure and under-disclosure, and support more consistent implementation across sectors.

21. In SAP's view, the guidance should keep the threshold focused on whether a program's output is a material factor in the outcome, rather than one input among many or a trivial computational step. Where a human decision-maker applies genuine independent judgment, including real capacity to assess, depart from or override an output in practice, the connection to the final decision is weaker. Examples would be especially useful in areas such as employment, access to essential services, financial services, education, housing and health-related decisions, where the significance of the effect on the individual is more readily understood. Illustrative examples, including in areas such as the use of generative AI to support decision-making, targeted service delivery, or customer segmentation and service eligibility decisions, would further assist entities in applying these concepts in practice.

Practical application

22. A further issue is the practical application of the obligation in real-world enterprise environments. In many cases, suppliers will not have complete visibility over how a customer ultimately configures and deploys a tool, while customers may not have access to detailed technical information about third-party systems they procure. The same is true where service providers themselves rely on upstream infrastructure, hosting and platform providers, and do not have full visibility into how those components operate in detail.

23. SAP considers that the guidance should acknowledge this operational reality. A simple one-provider, one-user model will not always reflect the way enterprise software and AI-enabled systems are deployed in practice.

24. This is particularly important in Australia, where privacy law does not map neatly onto controller and processor distinctions used in some other jurisdictions. International approaches nevertheless demonstrate the practical value of allocating primary responsibility to the entity that determines the purpose and practical use of personal information in a particular decision-making context. In that context, OAIC guidance has an

important role to play in reducing uncertainty and clarifying how the transparency obligation should be applied in multi-party environments.

25. SAP also notes that broad or unclear interpretations may generate substantial operational burden, including large volumes of highly granular customer requests for information that may not ultimately assist meaningful transparency for individuals. They may also increase the cost and complexity for APP entities of maintaining accurate transparency statements over time, including potential ongoing legal review to assess evolving use cases against the guidance. A more practical framework would support standardised, comprehensible disclosures and recognise legitimate limits relating to confidentiality and commercial sensitivity. The guidance should also avoid creating implied monitoring or oversight obligations for technology providers that do not have visibility into how customers deploy or configure tools after implementation, and instead allocate responsibility to the entities best placed to assess the relevant use case and make the required disclosure.

26. Over time, SAP will continue to support customers as expectations evolve in relation to automated decision-making transparency and responsible AI governance. For present purposes, however, SAP's main observation is that the guidance should give organisations a clear, proportionate and operationally workable basis for compliance.

Final observation

27. Overall, SAP supports the consultation and considers there is a strong foundation for practical guidance. The key refinement is to ensure the guidance remains focused on consequential uses of automated decision-making, aligns accountability with control, promotes meaningful transparency, and reflects the realities of enterprise technology environments.

28. SAP makes this submission as a long-standing participant in Australia's digital economy and as a company committed to safe and responsible AI governance. SAP would welcome guidance that protects individuals, strengthens trust, and supports responsible innovation and adoption of AI across the Australian economy.