25OIA1887

Inland Revenue
Te Tari Taake

14 March 2025

[redacted]

Dear [redacted]

Thank you for your request made under the Official Information Act 1982 (OIA), received on 3 February 2025. On 14 February 2025 you clarified sections 2.2 and 4.1 of your request. Your full request and clarification are attached as **Appendix A.**

Question One - AI Applications and Systems

Items 1.1, 1.2 Specific use of AI systems by Inland Revenue, and 1.3 AI technologies currently in use by Inland Revenue

Inland Revenue's business transformation programme has resulted in excellent digital capability and high-quality data across our broad customer base. The quality of data enables us to utilise AI to allow for further efficiencies and improve customer convenience. While extending the use of AI, Inland Revenue takes a benefit and value-focused approach. Our use of AI focuses on supporting decision-making in the management of tax/social policy compliance risk, intervention design and on increasing staff productivity.

The table on the following page details all current uses of AI within Inland Revenue, and whether the tools are for internal use, supporting staff, or supporting customers directly. The use of one predictive model is refused under section 18(c)(i) of the OIA, as making the requested information available would be contrary to the provisions of section 18(3) of the Tax Administration Act 1994 (TAA). The Commissioner of Inland Revenue is not required to disclose any item of revenue information if the release of the information would adversely affect the integrity of the tax system or prejudice the maintenance of the law.

Table 1: Uses of AI within Inland Revenue

Name	Description	Product Status	AI Technology	Use
ABBYY FineReader 16	Text recognition and document conversion tool, used to convert PDFs into excel.	In production	Optical Character Recognition	Staff Support
AI Futurist	Enables querying and summarisation of content.	Available	Generative	Staff Support
Āwhina Mai	Customised guidance for customers completing forms and reviewing guides on Inland Revenue's websites	Pilot	Generative	Supporting Customers Directly
Assurity Intelligence	Test scenario generation	Proof of concept	Generative	Staff Support
Copilot Chat (Bing/Browser)	AI-powered chat service	Available	Generative	Staff Support
Coveo	Inland Revenue's public websites and internal staff intranet and knowledge base use Coveo as a search platform. Coveo uses machine learning to continuously learn and improve from user searches and patterns to inform the best results to display.	In production	Machine Learning	Supporting Customers Directly and Staff Support
Creative desktop applications – Adobe Acrobat Pro DC **Adobe Photoshop** **Adobe Premiere Pro** **Adobe InDesign** **Adobe Captivate** **Adobe Illustrator** **Adobe Lightroom**	Understanding/making use of large data sets, creation or generation of new content primarily for training and marketing purposes.	Available	Generative Machine Learning	Staff Support

Name	Description	Product Status	AI Technology	Use
Adobe After Effects **Adobe Audition** **Adobe Media Encoder**				
Da Vinci AI	Provides forecasting for the voice channel, including long-term predictions and insights into how changes in one voice queue affect another.	Pilot	Machine Learning	Staff Support
Receipt, invoice, statement and tax/employer return review	Text recognition	In production	Optical Character Recognition	Internal Use
Dragon Naturally Speaking	Screen reader	In production	Machine Learning	Staff Support
DDoS protections **- AWS Shield** **- Azure DDoS Protection** **- Cloudflare** **- Magictransit** **- F5 BigIP** **- Oracle DDoS**	Inland Revenue uses a range of tools to prevent our systems and services from unexpected outages due to network attacks.	In production	Machine Learning	Internal Use
Email-user link prediction	Predict the similarity of email addresses and usernames such that Inland Revenue can identify the probability that they are controlled / used by the same real-world person.	In development	Machine Learning	Internal Use and Staff Support
Figma	Prototyping software that enables Inland Revenue to develop mock-ups of intended	In production	Machine Learning	Staff Support

Name	Description	Product Status	AI Technology	Use
	changes to products and services across both e-services and internal/external Inland Revenue websites			
Financial intelligence network detection	Links, matches and identifies multi-dimensional risks of Crypto asset users via operational and strategic visualisation	In production	Machine Learning	Staff Support
Genesys Agent Assist	Creates summaries of conversations with contact centre agents for post-call notes.	Pilot	Generative	Staff Support
Genesys Agent Copilot	Group of functions to support contact centre agents including: surfacing relevant knowledge content, recommended next best actions, and predicted subject code.	Proof-of-concept	Natural Language Understanding	Staff Support
GST integrity model	A predictive model to assess the risk of GST returns requesting refunds.	In production	Machine Learning	Internal Use
Graph Entity Resolution	Analyses and compares information held by Inland Revenue to external datasets provided by third parties to determine if records are referencing the same entity	In production	Machine Learning	Internal Use
Microsoft 365 Copilot	Copilot is integrated into the M365 suite of products and is designed to enhance staff productivity.	Pilot	Generative	Staff Support
Microsoft Defender	An enterprise-capable host protection solution that is integrated with a range of other Microsoft Apps, observes activity on devices for potential malicious behaviour.	In production	Machine Learning	Internal Use

Name	Description	Product Status	AI Technology	Use
Microsoft Teams Intelligent Recap	Provides recaps of Teams meetings.	Pilot	Generative	Staff Support
Microsoft power BI	Dashboard/reporting software that connects to multiple Azure services	Available	Machine Learning	Staff Support
Microsoft Purview	Portfolio of products that span Data governance, data security, and risk and compliance solutions.	In development	Machine Learning	Internal Use
Nuix	Digital forensics and analysis of information.	In production	Machine Learning Natural Language Processing	Staff Support
Overdue income tax return RIT prediction	Predicts residual income tax (RIT) on overdue returns.	In production	Machine Learning	Staff Support
Qualtrics	Analyses customer feedback from survey data.	In production	Machine Learning Natural Language Processing	Staff Support
Power Automate	Low Code solution that supports automating tasks.	In production	Machine Learning	Staff Support
Propensity to read letter or log-in to myIR	Helps Inland Revenue to select and use the right channels to communicate with customers.	In production	Machine Learning	Staff Support
Tableau Desktop	Data analytics and graphing tool used for analysing and visualising performance test results.	In production	Machine Learning Natural Language Processing	Internal Use

Name	Description	Product Status	AI Technology	Use
Viva Engage	Provides people with personalised content feeds and recommendations of communities to join or follow.	In production	Machine Learning	Internal Use
Viva Topics	Aggregates resources and expertise Inland Revenue holds on a range of topics.	Pilot	Machine Learning	Staff Support
Voice biometrics	Identity confirmation for contact centre calls.	In production	Biometrics	Supporting Customers Directly
Windows Hello for business	Authentication for Inland Revenue devices.	In production	Biometrics	Internal Use
Z scaler	Detection and classification of web traffic and websites.	In production	Machine Learning	Internal Use
ZoomText	Screen magnification software for accessibility purposes.	In production	Optical Character Recognition	Staff Support

Item 1.4 Please specify the types of AI technologies currently in use by IRD for international tax information exchange.

Inland Revenue does not use AI technology for international tax information exchange. Therefore, your request for this information is refused under section 18(g)(i) of the OIA, as the information is not held by Inland Revenue, and we do not believe it is held by another agency.

Item 1.5 Please specify the purposes for which IRD uses the following functions of AI systems:

Item 1.5a Web scraping or web-crawler applications

Inland Revenue does not use web scraping or web-crawler applications. Therefore, your request for this information is refused under section 18(g)(i) of the OIA.

Item 1.5b Social network analysis and 1.5c Transaction network analysis

Inland Revenue uses social and transaction network analysis to study the network of taxpayers. This helps us to infer, predict, and identify relationships and interlinked tax positions.

Item 1.5d Risk management systems

Inland Revenue does not currently use AI in risk management systems. Therefore, your request for the purposes for which Inland Revenue uses this function is refused under section 18(g)(i) of the OIA.

Question Two - Governance Framework and Decision-Making

Item 2.1 Please outline the governance framework for AI systems implemented by IRD, including:

a) safety measures from conception through monitoring

A risk assessment has been completed on how Inland Revenue will respond to the use of AI. The risks presented by the internal use of AI are generally able to be managed through the adaptation of existing controls to respond to AI. This includes use of security and privacy processes, initiative management and vendor management. Risk assessments are used as part of the use case evaluation, to ensure these areas are sufficiently covered, and to identify where controls are needed specific to a system or use case.

Inland Revenue has taken an approach to restrict the use of AI within our systems except within specific, approved tools and for approved use cases (as reflected in our AI Staff Use Policy, attached in **Appendix B**). Inland Revenue's user behaviour is being monitored by means of a weekly report and analysed by our Information Security Office (ISO) team.

We have a drafted a Preliminary Risk Assessment (PRA) template for assessing a potential AI Proof-of-Concept (PoC). All project leads proposing an AI PoC must complete the PRA that defines the risks and controls specifically for the proposed PoC, with a focus on:

- Security and Compliance: Evaluating the security of the model, identifying potential vulnerabilities and assessing if data would be leaving New Zealand/designated countries.

