



AGENDA ITEM: 7



DISCUSSION PAPER

Interim AI Governance Committee

Thursday, 6 November 2025

Title:	AI Technical Standard Gap Analysis
Purpose:	To discuss findings from a gap analysis conducted on the AI Technical Standard.

RECOMMENDATION:

That the Committee **DISCUSS** the AI Technical Standard Gap analysis, key issues identified and proposed next steps.

BACKGROUND:

On 31 July, the DTA released the '[Australian Government AI Technical Standard](#)' (the standard). It is designed to complement the 'Policy for the responsible use of AI in government', 'AI Assurance Framework', 'Voluntary AI Safety Standard', and 'AI Ethics Principles' by providing guidance for embedding AI in government systems and operations. The standard outlines best practice for the end-to-end design, development, deployment, and use of AI systems.

The standard emphasises agency integration into existing governance infrastructure, by adapting existing policies, frameworks and practices, rather than introducing new or duplicative processes. Compliance with the standard is currently voluntary and encouraged by the DTA.

The Deputy Director General requested that the Chief Risk Officer, in consultation with DTG, perform a gap analysis of IP Australia's current practices against the standard (IAIGC2025/3-A3). The gap analysis is available at **Attachment A**.

KEY ISSUES/CONSIDERATIONS:

The standard sets out 148 criteria, of which 90 are 'required' and 58 are 'recommended'. In the gap analysis, we determined that individual criteria are either:

- 'Satisfied'- IP Australia meets this criterion to a sufficient degree in all applicable instances
- 'Partially satisfied' - IP Australia meets this criterion to some degree, or only in some applicable instances
- 'Not satisfied' - IP Australia does not have practices ensuring this criterion could be met consistently.

Of the 90 'required' criteria, IP Australia satisfies 68, partially satisfies 17, and does not satisfy 3. The wording of 2 criteria were too ambiguous for consulted SMEs to decide.

Of the 58 'recommended' criteria, IP Australia satisfies 45, partially satisfies 9, and does not satisfy 2. The wording of 2 criteria were too ambiguous for consulted SMEs to decide.

Overall, we satisfy 76%, partially satisfy 18%, and do not satisfy 3% of criteria in the standard.

Criteria not satisfied

The 5 unsatisfied criteria (3 of which are 'required') relate to watermarking. This includes visual watermarking and metadata for generated media content, and to indicate when a user is interacting with an AI system. Currently, we do not watermark AI generated content, and do not have established policies or practices in place to support this. **s47C**

Partially satisfied criteria

The partially satisfied criteria are spread across various aspects of the standard and are not concentrated in a specific field or indicative of a particular technical gap. For many of these, we have practices in place that address the criteria, but only partially, as they are not standardised, mandatory, or formalised agency-wide approaches. **s47C**

Next steps

s47C

RISKS:

GOVERNANCE
Risk -s47C
Control - s47C

SENSITIVITIES: Nil.

CONSULTATION:

Role	Name	Feedback	Support
Chief Data Officer	s47F	Incorporated	Support
Director, Business Transformation and New Products		Incorporated	Support
Director, Artificial Intelligence and Frontier Technologies		Incorporated	Support
Director, Digital Delivery and Assurance		Incorporated	Support
IT Security Advisor (ITSA)		Incorporated	Support
Solution Architect, Strategy and Architecture		Incorporated	Support
Director, Data Engineering		Incorporated	Support
Director, Data and Information Governance		Incorporated	Support
Data Analyst, Data Front Door and Analytics		Incorporated	Support

COMMUNICATION:

The outcomes of this paper will be included in the minutes of the meeting and published on the intranet following committee approval at the next available committee meeting.

ATTACHMENT:

Attachment A – AI Technical Standard Gap Analysis

Brief prepared by: s47F, acting Assistant Director, and s47F, Risk Officer, Enterprise Risk
s47F, acting Director and Chief Risk Officer, Enterprise Risk

