24OIA2041

Inland Revenue
Te Tari Taake

8 May 2024

[redacted]

Dear [redacted]

Thank you for your request made under the Official Information Act 1982 (OIA), received on 3 April 2024. You requested the following:

Reports or materials held by your organisation:

1. *Setting out and/or summarising how and when your organisation has met, or intends to meet, the commitments to transparency and partnership set out in the Algorithm Charter for Aotearoa New Zealand.*
2. *Setting out or explaining how decisions made by the organisation are informed by algorithms. This may include without limitation "plain English" documentation of the algorithm/s, information about the data and processes involved, or published information about how data is collected, secured, and stored.*
3. *That demonstrate or detail how your organisation is delivering clear public benefits through Treaty of Waitangi commitments by embedding a te ao Māori perspective in the organisation's development and use of algorithms consistent with the Treaty of Waitangi.*
4. *That demonstrate or detail how your organisation has, or intends to, identify and consult with people, communities and groups who have an interest in algorithms, including Māori.*
5. *That describe how your organisation makes sure data is fit for purpose by identifying and managing bias.*
6. *Showing how privacy, ethics, and human rights are safeguarded by regular peer reviews of algorithms to assess for unintended consequences, and how the organisation acts on this information.*
7. *Setting out the nominated point of contact for public inquiries about algorithms -*

together with any internal policies, principles, rules, or guidelines that relate to the above matters.

On 9 April 2024, you clarified your request to:

For clarity, I am specifically interested in the information in the IRD's possession setting out how and when the IRD has discharged, or intends to discharge, all or any of the commitments made by it in becoming a signatory to the Algorithm Charter of New Zealand. The items listed in points 1-7 of your email below are the commitments set out in the Algorithm Charter of New Zealand.

I am requesting any evaluative reports, costing plans, execution or operational reports or plans, reports, internal policies, principles, rules, or guidelines specifically focussed on discharging any or all of the commitments listed in the Algorithm Charter of New Zealand as set out in points 1-7 below. This will include but is not limited to advice or reporting issued to staff, executives, officials, or ministers.

As the IRD made the commitments by becoming a signatory, it has been assumed that the IRD has a plan or an intention to meet the commitments and it is reporting against that plan. If the IRD has no such plan or intention and/or is not attempting to establish or measure whether it is discharging all or any of the commitments made by it in becoming a signatory to the Algorithm Charter of New Zealand, that would also be useful to know.

Information being released

Please note, some documents released do address more than one of the commitments in the Algorithm Charter for Aotearoa New Zealand (the Charter). To avoid repetition, I have listed these only under one question.

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) – making the requested information available would be contrary to the provisions of a specified enactment, namely Inland Revenue's confidentiality obligation in section 18 of the Tax Administration Act 1994 (TAA). 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.

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.

Information outside scope

The enclosed documents contain some information that is outside the scope of your request. This information has not been considered for release and has been withheld as 'not in scope'.

Introduction

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

Inland Revenue's priority is to make sure it uses algorithms in a way that considers its obligations under the Revenue Acts, the Privacy Act, the Charter and all other New Zealand 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.

Inland Revenue utilises many business rules which enable the automation of its business processes. These support its service delivery to both staff and tax/social policy customers. It has integrated rules and algorithms in the products that it has procured through vendors to support productivity and efficiency. IR also has machine learning algorithms which work across large data sets to provide advanced analytics, decision support and predictive modelling.

The Charter states the intention of the application and commitment to the Charter "is to take a particular focus on those algorithms that have a high risk of unintended consequences and/or have a significant impact if things do go wrong, particularly for vulnerable communities". This excludes "most of the many business rules that government agencies use every day to give effect to legislative requirements and for business-as-usual activities".

The majority of algorithms that IR uses are business rules that give effect to legislative requirements (tax and social policy) and for business-as-usual activities. These therefore fall outside the intention of the application of the Charter. I have considered the intention of the Charter, alongside the clarification received from you on 9 April 2024, and as such, I have excluded documents relating to these business rules and machine learning algorithms from the scope of your request.

Where the Charter does apply to IR's use of business rules and machine learning, is to support human decision making in the management of tax/social policy compliance risk and intervention design. In this regard, IR does use complex algorithms and machine learning algorithms alongside data that is collected under the appropriate legalisation. 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. These algorithms also fall under the purview of the supplied governance model, and business and technical processes. There is no automated deployment that could materially and adversely impact any one individual or community. Similarly, there is no existing Generative Artificial Intelligence use case for managing compliance risk.

Question One

Setting out and/or summarising how and when your organisation has met, or intends to meet, the commitments to transparency and partnership set out in the Algorithm Charter for Aotearoa New Zealand

IR's AI Strategic Consideration & Roadmap document (partially released) details how IR, has begun to, and intends to meet the commitments of the Charter.

Governance groups have been set up within IR to provide oversight and direction for how it uses algorithms. These include its Artificial Intelligence Oversight Group and its Artificial Intelligence Working Group. The members of both of these groups work with IR's commitments as a public sector organisation in strengthening the Māori-Crown relationship and integrating te Tiriti o Waitangi and Te Ao Māori concepts and perspectives into IR's work. IR have also established a Māori Data Governance and Sovereignty Steering Group which are stewarding IR's approach to Māori Data Governance and Sovereignty.

Inland Revenue also has the Data and Information Governance Authority (DIGA). The DIGA is an executive-level governance body which ensures right-sized, sustainable, and efficient data and information governance that delivers oversight, transparency, and accountability in IR's stewardship and use of data and information as assets.

Please find enclosed the following documents:

Item	Date	Document	Decision
1.	15/12/2023	Māori Data Governance and Sovereignty Steering Group: Terms of Reference	Released
2.	26/03/2024	Data and Information Governance Authority: Terms of Reference	Released
3.	08/02/2024	Artificial Intelligence Working Group: Terms of Reference	Released
4.	09/02/2024	Artificial Intelligence Oversight Group: Terms of Reference	Released
5.	29/02/2024	AI Strategic Considerations & Roadmap	Partially released

Question Two

Setting out or explaining how decisions made by the organisation are informed by algorithms. This may include without limitation "plain English" documentation of the algorithm/s, information about the data and processes involved, or published information about how data is collected, secured, and stored.

Please refer to my introductory statement for further detail.

Inland Revenue's staff must use its Artificial Intelligence staff use policy and Artificial Intelligence use case guidelines. These set out IRs approach to using algorithms safely and securely in the workplace. They help IR's staff to make good decisions and deliver effective and efficient services.

Please find enclosed the following documents:

Item	Date	Document	Decision
6.	19/12/2023	Get Involved: submitting an idea to the AI Working Group	Released

Question Three

That demonstrate or detail how your organisation is delivering clear public benefits through Treaty of Waitangi commitments by embedding a te ao Māori perspective in the organisation's development and use of algorithms consistent with the Treaty of Waitangi

Inland Revenue recognises the importance of collaboration with tangata whenua in the adoption of algorithms. Please refer to the AI Strategic Considerations & Roadmap for further detail, specifically pages 12-18, and 46-47.

Te Kāhui Tūhono are an internal team who guide IR to contribute to and support mana motuhake through the adoption of Māori indigenous approaches that are underpinned by kaupapa Māori and Te Ao Māori principles and values. They are focused on growing IR's Māori indigenous perspective of Māori Data Governance and Sovereignty within IR. They form part of IR's AI Working Group and engage with the AI Oversight Group. IR also has a Kaihautū Rautaki who provides advice and strategic thinking to the Executive Leadership Team, as well as across Te Tari Taake and the wider public sector, on issues of strategic importance to the Māori Crown relationship.

Engagement and collaboration with whānau Māori and Te Kāhui Tūhono are essential in understanding and embedding a te ao Māori perspective in the development and use of algorithms, being consistent with IR's commitments through Te Tiriti o Waitangi.

Question Four

That demonstrate or detail how your organisation has, or intends to, identify and consult with people, communities and groups who have an interest in algorithms, including Māori.

Please refer to the AI Strategic Considerations & Roadmap for further detail, specifically pages 14-16, 18, 23, 46-48 and 51.

Please find enclosed the following documents:

Item	Date	Document	Decision
7.	16/01/2024	Artificial Intelligence Community of Interest: Overview	Released
8.	09/10/2023	Excerpt from meeting minutes: Artificial Intelligence Oversight Group	Partially released
9.	13/11/2023	Excerpt from meeting minutes: Artificial Intelligence Working Group	Partially released
10.	17/11/2023	Excerpt from Key Messages: Data and Information Governance Authority	Partially released
11.	20/11/2023	Excerpt from meeting minutes: Artificial Intelligence Oversight Group	Partially released

Question Five

That describe how your organisation makes sure data is fit for purpose by identifying and managing bias.

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

Mandatory training is in place for all IR staff (permanent and fixed term) on unconscious bias. This training helps everyone understand, recognise and address unconscious bias.

Please refer to the AI Strategic Considerations & Roadmap for further detail, specifically pages 14-15 and 22-25.

I also enclose the IR Data and Information Quality Model which is a tool used by IR staff to determine the fitness for purpose of IR's data, information and knowledge resources and products.

Item	Date	Document	Decision
12.	11/05/2023	Data and Information Quality Model	Partially released

Question Six

Showing how privacy, ethics, and human rights are safeguarded by regular peer reviews of algorithms to assess for unintended consequences, and how the organisation acts on this information.

Please refer to my response under Question Five for further detail.

If IR were to pursue a use case involving algorithms which involved personal information, a privacy impact assessment would be required to identify any potential risks and develop plans to mitigate them.

Please find enclosed the following documents:

Item	Date	Document	Decision
13.		Privacy & Ethics Impact Assessment Template	Released
14.		Privacy-Threshold Assessment Template	Released
15.		Excerpt from Centre of Enterprise Data and Analytics: Intervention Campaign Process Map	Partially released

Question Seven

Setting out the nominated point of contact for public inquiries about algorithms.

Enquiries from news media can be emailed to mediaqueries@ird.govt.nz

All other enquiries can be emailed to OIA@ird.govt.nz

Part two of request

Evaluative reports, costing plans, execution or operational reports or plans, reports, internal policies, principles, rules, or guidelines specifically focussed on discharging any or all of the commitments listed in the Algorithm Charter of New Zealand.

Please find enclosed the following documents:

Item	Date	Document	Decision
16.	18/12/2023	Artificial Intelligence use case guidelines	Partially released
17.	19/12/2023	Customer facing staff guidance	Partially released
18.	18/12/2023	Artificial Intelligence (AI) Staff Use Policy	Partially released
19.	18/12/2023	Information Security Policy	Partially released
20.	19/12/2023	What is Artificial Intelligence	Partially released

Advice or reporting issued to staff, executives, officials, or ministers specifically focussed on discharging any or all of the commitments listed in the Algorithm Charter of New Zealand.

Please find enclosed the following documents:

Item	Date	Document	Decision
21.	16/07/2020	Excerpt from Status Report: Weekly update for the Minister of Revenue	Partially released
22.	04/08/2020	Advice to Data and Information Governance Authority: Update on Government Algorithm Charter	Released
23.	04/08/2020	Excerpt from Data and Information Governance Authority: Minutes	Partially released
24.	06/08/2020	Excerpt from Status Report: Weekly update for the Minister of Revenue	Partially released
25.	23/08/2021	Report to Minister of Revenue: Inland Revenue's contribution to advancing Government data outcomes	Partially released

Item	Date	Document	Decision
26.	10/05/2022	Advice to Data and Information Governance Authority: Update on the Government Data Strategy and Roadmap	Released
27.	28/07/2022	Email from Craig Jones (StatsNZ) to Mike Cunnington (Deputy Commissioner) and Tina MacLean (Intelligence Leader – Data)	Partially released
28.	21/03/2023	Email from Mary Craig to IR staff	Partially released
29.	07/08/2023	Excerpt from Artificial Intelligence Oversight Group: Minutes	Partially released
30.	19/12/2023	Using Artificial Intelligence at IR: authored by Makayla Stewart	Released
31.	31/01/2024	Advice to Artificial Intelligence Working Group: Publication of IR's transparency obligations under Algorithm Charter for Aotearoa NZ	Released

Information publicly available

Your request for the following documents is refused under section 18(d) of the OIA, as the information is publicly available:

Item	Date	Document	Website address
32.	04/10/2018	Inland Revenue's Tony Morris -let's bust the myth about robot tax investigators	Inland Revenue's Tony Morris – let's bust the myth about robot tax investigators (ird.govt.nz)
33.	17/11/2020	A new world of tax compliance	A new world of tax compliance (ird.govt.nz)
34.	28/04/2021	Open Data	Open Data (ird.govt.nz)
35.	01/07/2023	Initial advice on Generative Artificial Intelligence in the public service	Joint System Leads tactical guidance on public service use of GenAI (digital.govt.nz)
36.	01/07/2023	Generative Artificial Intelligence: System Leaders' guidance for use of gen-AI across the New Zealand Public Service	Joint System Leads Tactical Guidance on Public Service USe of GenAI Summary (digital.govt.nz)

Item	Date	Document	Website address
37.	18/12/2023	Inland Revenue Annual Report: page 196	Inland Revenue Annual Report, Te Tari Taake Pūrongo ā-Tau, 2022–23 (ird.govt.nz)
38.	19/12/2023	Algorithm Charter for Aotearoa New Zealand (2023)	Algorithm Charter for Aotearoa New Zealand (2023) (ird.govt.nz)
39.	26/01/2024	Information Sharing	Information sharing (ird.govt.nz)
40.	15/04/2024	Our use of algorithms	Our use of algorithms (ird.govt.nz)

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.