- Privacy: Assessing the use, creation or modification of Personally Identifiable Information (PII) and Sensitive Revenue Information (SRI) and ensuring effective anonymisation.
- Legal and Ethical Considerations: Evaluating potential infringements on copyright, intellectual property or other rights in AI/LLM operations, and considering the service provider's terms and conditions and the use of Inland Revenue data for training the model.

Upon completion, the PRA is considered by Inland Revenue's AI governance bodies alongside the overall PoC assessment, with the risk management process being an ongoing activity that requires the PoC teams to revisit steps as new information arises.

The criteria to identify acceptable and unacceptable risks in relation to AI is based on Inland Revenue's Enterprise Risk Model and the ISO/IEC 42001 AI – Management System.

Inland Revenue continues to review our existing risk management processes and update as required to take into account increasing use of AI by both Inland Revenue and our partners.

As an example of managing risks associated with access to and usage of data more generally, digital service providers must complete an external due diligence questionnaire (IR1266) to gain access to Inland Revenue's gateway services. This ensures their risk, security, privacy, and access management controls are thoroughly assessed. Since their software interacts with Inland Revenue's systems, which handle sensitive taxpayer information, it is crucial to verify the risk posture against Inland Revenue's standards. This assessment helps maintain the integrity and security of our systems.

Please find enclosed the following documents, listed in the table below, which provide additional detail. Where some information in the documents is withheld, the relevant withholding ground of the OIA is specified in the document. An explanation of the withholding grounds follows:

- Section 9(2)(a) – to protect the privacy of natural persons
- Section 18(c)(i) – where making the requested information available would be contrary to the provisions of a specified enactment, namely Inland Revenue's confidentiality obligations in section 18(1) of the Tax Administration Act 1994 (TAA). Further, section 18(3) (in conjunction with section 143C(1)(a) of the TAA) prevents the Commissioner of Inland Revenue from disclosing any item of revenue information if the release of the information would adversely affect the integrity of the tax system or prejudice the maintenance of the law.

Appendix B: Inland Revenue's AI Use guidelines

Item	Date	Document	Decision
1.	09/12/2024	Artificial Intelligence AI Staff Use Policy	Released with some redactions under sections 9(2)(a) and 18(c)(i).

Item	Date	Document	Decision
2.	31/01/2025	Artificial Intelligence Use Case Guidelines	Released with some redactions under section 18(c)(i).

Where information has been withheld, as required by section 9(1) of the OIA, I have considered whether the grounds for withholding the information requested is outweighed by the public interest. In this instance, I do not consider that to be the case.

Item 2.1b Coverage of lifecycle stages of AI systems

Inland Revenue has an initiative management framework in place which is supported by governance bodies of senior leaders representing each business group.

At a high level, all change initiatives at Inland Revenue follow the same stages of a lifecycle although the detailed steps may vary depending on the business group, the initiative and the delivery approach. The lifecycle uses a consistent process of Emerge, Discover, Design & Deliver and Release.

For AI specific initiatives, two additional governance bodies have been established, as detailed in my response to Item 2.1c below.

For each initiative, a business owner is identified, so ongoing support for the new change or product can be provided, the benefits and value delivered assessed, and lessons learned documented and actioned by this business owner.

Item 2.1c AI-related decision-making structures, specifically who determines which problems are addressed by AI systems and who makes strategic decisions about AI system implementation

Inland Revenue has existing governance of operations, change, strategy and investment who are the decision makers for implementing AI tools. Two additional governance bodies have been established to manage identifying and testing AI tools and solutions prior to going to these executive-level governance bodies.

Please find enclosed the following documents (**Appendix C**), listed in the table below, which provide further detail regarding the governance bodies:

Appendix C: Inland Revenue's governance bodies

Item	Date	Document	Decision
3.	15/08/2024	Artificial Intelligence Working Group: Terms of Reference	Released in full.
4.	24/01/2025	Artificial Intelligence Oversight Group: Terms of Reference	Released in full.

Item 2.1d how AI systems are procured (in-house or external)

Inland Revenue has the necessary technology platforms and capabilities required to support machine learning in-house. Inland Revenue can undertake this activity in a self-contained way and has established robust processes and governance to safely use machine learning with the necessary checks and balances.

For procurement of IT products, this is treated as any other procurement. Sufficient due diligence is done to ensure the optimal terms are contracted and Inland Revenue's requirements are met. Any AI inclusion would be approved in a business case and architectural/security review prior to the procurement taking place.

Item 2.1e how inter-departmental communication regarding AI is organised

Inland Revenue has ongoing engagement with other agencies via the Algorithm Charter Community of Interest for discussion of emerging AI issues. Inland Revenue participates in the Ministry of Business, Innovation, and Employment-led cross-agency AI working group under the direction of the Minister for Digitising Government, Hon Minister Collins, which is focused on digital policy, and also the Government AI Community of Practice. The Commissioner of Inland Revenue is a member of the Digital Government Leadership Group.

Through these communities, as well as connections already held between staff at different agencies, Inland Revenue is actively sharing knowledge and documentation on ways to guide the secure and effective adoption of AI. This includes our governance framework, strategic considerations, training modules, supporting documentation and product specific documents.

Item 2.2 Please outline legal and quality control frameworks implemented by IRD. I am specifically interested in AI system errors that arise from poor quality training data, security of data used by AI systems, security of infrastructure for storage of data used by/processed through AI systems and biases in AI systems and other errors that can harm people.

Item 2.2a who bears legal responsibility for AI system errors listed above

Responsibility for any legal liability that may arise from any AI system errors caused by the specified factors will vary depending on the specific error and circumstances involved (including the factual matrix, causes of action available, parties involved etc.).

For a hypothetical example, consider Information Privacy Principle (IPP) 5 of the Privacy Act 2020. It imposes an obligation on Inland Revenue to ensure that personal information held is protected by such security safeguards as are "reasonable in the circumstances to take" against unauthorised loss, access, use and disclosure (www.legislation.govt.nz). An unauthorised access to personal information, while an AI service provider is providing those services might, in theory, result in Inland Revenue's liability for a breach of IPP 5 (towards any individual concerned) as well as the provider's liability to Inland Revenue for breach of contractual terms. Any applicable contractual terms and conditions may vary between service providers.

Item 2.2b who conducts quality control that aims to mitigate AI system errors listed above (this relates to pre-purchase, testing, implementation, ongoing use and other stages of the use of AI tools by IRD)

Inland Revenue recognises the importance of testing all data and algorithms for quality and accuracy. We have highly skilled data and statistical scientists who employ proven methodologies and testing frameworks to ensure that any complex algorithm is rigorously tested before deployment.

Mandatory training is in place for all Inland Revenue staff (permanent and fixed term) on how to understand, recognise and address unconscious bias.

With regards to generative AI tools, I refer you to my response to Item 2.2d, and the membership of the AI Working Group released under Item 2.1c.

Item 2.2c Whether quality control that aims to mitigate AI system errors listed above is internal or external

Inland Revenue does not rely solely on quality control by vendors and undertakes our own quality controls (in addition to validating quality controls provided by product vendors).

Item 2.2d what components are included in quality control processes that aim to mitigate AI system errors listed above

Quality controls include but are not limited to: manual and automated testing methods, unit, integration and function testing by developers, independent testing by peers, user acceptance testing by business representatives contextual to the targeted use outcomes and feedback loops, and bugs and features management as part of the natural processes of managing technologies based on data.

Specifically for generative AI, Inland Revenue has implemented a process of review for all proposed AI PoC. This process requires AI models to undertake a variety of reviews with the first stage conducting an initial security review to determine the type of AI used in the model. The second stage of review requires members of each business capability to evaluate and assess any risks associated with the use of the tool. This includes but is not limited to any potential legal considerations of using the tool, any security concerns with the model, who will have access to it, and consider the potential impact to people and communities. From here, the tool and the proposed use is evaluated by the AI Oversight Group for final sign off for the PoC. Please note that a tool being approved for a PoC will undergo additional evaluation before being considered for further implementation.

Evaluation of the tool will consider the impact to the organisation and the number of employees who have access to the tool. It also includes considerations of 'human in the loop' capabilities in which any output generated by AI must undergo a review by a person. This is required to reduce the likelihood of bias, misinformation as a result of AI hallucinations, and to check the quality of the output provided by AI tools.

Item 2.2e How responsibilities related to quality control processes that aim to mitigate AI system errors listed above are delineated within IRD. Please specify the quality control function of each governance body or an individual participating in quality control processes that aim to mitigate AI system errors listed above.

Inland Revenue's data and information governance policies prescribe accountabilities for Inland Revenue's data management. This includes Inland Revenue's data quality model which specifies that data quality responsibility rests with everyone at Inland Revenue.