Brief cleared by: Jennifer Hutchinson, General Manager, Governance Group

Date: 27 October 2025

Technical Standard for AI in Government						IP Australia Technical Standards	
Lifecycle Stage	Statement #	Statement	Criterion #	Criterion	Designation	Current State	Rationale
Whole of AI lifecycle	1	Define the reference architecture	1	Identify a suitable operational model to design, develop, and deliver the system securely and efficiently	Recommended	Criterion Partially Satisfied	A suitable operational model can be identified in the high-level design document, the architecture decision document, and/or the architecture pattern document. While Strategy & Architecture provides support through application portfolio management, the whole-of-lifecycle governance is the responsibility of the relevant platform owner, meaning it is difficult to provide universal assurance.
Whole of AI lifecycle	1	Define the reference architecture	2	Consider the technology impacts of the operating model	Recommended	Criterion Satisfied	
Whole of AI lifecycle	1	Define the reference architecture	3	Consider technology hosting strategies	Recommended	Criterion Satisfied	The architecture design process considers technology impacts for all systems, including the impacts of AI. Many impacts of AI use cases are also reviewed and evaluated by SMEs as part of the risk assessment.
Whole of AI lifecycle	2	Define the reference architecture	4	Evaluate existing reference architectures	Required	Criterion Partially Satisfied	The architecture design process considers technology hosting strategies, supported by specific architecture principles and approved enterprise non-functional requirements (ENFRs). This process includes all variants of Cloud.
Whole of AI lifecycle	2	Define the reference architecture	5	Monitor emerging reference architectures to evaluate and update the AI system	Recommended	Criterion Partially Satisfied	Published AI reference architectures are evaluated in architectural design exercises. Currently, IPA does not maintain a repository of AI reference architectures but is planning to develop one.
Whole of AI lifecycle	3	Identify and build people capabilities	6	Identify and assign AI roles to ensure a diverse team of business and technology professionals with specialised skills	Required	Criterion Satisfied	Strategy & Architecture monitor emerging reference architectures in all architectural domains, including AI.
Whole of AI lifecycle	3	Identify and build people capabilities	7	Build and maintain AI capabilities by undertaking regular training and education of end users, staff, and stakeholders	Required	Criterion Partially Satisfied	IPA's AI development and governance brings together a range of technical professionals. For example, the AI/ADM Assessment Panel comprises SMEs from the Office of Legal Counsel, Data Governance, Strategy & Architecture, IT Security Advisor, and Enterprise Risk.
Whole of AI lifecycle	3	Identify and build people capabilities	8	Mitigate staff over reliance, under reliance, and aversion of AI	Recommended	Criterion Partially Satisfied	IPA has made information about the agency's approach to AI available to external stakeholders. Specialised staff who are involved in developing or governing AI are routinely seeking out further training and education. LXP units, UL courses, and COPs are available to staff, and AI capability is included in the ILS. However, there is no mandatory AI training, and no assurance that all staff using AI have undergone foundational AI education. Further, IPA does not currently have a targeted plan of accommodations to support staff with additional needs to transition to AI use.
Whole of AI lifecycle	4	Enable AI auditing	9	Provide end-to-end auditability	Required	Criterion Satisfied	As indicated above, staff have access to a range of training, including the Data Capability Hub, although no training is compulsory. Business areas have carried out a range of change management measures as AI becomes increasingly integrated. Various controls and quality review programs are in place to minimise over reliance on AI. However, IPA lacks extensive data on current staff competency with and perception of AI, and does not yet have a comprehensive plan on how to combat potential skill atrophy arising from AI reliance.
Whole of AI lifecycle	4	Enable AI auditing	10	Perform ongoing data-specific checks across the AI lifecycle	Required	Criterion Partially Satisfied	End-to-end auditability is mandatory for system approval in all architectural design exercises, including systems with AI components.
Whole of AI lifecycle	4	Enable AI auditing	11	Perform ongoing model-specific checks across the AI lifecycle	Required	Criterion Partially Satisfied	Performing ongoing data-specific checks is generally the responsibility of the platform owner in consultation with the relevant data steward's and data custodian's, making it difficult to provide universal assurance.
Whole of AI lifecycle	5	Provide explainability based on the use case	12	Explain the AI system and technology used, including the limitations and capabilities of the system	Required	Criterion Satisfied	Ongoing checks of AI components, including models, are generally performed throughout the technology lifecycle. However, this has not yet been formally standardised.
Whole of AI lifecycle	5	Provide explainability based on the use case	13	Explain outputs made by the AI system to end users	Recommended	Criterion Partially Satisfied	The AI/ADM governance process requires explanations of the systems and their limitations or associated risks. The high-level design template reviews 'assumptions, dependencies, and constraints' of a given system. Where applicable, further information will be captured in technical/system use case cards and governance plans.
Whole of AI lifecycle	5	Provide explainability based on the use case	14	Explain how data is used and shared by the AI system	Recommended	Criterion Satisfied	AI system outputs have been partially explained to end users by emphasising to staff the limitations of AI-generated responses and their enduring responsibility in decision-making. However, we have little data confirming staff understanding of AI outputs. IPA's public-facing tools are accompanied by disclaimers. Depending on the nature of a use case, further information would also be captured in a use case-specific Change Management Plan.
Whole of AI lifecycle	6	Manage system bias	15	Identify how bias could affect people, processes, data, and technologies involved in the AI system lifecycle	Required	Criterion Satisfied	How data is used and shared by an AI system is captured within and across the following governance documents, which are completed in combination as relevant: the CAAF or IPA AI/ADM risk assessment; Security Risk Management Plan or Security Risk Assessment; High-Level Design document; Privacy Threshold Assessment or Privacy Impact Assessment; and/or technical/system use case cards and governance plans.
Whole of AI lifecycle	6	Manage system bias	16	Assess the impact of bias on your use case	Required	Criterion Satisfied	The CAAF and the IPA AI/ADM risk assessment both require information about harm to individuals, groups, and communities, as well as potential discrimination and barriers to accessibility or inclusion. Where applicable, further information is captured in technical/system use case cards and governance plans. The approach is informed by IPA's Data Ethics guidelines.
Whole of AI lifecycle	6	Manage system bias	17	Manage identified bias across the AI system lifecycle	Required	Criterion Partially Satisfied	The CAAF and IPA risk assessment both require information about the potential for the use case to facilitate discrimination, stereotyping, or harm, and to influence decision-making. Multiple facets of bias and its impacts are captured through this evaluation. Further, the CAAF includes a more specific, granular consideration of use case impact on discrete stakeholder groups. Where applicable, information about bias or possible bias is also recorded in the technical/system use case cards and governance plans.
Whole of AI lifecycle	7	Apply version control practices	18	Apply version management practices to the end-to-end development lifecycle	Required	Criterion Satisfied	As above, bias is thoroughly considered in the various governance documents required for a use case, and these documents are then reviewed by multiple SMEs with diverse knowledge bases. Identified potential for bias must be managed with controls. However, human agents are involved in the use of AI, and while this often serves to mitigate bias, there is the risk of increasing automation bias in operational activity as AI is integrated into BAU, which must be combated with education and quality review.
Whole of AI lifecycle	7	Apply version control practices	19	Use metadata in version control to distinguish between production and non-production data, models, and code	Recommended	Criterion Satisfied	Version management practices are in place for the governance documents, including the use case card, risk assessment, and SRMP. Version management are also in place for technical records, as code versioning is utilised in GitHub and the AWS environment, through CodeCommit.
Whole of AI lifecycle	7	Apply version control practices	20	Use a version control toolset to improve usability for users	Recommended	Insufficient information from DTA	The Data Asset Inventory Playbook and associated data guidance documents set out the appropriate use and management of metadata, which applies to all AI systems. Continuous integration and continuous delivery pipelines are in place across all development environments.
Whole of AI lifecycle	7	Apply version control practices	21	Record version control information in audit logs	Recommended	Criterion Satisfied	DTG advised criteria unclear, could not provide a response.
Whole of AI lifecycle	8	Apply watermarking techniques	22	Apply visual watermarks and metadata to generated media content to provide transparency and provenance, including authorship	Required	Criterion Not Satisfied	Audit logging is a standard consideration in design exercises. Version control information can be stored if configured during the system delivery by the delivery team.
Whole of AI lifecycle	8	Apply watermarking techniques	23	Apply watermarks that are WCAG compatible where relevant	Required	Criterion Not Satisfied	DTG unable to confirm that IPA has watermarking policies or practices beyond standard Information Management Markers.
Whole of AI lifecycle	8	Apply watermarking techniques	24	Apply visual and accessible content to indicate when a user is interacting with an AI system	Required	Criterion Not Satisfied	DTG unable to confirm that IPA has watermarking policies or practices beyond standard Information Management Markers.
Whole of AI lifecycle	8	Apply watermarking techniques	25	For hidden watermarks, use watermarking tools based on the use case and content risk	Recommended	Criterion Not Satisfied	DTG unable to confirm that IPA has watermarking policies or practices beyond standard Information Management Markers.
Whole of AI lifecycle	8	Apply watermarking techniques	26	Assess watermarking risks and limitations	Recommended	Criterion Not Satisfied	DTG unable to confirm that IPA has watermarking policies or practices beyond standard Information Management Markers.