Publishing of OIA response

We intend to publish our response to your request on Inland Revenue's website (www.ird.govt.nz) as this information may be of interest to other members of the public. This letter, with your personal details removed, will 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 for your request.

Yours sincerely

Jay Harris
Chief Information Security Officer

Inland Revenue
Te Tari Taake

Māori Data Governance and Sovereignty Steering Group: Terms of Reference

The Māori Data Governance and Sovereignty Steering Group (the Steering Group) reports to the Data and information Governance Authority (DIGA) and provide any required updates or decisions as appropriate. The Chair represents the Group to the DIGA.

Purpose

The purpose of the Steering Group is to provide direction and guidance on the activities being undertaken at Te Tari Taake, in response to the publication or request for consultation on any external Māori Data Governance or Sovereignty documents. This can include but is not limited

Not in scope

The Steering Group is responsible for:

- stewarding Te Tari Taake's approach to Māori Data Governance and Sovereignty
- maintaining visibility of Te Tari Taake's approach to Māori Data Governance and Sovereignty and reporting to Executives as appropriate
- providing specialist thought leadership to shape organisation-wide strategic choices and decision making; assisting others to ensure policy, process and decision-making is optimal and complements sector-wide work
- taking an enterprise-wide view, collaborating across Te Tari Taake to identify and proactively manage workflow of the Project Team and capability needs related to Māori Data Governance and Sovereignty
- the decision to undertake internal or external consultation or engagement, relating to Māori Data Governance and Sovereignty
- providing direction and sequencing of IR activities

Membership

The Steering Group Members are:

Enterprise Design and Integrity
Enterprise Services
Policy & Regulatory Stewardship

The members may change as work in this space changes over time.

Roles and Responsibilities

The Steering Group recognises the cultural importance of data to Māori, and the need for culturally grounded models and practices of access, storage, protection and disposal.

The Chair

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

support of the Information Specialist.

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.

Steering Group Members

The Steering 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 Steering 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.

Attendees and observers

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

Panels

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

Logistics

Meetings

The Steering Group will meet 6 weekly 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.

[1] Version 1.1 2023.12.15

Agendas and papers

The agenda and papers are made available to Steering 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 80% 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 Governance Team can prepare minutes if required and provide them to the Steering Group for their following meeting. Within that following meeting, the Steering Group is asked to approve the minutes as an accurate record.

Review

The Steering Group will review its performance at approximately six-monthly intervals, to enable ongoing improvement, and also consider the need to continue the Steering Group.

Reviews will be recorded in the minutes.

Version Control

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

Inland Revenue
Te Tari Taake

Data and Information Governance Authority

Terms of Reference

IR's Executive-Level Governance System

The Executive-Level Governance System supports IR's Commissioner and Chief Executive with their stewardship of the revenue system and their responsibilities for IR's financial management, performance, sustainability, organisational health, and capability.

Our governance bodies steer us in the near and medium term, and they look out to the long term, ensuring sustainability and agility for the future. They ensure we successfully balance competing priorities, take up the right opportunities, and deliver for government and customers in the face of risks and challenges.

The Executive-Level Governance System evolves so that it continues to be fit for purpose. Governance roles, practices and processes iterate as needed to ensure the system and the bodies within it remain effective and responsive. The system is reviewed every six months, and any proposed changes to the structure are approved by the Commissioner.

Purpose

The Data and Information Governance Authority (DIGA) ensures right-sized, sustainable, and efficient data and information governance that delivers oversight, transparency, and accountability in IR's stewardship and use of data and information as assets.

The Authority is responsible for:

- Custodianship of Enterprise Risk 5 and ensuring appropriate practice management of identified risks and controls, in accordance with IR's risk settings
- Ensuring external assurance activities are completed annually by IR
  - Government Chief Privacy Officer's Privacy Maturity Assessment
  - Archives New Zealand's Information Management Maturity Assessment
  - Public Service Commission's Information Gathering and Public Trust Model Standards

Membership

Members are:

Business Unit	Role
Enterprise Services	Deputy Commissioner (Chair)
Enterprise Design and Integrity	Deputy Commissioner
	Chief Information Security Officer
	Privacy Officer

Not in scope

Business Unit	Role
	Enterprise Leader – Strategic Architecture
Policy & Regulatory Stewardship	Deputy Commissioner
Tax Counsel Office	Chief Tax Counsel
Customer & Compliance Services - Individuals	Deputy Commissioner

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.

Authority Members

The Authority 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.

All papers must be sponsored by an Authority member, who ensures the paper provides the right information and meets the standards required for robust discussion and decision making.

Information Governance Team Support

The Authority is supported by the Information Governance team. They are responsible for:

- ensuring the Authority follows good governance principles and practices
- working with the Chair to set the agenda, considering key enterprise issues and decisions required
- providing support, advice and quality assurance for Authority papers, with the mandate to refuse any that do not meet quality standards
- 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 Authority meetings to provide input as needed.

Observers may attend with approval from the Chair. Observers do not contribute to the Authority's discussion unless a member asks them to.

Logistics

Meetings

The Authority meets quarterly, and members are advised of the meeting dates set for the year.

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.

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.

Agendas and papers

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

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

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 Governance Team writes the minutes and provides them to the Authority for their following meeting. Within that following meeting, the Authority is asked to approve the minutes as an accurate record.

Review

To enable ongoing improvement, the Authority 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[1] for continuity.

[1] Version 2.0 2024.03.26

Artificial Intelligence (AI) Working Group Terms of Reference

The Artificial Intelligence Working Group (The Working Group) reports to the Artificial Intelligence 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 AI Oversight Group. The Working Group will help with co-ordination, cohesion, and prioritisation as directed by the AI 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)
	Technical Lead, Organisational Development
	Domain Principal, Technology Commercial
	Domain Principal, ES Planning, Design & Delivery
	Domain Specialist, Digital Product Experience
	Change Analyst, Change, Design & Enablement
	Domain Lead, Technology Experience
	Domain Lead, Digital Product Experience
Enterprise Design and Integrity	Domain Lead, Technology Architecture

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

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 AI 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 AI 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 2.0 2024.02.08

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)
	Enterprise Leader, Strategic Architecture
	Intelligence Leader, Centre for Enterprise Data & Analytics (CEDA)
	Chief Information Security Officer
	Domain Lead, Digital Ecosystem
	Strategic Advisor, Enterprise Design & Integrity
	Enterprise Leader, Strategic Portfolio Stewardship
	Kaihautū Rautaki, Enterprise Design & Integrity
Enterprise Services	Intelligence Leader - Enterprise Information and Knowledge (EI&K)
	Enterprise Leader, Technology Services
Customer & Compliance Services - Individuals	Domain Lead, Planning, Design & Delivery - Digital
Policy & Regulatory Stewardship (PaRS)	Policy Director

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 2.0 2024.02.09

Inland Revenue
Te Tari Taake

AI Strategic Considerations & Roadmap

Version: FINAL 29/01/24

[Not in scope

Table of Contents

Not in scope

[Not in scope

Executive Summary

At Inland Revenue we have a structure and governance in place for AI and are now looking to further develop strategic considerations and a roadmap for this capability. This is a key input to Inland Revenue's Enterprise Strategy and long-term planning.

We have established a strategic framework to guide the secure and effective adoption of artificial intelligence (AI) within Inland Revenue. Our AI strategy consideration pack is informed by comprehensive research and collaboration with experts, and it presents an approach structured around seven strategic pillars that resonate with our organisational goals, ensure fair user experiences, and maintain secure, transparent innovation.

Key Components:

- **Strategic Considerations:** The foundations of our AI strategy consideration pack are pillars that align with our organisational aims, champion inclusivity, fortify security, prepare infrastructure, and conform to regulatory standards.

- **Prioritisation and Roadmap:** We introduce a prioritisation model for AI initiatives, coupled with a phased roadmap. This framework provides direction for foundational and optional activities from the initial stages through to broader implementation.

- **Impact on Inland Revenue:** This framework allows us to carefully consider ways to integrate AI into IR's operations and people capabilities. AI has the potential to increase operational efficiency and enhance the customer experience, whilst supporting IR's broad roles as defined in our enterprise strategy.

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

[Not in scope

Strategic Alignment

Effectiveness & Efficiency | **Stewardship** | **Broader Contribution**

- **Effectiveness & Efficiency:** By integrating AI, we can optimise our resource utilisation and streamline operations, enhancing our decision-making capabilities and improving our overall effectiveness and efficiency through optimisation and automation.
- **Stewardship:** Stewardship is crucial as we implement AI, and with careful planning and strategic allocation of resources, we can improve our adaptability, safeguard existing systems, and maintain the integrity and the relevancy of the tax and social policy system.
- **Broader Contribution:** As we strive to make a broader contribution to New Zealand, integrating AI can enhance cross-departmental data processing and enhance our capabilities for efficient partnerships, fostering holistic insights and innovative solutions.

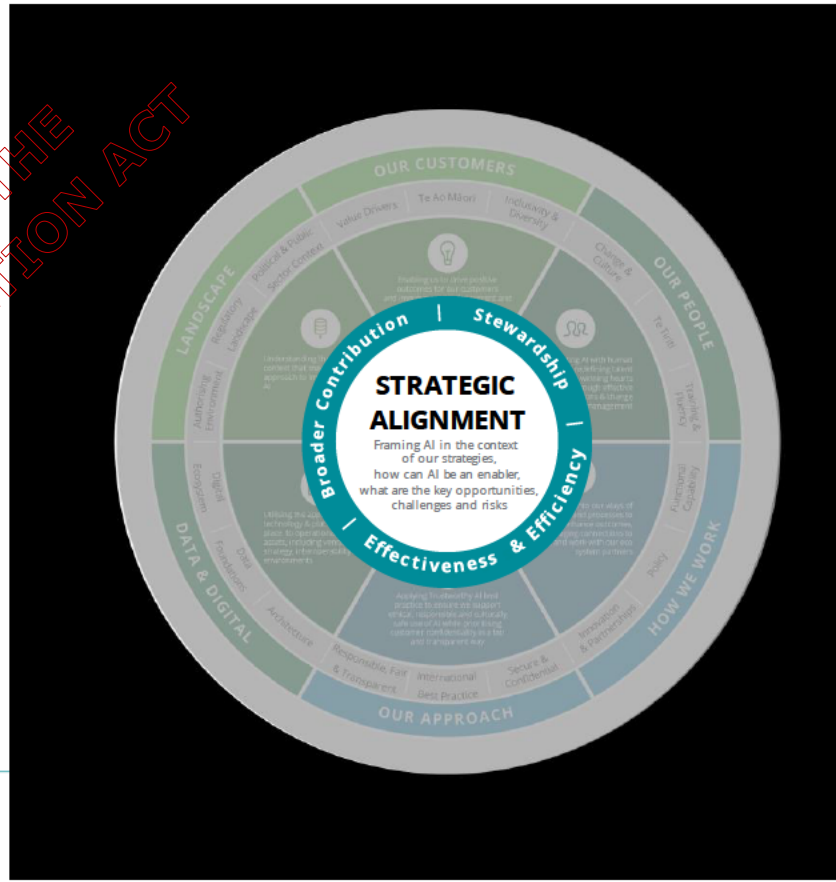

Inland Revenue
Te Tari Taake

[Not in scope

Strategic Alignment

STRATEGIC CONTEXT

Aligning the use of AI to accelerate the realisation of our organisation's aspiration and outcomes.

Our Enterprise Strategy is central to shaping the approach to AI. This strategy revolves around four key elements: our people, the driving force behind our operations; our customers, who guide our operational focus; supporting the Crown to be a better Treaty partner, promoting equitable services across Aotearoa New Zealand; and maintaining the integrity of the tax and social policy system, our underlying cornerstone. AI holds the potential to strengthen these pillars by streamlining processes, anticipating risks, and optimising public service delivery. However, as AI innovations and data flow escalate, stewardship becomes more critical. Falling behind in this rapidly evolving field could jeopardise our ability to effectively maintain our services. It's also crucial that te ao Māori principles are deeply woven into our strategy. By aligning AI with our objectives, we can strive towards promoting 'Oranga' via cross-government cooperation, robust social policy, and revenue generation. This combined strategic precision and deepened operational understanding could significantly elevate our performance.

Not in scope

1. Not in scope

[Not in scope

Our Customers

Value Drivers	Te Ao Māori	I

- **Value Drivers:** We aim to leverage AI to enhance efficiency, and service delivery, refocusing our resources on higher-value improved customer outcomes.
- **Te Ao Māori:** We seek to ensure that data sovereignty princi AI initiatives, recognising Tangata whenua as kaitiaki (guardi data and fostering oranga (health and well-being) through eq processes.
- **Inclusivity & Diversity:** Ensuring our AI strategy at Inland R prioritises accessibility and inclusiveness for all our customers e is key in delivering a diverse, equitable, and representative serv with the principle of manaakitanga in Te Pou o Te Tangata.