With regards to the specific AI system errors you have listed above, Inland Revenue has established an AI Working Group who will work together with a Project Team on any potential PoC, pilot or trial. The AI Working Group Terms of Reference, released under Item 2.1c, shows which areas of the business are represented (under 'Membership'), and therefore participating in the quality control processes you have listed in your request.

With regards to potential biases, as discussed in Item 2.2b, mandatory training is in place for all Inland Revenue staff (permanent and fixed term) on unconscious bias. This training helps everyone understand, recognise and address unconscious bias. Those staff participating in any pilot, trial or PoC, will regularly engage with the applicable Project Team and/or AI Working Group members to discuss their findings from the experience, including observations around any biases.

Item 2.3 Please outline an organisational structure for the use of AI systems by IRD, including whether there are
Item 2.2a reporting structures/a single structure for AI matters

As discussed in the response to Item 2.1, the AI Oversight Group and AI Working Group are responsible for the creation of AI policies, guidelines and processes to support Inland Revenue's approach to the implementation, adoption, and use of AI.

We constantly review our organisation and governance structure as the AI landscape continuously evolves, to ensure they are fit for purpose in supporting sound decision making.

The AI Oversight Group reports to the Enterprise Priorities and Performance Committee (an Executive Level governance body) to ensure the activities of the Oversight Group reflect the needs of Inland Revenue's business users and Inland Revenue's strategic direction for AI.

Each specific tool in use will have a Business Owner and a Product Administrator.

Inland Revenue's algorithm governance follows All-of-Government directives. We are committed to keeping customers safe and have signed up to the Algorithm Charter for Aotearoa New Zealand. Inland Revenue has operated under this Charter since July 2020. Inland Revenue continues to adapt this approach in line with Government expectations and technology advancements.

Item 2.3b training programs for staff using AI systems

In September 2024, seven e-learning modules were released to all staff; Introduction to AI, Gen AI 101, Value Capture, Trust, Ethics & Governance, Prompt Training, Risks of AI and GenAI, Threats of AI, and Data & Information Management. Select leaders (including our Executive leadership) received tailored fluency training to aid and assist decision making around the use of AI.

Product specific training is also provided to Inland Revenue staff who will be using a specific tool.

Wider professional development opportunities may be undertaken on an exception basis for those in specific roles.

Item 2.3c adaptation of hiring processes for AI-related skills

Inland Revenue recruits according to the skills, knowledge and attributes that are required to perform successfully in the role. One of the capabilities that is important to Inland Revenue is digital literacy. Depending on the role, some recruitment selection approaches may seek evidence of skills, knowledge and experience in using a range of different technologies to support their work.

Question Three - Benefits and Impact Assessment (including statistics)

Item 3.1 Please provide information about the measurable benefits of AI use to New Zealand tax administrators, including changes in query volumes, email correspondence reduction, efficiency improvements, cost reductions, audit resource allocation, tax collection efficiency, revenue impacts, other measurables, and

Item 3.2 Please provide information about the measurable benefits of AI use by IRD to New Zealand taxpayers, including acceleration of tax procedures and audits, increased refund procedure efficiency, legal certainty enhancement and other measurable benefits.

As shown in the response to Item 1.1, there are a wide range of AI systems used by Inland Revenue. Although each has its benefits, many do not demonstrate a measurable benefit such as those listed below (for example, protective security mechanisms, visual content creation, search abilities etc). As a result, I have focused my response on the measurable benefits that directly impact Inland Revenue's customers.

Annual Report

The change over time in query volumes (including call and email correspondence), cost and revenue can be found in Inland Revenue's annual reports. These can be found on Inland Revenue's website at: www.ird.govt.nz/about-us/publications/annual-corporate-reports/annual-report. The effect a specific AI tool has on these outcomes cannot always be identified, however further detail is provided below, where this link can be identified.

Traditional AI

Inland Revenue has integrated rules and algorithms in our products procured through vendors to support productivity and efficiency. We utilise many business rules which enable the automation of its business processes. These support our service delivery to both staff and tax/social policy customers. We also have machine learning algorithms which work across large data sets to provide advanced analytics, decision support and predictive modelling. All activities undertaken as a consequence of these algorithms are subject to human oversight and human decision-making, in both a technical and business context.

Each use of traditional AI is not measured individually. Therefore, I refer you to Inland Revenue's annual reports for more specific detail on our performance as a whole, year on year.

Optical Character Recognition (OCR) – Donation Tax Credits

OCR has been in use at Inland Revenue for many years to assist with processing specific high-volume returns (GST returns and employer returns). OCR was extended during our business transformation programme in March 2020 for donation tax credit claims, and in 2024 for FamilyBoost claims, to review submitted electronic images of individual receipts, invoices or statements to support the claims lodged. Use of OCR enables automated validation against other information held and/or pre-set conditions and rules, including whether or not a receipt may have alterations, erasures, or be incomplete in form.

For donation tax credits, where a receipt fails an automated validation check, it may be declined without need for manual review or, where there may be several exception reasons within the same claim, re-directed for manual review. The OCR has a high degree of accuracy and is regularly reviewed alongside system tolerances to ensure only non-qualifying receipts are identified and declined, and that receipts automatically approved are qualifying.

This has led to a significant reduction in manual interventions. Prior to using OCR for donation tax credit claims, all claims required manual intervention. With the use of OCR, 29% go through straight through processing. Claims submitted prior to the end of the tax year, which meet the business rules, can be automatically released when the tax year ends. This reduces the timeframes from 12 weeks to the same day.

In the 2023/24 financial year, OCR identified $30.3m (rounded) in ineligible donation receipts. Use of OCR enabled AI to do the heavy lifting, and Inland Revenue's compliance staff to undertake additional audits culminating in a further $2.3m in ineligible donation receipts.

Optical Character Recognition (OCR) – FamilyBoost

Inland Revenue is using OCR technology to convert images/PDFs to text, enabling key information to be extracted from Early Childhood Education (ECE) invoices and statements to match against the claim information submitted by customers. When all claim information matches and no other risks are triggered, FamilyBoost claims are processed automatically.

In the December 2024 quarter, almost 50% of FamilyBoost claims were automatically processed. Inland Revenue is continuing to work with customers to educate them on submitting quarterly

statements (if their ECE provides them) and inputting information correctly so their claims are more likely to be processed straight through. Inland Revenue are also working with the ECE sector including Student Management System (SMS) vendors, to ensure the ECE invoices and quarterly statements have all the required information that is structured in OCR 'friendly' format to increase the success of straight through processing. Automatically processed claims reduce the need for manual review, improving efficiency and resulting in faster payouts for customers.

Contact Centre

Inland Revenue uses the Genesys contact centre platform and recently completed a PoC using Genesys Cloud AI (a new tool that Genesys provides in their platform) which creates summaries of calls between a customer and Inland Revenue. This AI solution utilises call transcripts and generates a summary of the key points from the interaction with the customer which, once checked for completeness and correctness, can be copied to the customer record in Inland Revenue's operating system START. The AI generated summary is not retained. The PoC demonstrated an average 17% decrease in after-call work time, no degradation in quality of notes, and positive sentiment from Inland Revenue's staff who trialled this feature. The 'Conversation Summarisation' feature performed particularly well with general voice calls where time efficiencies were seen. More complex calls still saw a decrease in after-call work time, but to a lesser extent than general queues. Inland Revenue are considering wider roll out of this tool.

Microsoft Copilot Chat (Bing/Browser)

Microsoft Copilot chat (Bing/Browser) was trialled at Inland Revenue in 2024. 72% of users reported improved work quality and 93% of users reported time savings with an average of an hour saved per week. 22% of participants saved more than 2 hours per week. The top reported uses were summarising information, writing tasks, brainstorming, internet searches and technical assistance. Use of Copilot Chat was found to have a positive impact on the experience of Inland Revenue's people doing their work with 81% of users reporting a positive or somewhat positive impact on their day-to-day experience of doing their work.

This is now being deployed across the enterprise to appropriate roles and business groups, in a staged approach, with mandatory training before access is available.

GST integrity model

GST Integrity Manager uses a model to analyse patterns and predict behaviours in relation to GST compliance. It scans every GST credit return filed by customers and assesses its level of risk. Inland Revenue has improved the model this past year, enhancing its ability to identify high-risk returns. I refer you to Inland Revenue's Annual Reports, specifically data on '*Interventions that resulted in additional assessed revenue and recovered or disallowed expenditure*'.

Audit resource allocation

Inland Revenue utilises predictive modelling and traditional AI to support human decision-making in selection of customers for audit. I refer you to Inland Revenue's Annual Reports for

return on investment data, specifically under the section titled *Investigations*. Inland Revenue also recently published a media release which further details outcomes in audit and compliance (ird.govt.nz/media-releases/2024/full-throttle-on-compliance-work).