s47F

Design	9	Conduct pre-work	27	Define the problem to be solved, its context, intended use, and impacted stakeholders	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	This information is captured in project management plans, terms of reference, and other project artefacts. The CAAF requires a problem definition (section 2.1), a description of the use case purpose (section 2.2), a reflection on stakeholder impact (section 2.4), and its expected benefits (section 2.5).
Design	9	Conduct pre-work	28	Assess AI and non-AI alternatives	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF requires a discussion of non-AI alternatives (section 2.3). However, not all use cases complete the CAAF, and the IPA risk assessment does not have an equivalent section. Typically, this information is also considered in the high-level design documents, and is reviewed as part of some procurement requests.
Design	9	Conduct pre-work	29	Assess environmental impact and sustainability	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Both the CAAF and the IPA risk assessment require a discussion of possible environmental harms of the use case, in sections 3.4 and 3.3 respectively. However, it is unclear how increasing AI usage will affect IPA's commitment to achieving net zero, and Internal AI policies do not evaluate environmental impacts in-depth.
Design	9	Conduct pre-work	30	Perform cost analysis across all aspects of the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	IPA's Asset Management document gives guidelines for cost, and the Procurement process requires appropriate pricing to be considered. The agency has multiple governance measures in place that require review of initial and ongoing expenses for developing a system, and further information is also frequently captured in the business case, benefits profile, and by the Cloud Ops team.
Design	9	Conduct pre-work	31	Analyse how the use of AI will impact the solution and its delivery	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The use case card and CAAF/IPA risk assessment review the implications of integrating AI, identify the type of AI, and note relevant legislation. Project management plans, terms of reference, and other project artefacts also consider the system impacts.
Design	10	Adopt a human-centred approach	32	Identify human values requirements	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF and IPA risk assessment both require information about how individuals, communities, and groups will be affected by use of the AI. All AI integration is also governed by APS values. However, many use cases have not yet received feedback that may contribute towards further refining the systems.
Design	10	Adopt a human-centred approach	33	Establish a mechanism to inform users of AI interactions and output, as part of transparency	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	IPA has published an AI Transparency Statement, which is reviewed annually. Further, public-facing AI tools offered by IPA are marked with clear disclaimers.
Design	10	Adopt a human-centred approach	34	Design AI systems to be inclusive, ethical, and meet accessibility standards using appropriate mechanisms	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF and IPA risk assessment both require an assessment of the accessibility, inclusivity, and privacy of a use case, as well as its capacity to discriminate or harm. SMEs consider the given controls to ensure each use case meets agency standards. Agency staff have multiple points of contact for providing feedback during the pilot/initial phase of a tool. Further, feedback mechanisms are often built into production tools for users to give feedback at any time. However, it is unclear whether staff are confident providing feedback during the ongoing operational phase, and whether most staff can successfully navigate feedback mechanisms.
Design	10	Adopt a human-centred approach	35	Design feedback mechanisms	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF and IPA risk assessment both require product owners to identify any decisions made without human oversight and the potential for a use case to influence human decision-making, and the controls to manage these risks. Further, in many cases, the principal control for a risk is the inclusion of a human agent in the process.
Design	10	Adopt a human-centred approach	36	Define human oversight and control mechanisms	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	All solutions at the agency are business-driven. Development and delivery teams work closely with users throughout formulating the business case, the design process, and the proof-of-concept stage. Many production solutions are subject to ongoing revisions through user-driven, long-term continuous improvement programs.
Design	11	Design safety systematically	37	Involve users in the design process	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF and IPA risk assessment each require an analysis of potential harms to individuals, communities, groups, and the environment in sections 3.4 and 4 respectively. This assessment is endorsed by a range of SMEs and reviewed regularly thereafter.
Design	11	Design safety systematically	38	Analyse and assess harms	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	The CAAF and IPA risk assessment each require an analysis of the controls for potential harms to individuals, communities, groups, and the environment in sections 3.4 and 4 respectively. This assessment is endorsed by a range of SMEs and reviewed regularly. Controls are monitored and revised as necessary.
Design	11	Design safety systematically	39	Mitigate harms by embedding mechanisms for prevention, detection, and intervention	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	
Design	11	Design safety systematically	40	Design the system to allow calibration at deployment	Recommended	Insufficient information from DTA	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	DTG advised criteria unclear, could not provide a response.
Design	12	Define success criteria	41	Identify, assess, and select metrics appropriate to the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Project teams are expected to define their objectives and key results metrics in consultation with Business Delivery & Architecture during the initial stages. These are recorded in the project management plan and other project artefacts.
Design	12	Define success criteria	42	Reevaluate the selection of appropriate success metrics as the AI system moves through the AI lifecycle	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Project teams define their objectives and key results metrics during the early phases. These metrics are then monitored and reviewed through the proof-of-concept and pilot evaluations, and tools are reassessed after any model changes or retraining.
Design	12	Define success criteria	43	Continuously verify correctness of the metrics	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	AI use cases are subject to continuous benchmarking where appropriate, and tools are reassessed after any model changes or retraining. Additional metric monitoring and verification is used in some cases; for example, real-time performance dashboards. A number of presently built AI tools use data from the IP Data Platform. This is partially traceable using logs, but we are working to create practices around creating more traceability using the lineage tool that has been developed for the IPDP. For tools that are sourcing data from elsewhere, they are not currently traceable.
Data	13	Establish data supply chain management processes	44	Create and collect data for the AI system and identify the purpose for its use	Required	Criterion Partially Satisfied	Data Governance	IPA has guidelines for archiving and destroying data that align with the agency's legislative obligations under the Archives Act (1983). The process is outlined in Data and Information Governance Framework.
Data	13	Establish data supply chain management processes	45	Plan for data archival and destruction	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	A number of presently built AI tools use data from the IP Data Platform. This is partially traceable using logs, but we are working to create practices around creating more traceability using the lineage tool that has been developed for the IPDP. For tools that are sourcing data from elsewhere, they are not currently traceable.
Data	13	Establish data supply chain management processes	46	Analyse data for use by mapping the data supply chain and ensuring traceability	Recommended	Criterion Partially Satisfied	Data Governance	A number of data artefacts (Business Facts and Datamarts) have been built or are being built on the IP Data Platform for downstream data users. These are being leveraged for reporting and analytics use. However, some direct sourcing of data still occurs regularly.
Data	13	Establish data supply chain management processes	47	Implement practices to maintain and reuse data	Recommended	Criterion Partially Satisfied	Data Governance	Satisfied through the Data Asset Inventory and processes for allowing data access and sharing through Service Now. Access and retrieval are provided for under the Data and Information Governance Framework and the Data Asset Inventory Playbook. The Data Custodian Roles and Responsibilities, Usage and sharing falls under their purview. Data sharing procedures are also addressed in IP Australia Information Sharing Policy.
Data	14	Implement data orchestration processes	48	Implement processes to enable data access and retrieval, encompassing the sharing, archiving, and deletion of data	Required	Criterion Satisfied	Data Governance	There is currently an architecture decision which is reasonably well followed for the integration/ETL tools to use for different types of data orchestration and integration based on source and target domains for data.
Data	14	Implement data orchestration processes	49	Establish standard operating procedures for data orchestration	Recommended	Criterion Satisfied	Data Governance	Integration is done incrementally for the large majority of data workflows in the agency. The only exceptions are when the quantity of data is small enough to not worry about incremental loads and instead fall back to a full daily refresh.
Data	14	Implement data orchestration processes	50	Configure integration processes to integrate data in increments	Recommended	Criterion Satisfied	Data Governance	As discussed under item 49 - this is well governed based on the architecture decision around ETL tools.
Data	14	Implement data orchestration processes	51	Implement automation processes to orchestrate the reliable flow of data between systems and platforms	Recommended	Criterion Satisfied	Data Governance	Our data processes are tested before release and then validated continuously by data users.
Data	14	Implement data orchestration processes	52	Perform oversight and regular testing of task dependencies	Recommended	Criterion Satisfied	Data Governance	Data exchange processes (assuming this means external to the agency) are in a few different areas at the moment. There was a document prepared by DTG about all the ways data exits the organisation. There are a few processes but they are well documented, though we should aim to consolidate them.
Data	14	Implement data orchestration and feature engineering practices	53	Establish and maintain data exchange processes	Recommended	Criterion Partially Satisfied	Data Governance	All data cleansing is carried out by Data Engineering with approval from Data stewards. There is an established process for fixes and changes to data.
Data	15	Implement data transformation and feature engineering practices	54	Establish data cleansing procedures to manage any data issues	Recommended	Criterion Satisfied	Data Governance	IPA is guided by the Department of Finance's Data Maturity Assessment Tool. When required, information about data transformation processes is captured in the high-level design documents, the technical use case card, and other system documentation that is typically maintained in Confluence.
Data	15	Implement data transformation and feature engineering practices	55	Define data transformation processes to convert and optimise data for the AI system	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	When required, information about transformation between datasets is captured in the High-level design documents, the technical use case card, and other system documentation that is typically maintained in Confluence.
Data	15	Implement data transformation and feature engineering practices	56	Map the points where transformation occurs between datasets and across the AI system	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	
Data	15	Implement data transformation and feature engineering practices	57	Identify fit-for-purpose feature engineering techniques	Recommended			