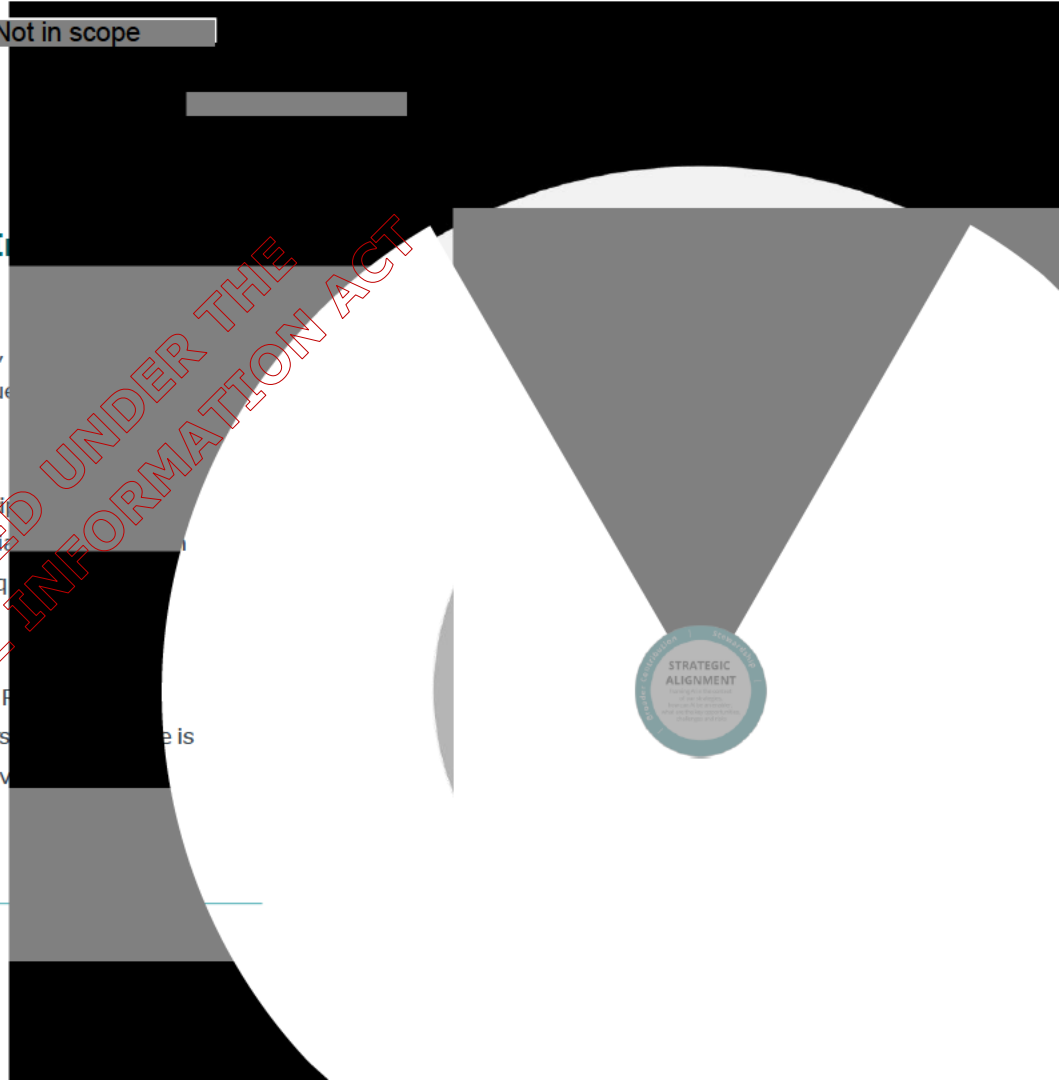

Inland Revenue
Te Tari Taake

Not in scope

[Not in scope]

Te Ao Māori

STRATEGIC CONTEXT

Adopting a Māori lens when considering how we use AI to better serve Aotearoa.

In order to foster oranga in all areas, it will be beneficial for us to think holistically when considering how to create enhanced outcomes for tangata whenua through the adoption of AI.

He rei ngā niho, he paraoa ngā kauae – One must have the right principles for large undertakings.

STAKEHOLDER FEEDBACK

AI has potential to counteract unconscious bias, and its implementation at IR will require careful consideration of data storage locations. Collaborative opportunities with Te Kāhui Tūhono need to be explored and any decisions must be viewed through a risk lens.

Interviewees: Anil, Mike, Brijesh, Mary, Cate.

OPPORTUNITIES & CHALLENGES

- The adoption of AI tools presents a unique opportunity to build more trust with Māori communities, providing principles of Māori data sovereignty are incorporated. Notably, there is a need to ensure Māori are kaitiaki of their own data and prioritise benefits back to the community.
- We have a key opportunity to think at the level of the broader tax system, and set out to reduce inequitable outcomes for Māori communities through AI implementation. We could explore opportunities to incorporate mātauranga Māori in AI systems through sustained community engagement.

RISKS

- Many AI applications require data to be stored in overseas data centers, which may conflict with principles of indigenous data sovereignty whereby Māori have the intrinsic right to control over their own data. Similarly, many Māori communities have expressed concern over the use of their language data to train Large Language Models for generative tasks without their consent, when they should be active decision-makers in the process [1].
- Potential for AI initiatives to perpetuate further harm, if insufficient consideration is given to te ao Māori. Particularly, care is needed to ensure that further bias is not encoded in AI models through the use of imbalanced datasets, and that generative AI tools do not misrepresent reo and tikanga in its outputs. Preventative steps should be communicated to the public.

KEY INSIGHTS

- Principles of Māori Data Sovereignty are designed to comprehensively cover all aspects of the collection, storage and use of Māori data. Notably, Te Mana Raraunga's Māori Data Governance model can be used to support IR efforts to foster oranga for current and future generations [2]. Critically, data is considered to be taonga and should be collected respectfully in ways that prioritise Māori needs, fostering kotahitanga (collective benefit).
- Data storage considerations include ensuring Māori control as well as sufficient privacy and security measures, ensuring that tangata whenua are kaitiaki over their own data. Finally, the use of Māori data should benefit the Māori community, with the aim of reducing inequities over time and putting protection in place to prevent future harm.
- Work in this space should be done in conjunction with Te Kāhui Tūhono within IR.

ROADMAP CONSIDERATIONS

Partnership & Engagement with Māori

1. *Indigenous groups in NZ, US fear AI colonisation | Reuters*
2. *Te Kāhui Raraunga*

[Not in scope]

Inclusivity & Diversity

STRATEGIC CONTEXT

Ensuring accessibility of our AI tools for the organisation and customers, and increasing inclusive outcomes.

This consideration will be of utmost importance when incorporating AI into Te Pou o Te Tangata, or how our enterprise strategy is delivered. It aims to ensure that the entirety of our diverse customer base and workforce is included on the journey of AI implementation, and that accessibility is prioritised.

STAKEHOLDER FEEDBACK

As technology becomes more integrated, efforts must be made to prevent exclusion of individuals who lack access, skills, motivation, or trust in digital capabilities.

Interviewee: Anil.

OPPORTUNITIES & CHALLENGES

- AI, and specifically GenAI, is poised to help users more easily access relevant information, providing a key opportunity to improve the inclusivity of our channels [1]. To activate the full gains of AI tools, we should ensure that services are fully accessible, and carefully considered to all internal and external end-users, considering the diversity of the taxpayer base, providing multilingual options and support for those with disabilities.
- Seeking diverse perspectives (manaakitanga) to inform the work we do, a key pillar of Te Pou o Te Tangata, will ensure our AI initiatives benefit all of our people. Internal workshops, with a team of staff from all backgrounds and levels of the organisation, will notably help to eliminate possible biases in AI algorithms from early in the design process, as any potential blind spots will quickly be identified by such a team [2].

RISKS

- Without sufficient engagement and co-design, there is a risk that AI solutions will only be beneficial for certain segments of our workforce and customers. This connects to the broader issue of digital exclusion, whereby certain groups may be less able to use digital tools (including those built with AI) due to a lack of experience with technology, among other factors [3]. This risk is particularly pronounced in our organisation's case, as tax administration services need to be fully accessible to the wider population.
- Particular consideration should be given to upskilling and engaging Māori, as is further discussed in the foundations segment of this strategic framework. Otherwise, there is a risk that AI tools will perpetuate harm for these communities rather than being part of the solution.

KEY INSIGHTS

- People living with disabilities have expressed a great deal of interest in how the use of AI is poised to make their lives easier. Indeed, there are many potential benefits of use cases such as a Generative AI assistant, for example, to help customers more easily navigate through and interpret material when interacting with our organisation [1].
- Inclusivity by design can be implemented in AI tools through a collaborative engagement process. Notably, consulting a diverse range of end-users will be critical to drive innovation and development, including hearing about their previous experiences of using AI applications and taking suggestions for improvement on board. Specifically engaging with members of the neurodivergent community can, for example, help our designers simplify content in ways that may seem small but make the world of difference when it comes to accessibility [4].
- Inclusive design can also be enabled through technical measures that target accessibility, such as speech-to-text capability or keyboard navigation. These are most effective when built in early on in the process, so that they are central to the design of a particular solution rather than worked in just before production [1].

ROADMAP CONSIDERATIONS

Use Case Proposal & Prioritisation Framework

Partnership & Engagement with Māori

1. Designing GenAI for People with Disabilities (hbr.org)
2. us-ai-institute-state-of-ai-fifth-edition.pdf (deloitte.com)
3. OECD State of the art | (oecd-ilibrary.org)
4. Inclusive Design (deloittedigital.com.au)

Related Resource 1: DI_CDO-Playbook-2023 (deloitte.com)
Related Resource 2: Entering the world of tax (deloitte.com)
Related Resource 3: Scaling AI in government (deloitte.com)

Not in scope

Our People

Not in scope

Not in scope

- **Te Tiriti:** We recognise the importance of u[illegible]ommitments to Māori under Te Tiriti o Waitangi within our AI strat[illegible]edging the potential of AI to enhance our contributions to oranga[illegible]nga whānau and oranga whenua. This is achieved through a [illegible]agement and partnership with Māori, ensuring the fair an[illegible]se of AI tools in tax services that respect and incorporate mātau[illegible]

Not in scope

Not in scope

[Not in scope

Te Tiriti

STRATEGIC CONTEXT

Ensuring we support the Crown to deliver on our commitments to Māori under te Tiriti o Waitangi. In the context of AI, this aligns with commitments made as a signatory to the Algorithm Charter.

We aim to move closer to our aspirations of improving oranga by becoming te Tiriti-based, establishing this consideration as a crucial foundation for our approach to AI.

STAKEHOLDER FEEDBACK

We should formulate our AI strategy with particular emphasis on working in alignment with Te Kāhui Tūhono, as this would promote the values embedded in Te Tiriti.
Interviewee: Brijesh.

OPPORTUNITIES & CHALLENGES

- The contribution to oranga tāngata, oranga whānau and oranga whenua we have made has the potential to be greatly enhanced by AI initiatives that build on a meaningful, mutually beneficial partnership with Māori. There is an opportunity to engage with iwi Māori to raise awareness on the potential benefits of using AI for tax purposes, as well as to foster collaborative co-design, agreeing on system parameters of relevant AI tools and creating space in which mātauranga Māori can be shared.

RISKS

- Potential for tangata whenua to have a low appetite to engage with customer-facing AI solutions, due to a perception that AI will impact their right to control how services are delivered to them and how their information is utilised. This highlights the importance of transparency of AI tools and accountability, ensuring the AI tools are seen as trustworthy by both the organisation and the community.
- There is a risk that, if partnership with Māori is seen as a one-time engagement process, the full benefit of our AI initiatives may not be derived for Māori communities. Instead, collaboration with tangata whenua should be seen as a foundational element of our AI framework that will enable its broader success. As AI tools continue to evolve, the connections to Te Tiriti will also change and need to be considered on an ongoing basis. Specifically, fairness and transparency should be prioritised to foster equitable outcomes and align with the associated equity principle of te Tiriti.

KEY INSIGHTS

The key principles of te tiriti are highly relevant when implementing AI initiatives. The protection of Māori interests and prevention of future harm should be front of mind, alongside the incorporation of mātauranga Māori in all initiatives. This is supported by strong partnership with Māori, and a willingness to prioritise participation and engagement with tangata whenua. Many Māori researchers are already involved in determining how AI can be made mana-enhancing and beneficial for Māori, and several academic workshops have been held across the country [1]. These efforts have been supported by Te Hiku Media, among others. These organisations are partnering to implement AI-enabled linguistic resources for te reo Māori, in collaboration with local iwi [2]. This ongoing initiative has received government funding [3] and provides a practical example of how effective partnership, underpinned by te Tiriti, can provide beneficial outcomes for Māori.

- When planning an AI initiative, the three voices framework can help with the synthesis of information from different knowledge sources across complex systems. Specifically, the voice of intent (legislation), voice of expertise (research) and voice of experience (community perspective) all come together to form a holistic view of how collective benefit can be derived for all involved [4].