Tax collection efficiency

Traditional AI is utilised by Inland Revenue to support the automation of deduction notices, reducing the need for manual intervention in many instances. Predictive modelling is used to inform Inland Revenue leaders when they are selecting customers requiring manual compliance work. The details of measurement and benefits relating to the use of predictive modelling is withheld under section 18(c)(i) of the OIA, as making the requested information available would be contrary to the provisions of a specified enactment, namely Inland Revenue's Tax Administration Act 1994. Disclosure of this information does not fall within any of the exceptions to the confidentiality obligation listed in sections 18D to 18J of the TAA. I refer you to Inland Revenue's Annual Reports, specifically the section *Management of Debt and Unfiled Returns*. This provides broader performance measurement results which in part, reflect the use of this modelling.

Question Four - Data Protection and Management

Item 4.1 Please provide information about taxpayers' data collection and processing by IRD, including

Item 4.1a broad high-level list of types of data that IR collects directly and through third parties

As Inland Revenue works through potential AI-use opportunities, we will continue to ensure that privacy and security are prioritised, and that customer data remains safe and secure. Our internal guidelines state that no customer data should ever be used as an input to any external and/or publicly available AI solutions. No Inland Revenue customer data leaves Inland Revenue for the purposes of any AI unit training.

Our priority is to make sure we use AI in a way that is compliant with our obligations under the Revenue Acts, the Privacy Act, the Tax Administration Act, and the Algorithm Charter of Aotearoa and all other New Zealand Government authoritative guidance.

Inland Revenue collects information for the administration or enforcement of the Revenue Acts and for other functions conferred on the Commissioner. This includes the registration, assessment and processing of tax and social policy obligations and entitlements.

Inland Revenue collects family details, financial information and supporting disclosures directly from customers to support their self-assessment obligations. Indirect data collection examples include information provided from banks and employers that allows Inland Revenue to prepopulate Individual Income Tax Assessments.

The types of information that we collect under Memoranda of Understanding, and other types of long-term information sharing agreements, can be found on Inland Revenue's website at: ird.govt.nz/about-us/information-sharing

Inland Revenue also collects bulk datasets for limited purposes such as:

- targeted tax and social policy compliance actions; and
- to help customers receive the right entitlements.

Further information regarding bulk datasets can be found on Inland Revenue's website at: ird.govt.nz/about-us/information-sharing/external-dataset-framework

Item 4.1b a detailed list of data sources utilised in the development of AI systems (i.e. source of data used to train AI systems) and utilised by AI systems (i.e. source of input data that is processed/analysed with the use of AI systems)

A number of data sources are used in the training of AI systems (machine learning, predictive analysis) as per the response to Item 4.1a. These include:

- Customer demographic information,
- Financial and tax information,
- Corporate relationships and ownership structures,
- Asset purchases and ownership,
- Inland Revenue's compliance history, and
- Customer interaction and survey response data.

Please note, this information is not used to train generative AI models.

Item 4.1d use of AI for systematic data collection from taxpayers or third parties, including through e-invoicing systems, online cash registers, online PAYE processes, online platform reporting, and other online tools

Inland Revenue does not currently use AI for systematic data collection from taxpayers. Therefore, your request for the purposes for which Inland Revenue uses this function is refused under section 18(g)(i) of the OIA.

Question Five - International Cooperation

5.1 Please detail New Zealand's current cooperation with international or regional organisations regarding AI development and implementation in tax administration.

Looking ahead to the future of tax administration, as defined by the OECD, a key part of Inland Revenue's thinking is focused on working with its key strategic partners and vendors (oecd.org/en/topics/policy-issues/tax-administration).

Inland Revenue prioritises international engagement on AI, specifically on international rules and norms. Inland Revenue is part of the OECD tax administration "cluster" on AI. This cluster of countries aims to enhance the trustworthy use of AI through international collaboration, developing a Trustworthy AI framework, and sharing best practices.

The project targets several impacts including reducing taxpayer burden, optimising internal operations, increasing AI trust, and prioritising AI investments. New Zealand supports these activities to understand AI deployment in different jurisdictions and utilise the generated assets.

Inland Revenue will continue to be involved in the next phase of the project, with a focus on learning and evaluating appropriate Trustworthy AI frameworks in the context of tax administration.

Right of review

If you disagree with my decision on your OIA request, you can ask an Inland Revenue review officer to review my decision. To ask for an internal review, please email the Commissioner of Inland Revenue at: commissionerscorrespondence@ird.govt.nz.

Alternatively, under section 28(3) of the OIA, you have the right to ask the Ombudsman to investigate and review my decision. You can contact the office of the Ombudsman by email at: info@ombudsman.parliament.nz.

If you choose to have an internal review, you can still ask the Ombudsman for a review.

Publishing of OIA response

We intend to publish our response to your request on Inland Revenue's website (ird.govt.nz) as this information may be of interest to other members of the public. This letter, with your personal details removed, may be published in its entirety. Publishing responses increases the availability of information to the public and is consistent with the OIA's purpose of enabling more effective participation in the making and administration of laws and policies and promoting the accountability of officials.

Thank you again for your request.

Yours sincerely

Cate Robertson
Enterprise Leader, Strategic Architecture

Appendix A: Your full request and clarification of your request

I am writing pursuant to the Official Information Act 1982 to request information regarding Inland Revenue's implementation, governance, and operational use of artificial intelligence systems.

The scope of this request encompasses the following areas:

1. *AI Applications and Systems*
    - 1.1 *Please provide information detailing the specific use of AI systems by IRD for internal purposes, including:*
        - a. *research*
        - b. *correspondence with taxpayers*
        - c. *taxpayer services (virtual conversational assistants, chatbots, personalised services, nudges etc.)*
        - d. *policy analysis (e.g. data-driven insights, scenario modelling)*
    - 1.2 *Please provide information detailing the specific use of AI systems by IRD for external purposes, including:*
        - a. *automated data collection and matching processes*
        - b. *risk assessment and compliance monitoring systems*
        - c. *fraud detection mechanisms*
        - d. *audit selection procedures*
        - e. *taxpayer profiling capabilities*
        - f. *predictive analytics applications*
    - 1.3 *List all AI technologies currently in use by IRD, specifically:*
        - a. *machine learning (ML)*
        - b. *optical character recognition (OCR)*
        - c. *natural language processing (NLP)*
        - d. *any other AI technologies*
    - 1.4 *Please specify the types of AI technologies currently in use by IRD for international tax information exchange.*
    - 1.5 *Please specify the purposes for which IRD uses the following functions of AI systems:*
        - a. *web-scraping or web-crawler applications*
        - b. *social network analysis (SNA)*
        - c. *transaction network analysis (TNA)*
        - d. *risk management systems (RMS)*

2 Governance Framework and Decision-Making

2.1 Please outline the governance framework for AI systems implemented by IRD, including:

a. safety measures from conception through monitoring

b. coverage of lifecycle stages of AI systems

c. AI-related decision-making structures, specifically:

i. who determines which problems are addressed by AI systems

ii. who makes strategic decisions about AI system implementation

d. how AI systems are procured (in-house or external)

e. how inter-departmental communication regarding AI is organised

2.2 Please outline legal and quality control frameworks implemented by IRD, including:

a. who bears legal responsibility for AI system errors

b. who conducts quality control

c. whether quality control is internal or external

d. what components are included in quality control processes

e. how responsibilities are delineated between decision-makers

2.3 Please outline an organisational structure for the use of AI systems by IRD, including whether there are:

a. reporting structures/a single structure for AI matters

b. training programs for staff using AI systems

c. adaptation of hiring processes for AI-related skills

3 Benefits and Impact Assessment

3.1 Please provide information about the measurable benefits of AI use to New Zealand tax administrators, including:

a. changes in query volumes

b. email correspondence reduction

c. efficiency improvements

d. cost reductions

e. audit resource allocation

f. tax collection efficiency

g. revenue impacts

h. other measurable

3.2 Please provide information about the measurable benefits of AI use by IRD to New Zealand taxpayers, including:

a. acceleration of tax procedures and audits,

b. increased refund procedure efficiency

c. legal certainty enhancement

d. other measurable benefits

3.3 Please support your answers to questions 3.1 and 3.2 with available statistics or refer to publicly available sources of such statistics

4 Data Protection and Management

4.1 Please provide information about taxpayers' data collection and processing by IRD, including:

a. types of data collected

b. data sources utilised

c. use of AI for data matching

d. use of AI for systematic data collection from taxpayers or third parties, including through e-invoicing systems, online cash registers, online PAYE processes, online platform reporting, and other online tools

5 International Cooperation

5.1 Please detail New Zealand's current cooperation with international or regional organisations regarding AI development and implementation in tax administration.

On 14 March 2024, you clarified the following aspects of your request:

An "AI system" refers to any AI system that is used by IRD to perform its functions. "AI system errors" refers to errors in an AI system that is used by IRD to perform its functions.