		Implement data transformation and feature engineering practices	58	Apply consistent data transformation and feature engineering methods to support data reuse	Recommended			s47F	Under IPA's Performance Reporting Framework, owners must implement appropriate data quality control and assurance processes. The quality assessment criteria for data are generally use case-specific and depend on defined best practice for the use case, as well as the acceptable risk levels.
Data	16	Ensure data quality is acceptable	59	Define quality assessment criteria for the data used in the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Under IPA's Data Custodian Roles and Responsibilities document, monitoring and upkeeping data quality is one of the custodian's duties. The data profiling activities and remediation measures are generally use case-specific and depend on defined best practice for the use case, as well as the acceptable risk levels.
Data	16	Ensure data quality is acceptable	60	Implement data profiling activities and remediate any data quality issues	Recommended	Criterion Satisfied	Technologies; Enterprise Risk		
Data	16	Ensure data quality is acceptable	61	Define processes for labelling data and managing the quality of data labels	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Under the Information Custody Policy, data has several key aspects labelled and maintained in Confidence. The Data Asset Inventory also records various metadata. The labelling of data and quality management thereof are generally use case-specific and depend on defined best practice for the use case, as well as the acceptable risk levels.
Data	17	Validate and select data	62	Perform data validation activities to ensure data meets the requirements for the system's purpose	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		IPA's data validation activities are generally use case-specific and depend on the system requirements and the system's intended purpose, as well as defined best practice. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	17	Validate and select data	63	Select data for use that is aligned with the purpose of the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		IPA's data selection practices are generally use case-specific and depend on the intended purpose of the AI system and associated defined best practice, as well as the acceptable risk levels. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	18	Enable data fusion, integration and sharing	64	Analyse data fusion and integration requirements	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for analysing data fusion and integration requirements is generally use case-specific. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	18	Enable data fusion, integration and sharing	65	Establish an approach to data fusion and integration	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for data fusion and integration is generally use case-specific, guided by the Enterprise Pattern for Data Integration. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	18	Enable data fusion, integration and sharing	66	Identify data sharing arrangements and processes to maintain consistency	Recommended	Criterion Satisfied	Data Governance		We have policies for what data we can share and when (including IP Australia Information Sharing Policy). In addition there is information in the data asset inventory about whether we can share specific datasets.
Data	19	Establish the model and context dataset	67	Measure how representative the model dataset is	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for how to measure the representativeness of model datasets is considered in relation to a specific use case. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	19	Establish the model and context dataset	68	Separate the model training dataset from the validation and testing datasets	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for how to manage model datasets is considered in relation to a specific use case. The information is captured in the system documentation, which is typically maintained in Confidence.
Data	19	Establish the model and context dataset	69	Manage bias in the data	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for how to manage bias in data is considered in relation to a specific use case. IPA is guided by the APS Data Ethics Framework and our Internal Data and Information Governance Framework. The bias management information is captured in the system documentation, which is typically maintained in Confidence.
Data	19	Establish the model and context dataset	70	For generative AI, build reference or contextual datasets to improve the quality of AI outputs	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for building reference/contextual datasets to improve AI output quality is considered in relation to a specific use case. The information is captured in the system documentation, which is typically maintained in Confidence.
Train	20	Plan the model architecture	71	Establish success criteria that cover any AI training and operational limitations for infrastructure and costs	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Project artefacts such as the business case, project management plan, and benefits profile capture the success criteria for AI training and operational limitations. Costs are outlined in the first instance when the project is being approved and carefully monitored thereafter.
Train	20	Plan the model architecture	72	Define a model architecture for the use case suitable to the data and AI system operation	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for defining suitable model architecture is considered in relation to a specific use case. The information is captured in the high-level design and other system documentation, which is typically maintained in Confidence.
Train	20	Plan the model architecture	73	Select algorithms aligned with the purpose of the AI system and the available data	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for selecting appropriate algorithms is considered in relation to a specific use case and data. The information is captured in the system documentation, which is typically maintained in Confidence.
Train	20	Plan the model architecture	74	Set training boundaries in relation to any infrastructure, performance, and cost limitations	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for training boundaries is considered in relation to a specific use case and data. The information is captured in the system documentation, which is typically maintained in Confidence.
Train	20	Plan the model architecture	75	Start small, scale gradually	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for scaling projects is considered in relation to a specific use case and data. The relevant information is captured in the system documentation, which is typically maintained in Confidence. Projects at IPA are subject to extensive development and testing before any widespread rollout or expansion.
Train	21	Establish training environment	76	Establish compute resources and infrastructure for the training environment	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Best practice for establishing compute resources and infrastructure in the training environment is considered in relation to a specific use case and data. The relevant information is captured in the system documentation, which is typically maintained in Confidence.
Train	21	Establish training environment	77	Secure the infrastructure	Required		IT Security Advisor; Enterprise Risk		
Train	21	Establish training environment	78	Reuse available approved AI modelling frameworks, libraries, and tools	Recommended	Criterion Satisfied	Strategy & Architecture; Enterprise Risk		IPA's architectural principles consider reuse of available approved AI modelling frameworks, libraries, and tools.
Train	22	Implement model creation, tuning, and grounding	79	Set assessment criteria for the AI models, with respect to pre-defined metrics for the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Project teams are expected to define their objectives and key results metrics, including setting their assessment criteria, in consultation with Business Delivery & Architecture during the initial stages. These are recorded in the project management plan and other project artefacts.
Train	22	Implement model creation, tuning, and grounding	80	Identify and address situations where AI outputs should not be provided	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		IPA's Internal AI Governance Policy provides clear ethical guidelines and outlines appropriate uses for AI. The AIAADM governance process evaluates each AI use case, ensuring that the tool is applied to suitable situations and that human controls are in place as necessary.
Train	22	Implement model creation, tuning, and grounding	81	Apply considerations for reusing existing agency models, off-the-shelf, and pre-trained models	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		IPA's architectural principles consider the reuse of existing models and pre-trained models.
Train	22	Implement model creation, tuning, and grounding	82	Create or fine-tune models optimised for target domain environment	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to create or fine-tune models optimised for the target domain environment.
Train	22	Implement model creation, tuning, and grounding	83	Create and train using multiple model architectures and learning strategies	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to create and train using multiple model architectures and learning strategies.
Train	23	Validate, assess, and update model	84	Set techniques to validate AI trained models	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to set techniques to validate AI trained models.
Train	23	Validate, assess, and update model	85	Evaluate the model against training boundaries	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to evaluate the model against training boundaries.
Train	23	Validate, assess, and update model	86	Evaluate the model for bias, implement and test bias mitigations	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to evaluate the model for bias, implement and test bias mitigations.
Train	23	Validate, assess, and update model	87	Identify relevant model refinement methods	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		It is considered a standard data science practice implemented for all IPA projects to identify relevant model refinement methods.
Train	24	Selected training models	88	Assess a pool of trained models against acceptance metrics to select a model for the AI system	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Early-stage research and development and proof-of-concept processes involve considering and assessing a pool of trained models against acceptance metrics to select a model for the AI system.
Train	25	Implement continuous improvement frameworks	89	Establish interface tools and feedback channels for machines and humans	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk		Early stage research and development and proof-of-concept processes involve considering and establishing interface tools and feedback channels for machines and humans.