ROADMAP CONSIDERATIONS

Te Tiriti Alignment

1. *Māori Speech Hui 2021 | Speech research @ UoA (auckland.ac.nz), Wānanga for Māori Artificial Intelligence: University of Waikato*
2. *Papa Reo (tehiku.nz)*
3. *Te Hiku Media Awarded $13M*
4. *Three Voices Infographic (Deloitte.com)*

Related Resource 1: Ngā mātāpono o te Tiriti o Waitangi (teara.govt.nz)

Not in scope

Not in scope

Not in scope

[Not in scope

Policy

Not in scope

Not in scope

KEY INSIGHTS

- Effective governance structures are essential to instill trust and confidence in AI systems, especially in addressing issues of bias and discrimination [2].
- Establishing a centre of excellence or similar structures can help in developing best practices, sharing knowledge, and ensuring the quality of AI deployments [2]
- Documenting and enforcing machine learning operations (MLOps) is crucial for the ethical and effective deployment of AI, as it facilitates addressing any issues with AI models. [5]
- Embracing new operational models and processes is vital for leveraging AI in a way that drives sustained quality, innovation and value creation[5].

RISKS

- Policies need to account for the risks associated with data privacy and protection, especially considering regulations which place restriction on automated decision making and profiling. [3]
- Operational policy must include mechanisms to continuously evaluate and mitigate biases in AI algorithms and secure sensitive financial data against unauthorised access or breaches. [4]
- There is a need for clear documentation and enforceable processes to address ethical risks, such as bias or misuse of AI, underlining the importance of robust governance structures.[2]

Not in scope

1. *Building the digital revenue agency of the future*
2. *DI crafting an AI strategy for govt leaders*
3. *Ddtl tax artificial intelligence in tax*
4. *Us Deloitte the implications of generative*
5. *State of ai for government*

[Not in scope

Innovation & Partnerships

Not in scope

Not in scope

KEY INSIGHTS

Not in scope

- Ethical frameworks and toolkits are essential in AI implementation, promoting privacy, reducing bias, and ensuring diverse and inclusive design teams [4].

Not in scope

RISKS

Not in scope

- Collaborations, especially those involving data sharing and AI implementation, must navigate ethical frameworks and privacy concerns, ensuring adherence to standards and mitigating bias [4][5].

Not in scope

Not in scope

Not in scope

Not in scope

4. Crafting an AI strategy for government leaders | Deloitte Insights

5. DI_building-the digital-revenue-agency.pdf (deloitte.com)

Not in scope

[Not in scope

Our Approach

Responsible, Fair & Transparent | **International Best Practice** | **…al**

- **Responsible, Fair & Transparent:** We are committed to d… operating our AI systems in a socially responsible way, prov… …qual treatment for all and maintaining a level of transparency tha… understand and trust.
- **Secure & Confidential:** As the custodians of New Zealand'… we must prioritise implementing Trustworthy AI principles to… robust security and confidentiality measures, which are intri… …ng trust and confidence.
- **International Best Practice:** Leveraging insights from Int… standards and collaborating with international bodies will en… AI-related frameworks well-tailored to New Zealand's contex… efficient learning transfer and enhancing organisational tran… Revenue.

Inland Revenue
Te Tari Taake

[Not in scope

Responsible, Fair & Transparent

STRATEGIC CONTEXT

A socio-technical ecosystem enabling AI that is ethical, lawful and technically robust. It is achieved through governance of AI risks across people, processes and technology – in a systematic fashion, leveraging an impact-based tiering approach.

As the adoption of AI grows, so too does the need for fostering trust and accountability. Compared to private sector organisations, government agencies face additional legal and risk constraints when it comes to AI adoption. It is imperative to ensure everyone can understand and scrutinise how their data is being used by us, how AI decisions are being made, and who is responsible for them.

This is consistent with our Stewardship role – as set out in the Enterprise Strategy, we have a responsibility to ensure that any AI solutions implemented are trustworthy to be fit for purpose today and tomorrow.

STAKEHOLDER FEEDBACK

We must be responsible, fair and transparent in our application of AI. This involves effective risk management against malicious actors and bias, maintaining the integrity of the tax system, and ensuring rigorous testing and explainability of AI models. The credibility and trustworthiness of the AI system is crucial, hence, security, anonymity regarding Personal Identifiable Information (PII), and an absence of bias in practice are paramount. Also, engaging experts and improving knowledge, especially among decision-makers, on ethics and data governance is essential to avoid potential public relations issues related to AI use. Interviewees: James, Cate, Anil, Tina, Lisa.

OPPORTUNITIES & CHALLENGES

- We have the opportunity to introduce guardrails embedding Trustworthy AI principles early in AI journey, ensuring AI systems are responsible, reliable, fair and transparent.
- Responsible: Create and operate the technology in a socially responsible manner, with clear accountability of who is responsible for decisions made using AI systems.
- Robust & Reliable: Confirm that AI functions properly beyond pilot stage, producing consistently accurate, relevant outputs.
- Fair: Guard against illegal and unethical discrimination, ensuring equitable treatment of all.
- Transparent: Help users (IR employees/ the public) understand how their data can be used and how AI systems make decisions.

RISKS

Most organisations are still grappling with the risks associated with traditional AI, and Generative AI brings renewed attention to these. We should seek a balanced view of value creation opportunities with risks involved. Examples include:

- Generation of misleading or false content, potentially causing confusion or harm to users who blindly rely on the generated output – e.g. our employees using Generative AI knowledge tools without proper training.
- Users may not be aware that specific IR content was machine-generated. We must communicate how the system works and build transparency and trust.
- Bias in the underlying data is a risk that can be amplified when AI models are trained on them – e.g. public-facing AI assistant less able to answer questions from certain socio-economic groups due to lack of past examples, perpetuating barriers to use.

Not in scope

Not in scope

Not in scope

[Not in scope

Secure & Confidential

STRATEGIC CONTEXT

To maintain trust and confidence, it is paramount in today's data-driven landscape to ensure the security and confidentiality of AI.

It is important to implement robust measures to safeguard sensitive information and protect against risks, including data breaches, that may cause physical and/or digital harm. Privacy must also be respected, and consumer and employee data should not be used beyond its intended and stated use.

This is also consistent with our Effective and Efficient role – as set out in the Enterprise Strategy, we have a responsibility to use our knowledge, resources and capabilities wisely.

STAKEHOLDER FEEDBACK

We should prioritise risk management, information security and elimination of bias, to preserve the integrity of the tax system. Ensuring model explainability, stringent testing and maintaining the anonymity and security of Personally Identifiable Information (PII) is crucial. Developing deep expertise and building capability relating to ethics and data governance are important points. It is vital to accommodate the evolving advances yet ensuring sensitive data is securely held for the credibility of both the tax system and the organisation.
Interviewees: James, Anil, Tina, Lisa.

OPPORTUNITIES & CHALLENGES

- We have the opportunity to introduce guardrails embedding Trustworthy AI principles early in our AI journey, ensuring AI systems are secure and confidential.
- Preserving Privacy: Train AI models on representative data without compromising sensitive sources of training data.
- Safe & Secure: Ensure that the technology, and the data that feeds into it, is protected from risks that may cause individual and/or collective physical, emotional environmental, and/or digital harm.

Not in scope

KEY INSIGHTS

- There is increased guidance from the Privacy Commissioner on Artificial Intelligence and the Information Privacy Principles [1]. These may provide a useful starting place in understanding our security and confidentiality needs.
- Additionally, global frameworks may be leveraged. Most notably, the NIST AI Risk Management Framework which has been adopted by other agencies in the New Zealand public sector [2]. The recently published ISO42001:2023 AIMS can be leveraged alongside NIST to develop a robust AI risk and controls library in the near term, and work towards compliance in the longer term [3].
- To ensure effective and efficient leadership, an IR specific approach should be taken to ensure confidentiality and integrity is maintained. This could be achieved through tools such as the AI Impact Assessment (tailored to our needs), and MLOps (tailored to the development systems in play at IR).

ROADMAP CONSIDERATIONS

AI Risk Framework | AI Risk Assessment

AI Governance Framework | AI Governance Pilot

1. AI-and-the-Information-Privacy-Principles.pdf
2. AI Risk Management Framework | NIST
3. ISO 42001:2023 AIMS Standard

[Not in scope

International Best Practice

STRATEGIC CONTEXT

Using global context to inform our navigation of AI-related risks as New Zealand's tax authority.

International standards will set a benchmark that we may aim to meet or surpass when defining our approach to AI. This consideration feeds into Te Pou o te Tangata, in that it involves seeking diverse perspectives to inform the work we do.

STAKEHOLDER FEEDBACK

Tax Admin 3.0 is a key area of focus

OPPORTUNITIES & CHALLENGES

- There may be an opportunity to partner with an international body to co-design AI-related frameworks, building on existing standards and ensuring relevance to the specific New Zealand context. In 2020, the New Zealand government partnered with experts from the World Economic Forum in a pilot project to design a fit-for-purpose governance framework, producing recommendations to be fed into future work [1]. Any further collaboration in this space has the potential to be highly beneficial to IR, as international best practices can be directly incorporated into local strategy.
- Learning and sharing enables efficient learning transfers as well as improves organisational transparency. We could leverage this to solidify relationships and capability

Not in scope

Not in scope

1. WEF_Reimagining_Regulation_Age_AI_2020.pdf (weforum.org)
2. White House AI Executive Order | Brookings
3. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
4. AI-Report-Online.pdf (oecd-opsi.org)
5. Tax Administration_3.0 (oecd-ilibrary.org)

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

[Not in scope

Wider Regulatory Landscape

Not in scope

Not in scope

OPPORTUNITIES & CHALLENGES

- In line with our target outcomes, the Algorithm Charter recognises the power of AI to help public sector organisations "deliver services that are more effective and efficient" [1]. Due to our broad ranging public role, any AI use cases have the potential to significantly impact wellbeing, whether intentionally or not. Applying the charter's recommendations around Trustworthy AI and Te Tiriti will enable us to mitigate any risks as they arise.

Not in scope

RISKS

- Extra care is required when using third-party services that integrate AI. These might be subject to international regulations surrounding AI and privacy, particularly if their service uses international data sharing. As such, procurement processes should be strengthened to ensure provider use of AI is compliant with all applicable standards [3]. We should seek transparency and control over how providers are implementing AI, to ensure that any exposure of personal information or associated harm is prevented.
- Any deviation from applicable standards for AI use could result in severe reputational damage for our organisation due to its wide-ranging role as steward of New Zealand's tax system. Strong transparency and human oversight are especially critical for us, as inaccurate AI output could jeopardise public trust.

KEY INSIGHTS

- Regulatory compliance should be future-focussed - not only appreciating today's legislation but also proactively anticipating policies to follow. We should monitor for changes in regulatory policy that could impact our AI initiatives, and particularly for the introduction of any AI-related legislation. Thus far, pathways to AI regulation have been similar internationally, with countries moving from investigating AI capabilities to actively growing the industry [4]. As many governments begin looking at how to shape and regulate AI development, policies may diverge based on local factors. We should watch this regulatory landscape carefully, to pre-emptively ensure compliance.
- We should develop particularly robust procedures around the use of Generative AI. In addition to the privacy protections outlined in the Privacy Act 2020, specific Generative AI guidelines the Privacy Commissioner earlier this year focus on the importance of human oversight, including validation to ensure accuracy and confidentiality, as well as feedback mechanisms to enable improvement [5].