I am specifically interested in AI system errors that arise from:

- *poor-quality training data (i.e. data that is not clean, incomplete, inaccurate, or unrepresentative);*
- *security of data used by AI systems;*
- *security of infrastructure for storage of data used by/processed through AI systems (specifically, I want to know whether all data used by AI systems is stored and processed in New Zealand or not);*
- *biases in AI systems and other errors that can harm people (e.g. false positives in detection of tax non-compliance/fraud or false negatives in assessments of various eligibility criteria).*

Section 2.2 was amended as follows:

Please outline legal and quality control frameworks implemented by IRD, including:

a. *Who bears legal responsibility for AI system errors listed above*
b. *Who conducts quality control that aims to mitigate AI system errors listed above (this relates to pre-purchase, testing, implementation, ongoing use and other stages of the use of AI tools by IRD)*
c. *Whether quality control that aims to mitigate AI system errors listed above is internal or external*
d. *What components are included in quality control processes that aim to mitigate AI system errors listed above*
e. *How responsibilities related to quality control processes that aim to mitigate AI system errors listed above are delineated within IRD. Please specify the quality control function of each governance body or an individual participating in quality control processes that aim to mitigate AI system errors listed above.*

Section 4.1 was amended as follows:

a) *Please provide a broad high-level list of **types of data** that Inland Revenue collects directly and through third parties.*
b) *Please provide a detailed list of **data sources**:*
   - *utilised in the development of AI systems (i.e. sources of data used to train AI systems), and*
   - *utilised by AI system (i.e. sources of input data that is processed/analysed with the use of AI systems).*

Inland Revenue
Te Tari Taake

Ka whakamahi nga kaimahi AI i nga kaupapa here

Artificial Intelligence (AI) staff use policy

This policy sets out the requirements for using AI products and services for staff at Inland Revenue

Note: A glossary of terms is available in this policy.

Why we have this policy

This policy sets out Inland Revenue's approach to safely and securely look to use AI in the workplace, to help make good decisions and deliver services that are more effective and efficient.

It is a priority for IR to ensure that AI is adopted in a way that considers not only our obligations under the Revenue Acts and Privacy Act but also under the Algorithm Charter for Aotearoa New Zealand and any other NZ Government authoritative guidance.

This includes embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of Te Tiriti o Waitangi.

In addition to these absolute requirements, as a leading public agency, we must also robustly assess the security, privacy, operational and technical impacts that any new AI solution or use case brings, and then the associated transparency of this to the public.

This is a rapidly evolving area, and this policy is subject to change.

Who this policy applies to

This policy applies to all *users* engaging in the use of approved AI products and services for an Inland Revenue purpose and are:

- An Inland Revenue employee
- A contingent worker (consultant, contractor, subcontractor or otherwise) working for Inland Revenue with access to Inland Revenue systems and information

Your responsibility

As someone working for Inland Revenue, you are responsible for understanding and following this policy. This means: **Compliance Measurement**

Inland Revenue may verify compliance to this policy through various methods, including, but not limited to, business tool reports, internal and external audits, and feedback to the policy owner. **Non-Compliance**

For IR employees, if any possible compliance issue is identified appropriate action will be taken, including referral to IR's Integrity team for consideration under Inland Revenue's Code of Conduct – Tikanga Whanonga.

For non-employees (contingent workers/contractors/suppliers) any possible compliance issues would be escalated and managed consistent with written agreements/contractual arrangements.

Exemptions

Any exemption to the policy must go through the required process. See the Exemptions Standard for more information.

Our policy

1. Principles

To help guide our thinking and behaviour, Inland Revenue has defined some key principles to apply to AI use cases. Where these principles are not able to be met, it may present additional risks to Inland Revenue that need to be considered and appropriate action taken.

- **Transparency** - we will maintain transparency by clearly explaining how decisions are informed by algorithms and where algorithms have generated content. This applies both internally, promoting transparency within our working teams, and with our customers and partners around how we will utilize these tools.
- **Human oversight –** we will retain human oversight to assess for unintended consequences and act on this information. This includes understanding limitations and identifying and managing bias.
- **Partnership –** we will embed a Te Ao Māori perspectives in the development and use of algorithms consistent with the principles of Te Tiriti o Waitangi.
- **Ethics –** public trust in how Inland Revenue manages data makes us all data stewards. Our AI systems should respect human rights, diversity and the autonomy of individuals and not result in discrimination against individuals, communities or groups.
- **Integrity -** unauthorised access, misuse or security incidents involving the use of AI tools/solutions will be reported to Information Security.
- **Robust testing –** we will use a robust testing and review process to evaluate the use of new AI systems and tools or updates to existing tools that have adopted increased functionality.

2. Adoption of AI solutions

Any new proposed use case for an AI or Large Language Model (LLM) tool, AI integrated tool, or new use of an AI tool should be escalated to Inland Revenue's AI Oversight Group and AI Working Group (contact s 18(c)(i)) with a use case submission form completed for new tools. This ensures adequate testing, review and consideration is applied to all use cases, and prevents effort duplication and risk.

Existing approved workplace technologies that use AI do not require new consideration or exemption. Some examples are listed in the glossary. If you are unsure about the current approval status of a system for use with your business unit or information, please contact s 18(c)(i) .

3. Creation of content

Users Must Not:

- Use any AI/LLM or other intelligent tool that is not approved for use or use an approved tool/solution for an unapproved use case.
- Use any information classified as Sensitive or above with an AI/LLM solution.
- Intentionally generate, or use AI/LLMs to create any misleading, illegal, discriminatory, defamatory or offensive content.
- Process or use in AI/LLM solutions, any information that is reasonably capable of being used to identify an Inland Revenue customer (or Inland Revenue staff) without express approval and consideration from the AI Oversight Group.
- Process or use commercially In-Confidence information without express approval and consideration from the AI Oversight Group.
- Provide automated financial, legal, tax advice, guidance, or AI/LLM information to another party (either internal or external) without human involvement.
- Infringe upon copyright or other right of use in operation of AI/LLMs.

Users Must:

- Consistently review and confirm the accuracy of any generated AI/LLM output (including text, audio, visual or other) or intelligent system output. This includes checking code or technical information. Robust technical testing must be performed to ensure this type of content is free from security issues.
- Review intelligent system output for potential bias, ethical concerns, and unintended messaging.
- Clearly state through visual, verbal or written indicators/mechanisms the use of AI/LLM and other intelligent solutions in the production or contribution to content.
- Consider IR's Enterprise Risk policy and framework to ensure a wider risk consideration is given for the use.
- Seek the approval of the AI Oversight Group for proposed use cases of unapproved AI/LLM use cases or solutions.
- Acknowledge that any information produced in part or whole by an AI system will be attributable to the Inland Revenue staff working with that information, and ultimately Inland Revenue will be responsible for the quality and outcomes of that information.

4. Disposal of content

The retention and disposal of any AI information or knowledge will be handled in conjunction with Inland Revenue current retention and disposal rules and processes.

Roles and responsibilities

The table below contains roles and responsibilities for functions and users to deliver effective AI management at Inland Revenue.

Role	Responsibility
Executive Leadership Team	Reviews Inland Revenue's approach and use of AI as part of Governance Board meetings
AI Oversight Group	• Provides oversight and direction for AI systems and use at Inland Revenue. • Will act as the key point for AI related work and proposals. • The approval group for proposed AI/LLM use cases or solutions.
AI Working Group	• Takes direction from the AI Oversight group. • Facilitates open discussion of possible current and future AI use.
CISO Team (including SecOps, Integrity, and Assurance)	• Responds to any report of abuse, misuse, or non-compliance with this policy. Monitor user behaviours with AI tooling • Educate users on the capabilities and acceptable use of emerging technologies
Procurement team	Ensures appropriate checks are undertaken and understanding is held of vendors where AI is in the services provided. E.g.: conduct market research on vendors and offerings, covering privacy, security and ethical risks.
Business and Technical owners of AI/LLM and intelligent solutions and responsible people	• Ensure consistent and fit for purpose assessments of AI/LLM and AI integrated solutions, including where these capabilities may be integrated into existing products. • Govern and maintain an awareness of the use of AI/LLM though Inland Revenue's vendors and partners, and also Inland Revenue's fourth and fifth parties. • Govern and maintain relevant / useful guidance for staff. • Consider the new high level risks Inland Revenue may be exposed to and manage the risk as per Inland Revenue's Enterprise Risk policy and framework. • Consider the carbon and environmental impact of Inland Revenue consumption of AI/LLM solutions. • Where satisfactory controls are not known or available to govern the use of this information. This includes limiting the collection of Inland Revenue information by AI tools and products as much as possible, including submitted query and returned answers from the AI vendor.

Inland Revenue employees and contingent workers (including consultants, contractors, subcontractors or otherwise)	• Adhere to the AI policy and guidelines. • Take accountability and responsibility for all AI generated or intelligently supported content or information. • Report any abuse, misuse or non-compliance with this policy to the Inland Revenue Information Security Team.