Train	25	Implement continuous improvement frameworks	90	Perform model version control	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Early-stage research and development and proof-of-concept processes involve considering and performing model version control
Evaluate	26	Adapt strategies and practices for AI systems	91	Mitigate bias in the testing process	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	It is considered a standard data science practice implemented for all IPA projects to mitigate bias in the testing process
Evaluate	26	Adapt strategies and practices for AI systems	92	Define test criteria approaches	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	It is considered a standard data science practice implemented for all IPA projects to define test criteria approaches
Evaluate	26	Adapt strategies and practices for AI systems	93	Define how test coverage will be measured	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	It is considered a standard data science practice implemented for all IPA projects to define how test coverage will be measured
Evaluate	26	Adapt strategies and practices for AI systems	94	Define a strategy to ensure test adequacy	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Strategies to ensure test adequacy are captured in use case-specific test plans and corporate testing processes
Evaluate	27	Test for specified behaviour	95	Undertake human verification of test design and implementation for correctness, consistency, and completeness	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Human verification of test design and implementation for correctness, consistency, and completeness is undertaken and captured in use case-specific test plans and corporate testing processes
Evaluate	27	Test for specified behaviour	96	Conduct functional performance testing to verify the correctness of the AI System Under Test (SUT) as per the pre-defined metrics	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Functional performance testing to verify correctness per defined metrics is conducted and captured in use case-specific test plans and corporate testing processes
Evaluate	27	Test for specified behaviour	97	Perform controllability testing to verify human oversight and control, and system control requirements	Required	Insufficient information from DTA	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	DTG advised criteria unclear, could not provide a response
Evaluate	27	Test for specified behaviour	98	Perform explainability and transparency testing as per the requirements	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	How explainability and transparency testing are performed specific to the use case, and relevant information is captured in the use case governance plan
Evaluate	27	Test for specified behaviour	99	Perform calibration testing as per the requirements	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	It is considered a standard data science practice implemented for all IPA projects to perform calibration testing
Evaluate	27	Test for specified behaviour	100	Perform logging tests as per the requirements	Required	Criterion Satisfied	Digital Delivery & Assurance; IT Security Advisor; Enterprise Risk	The Digital Delivery & Assurance corporate testing team conducts testing that addresses project requirements, including testing of edge cases and cases created by the team For each AI use case, the team identifies which events will be logged and determines whether alerts are required These alerts are then tested through deliberate triggering
Evaluate	28	Test for safety, robustness, and reliability	101	Test the computational performance of the system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Testing the computational performance of the system is considered a standard system development practice and implemented for all IPA projects
Evaluate	28	Test for safety, robustness, and reliability	102	Test safety measures through negative testing methods, failure testing, and fault injection	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; IT Security Advisor; Enterprise Risk	Testing safety measures through negative testing methods, failure testing, and fault injection is captured in use case-specific test plans and corporate testing processes
Evaluate	28	Test for safety, robustness, and reliability	103	Test reliability of the AI output, through stress testing over an extended period, simulating edge cases, and operating under extreme conditions	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; IT Security Advisor; Enterprise Risk	Testing the reliability of AI output through extended stress testing, simulating edge cases, and operation under extreme conditions is captured in use case-specific test plans and corporate testing processes This demonstrates whether edge cases can perform malicious activity, which can then facilitate the creation of preventative measures
Evaluate	28	Test for safety, robustness, and reliability	104	Undertake adversarial testing (red team testing), attempting to break security and privacy measures to identify weaknesses	Recommended	Criterion Satisfied	IT Security Advisor; Enterprise Risk	Penetration testing is used as required to assess and maintain the security and privacy of AI use cases
Evaluate	29	Test for conformance and compliance	105	Verify compliance with relevant policies, frameworks, and legislation	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	AI use cases are required complete a use case card and risk assessment, which serve to implement and facilitate compliance with internal and whole-of-government policies, frameworks, and legislation These governance documents are reviewed by a range of SMEs, including legal SMEs
Evaluate	29	Test for conformance and compliance	106	Verify conformance against organisation and industry-specific coding standards	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	How conformance against organisation and industry-specific coding standards is verified is currently specific to the project team IPA is investigating coding standards
Evaluate	29	Test for conformance and compliance	107	Perform vulnerability testing to identify any well-known vulnerabilities	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Security Operations perform vulnerability testing to identify any well-known vulnerabilities
Evaluate	30	Test for intended and unintended consequences	108	Perform user acceptance testing (UAT) and scenario testing, validating the system with a diversity of end-users in their operating contexts and real-world scenarios	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Whenever a change will affect the functionality of a tool or service, IPA engages in user acceptance testing and scenario testing to validate the system with end-users in their operating contexts
Evaluate	30	Test for intended and unintended consequences	109	Perform robust regression testing to mitigate the heightened risk of escaped defects resulting from changes, such as a step change in parameters	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Whenever a change will affect the functionality of a tool or service, IPA engages in regression testing to mitigate the risk of emergent defects
Integrate	31	Undertake integration planning	110	Ensure the AI system meets architecture and operational requirements with the Australian Government Security Authority to Operate (SATO)	Recommended	Criterion Satisfied	IT Security Advisor; Enterprise Risk	As part of the Authority to Operate, IPA requires AI use cases to complete a Security Risk Assessment with recommended treatments This aligns with the Australian Government Security Authority to Operate (SATO)
Integrate	31	Undertake integration planning	111	Identify suitable tests for integration with the operational environment, systems, and data	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Continuous integration and continuous delivery pipelines are in place across all development environments Further, corporate testing is conducting across all environments
Integrate	32	Manage integration as a continuous practice	112	Apply secure and auditable continuous integration practices for AI systems	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Continuous integration and continuous delivery pipelines are in place across all development environments Further, corporate testing is conducting across all environments
Deploy	33	Create business continuity plans	113	Develop plans to ensure critical systems remain operational during disruptions	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk	Our Enterprise Resilience Policy and Framework guides operational response to disruptions Disaster recovery, recoverability, and business continuity are all standard considerations during the architecture design process, although the specifics may vary depending on the nature of the use case
Deploy	34	Configure a staging environment	114	Ensure the staging environment mirrors the production environment in configurations, libraries, and dependencies for consistency and predictability	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Continuous integration and continuous delivery pipelines are in place across all development environments Further, corporate testing is conducting across all environments
Deploy	34	Configure a staging environment	115	Measure the performance of the AI system in the staging environment against predefined metrics	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Generally, the predefined metrics that measure the performance of the AI system are use case-specific and arise from a considered best practice approach The performance information is captured in the system documentation, which is typically maintained in Confluence
Deploy	34	Configure a staging environment	116	Ensure deployment strategies include monitoring for AI specific metrics, such as inference latency and AI output accuracy	Recommended	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Generally, the AI-specific monitoring integrated into deployment strategies is use case-specific and arise from a considered best practice approach The deployment information is captured in the system documentation, which is typically maintained in Confluence
Deploy	35	Deploy to a production environment	117	Apply strategies for phased roll-out	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Strategies for a phased roll-out are applied when required This process is guided by IP Australia's Project & Program Management Framework, and documented in the project management plan, as well as other relevant project artefacts
Deploy	35	Deploy to a production environment	118	Apply readiness verification, assurance checks and change management practices for the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk	Readiness verification, assurance checks, and change management practices are specific to the use case, and are documented in the project management plan
Deploy	35	Deploy to a production environment	119	Apply strategies for limiting service interruptions	Recommended	Criterion Satisfied	Digital Delivery & Assurance; Enterprise Risk	Most releases are deployed without outages For deployments that require outages, IPA has defined specific maintenance windows and communicated these to stakeholders to limit service disruptions
Deploy	36	Implement rollout and safe rollback mechanisms	120	Define a comprehensive rollout and rollback strategy	Recommended	Criterion Satisfied	Digital Delivery & Assurance; Enterprise Risk	The high-level design template asks for a description of transitional phases and for a change roadmap to be included Further, a backup plan, either rollback or roll-forward, must be included for each change, detailing the preferred approach
Deploy	36	Implement rollout and safe rollback mechanisms	121	Implement load balancing and traffic shifting methods for system rollout	Recommended	Criterion Partially Satisfied	DTG	This is team specific Part of current AWS deployments

Deploy	36	Implement rollout and safe rollback mechanisms	122	Conduct regular health checks, readiness, and startup probes to verify stability and performance on the deployment environment	Recommended	Criterion Satisfied	Digital Delivery & Assurance; Enterprise Risk
Deploy	36	Implement rollout and safe rollback mechanisms	123	Implement rollback mechanisms to revert to the last stable version in case of failure	Recommended	Criterion Satisfied	Digital Delivery & Assurance; Enterprise Risk
Monitor	37	Establish monitoring framework	124	Define reporting requirements	Recommended	Criterion Satisfied	Enterprise Risk
Monitor	37	Establish monitoring framework	125	Define alerting requirements	Recommended	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Monitor	37	Establish monitoring framework	126	Implement monitoring tools	Recommended	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Monitor	37	Establish monitoring framework	127	Implement feedback loop to ensure that insights from monitoring are fed back into the development and improvement of the AI system	Recommended	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	128	Test periodically after deployment and have a clear framework to manage any issues	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	129	Monitor the system as agreed and specified in its operating procedures	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	130	Monitor performance and AI drift as per pre-defined metrics	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	131	Monitor health of the system and infrastructure	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	132	Monitor safety	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	133	Monitor reliability metrics and mechanisms	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	134	Monitor human-machine collaboration	Required	Insufficient information from DTA	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	135	Monitor for unintended consequences	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	136	Monitor transparency and explainability	Required	Criterion Partially Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	137	Monitor costs	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	138	Monitor security	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	38	Undertake ongoing testing and monitoring	139	Monitor compliance of the AI system	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	39	Establish incident resolution processes	140	Define incident handling processes	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Monitor	39	Establish incident resolution processes	141	Implement corrective and preventive actions for incidents	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk
Decommission	40	Create a decommissioning plan	142	Define the scope of decommissioning activities	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	40	Create a decommissioning plan	143	Conduct an impact analysis of decommissioning the target AI system	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	40	Create a decommissioning plan	144	Proactively communicate system retirement	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	41	Shut down the AI system	145	Retain AI system compliance records	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	41	Shut down the AI system	146	Disable computing resources or components specifically dedicated to the AI system	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	41	Shut down the AI system	147	Securely decommission or repurpose all computing resources specifically dedicated to the AI system, including individual and shared components	Required	Criterion Satisfied	Strategy & Architecture; Enterprise Risk
Decommission	42	Finalise documentation and reporting	148	Finalise decommissioning information and update organisational documentation	Required	Criterion Satisfied	Business Transformation & New Products; AI and Frontier Technologies; Enterprise Risk