ROADMAP CONSIDERATIONS

AI Standards

AI Governance Standard Operating Procedures

1. *Algorithm-Charter-2020_Final-English-1.pdf (data.govt.nz)*
3. *Interim Generative AI guidance | NZ Digital government*
4. *AI regulation | Deloitte Insights*
5. *Privacy Commissioner outlines expectations around AI use*

Not in scope

Related Resources 1: Privacy Act 2020 No 31 (as at 01 November 2023)

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Roadmap for AI

Inland Revenue
Te Tari Taake

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Roadmap notes – Foundations – Our Customers

Activity	JIRA Key	Description	Phase	Status	Dependencies
Not in scope					
Partnership and Engagement with Māori		Collaborate with iwi Māori to enable mana-enhancing usage of AI tools and techniques in ways that uphold the principles of te Tiriti, Māori Data Sovereignty, and te Ao Māori more broadly. This will help to cement trust in these tools so that they are widely used to benefit tangata whenua. Building blocks: -Work with internal team, Te Kāhui Tūhono, to design engagement process, ensuring a wide range of communities are consulted and considering the three voices framework if desired -Undertake discussions including with Te Kāhui Tūhono -Implement recommendations and design principles where relevant -Plan further engagement so that Māori feedback can be incorporated on an ongoing basis			• Te Tiriti Alignment

Our Customers

[Not in scope

Roadmap notes – Foundations – Our People

Activity	JIRA Key	Description	Phase	Status	Dependencies
Not in scope					
Our Peop					
Te Tiriti Alignment		Consider how AI implementation at IR can best uphold the principles of Te Tiriti o Waitangi. Building blocks: -Consult with tangata whenua, system leaders and across government to understand diverse perspectives -Implement findings		To do	

[Not in scope

Roadmap notes – Foundations – How We Work

	Activity	JIRA Key	Description	Phase	Status	Dependencies
How we work	Not in scope					
	Cross-agency initiatives	AIW-33	Consider how IR can uplift its own AI capability and prioritise innovation by partnering with other agencies. Building blocks: -Participate in AoGP rogramme of Work and related activities -Undertake further environment scanning to identify ongoing opportunities for collaboration		In Progress	• Funding
	Not in scope					

Not in scope

Not in scope

[Not in scope

Roadmap notes – Foundations – Landscape

	Activity	JIRA Key	Description	Phase	Status	Dependencies
Landscape	**Public Engagement**		Engage with diverse groups from New Zealand's taxpayer base to ensure that the principles of IR's approach to customer-facing AI use cases have adequate social license to be implemented and drive benefits for the population. Building blocks: -Design engagement process, ensuring a wide range of individuals will be consulted. As part of this, ensure transparent communication in customers channels surrounding IR's use of AI -Undertake discussions -Review findings -Implement recommendations		To do	
	Not in scope					

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

Not in scope

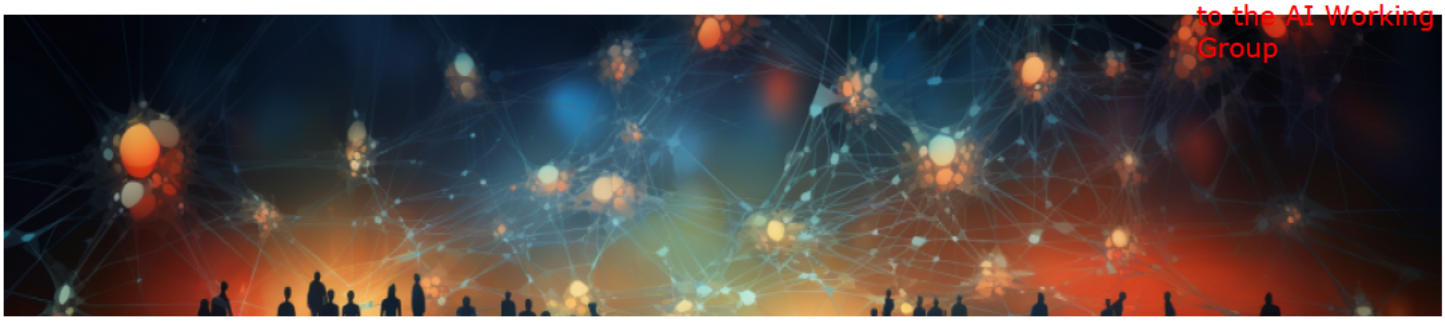

Get involved

Makayla Stewart
Change Analyst (L2)

Here is how you can get involved in the world of Artificial Intelligence.

Submitting an idea to the AI Working Group

Prior to making a submission, please ensure you have read through the AI strategic considerations and roadmap.

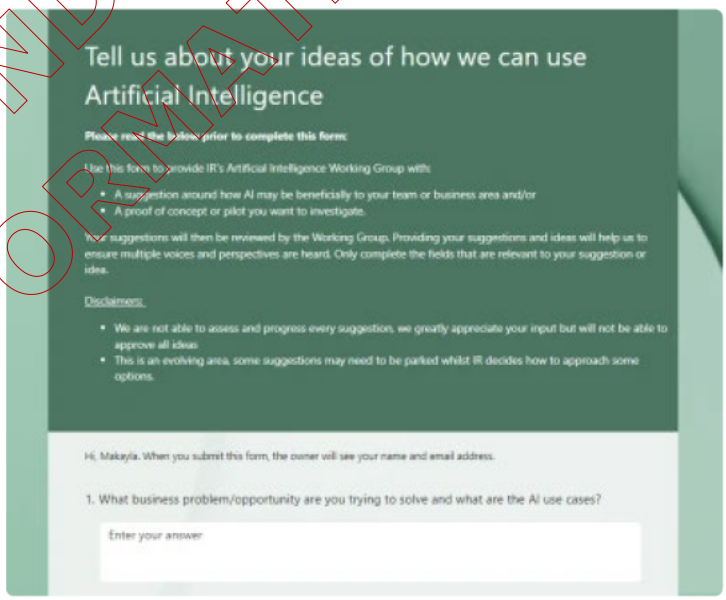

Proof of Concept or Pilot

Proofs of concept, pilots and trials must be raised to the AI working group for review prior to being approved by the AI Oversight Group.

Please complete the below form with as much detail as possible and we will be in contact.

Tell us about your ideas of how we can use Artificial Intelligence

An opportunity you've identified

If you have an opportunity you want to let the AI Working Group know about that you see for yourself, your team/business area or for Inland

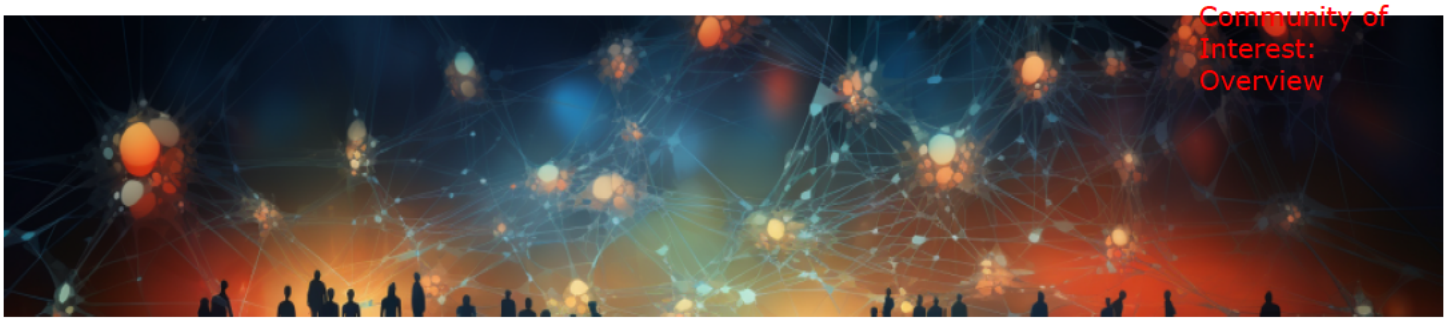

Artificial Intelligence Community of Interest

Makayla Stewart
Change Analyst (L2)

The Artificial Intelligence Community of Interest is focused on connecting like-minded people to grow our collective knowledge about Artificial Intelligence.

The community of interest is open to everyone at IR, we have both an active Teams site and Viva Engage were members are encouraged to share items of interest with the community.

Email AI@ird.govt.nz to join now!

Overview

Purpose

The community is capability and knowledge focused, the purpose is to connect our people who are interested in or passionate about AI to discuss the latest developments, how we are using the technology and where we can see opportunities in our organisation.

Responsibilities

The community holds a monthly meeting that is run by a facilitator, during these meetings we discuss:

- The latest developments in AI
- Updates from the AI working group and AI oversight group
- Opportunities we can see within our organisation for our people, customers and/or business partners
- Challenges or risks we can see within our organisation

Guest speakers both internally and externally are invited to our meeting to discuss their journey implementing and adopting this technology.

Artificial Intelligence Oversight Group

Inland Revenue
Te Tari Taake

Meeting Minutes for 9 October 2023, 1:30 – 2:30 pm

Core members	Brijesh John (Chair), Anil Srinivasa, Cate Robertson, Daniel Blank, Jay Harris, Jesse Thwaites, Malcolm Breadmore, Phil Whittington, Ron Grindle, Scott McCallum, Tina MacLean
Standing members	Conrad Bace, Graham Poppelwell, Vanessa Johnson, Virginia Flaus
Attendees	Tina McCaffrey, Ryan Hamilton, Makayla Stewart, Tanya Williams
Apologies	Underlined above

1. Welcome and Karakia

The Chair opened the meeting with a karakia.

2. Approval of previous meeting's minutes

The Oversight Group had no objections to the minutes of the 18 September 2023 meeting and were **accepted**.

There were no current open action items to discuss.

3. Aotearoa AI Summit. INFORM

Presenters: Tina McCaffrey, Service Owner, PD&D – CX / UX Design, Ryan Hamilton, Customer Experience Designer, PD&D – CX / UX Design, PD&D – CX / UX Design

An inform item providing insights on the AI summit run by the AI Forum in New Zealand on 25/09/2023 the presenters attended. The summit focused primarily on generative AI and with a particular emphasis on what we're doing or not doing in New Zealand. Talks were from both private sector companies and from the NZ public sector.

Some of the key themes from the presentations and discussions were:

- New Zealand rates very low on productivity measures and AI and technology is seen as one of the important ways NZ could bridge that productivity gap
- AI usage in NZ is low compared to other countries and with a tendency of being apprehensive about AI technologies
- Building of AI from a Maori perspective requires engagement and genuine interest in partnerships for the long term. Data is a taonga and its use must be in an appropriate cultural and ethical way.

- Trust – transparency on use so New Zealanders are aware of AI usage and agree with it. Establishment of the Centre for Data Ethics and Innovation.

The AI Oversight Group:

- **Noted** that outtakes presented to the group from the recently held AI Forum.

Not in scope

Not in scope

Excerpt from meeting minutes: Artificial Intelligence Working Group

Artificial Intelligence (AI) Working Group

Inland Revenue
Te Tari Taake

Meeting Minutes for 13 November 2023, 1:30 – 2:30 pm

Members	Graham Poppelwell (Chair), Aidan Roberts, Alex Steel, Brandon Sloan, Brent Jarnell, Brijesh John, Chris Hourigan, Dave Rowley, David Robinson, Dawn Swan, Erin Dyson, Jacinda Hughes, Jane Smathers, Jason Ratima, Jess Wawatai, Jo McGregor, Jos Crasborn, Kevin McArtney, Makayla Stewart, Prajakta Panse, Ryan Hamilton, Souradeep Gupta, Teressa Dillon, Tim Crook, Tina McCaffrey, Vanessa van der Schraft, William MacKay, Yolanda Wilke
Attendees	Tanya Williams, Sally Krogh, Kate Yong Deloitte: Roger Lee, Pieta Brown, Amanda Williamson
Apologies	Underlined above

Not in scope

Not in scope]

Not in scope

7. External relationship activities. INFORM

Presenters: Graham Poppelwell, Domain Lead, Information Governance & Sharing; Makayla Stewart, Change Analyst, Change, Design & Enablement; Brijesh John, Domain Lead - Technology Architecture.

Interim Centre for Data Ethics and Innovation (ICDEI)

The Interim Centre for Data Ethics and Innovation was established in September 2023 at the Artificial Summit. An invitation came through the Algorithm Charter committee seeking people to join workshops around ethics. The Domain Lead (Information Governance & Sharing) shared an open invite to interested members of the Working Group.

Algorithm Charter Cross-Agency Community

IR is hosting the next meeting on 6 December 2023 in Asteron Centre. The last meeting was hosted by NZ Police with a good presentation around how they work from a governance perspective. ICDEI will be invited to present as per the draft agenda that will be shared with the group.

AI Community of Interest and Auckland Council

The AI community of interest group hosted speakers from Auckland Council (AC) to talk about the intelligent solutions they are using to solve business problems. They discussed the way AC worked with businesses to identify the issues and challenges they were facing and then working with them for a solution. AC discussed scenarios of AI reducing working hours for better work life balance and retention of key people. The session recording is available for viewing by contacting the Change Analyst (Change, Design & Enablement).

Cross-agency Survey from the Deputy Chief Digital Officer

The presenters and the Technical Lead (Information Governance) have completed this. The survey covered areas like the kind of governance in place, its usefulness and the appetite for AI. There are opportunities to influence the programme and establish common capabilities. There will be an upcoming meeting with more updates from other public sector agencies starting their AI journey.

Not in scope

Not in scope

Excerpt from Key Messages: Data and Information Governance Authority

Inland Revenue
Te Tari Taake

Executive Level Governance

Date: 17 November 2023

To: Executive Leadership Team and Senior Leaders

From: Information Governance

Key messages from the 8 November 2023 Data and Information Governance Authority (DIGA) meeting

Not in scope

Not in scope

Artificial Information (AI) Governance

IR will be hosting the upcoming AI Charter cross-agency meeting in-house at Asteron on 6 December. Guests will also include representatives from the recently established Interim Centre for Data Ethics and Innovation who are within the office of the Government Chief Data Steward.

The Artificial Intelligence Te Mātāwai site is going through final review before go-live later in November. IR-wide communications are planned for its launch via Featured News.