Glossary

Term	Meaning and example
Generative AI	A generative AI is a direct interaction model which describe how the user interacts with the solution. It is often used in a question-answer format (such as ChatGPT/OpenAI or Microsoft Co-pilot models). However, GenAI models can generate other forms of media such as images, video, or audio. These are solutions capable of generating entirely new content and may leverage large language models. Examples: ChatGPT, Bing Discover, and Microsoft Co-pilot solutions.
Large Language Model (LLM)	An LLM is where a traditional machine learning algorithm may seek to find connections between attributes (such as the loose connection between height and shoe size). An LLM tries to understand how people talk about the words in a given query, not learning or understanding data "relationships" but merely predicting the words most likely to answer the query.
AI Integrated Solutions	These are solutions that already exist in the Inland Revenue workplace, to which the providing vendor has integrated an AI component. This may be very visible (like a chatbot or assistant) or more subtle, such as a ranking algorithm (most famously used in search engines). Inland Revenue, in many cases, does not have an option to prevent or control the AI component of these solutions, and increasingly will be included in all software tools. System specific guidelines may need to be developed based on the specific tool and use case. Examples: Grammar and spelling checkers, GitHub Co-pilot, Viva answers, Siri, Bing GPT, Spotify and Netflix recommendations.

Machine Learning (ML)	An algorithm is provided sample training data and "learns" relationships in that data. ML can be thought of as the grey "semi-intelligent" space between business rules (below) and LLM (above). In an ML solution we should have more control of the outcome, being able to interrogate the training data and clearly show relationships / paths from input to output (something that is unlikely to be possible with vendor AI solutions and LLM's). Examples: Digital Intelligence Platform (DIP) data investigations, Website and Haukāinga search, Viva insights.
Business Rules	These are the simplest kinds of computational intelligence. This includes decision trees (e.g.: if value X is above 100, then go to this step), simple coding logic such as if/else statements and for/while/until loops. Inland Revenue heavily utilities business rules within START (e.g. calculators, leave and timesheet delegations, customer decision trees, Excel Macros) to help sort and process customer requests.

Finding out more

This policy should be read in conjunction with:

- Inland Revenue's AI Use case guidelines and the wider supporting AI material on the Te Mātāwai page:
- Artificial Intelligence - Home (sharepoint.com)

IR has a number of additional relevant resources to support this Policy including:

- End user policy
- ICT security policy
- Data information policy
- Use of business tools policy
- Risk management policy
- Online services standard
- Software management standard
- Internet use standard
- Information classification guide

Document control	Version 1.1

Approved by	s 9(2)(a)
Review dates	09-Dec-2024 | 15:35 NZDT Date approved: Next review: Six months from approval date
Policy owner	Jay Harris – Chief Information Security Officer
Policy contact	s 18(c)(i)

Version History

Version	Date	Sections amended	Summary of amendment
1.1	18/12/23	All	Document created and approved
1.2	20/08/24	All	Amendments to the phrasing of content and the policy template. Added links to relevant sites.
1.3	5/02/2025	Disposal of Information	Updated link as previous link was throwing a 404 error.

GUIDELINE

Artificial Intelligence Use Case Guidelines

These guidelines set out the requirements when using AI technologies in existing approved workplace technologies at IR and provides guidance for submissions of a potential AI use case at IR.

These guidelines support IR's Artificial Intelligence Staff Use Policy and are to be read in conjunction with IR's information classification and handling guidelines.

1. Contents

2. AI Use Case Guidelines

IR has developed an Artificial Intelligence (AI) Staff Use Policy outlining permitted use and restrictions on artificially intelligent systems that **ALL STAFF** at IR must follow.

***Important note**: IR staff **ARE NOT** permitted to use AI for work purposes or on work devices outside the use cases specifically referenced as approved in the AI application register.*

If a use case falls outside of this guidance or is not approved for the level of information classification needed, approval must be sought from the AI Oversight Group. See Submitting an AI use case for more details.

The tables below are designed to provide clarity around what types of information can be used

with different AI solutions at IR. It is impossible to represent all use cases within IR in a single view, as such this document will be regularly updated as IR evolves its understanding of this landscape.

It is important to understand IR's information classification levels prior to using these guidelines.

2.1 Table 1: Types of AI mapped to IR's classification labels

This view has mapped IR's definition of each type of AI (noted in "Appendix 1" below) to our standard information classification labels. This is provided to staff to help you understand what types of information (such as documents, data or images) can be used with different kinds of solutions.

Note: Comprehensive definitions can be found in terms and definitions.

s 18(c)(i)

2.2 Table 2: Use cases mapped to IR examples.

This view has mapped a range of IR use cases and visually shows what kind of AI solutions these use cases can leverage. This is not an exhaustive list and requests for additional rows and any

questions can be made to AI@ird.govt.nz. Compare your use case against this table to help verify if it is currently permitted.

All rows coloured red with an "X" are not approved for that form of AI, while rows coloured green with a tick, are permitted. Note that this is not an exhaustive list of use cases and aims to provide an overview of commonly encountered forms of AI.

s 18(c)(i)