IPA has implemented automated testing suites to verify stability and performance in the deployment environment. These suites run scheduled checks daily, and where a test fails, nominated individuals are notified to investigate the issue

IPA requires specific backup plans to be recorded against each change record, detailing steps on the preferred backup method. This means that measures can be implemented to reestablish a stable version in case of failure

IPA's *ITS07 Monitoring and Reporting Policy* outlines reporting requirements. Further, the agency's risk-scaled governance defines the internal reporting and escalation of issues, including the identification of risk stewards. All use cases are reviewed by a diverse Panel of SMEs. Per the *Responsible use of AI in government policy*, any high-risk use cases are reported to the DTA. IPA's Continuous Monitoring Policy notes that a Continuous Monitoring Plan defines the alerting requirements for a system. Alerting is a standard consideration during the architecture design process, and varies according to business needs and the nature of a use case.

IPA uses Splunk and Dynatrace as monitoring tools, has security logging, and restricts/tracks access

IPA establishes tool-specific alerts and environment alerts, which are subsequently monitored. These alerts are tailored to the use case. However, these loops are only established as required, and are not implemented universally across all use cases.

Testing after deployment and managing issues in alignment with given frameworks is generally the responsibility of the platform owner. The platform owner can seek support from a range of business areas and SMEs.

Monitoring the system as specified in its operating procedures is generally the responsibility of the platform owner. The platform owner is guided and informed by IPA's Continuous Monitoring Policy.

Best practice for monitoring performance and AI drift is considered in relation to a specific use case and its metrics. The information is captured in the system documentation, which is typically maintained in Confluence.

Human agents can report any apparent issues with system and infrastructure. Splunk and/or Dynatrace monitoring can be implemented as required. The Continuous Monitoring Policy indicates appropriate timeframe responses to exploitable system and infrastructure issues.

The CAAF and IPA risk assessments analyse the potential harms of the use case and document the instituted controls. Further, IPA has an Agency Security Advisor to whom any incident that impacts staff safety is reported, as well as robust Workplace Health and Safety policies and IT security policies that create a safe environment. However, there is scope to extend and expand these safety measures in response to feedback.

Monitoring reliability metrics and measurements is generally the responsibility of the platform owner. Splunk and/or Dynatrace monitoring can be implemented as required.

DTG advised criteria unclear, could not provide a response

AI use cases are reviewed regularly, with the review cycle being relative to their risk level, with low-risk use cases being reviewed annually. This review would provide a platform for any unintended consequences to be considered. Measures could be subsequently detailed in a governance plan as necessary. However, there is no formal protocol in place to monitor specifically for unintended consequences.

The transparency and explainability requirements of all AI at IPA is outlined in the agency's AI Transparency Statement. Further transparency and explainability requirements could be introduced into the use case governance plan and be initiated as necessary.

The costs of a given AI system are outlined in the business case and/or benefits profile of the Project Management Plan. The ongoing costs are monitored in Cloud FinOps, and the Financial Business Partner from FP&G monitors finances in consultation with product owners. Cost management is further governed by existing financial assurance mechanisms.

This is governed by the *ITS07 Monitoring and Reporting Policy* under the Security Policy. Further, per the Continuous Monitoring Policy, the Cyber Security Operations team use automated and manual processes to identify any vulnerabilities.

AI use cases are reviewed by the AI/ADM Panel regularly, with the review cycle being relative to their risk level, with low-risk use cases being reviewed annually. This review assesses many aspects of the AI use case, including compliance with internal policies and whole-of-government policies.

Defining AI-related incident handling processes is the responsibility of the Service Centre under the Infrastructure & Technology Operations. These processes are governed by the Protective Security Incident Management Procedure and the corresponding Incident Management Template and workflow.

Management of AI-related incidents, including corrective and preventive actions, are governed by the Enterprise Resilience Policy and Framework and the ICT Disaster Recovery Strategy. All Tier 1 and Tier 2 incidents require a formal debrief and a Post Incident Review.

Business Delivery & Architecture manages decommissioning of assets. The Digital Delivery & Assurance team oversees the decommissioning processes and ensures compliance.

Impact analysis is part of the decommissioning process. The Digital Delivery and Assurance team oversees the decommissioning processes and ensures compliance.

It is standard practice at IPA to communicate clearly to staff when products or systems are being retired. Staff are provided with warnings, reminders, and guides on how to take any necessary steps to transition to a new system.

Records are kept as required under the Archives Act 1983 and are a responsibility under records management and information governance. Decommissioned AI use cases remain in the AI/ADM Inventory, with their essential documentation still viewable. The Information Management Policy outlines IPA's record destruction guidelines.

Disabling computing resources and components previously dedicated to the AI system is a standard part of the decommissioning process. The Digital Delivery and Assurance team oversees the decommissioning processes and ensures compliance.

Securely decommissioning and repurposing computing resources previously dedicated to the AI system, including individual and shared components, is a standard part of the decommissioning process. The Digital Delivery and Assurance team oversees the decommissioning processes and ensures compliance.

When an AI use case is decommissioned, the end-of-lifecycle process discussed above terminates with the AI/ADM Inventory records being updated to reflect the date and reason of the use case being decommissioned.



Preliminary Compliance assessment

Policy for responsible use of AI in Government

(version 2.0)

Purpose

Assess current compliance and determine action required to reach full compliance with all mandatory requirements of the updated *Policy for responsible use of AI in Government*, released in December 2025.

Proposed actions – pending detailed governance review

Action	Responsible	Requirement Reference
Due 15 June 2026		
Strategic position on AI Adoption	PTE project #3 / DDG	#2 – strategic position on AI adoption
Share AI & ADM Inventory with DTA (AI Use cases only)	GG	#5 – Internal AI use case register #16 – Internal AI use case register
Update AI Governance Policy, templates & tools	GG	#4 – Accountable use case owners #6 – Operationalise the responsible use of AI #8 – AI use case impact assessment #15 – Accountable Use Case Owners #16 – Internal use case register
Due 15 December 2026		
Complete entry of all existing use cases into the AI & ADM Inventory	GG	#5 Internal AI use case register
Establish pathways for staff and public to report AI safety concerns	GG & DTG	#6 Operationalise the responsible use of AI
Establish clear processes to address AI incidents aligned to their ICT incident management approach	GG & DTG	#6 Operationalise the responsible use of AI
Mandatory training for all staff on responsible AI use	HR	#6 Operationalise the responsible use of AI
Due 30 April 2027		
Assessment of existing use-cases against updated criteria to determine if they are “in-scope” of whole of government policy. For AI use cases that are in-scope, agencies must conduct an AI use case impact assessment.	GG, in collaboration with product owners & AI/ADM Assessment panel	#4 – Accountable use case owners
Due date TBC		
Establish process to ensure DDG (Accountable AI Officer) has visibility of the AI & ADM Inventory	GG	#15 – Accountable Use Case Owners

Assessment current as at 9 February 2026.