Not in scope

Excerpt from meeting minutes: Artificial Intelligence Oversight Group

Artificial Intelligence Oversight Group

Inland Revenue
Te Tari Taake

Meeting Minutes for 20 November 2023, 1:30 – 2:20 pm

Core members	Brijesh John (Chair), Anil Srinivasa, Cate Robertson, Craig Thomas, Daniel Blank, Jay Harris, Jesse Thwaites, Malcolm Breadmore, Phil Whittington, Ron Grindle, Scott McCallum, Tina MacLean
Standing members	Conrad Bace, Graham Poppelwell, Vanessa Johnson, Virginia Flaus
Attendees	Teressa Dillon, Aiden Roberts, Vanessa van der Schraft, Lucy Cording, Tanya Williams, Kate Yong Deloitte: Roger Lee, Pieta Brown, Amanda Williamson
Apologies	Underlined above

Not in scope

Not in scope

7. General Business. INFORM

Presenter: Graham Poppelwell, Domain Lead, Information Governance & Sharing

Interim Centre for Data Ethics and Innovation
There is an open invite to members of the Oversight Group for a workshop.

Community of Practice for the AI Charter
This brings together government agencies to talk on an informal basis. IR is hosting the next meeting on 6 December 2023 in Asteron Centre.

Data and Information Quality Model

This document is a Model. Models support or expand upon a related Standard by outlining an explicit set of technical requirements. This Model should be read in conjunction with the Data and Information Policy which outlines Inland Revenue's approach to the stewardship of data, information and knowledge, and the Data and Information Quality Standard which establishes high-level requirements for the quality of our data, information, and knowledge. The Data and Information Quality Standard is being drafted at the time this Model was approved.

Contents

Why we have this Model

The Model provides guidance for anyone at Inland Revenue who is interested in understanding, communicating, measuring, and improving the quality of data, information, and knowledge within Inland Revenue. The Model includes definitions of the most relevant quality concepts and high-level guidance about how to apply these concepts in practice.

The Model is positioned as a Tier 3 instrument in the Data and Information Governance Instruments Framework – see the Data and Information Governance Standard for further information about the Framework.

Tier 1: Explains why we are going to do what we say we are going to do.

Tier 2: Defines what we are going to do.

Tier 3: Tells the business how to do something.

Tier 4: Are the tools we work with to ensure the things in Tier 1-3 documents can be done.

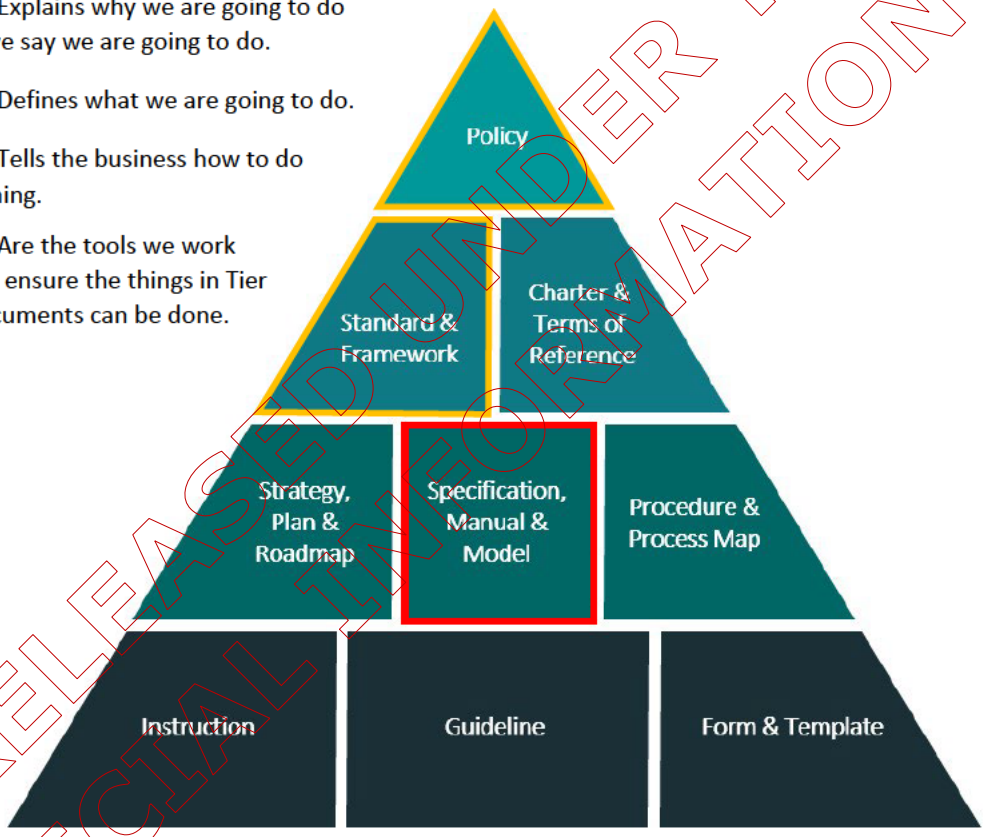

Data and Information Governance Instruments Framework

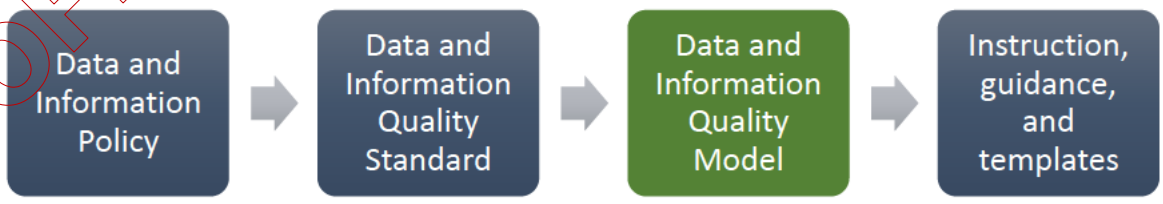

Additional guidance

The Model will be progressively supplemented by additional guidelines, instructions, templates, and other material developed to support its practical application.

What we mean by data and information quality

Data and Information Quality is defined **as how fit-for-purpose data, information, and knowledge resources and products are.**

This definition requires determination of **fitness** and **purpose**.

- Purpose is based on business context which determines quality requirements.
- Fitness requires determination of measurable quality characteristics and the comparison of these measurements to required quality characteristics.

There may be multiple purposes for use of the same data and information.

For example: Service delivery activities might require some data to be current at the time it is used to engage with customers, whereas for some analytical activities the same data could be considered usable if it is far less current.

Data and Information Quality is not limited to only characteristics of data and information itself but also the origin, supply chain, and current access to and use of the data and information.

For example: s 18(c)(i) In this case the intrinsic data quality characteristics might meet quality requirements but the data still considered not immediately usable.

Understanding the purpose for which data, information, and knowledge was collected or generated is essential to determining its fitness for use.

For example: A customer address might have been collected to be used for correspondence. The address data might be accurate, current, and complete. However, this data might not be usable for an alternate purpose, perhaps to determine if a customer is resident in New Zealand or not.

Why this Model is important

Not in scope

This Model presents a consistent and comprehensive means for describing and defining data, information, and knowledge quality requirements and characteristics of data, information, and knowledge collections. This consistency and comprehensiveness results in greater clarity about what quality of data, information, and knowledge is required for specific purposes.

The Model also presents the collection of elements necessary to support the management of the quality of data, information, and knowledge. This includes what good measurement of quality includes, and what different types of quality treatments and controls provide.

Collectively, the elements of the Model, when applied as intended, result in reduced risk and cost, and increased business confidence and likelihood of achieving business objectives.

Not in scope

What this Model applies to

The Model applies to all data, information, and knowledge stewarded by Inland Revenue.

What the requirements of this Model are

The mandatory requirements of the Model are:

1. The adoption of the Model across Inland Revenue for qualifying and specifying data, information, and knowledge quality. Adoption should align with opportunity to do so. It is not expected that adoption of the Model be immediate nor without consideration of the impacts and resourcing associated with its adoption.

Not in scope

Not in scope

The Data and Information Quality Model

The Model outlines the minimum elements that need to come together to describe:

- The quality of data, information, and knowledge, including its measurement.
- The quality requirements for data, information, and knowledge.
- Considerations for treating and controlling data, information, and knowledge quality.

The Model can be conceptually visualised as:

Data and Information Quality Model

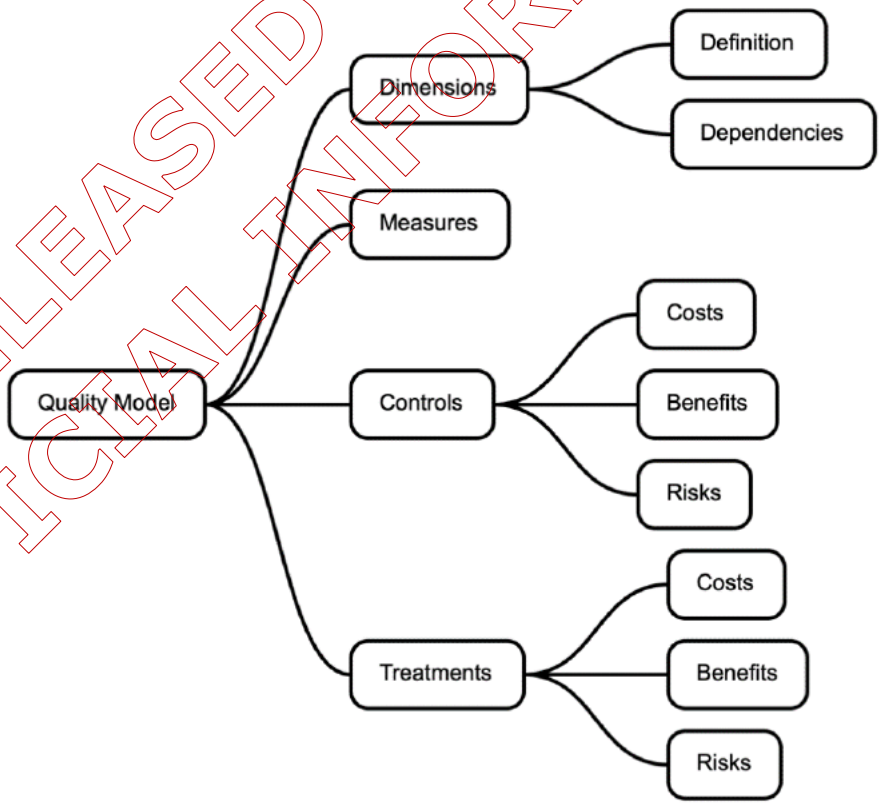

The elements of the Model are defined in the following table.

Element	Definition
Dimensions	Dimensions are representation of concepts that underpin quality. Some of these concepts are specific or intrinsic to the data, information, and knowledge, and some concepts are about how the environment that data, information, and knowledge exists in can affect the quality of data information, and knowledge. Examples of intrinsic quality concepts are accuracy, completeness, consistency, and currency. Examples of non-intrinsic quality concepts are availability and credibility.
Measures	Measures are defined ways to assess the extent of an aspect of quality. Measures can be quantitative or qualitative in their basis. They can be coarse or precise in their degree of assessment. For example: determining the currency of data could be measured by comparing the last time the data was updated compared to now. Thresholds are defined measure values that represent something material in practice. For example: data might be considered current enough for a purpose if it was last updated within 30 days of its use. Measures must be clear and unambiguous in their definition so implementations of measures can be trusted across multiple implementations. In practice, it is necessary to have robust governance and management around quality measures.
Controls	Controls are mechanisms that reduce quality issues arising. Controls can be applied prior to a point in a data, information, or knowledge supply chain where the quality matters most. These are preventative controls. Controls can also determine that a quality problem exists after the fact and inform roles or processes that action is required to avoid the quality problem increasing. In this regard, controls manage specific types of risk. Controls can also introduce risks and this needs to be considered when deciding to implement controls. For example: a field on an input form could examine the data input by a person and if the data doesn't meet a specific rule the data could be rejected before being accepted by the information system behind the form. On first glance this appears an excellent change to apply to reduce receiving

Element	Definition
	poor quality data. However, it is possible a person does not have data to input that conforms exactly to the fields validation rule yet it is good enough for a person to inspect and make sense of once it is in the information system. Rejecting data using validations rules can annoy people trying to provide the best information they have. As a result a person might give up and not provide any information, essentially an "own goal" on the part of the form designer. Controls must have control owners. These are roles responsible for ensuring controls are reliably in place and are being routinely assured for their effectiveness. Implementing and maintaining controls comes at a cost and with benefits. Ideally, the cost and benefit of a control is well understood and this contributes to the degree of investment in it.
Treatments	Treatments are activities carried out or mechanisms implemented that intentionally affect the quality of data, information, and knowledge. Some treatments are also controls. Like controls, treatments come at a cost and with benefits, and these need to be understood to determine the appropriate level of ongoing investment in a treatment. And like controls, treatments can also introduce risks and these risks need to be considered when deciding on and applying treatments. For example: It may have been determined that a collection of data, say customer correspondence addresses, can have its quality improved if it is combined with another collection of data, perhaps to add missing address details. It is possible the address details that are to be added, perhaps customer physical addresses, do not represent the same thing. The result of combining these two collections of data could be a trusted collection of misinformation. There is another risk: that data is poorly combined and data not intended to be updated becomes updated and reduces the quality. In both cases, robust due diligence and recovery options can mitigate the risks to a degree.

This Model does not dwell on measures, controls, and treatments. These should be defined in subsequent documents such as Specifications.

However, the quality dimensions are covered in detail in this Model as they are generic and underpin measures, controls, and treatments.