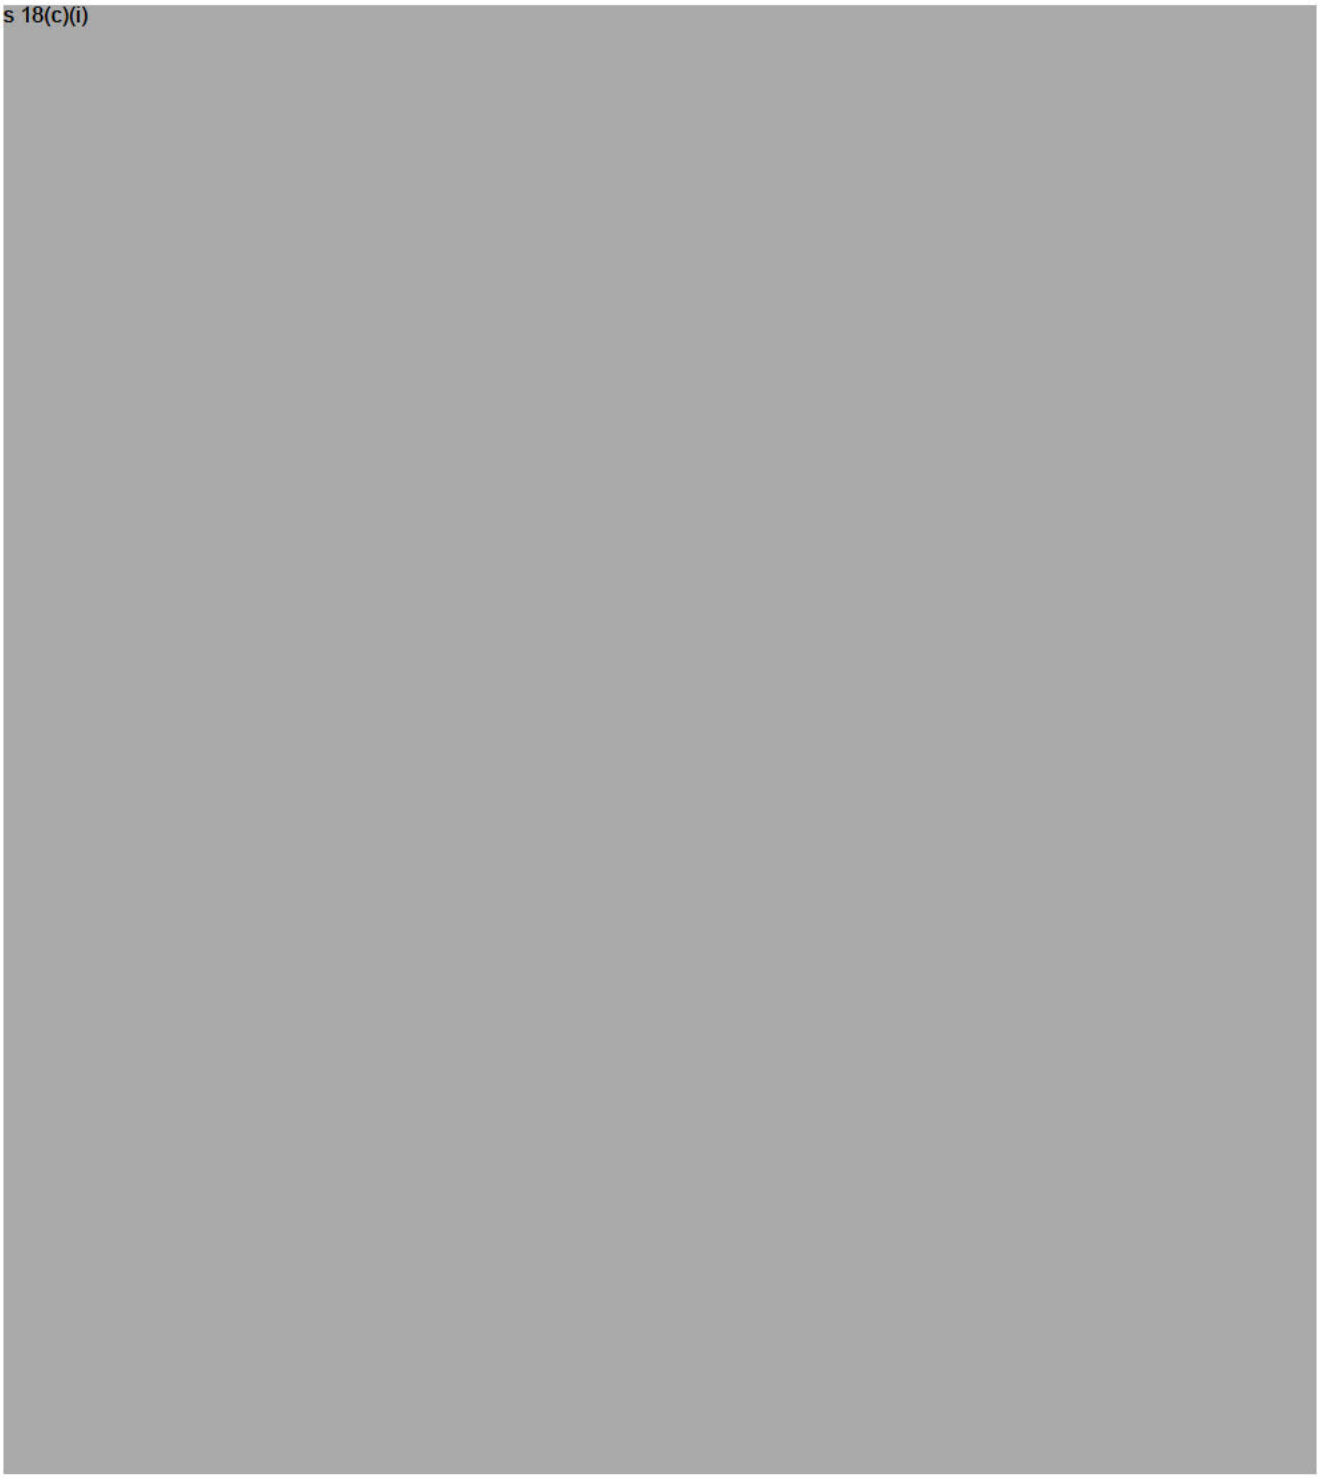

s 18(c)(i)

3. Responsibilities

The following establishes the broad accountabilities and responsibilities of the key internal stakeholders applicable to this Guideline.

Who	What
All Inland Revenue employees and contracted individuals	**MUST** follow these guidelines

4. Terms and definitions

The following definitions apply throughout this document:

Artificial Intelligence (AI)	Artificial Intelligence refers to the ability for computer systems to perform tasks that typically require human forms of intelligence. This includes tasks such as learning, perceiving, reasoning and planning, which allows machines to complete decision-making tasks and simulate thought adjacent to that of a human.
Generative AI (GenAI)	A generative AI is a direct interaction model which describes how the user interacts with the solution. It is often used in a question-answer format (such as ChatGPT/OpenAI or Microsoft Co-pilot models). However, GenAI models can generate other forms of media such as images, video, or audio. These are solutions capable of generating entirely new content and may leverage large language models.
Large Language Model (LLM)	An LLM is where a traditional machine learning algorithm may seek to find connections between attributes (such as the loose connection between height and shoe size). An LLM tries to understand how people talk about the words in a given query, not learning, or understanding data "relationships" but merely predicting the words most likely to answer the query.
AI Integrated Solutions	These are solutions that already exist in the Inland Revenue workplace, to which the providing vendor has integrated an AI component. This may be very visible (like a chatbot or assistant) or more subtle, such as a ranking algorithm (most famously used in search engines). Inland Revenue, in many cases, does not have an option to prevent or control the AI component of these solutions, and increasingly will be included in all

	software tools. System specific guidelines may need to be developed based on the specific tool and use case. Examples: Grammar and spelling checkers, GitHub Co-pilot, Viva answers, Siri, Bing GPT, Spotify and Netflix recommendations.
Machine Learning (ML)	An algorithm is provided sample training data and "learns" relationships in that data (which can be done in a number of ways). ML can be thought of as the grey "semi-intelligent" space between business rules (below) and LLM (above). In an ML solution we should have more control of the outcome, being able to interrogate the training data and clearly show relationships / paths from input to output (something that is unlikely to be possible with vendor AI solutions and LLM's). Examples: Digital Intelligence Platform (DIP) data investigations, Website and Haukāinga search, Viva insights
Business Rules	These are the simplest kinds of computational intelligence. This includes decision trees (e.g.: if value X is above 100, then go to this step), simple coding logic such as if/else statements and for/while/until loops. Inland Revenue heavily utilities business rules within START to help sort and process customer requests. Examples: START calculators, leave and timesheet delegations, Excel macros, available to customer decision trees.

5. Where to find out more

This guideline will be available to all employees, contractors and consultants. The current version of this Guideline is in the Governance Documents Centre. If you require any further information, please contact the Document Owner (refer Section 6).

6. Document information

Current version	31 January 2025
First released	18 December 2023
Last updated	31 January 2025
Review frequency	Yearly or as needed
Review before	31 January 2026

Business function owner	Jay Harris
Document owner	Graham Poppelwell
Audience	All IR employees, contractors and consultants

6.1 Document amendment history

Version	date	Sections amended	Summary of amendment
1.1	18/12/23	All	Document created and approved
1.2	20/08/24	All	Amendments to the phrasing of content and the policy template. Made the colours in the table more accessible/higher contrast. Updated definitions.
1.3	21/10/2024	All	Updates to the document based on the comments from the AI sub-working group and amendments to the wording of in-confidence budget sensitive wording.
1.4	31/01/2025	All	Content updated to guidelines template. Added in AI agents to table 1. Added new content in Table 2 for approaching an external's request to use AI call transcriber or assistance.

7. Related documents

Title	Document
AI at IR	Using Artificial Intelligence at IR
Government Joint Systems Leads	Joint System Leads advice on Public Service Use of GenAI Advice Document
Government Joint Systems Leads	Joint System Leader tactical guidance on public service use of GenAI summary

Artificial Intelligence (AI) Working Group Terms of Reference

The Artificial Intelligence Working Group (The Working Group) reports to the Artificial Intelligence Oversight Group (The Oversight Group) to ensure the activities of the Working Group reflect the needs of IR's business users and IR's strategic direction for AI.

Purpose

The purpose of the Working Group is to provide an operational response to governance direction and outcomes from the Oversight Group. The Working Group helps with co-ordination, cohesion, and prioritisation as directed by the Oversight Group.

The Working Group is responsible for:

- defining, reviewing, and implementing governance instruments in line with IR's corporate strategy (such as roadmaps, frameworks, policies)
- managing identified AI issues and coordinating progression and delivery of outcomes
- identifying threats and opportunities for future development, enhancements, and expansion
- creation of communications to inform and educate stakeholders about AI
- maintaining high-quality working relationships with complementary forums within and outside IR (including the AI Community of Practice).

Membership

Members are:

Business Unit	Role
Enterprise Services	Domain Lead, Information Governance & Sharing (Chair)
	Analyst, People & Workplace Analytics (non-voting)
	Change Analyst, Change, Design & Enablement
	Domain Lead, Enterprise Information & Knowledge
	Domain Lead, Technology Experience
	Domain Lead, Digital Product Experience
	Domain Principal, ES Planning, Design & Delivery
	Domain Principal, Organisational Development
	Domain Principal, People Operations & Policy (non-voting)

	Domain Principal, Technology Commercial x2
	Domain Specialist, Organisational Development (non-voting)
	Domain Specialist, Organisational Learning (non-voting)
	Technical Lead, Organisational Development
	Technology Specialist, Digital Product Experience
Enterprise Design and Integrity	Architect, Technology Architecture x2
	Change Analyst, Digital Ecosystem
	Domain Lead, Business Architecture
	Domain Lead, Information Security Office - Internal
	Domain Lead, Technology Architecture
	Domain Lead, Testing Central
	Domain Principal, Centre for Enterprise Data and Analytics
	Domain Principal, Intelligence
	Domain Principal, Internal Assurance (non-voting)
	Domain Principal, Strategic Portfolio Stewardship
	Domain Principal, Technology Commercial
	Kaihautū, Te Kāhui Tūhono (non-voting)
	Kaitohutohu, Te Kāhui Tūhono
	Privacy Officer
	Project Manager, Strategic Portfolio Stewardship
	Project Team Leader, Information Security Office
	Technology Specialist, Information Security Office
Policy and Regulatory Stewardship	Policy Advisor, Māori Perspectives
Customer and Compliance Services	Customer Experience Designer, PD&D - CX/UX Design (non-voting)
	Domain Lead, Insights & Configuration
	Group Lead, Individuals
	Service Owner, PD&D - CX/UX Design

	Service Owner, Pou Here Tangata IPV - Programme
	Technical Lead, Legal Services
Tax Counsel Office	Information Specialist, TCO – Knowledge & Development (non-voting)

The Chair approves proposed changes to the Working Group membership.