Policy for responsible use of AI in Government - Mandatory requirements

KEY

Compliant	Existing policies/practices meet the requirement
Action required	Change required to ensure compliance

#	Requirement		Current state	Compliant/Action required	Required action	Responsible	Due
1	AI Transparency statement Standard for transparency statements.	Agencies must make a publicly available statement outlining their approach to AI adoption and use, as prescribed under the Standard for transparency statements .	Published here Accountability and reporting IP Australia 27 Feb 2025	Compliant	Continue with scheduled annual review. Refer Agenda Item 5.	GG	27 February 2026
		The statement must be reviewed and updated annually or sooner, should the agency make significant changes to its approach to AI.					
		Agencies must notify the DTA when they publish and make any changes to their AI transparency statement by emailing ai@dtg.gov.au .	DTA notified when original published.	Compliant	Continue with planned approach to notify DTA when annual review complete.	GG	27 February 2026
2	Strategic position on AI adoption	Agencies must develop a strategic position on AI adoption within 6 months of this policy taking effect. This position is to emphasise how AI opportunities can be identified and embraced by the agency.	Under development as part of P+T+E project 3 led by s47F , working closely with the CEG team on the communication elements.	ACTION REQUIRED	Develop a strategic position on AI adoption.	P+T+E project 3	15 June 2026
		Agencies must communicate their strategic position on AI to give staff clear direction on AI adoption.		ACTION REQUIRED	Communicate strategic position to staff.	P+T+E project 3	15 June 2026
3	Accountable officials	Agencies must designate accountable official(s) to take accountability for implementing this policy.	Margaret Tregurtha designated as accountable official.	Compliant	nil	nil	nil
		Agencies must follow the Standard for accountability when designating	See Standard for accountability assessment below	ACTION REQUIRED	See Standard for accountability assessment below		

#	Requirement		Current state	Compliant/Action required	Required action	Responsible	Due
		accountable official(s) and implementing this requirement. The responsibilities of accountable officials are set in the standard.					
		Agencies must notify the DTA when they designate and make any changes to their accountable official(s) by emailing ai@dtg.gov.au .	DTA notified when Margaret originally designated as our accountable official.	Compliant	nil	GG	nil
4	Accountable use case owners	Agencies must designate an accountable use case owner for each in-scope AI use case within 12 months of this policy taking effect.	In progress. Captured in the use case cards and AI/ADM Inventory.	ACTION REQUIRED	Complete use case cards for historic use cases to determine appropriate use case owners.	GG, in collaboration with product owners	15 December 2026
		Accountable official(s) are to maintain a register of accountable use case owners.			Update AI/ADM Governance policy & template language to make consistent with this policy. Various language including 'decision maker' and 'risk steward' have been used.		
		Agencies must follow the Standard for accountability when implementing this requirement. The responsibilities of accountable use case owners are set in the standard.	See Standard for accountability assessment below	ACTION REQUIRED	See Standard for accountability assessment below		
5	Internal AI use case register	Agencies must create a register of in-scope AI use cases to enable accountable official(s) to record accountable use case owners within 12 months of this policy taking effect.	Register has been created (AI/ADM Inventory) but is not yet fully populated for historic use cases.	ACTION REQUIRED	Continue data entry of stocktake into inventory. Complete use case cards for historic use cases.	GG	15 December 2026
		Agencies must share the register with the DTA every 6 months, commencing from when they create the register to meet the above requirement.			Share inventory with DTA.	GG	15 June 2026

#	Requirement		Current state	Compliant/Action required	Required action	Responsible	Due
6	Operationalise the responsible use of AI	Agencies must establish an approach to embed responsible AI practices within 12 months of this policy taking effect. This may vary according to the scale and scope of agency AI use. At a minimum, the approach will provide an agency with:	AI & ADM Governance policies processes established and implemented.	ACTION REQUIRED	Updates to policy required, as detailed in this assessment.	GG	15 December 2026
		a process for adopting AI use cases in line with the implemented actions of this policy, as well as the agency's enterprise risk management and governance approach.					
		a way to inform staff who are designing and implementing AI use cases about Australia's AI Ethics Principles.	Available to staff via intranet and referenced in AI & ADM Governance policies.	ACTION REQUIRED	Make more explicit via standardised comms through the risk assessment process.	GG	
		a pathway for staff to report AI safety concerns, including AI incidents.	General security incident response capability/procedure, including service now issue reporting capability.	ACTION REQUIRED	Review established capability to ensure AI coverage and fit for purpose.	GG/DTG	
		pathways for the public to report AI safety concerns, appropriate to the agency's AI use.	Standard customer complaint process. Existing pathway for public to report vulnerabilities.	ACTION REQUIRED	Review established capability to ensure AI coverage and fit for purpose.	GG/DTG	
		clear processes to address AI incidents aligned to their ICT incident management approach - incident remediation must be overseen by an appropriate governance body or senior executive and should be undertaken in line with any other legal obligations.	Cyber security incident response plan & playbooks in place.	ACTION REQUIRED	Review established capability to ensure AI coverage and fit for purpose.	GG/DTG	
7	Staff training on AI	Agencies must implement mandatory training for all staff on responsible AI use within 12 months of this policy taking effect.	In progress.		Implement mandatory training on responsible AI use.	HR	15 December 2026

#	Requirement		Current state	Compliant/Action required	Required action	Responsible	Due
8	AI use case impact assessment	Agencies must assess all new AI use cases against the in-scope criteria (Appendix C) to determine if they are in scope of the policy. The assessment must be documented and take place during the design phase while developing requirements. Agencies must begin AI use case assessments within 12 months of this policy taking effect. For existing use cases not yet assessed, agencies must determine whether they are in scope of this policy and apply all relevant policy actions by 30 April 2027.	AI use cases assessments against 'in scope criteria' are underway.	ACTION REQUIRED	Complete historic assessments.	GG, in collaboration with product owners.	30 April 2027
		For AI use cases that are in-scope, agencies must conduct an AI use case impact assessment. Agencies are to commence an assessment at the design stage. Before the solution is deployed, agencies must finalise the assessment and apply any agreed risk treatments.	AI use case impact assessments underway for in-scope use cases.	ACTION REQUIRED	Update CAAF templates to new impact assessment template for all use cases moving forward (not impact assessments already completed).	GG	nil
		Agencies must be able to revise their internal process in response to any impact assessment tool updates.	We can revise internal processes as required.	Compliant	nil	GG	nil
		Agencies must add each in-scope AI use case to their internal register of AI use cases and update it as required. Include risk rating and accountable use case owner changes. When deploying an in-scope AI use case, agencies must : - regularly monitor and evaluate their use case to ensure it is operating as intended and that risks are being effectively managed. - re-validate the AI use case impact assessment by checking its accuracy and updating it when there is a	All use cases (in or out of scope) are added to our AI/ ADM inventory. All use cases are regularly monitored and impact/risk assessments are reviewed at least annually.	ACTION REQUIRED	Clarify procedure for reviewing use cases.	GG	Nil

#	Requirement	Current state	Compliant/Action required	Required action	Responsible	Due
	material change in the use case scope, usage or operation.					
	If an agency determines their in-scope AI use case has an inherent high-risk rating when completing an AI use case impact assessment, they must: report the use case to the agency accountable official with the reasons for the inherent high-risk rating, proposed mitigations and residual risks govern the use case through a designated board or a senior executive, whichever is appropriate for the size and scope of the agency. Once an agency has decided to deploy the use case, they must: report the use case to the DTA through the accountable official, see the Standard for accountability establish a system to regularly review the use case every 12 months at a minimum. The review must report to the relevant governing board or senior executive on whether the use case is operating as intended and that risks are being effectively managed. The review must also consider the AI use case impact assessment and revisions to it, if required.	All high risk use cases must be signed off by the DDG or DG. Current policies require that all high risk use cases are reported to DTA. High risk use cases are required to be reviewed every 3 months.	ACTION REQUIRED	Clarify in AI & ADM Governance Policies that high risk use cases must be reported to AI Accountable Official (DDG). Emphasise requirements to record mitigations/treatments/controls for high risks. Clarify ongoing governance requirements for high risk use cases. Clarify procedure for review of use cases.	GG	?
	Appendix C: In-scope AI use cases criteria The use, misuse or failure of AI could lead to more than insignificant harm to individuals, communities, organisations, the environment or the collective rights of cultural groups including First Nations peoples.	In scope criteria from pilot CAAF currently being used: The estimated whole of life cost of the project or service incorporating the use of AI is more than \$10 million.	ACTION REQUIRED	Update use case card template to reflect new 'in scope' criteria. Reassess existing use cases to determine if they are now 'covered'	GG	TBC