Quality dimensions

Quality dimensions represent quality concepts. The Model organises these dimensions into two collections that characterise the usability and the reliability of data, information, and knowledge and this is illustrated below. Also included in the illustration are two groupings of quality dimensions: Trustworthiness groups Credibility and Authenticity, and Timeliness groups Currency and Latency. Trustworthiness and Timeliness are not themselves quality dimensions and are included to improve understandability of the Model.

The quality dimension dependencies are included in the Appendix.

Data and Information Quality Model – Quality Dimensions

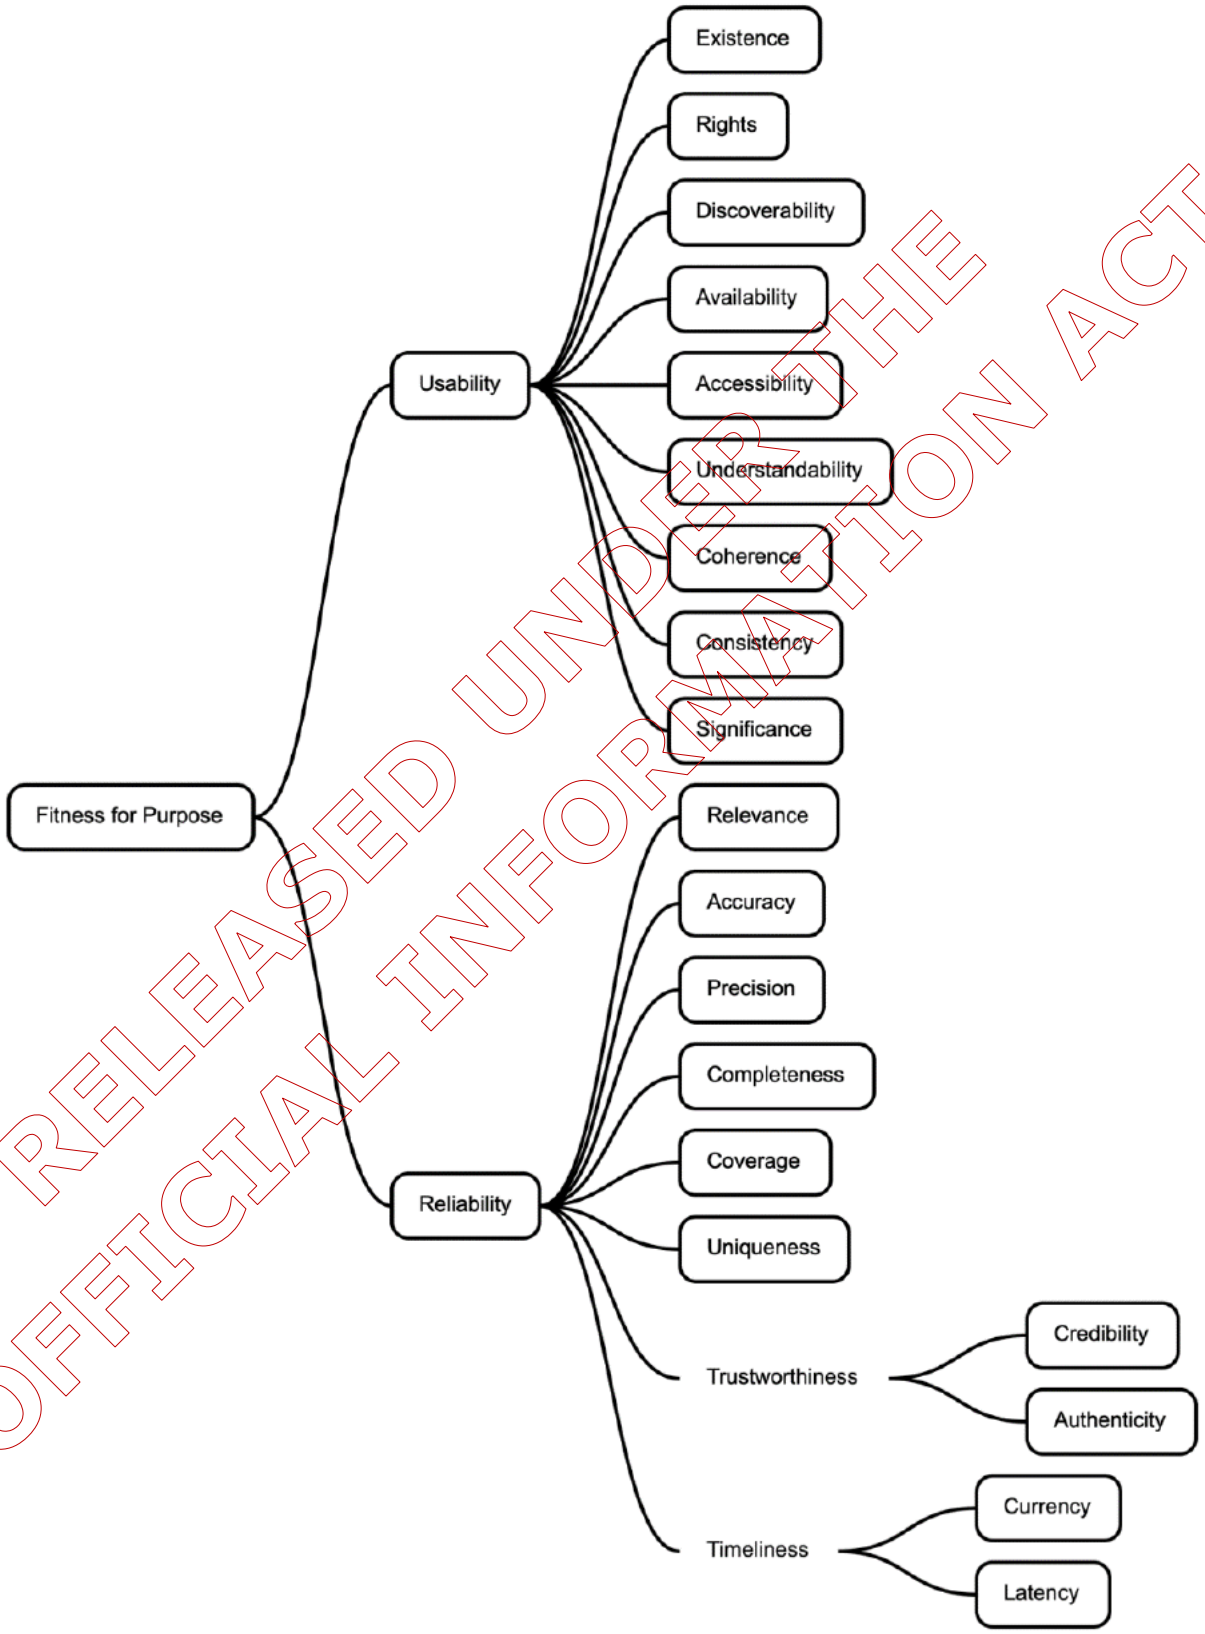

High-level quality dimensions

Dimension	Definition
Fitness for purpose	**The degree to which data and information meets specific quality requirements.** Fitness for purpose is context specific, meaning the purpose to which data or information is to be put is fundamental to determining if it is of suitable quality. Furthermore, quality requirements must be determined and specified in a way that allows for the quality characteristics of data and information to be measured and compared to the quality requirements. A consequence of this is the necessity to be clear about the purpose to which data and information is to be put, and to specify quality requirements that are purpose specific and that can be measured on an ongoing basis. Fitness for purpose often involves measuring quality across multiple quality dimensions that span Usability and Reliability.
Usability	**The degree to which a person or system can process data and information efficiently and effectively for a specific purpose and context.** Usability is a consolidation of several quality concepts that can be defined as quality dimensions. For this Model those dimensions are: Existence, Rights, Discoverability, Availability, Accessibility, Understandability, Coherence, Consistency, and Significance.
Reliability	**The degree to which a person or system can trust the integrity of data and information for a specific purpose and context.** Reliability is a consolidation of several quality concepts that can be defined as quality dimensions. For this Model those dimensions are: Relevance, Accuracy, Precision, Completeness, Coverage, Uniqueness, Credibility, Authenticity, Currency, and Latency.

Usability dimensions

Dimension	Definition
Existence	**Data and information has either been generated or collected.** This is the most fundamental quality dimension, upon which all others depend. An existence issue occurs when Inland Revenue does not possess nor have access to content addressing the subject matter. An existence issue does not arise where one part of Inland Revenue holds data, information, or knowledge, but it cannot be accessed by others who may get value from it. In this case the quality issue may be a matter of rights, discoverability, and/or availability, and should be treated accordingly.
Rights	**A person, organisation, or system has the rights to access the intended data and information.** This dimension describes whether a person, organisation, or system has the right to access and make use of the data, information, or knowledge for a

Dimension	Definition
	particular purpose. Rights encompass legal authority, privacy, and ethical considerations. This dimension should not be mistaken for Availability, describing the ability to access data, information, and knowledge. A person may have physical access to data, information, or knowledge, but not the right to use that information, or vice versa.
Discoverability	**A person or system can identify that data and information exists.** A discoverability issue means people or systems cannot find specific data, information, or knowledge when they need it. This is different from Availability, being unable to access the data and information. Discoverability and availability issues often co-exist, for example if data or information is stored in private repository the person or system cannot access.
Availability	**A person or system can access the intended data and information when needed.** Availability concerns a person's or system's ability to access data and information when it is needed, independent of their rights to access and use the data and information. If a person or system cannot access the data and information due to it being encrypted or protected by digital rights management technology then it is essentially not available.
Accessibility	**How processible data and information is via a person's senses or a system's sensors.** Accessibility recognises that data, information, and knowledge must be able to be processed to be useful. For example: a date could be incorrectly constructed and consequently not be able to be interpreted by software or people. Another example is information presented for direct consumption by people yet the information is presented too small to be read by an audience, or as a combination of colours that some people cannot visually differentiate. Accessibility should not be confused with Availability. In more general use of the term "accessibility" it can be a synonym for availability. Availability as used in this Model is about data, information, or knowledge being "within reach" when it is wanted.
Understandability	**Data and information can be comprehended without ambiguity.** Understandability describes a person's or system's ability to consume and make sense of data and information so it can be they can be used appropriately. This is dependent on both the information supplied with and about data and information, and on a consumer's knowledge and ability to interpret this accompanying information.

Dimension	Definition
	The accompanying information is often referred to as metadata and comes in a few conventionally classified types: • **Descriptive metadata**: Descriptive information about data or information. It provides the user with enough information to understand what the data or information is about. It can be used for discovery and identification of data and information. • **Structural metadata**: Information about the container data or information is packaged in. Examples of containers are files for documents, and database tables and columns for data. • **Administrative metadata**: Information to help manage the use of data and information. It permits a person, organisation, or system to understand the permitted and disallowed uses of data and information, and any obligations regarding stewardship of the data and information. • **Reference metadata**: Information about the contents and quality of data and information. It permits the user to understand whether the information resource is fit for purpose. It includes the value and risks of the information resource. • **Process metadata**: Information that details the processes that have handled data and information resource during its lifetime. It permits the user to understand where the resource has come from, who has used it, and how it has changed, often referred to as lineage or provenance. It includes the creation, transformation, migration, and disposal details.
Coherence	**Data and information is reconcilable with what it represents.** Coherence reflects the degree to which data and information are logically connected and mutually consistent. This means that people and systems can know that the same concept or item is represented by specific references. For example, the same term should not be used to refer to different things without this being clearly communicated, and that different terms should not be used to refer to the same thing without this being clearly communicated. Coherence is particularly important when wanting to link multiple datasets. If the columns/fields being linked don't represent the same thing then the columns/fields cannot be used this way and potentially reducing opportunity for significant value. In practice, references and terms change over time, location, and business context and this cannot be avoided. However, it can be minimised through explicit management of data and information design and accompanying metadata.
Consistency	**The degree to which a collection of data or information has no internal contradictions.**

Dimension	Definition
	Consistency applies to the multiple instances of a data, information, or knowledge item. Specifically, that the multiple instances of an item are aligned with each other. For example, a customer's date of birth might be stored in multiple locations and should be the same value in all those locations, or consistent across all those locations.
Significance	**The degree to which a collection of data or information has the appropriate volume for a context.** Significance applies to a collection of items. Significance is about whether there is a sufficient number of items to support a purpose. For example: There only needs to be one set of instructions about an activity, and there being more than one set of instructions intended to inform the same activity could be detrimental. In the case of data, the number of customers in a sample of customers with particular characteristics might need to be at least 20,000 to reliably infer anything about a larger population of customers.

Reliability dimensions

Dimension	Definition
Relevance	**Data and information is valuable for a particular purpose and context.** Relevance is about the reader's ability to determine the specific context content should be applied. This context could, for example, be a specific task, process, enquiry, decision, or project. An example of a potential relevance issue is an article titled "Cease a repaid loan account". From the title alone, it is unclear whether the article relates to a Student Loan or a Small Business Cashflow Loan. The reader would need additional information to make a determination, such as metadata or the collection in which the content is published.
Accuracy	**How closely data and information represents something (includes bias).** Accuracy is about close a representation of something is to what it represents. The closeness of representation might be valid for a particular purpose and at a particular point in time; in other words, accuracy is not absolute.
Precision	**How specific the data and information are in their representation.** Precision is about the level of specificity or detail data and information includes. For example: A person born on 18 September 1968 could be described on 27 April 2023 to be 54 years old, or between 50 and 55 years of age. Both are accurate but the former has a higher precision having a granularity of one year as opposed to five years.