Roles and responsibilities

Chair

The Chair is responsible for the overall direction of the meeting. They set the agenda, with the support of the Information Governance Team.

The Chair represents the Working Group to the Oversight Group.

Working Group Members

The Working Group members, or the people formally acting for them, are expected to attend every meeting. Where this is not possible, members are encouraged to send delegates in their place. Only people who are formally acting for members count toward quorum.

The Working Group members commit to:

- demonstrating the public service principles and values
- demonstrating Te Pou o te Tangata - how we do things at IR: Whanaungatanga, Manaakitanga and Mahi Tika
- working with IR's commitments as a public sector organisation in mind, including IR's commitment to strengthening the Māori-Crown relationship and to integrating te Tiriti o Waitangi and Māori concepts and perspectives into IR's work
- making sure risks, issues and challenges are brought into the open and explored
- welcoming different points of view and frank, robust discussion
- being clear when allocating responsibility
- collectively owning decisions made.

Information Governance Team Support

The Working Group is supported by an Information Specialist, from the Information Governance team. They are responsible for:

- ensuring the Working Group follows good governance principles and practices
- working with the Chair to set the agenda, considering key AI issues and decisions required
- providing support, advice and quality assurance for papers and speakers attending meetings
- ensuring that minutes reflect decisions and key discussion points, and that decisions are communicated to the people who need to know.

Attendees and observers

At the Chair's discretion, people are invited to attend Working Group meetings to provide input as needed.

Logistics

Meetings

The Working Group will meet monthly or as decided by the Oversight Group. Additional meetings may be scheduled when needed, and members are notified of these as early as possible. All procedures, rules and practices for regular meetings stay in place for additional meetings.

Agenda and papers

The agenda and papers are made available to Working Group members four working days before the meeting. The Chair decides whether to accept late agenda items and papers on the Information Specialist's recommendation.

Papers should be on the governance template (a choice of Word or PowerPoint) and should include a purpose statement and recommendations.

In some circumstances it may be necessary for papers to be circulated for feedback and decisions outside of meetings. The Chair's agreement is required for these 'out-of-cycle items.

Quorum

A quorum of 50% of members, including the Chair, is required for decisions to be made. If there is no quorum, the Chair decides whether to reschedule the meeting.

Minutes of meetings

The Information Specialist writes the minutes and provides them to the Chair for their following meeting. Within that following meeting, the Working Group is asked to approve the minutes as an accurate record.

Review

To enable ongoing improvement, the Working Group will review its performance at approximately six-monthly intervals. Reviews will be recorded in the minutes.

Version Control

The most recent document will be included in the footnote for continuity.[1]

[1] Version 2.1 2024.08.15

Inland Revenue
Te Tari Taake

Artificial Intelligence (AI) Oversight Group: Terms of Reference

The Artificial Intelligence Oversight Group (The Oversight Group) reports to the Enterprise Priorities and Performance Committee (EPPC) to ensure the activities of the Oversight Group reflect the needs of Inland Revenue's business users and Inland Revenue's strategic direction for AI. The Chair represents the Oversight Group to the EPPC or the appropriate Executive Level Governance body.

It is a priority for IR to ensure that AI is adopted in a way that considers not only our obligations under the Revenue Acts and Privacy Act but also under the Algorithm Charter for Aotearoa New Zealand (of which IR is a signatory) and any other NZ Government authoritative guidance.

This includes embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi.

Purpose

The purpose of the Oversight Group is to steer IR's AI direction and develop IR's capability to respond to AI threats and opportunities in alignment with IR's strategic goals. The term AI encompasses solutions such as generative AI/Large Language Models, AI integrated tools, Machine Learning and Business Rules.

The AI Oversight Group approves:

- Proof of Concepts and Proof of Technologies relating to AI
- all organization wide communications relating to AI
- all governance processes and instruments relating to AI (such as roadmaps, frameworks, policies)
- all proposed AI/LLM use cases or solutions.

The AI Oversight Group is responsible for:

- ensuring IR is well informed on industry AI developments and advancements
- assisting in the co-ordination, cohesion, strategic fit and prioritisation of AI initiatives
- creating and maintaining the organisational understanding and use of AI related services and products
- contributing to All of Government AI initiatives as required
- building and maintaining public and private sector relationships.

Membership

Members are:

Business Unit	Role
Enterprise Design & Integrity	Domain Lead, Technology Architecture (Chair)
	Architect, Technology Architecture (non-voting)
	Chief Information Security Officer
	Corporate Solicitor, Corporate Legal (non-voting)
	Domain Lead, Digital Ecosystem
	Enterprise Leader, Strategic Architecture
	Enterprise Leader, Strategic Portfolio Stewardship
	Enterprise Leader, Data, Analytics & Insights
	Kaihautū Rautaki, Enterprise Design & Integrity
	Service Leader, Integrity & Internal Assurance (non-voting)
	Strategic Advisor, Enterprise Design & Integrity
Enterprise Services	Domain Lead, Information Governance & Sharing (non-voting)
	Enterprise Leader, People & Workplace Services
	Enterprise Leader, Technology Services
	Intelligence Leader - Enterprise Information and Knowledge (EI&K)
Customer & Compliance Services - Individuals	Domain Lead, Planning, Design & Delivery - Digital
Policy & Regulatory Stewardship (PaRS)	Policy Director
Tax Counsel Office	Technical Lead, TCO – Knowledge & Development (non-voting)

The Chair approves proposed changes to the Oversight Group membership.

Roles and Responsibilities

The Chair

The Chair is responsible for the overall direction of the meeting. They set the agenda, with the support of the Information Governance Team.

The aim in meetings is to achieve consensus, and the principles of collective responsibility apply. The Chair may, however, bring discussion to an end and determine a position.

Oversight Group Members

The Oversight Group members, or the people formally acting for them, are expected to attend every meeting. Where this is not possible, members are encouraged to send delegates in their place. Only people who are formally acting for members count toward quorum.

The members commit to:

- demonstrating the public service principles and values
- demonstrating Te Pou o te Tangata - how we do things at IR: Whanaungatanga, Manaakitanga and Mahi Tika
- working with IR's commitments as a public sector organisation in mind, including IR's commitment to strengthening the Maori-Crown relationship and to integrating te Tiriti o Waitangi and Māori concepts and perspectives into IR's work
- making sure risks, issues and challenges are brought into the open and explored
- welcoming different points of view and frank, robust discussion
- being clear when allocating responsibility and authority
- collectively owning decisions made
- operating with an agnostic system, platform, and information form perspective.

Information Governance Team Support

The Oversight Group is supported by the Information Governance team.

They are responsible for:

- ensuring the Oversight Group follows good governance principles and practices
- working with the Chair to set the agenda, considering key issues and decisions required
- providing support, advice and quality assurance for papers and speakers attending meetings
- ensuring that minutes reflect decisions and key discussion points, and that decisions are communicated to the people who need to know.

Attendees and observers

At the Chair's discretion, people are invited to attend Oversight Group meetings to provide input as needed.

Panels

The Oversight Group may form Panels of members to focus on specific initiatives or areas of activity. The Chair of the Oversight Group will appoint the Panel Lead. Panel Leads must be voting members of the Oversight Group. Panels will report on their progress to the Oversight Group at each Oversight Group meeting or as appropriate. Panels have no authority or responsibilities outside the Oversight group.

Logistics

Meetings

The Oversight Group will meet monthly or as decided by the Chair. Additional meetings may be scheduled when needed, and members are notified of these as early as possible. All procedures, rules and practices for regular meetings stay in place for additional meetings.

Agendas and papers

The agenda and papers are made available to Oversight Group members four working days before the meeting. The Chair decides whether to accept late agenda items and papers on the Information Specialist's recommendation.

Papers should be on the governance template (a choice of Word or PowerPoint) and should include a purpose statement and recommendations.

In some circumstances it may be necessary for papers to be circulated for feedback and decisions outside of meetings. The Chair's agreement is required for these 'out-of-cycle items.

Quorum

A quorum of 67% of members including the Chair, is required for decisions to be made. If there is no quorum, the Chair decides whether to reschedule the meeting.

Minutes of meetings

The Information Specialist writes the minutes and provides them to the Chair for their following meeting. Within that following meeting, the Oversight Group is asked to approve the minutes as an accurate record.

Review

To enable ongoing improvement, the Oversight Group will review its performance at approximately six-monthly intervals. Reviews will be recorded in the minutes.

Version Control

The most recent document will be included in the footnote for continuity.[1]

[1] Version 2.2 2025.01.24