#	Requirement		Current state	Compliant/Action required	Required action	Responsible	Due
		- The use of AI will materially influence administrative decisions that affect individuals, communities, organisations, the environment or the collective rights of cultural groups including First Nations peoples. - It is possible the public will directly interact with, or be significantly impacted by, the AI or its outputs without human review. - The AI is designed to use personal or sensitive data¹ or security classified information². - It is deemed an elevated risk AI use case as directed by the DTA.	- It is possible the use of AI will lead to more than insignificant harm to individuals, communities, organisations or the environment. - The use of AI will materially influence decision-making that affects individuals, communities, organisations or the environment. - It is possible the AI will either directly interact with the public or produce output that are not subject to human review prior to publication.		under WoG Impact assessment.		

Standard for AI transparency statements – Mandatory requirements (Due 27 Feb 2026)

#	Requirement		Current state	Compliant/Action required	Required action	Responsible
10	Minimum information	the intentions behind why the agency uses AI or is considering its adoption	Included.	Compliant	Nil	GG/DTG
		classification of AI use according to usage patterns and domains, as listed at Attachment A	High level summary of current state as at 27 Feb 2025	ACTION REQUIRED	Update according to info in the Inventory	GG
		classification of use where the public may directly interact with, or be significantly impacted by, AI or its outputs without human review	High level summary of current state as at 27 Feb 2025	ACTION REQUIRED	Update according to info in the Inventory	GG
		measures to monitor the effectiveness of deployed AI systems and protect the public against negative impacts	High level summary of current state as at 27 Feb 2025	ACTION REQUIRED	Update according to current policies and procedures	GG/DTG
		overview of compliance with the requirements under the Policy for responsible use of AI in government	Correct as per old policy	ACTION REQUIRED	Update according to this assessment	GG
		compliance with applicable legislation and regulation	High level summary of current state as at 27 Feb 2025	ACTION REQUIRED	Update according to current policies and procedures	GG
		when the statement was most recently updated.	n/a	Compliant	Continue with existing planned approach.	GG
		They must also provide or direct to a contact email for further public enquiries.	For any enquires regarding IP Australia’s use of AI, please contact MDB-Enterprise-Risk@ipaustalia.gov.au	ACTION REQUIRED	Update to new AI/ADM inbox	GG
11	Format	Statements must use clear, plain language that avoids technical jargon and is consistent with the Australian Government Style Manual.	Statement is clear and generally implements plain language consistent with the style manual.	ACTION REQUIRED	Review final updated statement to ensure consistency with this requirement	GG
12	Publication	Agencies must publish transparency statements on their public facing website.	n/a	Compliant	Continue with existing planned approach.	GG
		Agencies must send the DTA a link to the statement when it is published or updated by emailing ai@dtg.gov.au .	n/a	Compliant	Continue with existing planned approach.	GG
13	Review	Transparency statements must be reviewed and updated at these junctures: at least once a year when making a significant change to the agency’s approach to AI when any new factor materially impacts the existing statement’s accuracy.	“This statement will be updated at least once a year, or sooner if we make any significant change to our approach to AI, or if any new factor materially impacts the existing statement’s accuracy”.	Compliant	Continue with existing planned approach.	nil

Standard for accountability – mandatory requirements

#	Requirement		Current state	Compliant/Action required	Required action	Responsible
14	Accountable officials	Agencies must designate accountability for implementing the policy to accountable official(s) (AOs), who must: - be accountable for implementation of the policy within their agencies - notify the Digital Transformation Agency (DTA) where the agency has identified a new high-risk use case by emailing ai@dta.gov.au - be a contact point for whole-of-government AI coordination and respond to DTA requests for information - engage in whole-of-government AI forums and processes - keep up to date with changing requirements as they evolve over time.	Margaret designated as Accountable official.	Compliant	Nil	n/a
		Agencies must notify the DTA of AO selection including the contact details of all their AOs at initial selection and when the accountable roles change. Notify DTA by emailing ai@dta.gov.au .	DTA notified with first transparency statement on 27 Feb 2025when Margaret was first designated.	Compliant	Nil	n/a
		In the event their agency has decided to deploy a new use case with an inherent high risk rating , AOs must notify the DTA, by emailing ai@dta.gov.au . AOs must also notify the DTA when an existing AI use case has been re-assessed as having an inherent high risk, or when use case is no longer high risk.	No high risk use cases identified so far.	Compliant	Nil	n/a
		AOs must participate in, or nominate a delegate for, whole-of-government forums and processes which support collaboration and coordination on current and emerging AI issues. These forums will be communicated to AOs as they emerge.	Participation ongoing.	Compliant	Nil	n/a
15	Accountable Use Case Owners	Accountable use case owners must : ensure their use case is registered with the agency's AO	All use cases are 'registered' with the AO via the Inventory.	ACTION REQUIRED	Establish process to ensure Margaret has visibility of the inventory.	GG

#	Requirement	Current state	Compliant/Action required	Required action	Responsible	
		- be accountable for applying the actions under the AI use case impact assessment requirements. These include: - conducting an AI use case impact assessment - regularly monitoring and evaluating the use case - re-validating the assessment when required - applying the high-risk use case actions if their use case has an inherent high-risk rating.	Roles and responsibilities of 'risk stewards' (use case owners) documented in AI & ADM Governance Policies.		Update use case owner approval minute to clarify these responsibilities. Update use case owner approval minute to clarify that even low risk use cases require use case owner sign off acknowledging their responsibilities as a use case owner.	
		Accountable use case owners must also be familiar with Australia's AI Ethics Principles, the Australian Government Impact Assessment Tool and the policy.	General references to these materials in policy documents.	ACTION REQUIRED	Update use case owner approval minute to clarify this responsibility.	GG
		Accountable use case owners of high-risk use cases must have the ability to identify and manage risks and emerging issues.	Accountable use case owner of all high-risk use cases is DDG, Margaret Tregurtha, AI Accountable Official.	Compliant	Nil	n/a
16	Internal use case register	The register must include the following fields: use case name - agency identifier (reference number) - description - AI technology type - lifecycle stage - use of Technical standard for government's use of artificial intelligence - domain - usage pattern - accountable use case owner - what criteria the use case met to be in scope of the policy - inherent risk rating - residual risk rating	AI/ADM inventory current includes the following fields: - Description - AI technology type - Domain - Usage pattern - accountable use case owner - what criteria the use case met to be in scope of the policy - Inherent risk rating - Date of last review	ACTION REQUIRED	Update use case cards to also capture: - agency identifier (reference number) - use of Technical standard for government's use of artificial intelligence - Residual risk rating - Date for next review	GG

#	Requirement	Current state	Compliant/Action required	Required action	Responsible
	- the date when the AI impact assessment was last updated (if applicable). For AI use cases with an inherent high-risk rating, the register must also include: - the last date of review - date for next review.				
	Agencies must share the register with the DTA every 6 months, commencing from when they create the register to meet the policy requirement. They can share the register by emailing ai@dtg.gov.au	Not started	ACTION REQUIRED	Share in-scope AI items from AI & ADM inventory with DTA	GG