Dimension	Definition
	Importantly, Precision does not in any way speak to Accuracy. Data could be highly precise yet wrong.
Completeness	**Data and information content covers all the relevant aspects.** Completeness can be well explained with an example. Consider a customer physical address which has the address elements of dwelling name, dwelling number, street number, street name, suburb, town/city, region/state, and country. It might be defined for a particular purpose that some specific elements must be populated to be considered a complete address. Importantly, an address being complete does not mean the address elements are accurate or well-formed.
Coverage	**The degree to which a collection describes the expected scenarios within a context.** Coverage applies to a collection of data, information, or knowledge. It is similar to Completeness, except that it applies at a broader scale, spanning multiple records or collections. Coverage is about whether a collection of data, information, or knowledge is sufficiently inclusive for a particular purpose at a point in time.
Uniqueness	**The degree to which a set contains only one instance of the same object.** Uniqueness applies to a collection of data, information, or knowledge. It refers to the incidence of something within a collection. For example: a customer might be represented with a system once, or more than once. If a customer should be represented only once but is represented twice then a Uniqueness problem exists. Such a problem could be exacerbated if some transactions were associated with one of the customer records and some transactions were associated with the other customer record. This situation would be frustrating for a customer, and potentially reputationally damaging for the organisation responsible for the integrity of the customer records.
Credibility	**The degree to which data and information is considered trustworthy.** Credibility is an aspect of timeliness. It is the subjective counterpart of Authenticity. Credibility is about how true the data, information, or knowledge is considered to be by the party accessing it. Credibility can be influenced by the consumer's own knowledge. If it fits their existing understanding and/or biases, it is more likely to be accepted as credible than if it clashes with their current knowledge.
Authenticity	**Data and information are regarded to be a true representation.** Authenticity is an aspect of trustworthiness. It is the objective counterpart of credibility. Authenticity is more concerned with the source of the data, information, or knowledge than with whether it "rings true". If the presented data,

Dimension	Definition
	information, or knowledge comes from a trusted and reliable source, then it should be regarded as authentic.
Currency	**Data and information are up to date at a specific point in time and use.** Currency is an aspect of timeliness and represents the version of data, information, and knowledge. For some purposes the most recent version might be required. For some other purposes a history of versions might be required. For example: knowing the version of legislation that was in effect at a point in time is important and a Currency matter.
Latency	**Time between data or information existing and being available to a person or system.** Latency is about the delay between data, information, or knowledge about something existing and it being available to be used. Latency is an aspect of timeliness but should not be confused with Currency which is another aspect of timeliness. For example: when a customer goes into a bank branch and completes a deposit form and hands this to a teller with the funds to be deposited, the form is date and time stamped. However, the deposit details might not be immediately entered into the banking system, and instead is input as part of a larger batch at the end of the day. The customer will not see the deposit of funds in their account until their account details have been updated. It is important to understand if there is material latency in the data, information, or knowledge being accessed as the latency may affect its usefulness.

Practical application of the Model

This Model informs the definition of data, information, and knowledge quality measures, treatments, and controls. These measures, treatments, and controls can be brought together in various Specification documents, or similar formal definitions.

To realise the potential value of these definitions they need to be operationalised in the forms of quality measurement components and configuration, executed treatments, and implemented and assured controls.

> For example: When Inland Revenue supplies data to another organisation, that data should be checked by automated processes once it is produced and before it leaves Inland Revenue. The quality checks that are executed would be determined by both Inland Revenue who best understand the data, and quality checks agreed to by Inland Revenue and the consuming organisation. These agreed quality checks would be itemised and defined in an operational agreement of some kind and held by both organisations. These quality checks would be included because it is understood by all involved parties that data falling outside of the quality requirements these checks

assess might result in incorrect decisions and actions being undertaken by the consuming party. This in turn could result in detrimental impacts on people and organisations.

The results produced by running these quality checks should accompany the data in a machine-processible quality report.

Any problems found during the automated checks should be investigated and resolved, then the data reproduced and supplied to the consuming organisation.

Upon arrival at the consuming organisation, the data should have automated quality checks executed across it and the results from the sent and received quality checks compared. Any quality exceptions should be investigated and might result in Inland Revenue being contacted as part of an investigation.

To achieve all the above, effectively and efficiently, requires all involved parties to understand the value of implementing quality checks - this speaks to the purpose for which the data is being used, its value, and risks arising from the data being not fit for the purpose. The above also requires all involved parties to agree on a set of quality measures that implement the quality concepts outlined in this Model as Quality Dimensions. The implemented quality checks are quality controls, and these have owners, costs, and risks associated with them. When a quality problem is identified, treatments will be decided and executed – these treatments have owners, costs, and risks.

Importantly, this Model provides the consistent and comprehensive definition that shapes data, information, and knowledge quality measures, treatments, and controls.

Responsibilities

Responsibilities for the management of the Model

Role	Responsibilities
Business Owner	Accountable for ensuring the Model is fit for purpose, is broadly known to exist, and is actively being used when and where it should be.
Document Owner	Responsible for ensuring the Model is fit for purpose, is broadly known to exist, is actively being used when and where it should be, maintaining the Model, and being the point of contact regarding risks and issues about the Model.

Responsibilities for the adoption of the Model

Role	Responsibilities
All Inland Revenue employees, contractors, and consultants	Responsible for ensuring they use the Model when it is appropriate to do so and according to their responsibilities, and reporting risk and issues about the Model to the Document Owner.

Appendix

Terms and definitions

Term	Definition	Source

The origins of the Model

The Model is largely a consolidation of elements that are globally well accepted as foundational to good data, information, and knowledge quality management. Those elements are:

- Quality dimensions
- Quality measures
- Quality treatments
- Quality controls

The quality dimensions are a consolidation of quality concepts from several sources, including:

- Inland Revenue's Data Science Team
- New Zealand Government Web Accessibility Standard 1.1 (incorporates W3C Web Content Accessibility Guidelines 2.1)
- Quality Framework for OECD Statistical Activities
- Conformed Dimensions of Data Quality
- ISO Standard 8000-8:2015 Data Quality Part 8: Information and data quality: Concepts and measuring
- ISO Standard 25012:2018 Software engineering - Software product Quality Requirements and Evaluation (SQuaRE) - Data quality model

The Rights quality dimension is a novel addition to the Model and recognises the real-world usability consideration regarding data, information, and knowledge. See the Appendix section **How the Quality Dimensions depend on each other** to understand the precedence Rights enacts across almost all quality dimensions. The Rights quality dimension definition is reproduced below:

> **A person, organisation, or system has the rights to access the intended data and information.**
>
> This dimension describes whether a person, organisation, or system has the right to access and make use of the data, information, or knowledge for a particular purpose. Rights encompass legal authority, privacy, and ethical considerations.
>
> This dimension should not be mistaken for Availability, describing the ability to access data, information, and knowledge. A person may have physical access to data, information, or knowledge, but not the right to use that information, or vice versa.

The draft of this Model was consulted on with over 130 Inland Revenue people during 2021, covering it application to data, information, and knowledge. The Data and Information Quality Working Group endorsed the interim use of the Model for optional adoption and evolution.

How the Quality Dimensions depend on each other

Some quality concepts are dependent on other quality concepts, and this can be reflected in the quality dimensions. For example, data or information must exist before it can be discoverable, so discoverability depends on existence.

Dimension	Existence	Rights	Discoverability	Availability	Accessibility	Understandability	Coherence	Relevance	Consistency	Significance	Accuracy	Precision	Completeness	Coverage	Uniqueness	Credibility	Authenticity	Currency	Latency
Existence																			
Rights	✓																		
Discoverability	✓	✓																	
Availability	✓	✓	✓																
Accessibility	✓	✓	✓	✓															
Understandability	✓	✓	✓	✓	✓														
Coherence	✓	✓	✓	✓	✓	✓													
Relevance	✓	✓	✓	✓	✓	✓													
Consistency	✓	✓	✓	✓	✓														
Significance	✓	✓	✓	✓	✓														
Accuracy	✓	✓	✓	✓	✓														
Precision	✓	✓	✓	✓	✓														
Completeness	✓	✓	✓	✓	✓														
Coverage	✓	✓	✓	✓	✓														
Uniqueness	✓	✓	✓	✓	✓														
Credibility	✓	✓	✓	✓	✓	✓		✓	✓										
Authenticity	✓	✓	✓	✓	✓														
Currency	✓	✓	✓	✓	✓														
Latency	✓	✓	✓	✓	✓														

When the Model is effective, reviewed, and what it replaces

Not in scope

Authorisation

This Model is approved by the Data and Information Quality Working Group, and authorised by the Intel Leader Enterprise Information and Knowledge

s 9(2)(a)

Date: 9th November 2023 Signature:

Document information

Current version	
Approval date	Data and Information Quality Working Group meeting 11 May 2023
Authorisation date	11 May 2023
Effective date	11 May 2023
Last updated	11 May 2023
Review frequency	2 years
Review date	11 May 2025
Approval Party	Data and Information Quality Working Group
Business Owner	Leader Enterprise Information and Knowledge
Document Owner	Domain Lead Governance and Information Sharing, Enterprise Information and Knowledge
Content developer	Advisor – Data and Information Governance, Enterprise Information and Knowledge
Audience	All Inland Revenue employees, contractors, and consultants.

Document amendment history

Version	Date	Sections amended	Summary of amendment
0.5	7 May 2023	All	Final as approved.

Inland Revenue
Te Tari Taake

Privacy & Ethics Impact Assessment

Name of Initiative

Insert date

Author

Version

About this Document

The purpose of this document is to fully consider all privacy risks this proposal raises and how they will be mitigated. If you have completed the Privacy Threshold Assessment, information from that document can be included. The Privacy & Ethics Impact Assessment (PEIA) will answer the following questions:

- Does this proposal fully comply with the Privacy Act principles?
- Have all risks been identified, and mitigations proposed?
- Are we satisfied personal information is appropriately managed that the proposal can proceed?

Document version

Version	Date	Section	Page	Description of change	By

Document contributors

The following Inland Revenue business groups have been consulted on the project.

Name	Role	Business Group

Document sign-off

Name	Role & Business Group	Sign off date

Contents

Introduction

1.1 Project summary

Describe the project, background to the initiative and what it intends to achieve. For instance:

- *explain context and reasons why the project is being undertaken*
- *any assumptions underlying the assessment*
- *whether the project is a one-off initiative or part of on-going business development*
- *benefits to business*
- *justification for collecting or using personal information in the project*

Note: *this summary can be copied from the Brief Privacy Analysis document if it has already been completed.*

1.2 Scope of the assessment

Describe what the PIA covers and what it doesn't cover. For example:

- What parts of IR, which project or system is included?
- What are the information-management processes that the PIA will consider? (ie. Collection, storage, use, access, retention or disposal of personal information)
- What assumptions have been made?

1.3 Glossary

Term	Meaning

2. Personal Information

2.1 Information to be used

Identify and describe the personal information to be used in the project, ie full name, phone number, financial details, family relationships.

Type of personal information	Source of information	Purpose of using the information
e.g. full name		

2.2 Authority and access

Explain if this information is already held by IR, if new personal information will be collected, under what authority, and who will have access to information.

2.3 Information Flows

Document the flow of personal information to clearly illustrate how data is collected, how it circulates internally and how it is disseminated beyond IR. If relevant, show **current** and **future** state so the differences are visible at a glance.

3. Privacy Assessment

3.1 Privacy principles and responses

The following table lists the relevant excerpts of the Information Privacy Principles1 (IPP) and responses to each for this particular project.

Note: adapt your approach to the issues at stake. In some cases, the emphasis will be on only one or two issues. If an IPP is not relevant to the project, say so but you should at least consider each principle.

IPP 1 Purpose of collection

Only collect personal information if you really need it:

- *collection must be for a lawful purpose connected with a function or activity of IR; and*
- *collection must be necessary for that purpose*

If the lawful purpose for which personal information about an individual is collected does not require the collection of an individual's identifying information, then don't collect identifying information.

Summary of how IR will comply:
Explain what personal information is being collected and why IR needs it. How will it enable IR to operate? Are you only collecting what you actually need?

If you have identified any risk to complying with principle 1, complete the Table of risks and mitigations.

[1] Office of the Privacy Commission | Information Privacy Principles - http://privacy.org.nz/news-and-publications/guidance-notes//information-privacy-principles

IPP 2 Source of personal information

Get it directly from the person concerned wherever possible. You can collect from another source if:

- *the information is publicly available*
- *the individual concerned has authorised collection from someone else*
- *non-compliance would not prejudice the interests of the individual concerned*
- *it's necessary to collect the information to uphold or enforce the law or protect public revenue*
- *compliance would prejudice the purposes of the collection*
- *compliance is not reasonably practicable*

Summary of how IR will comply:
Where is IR getting the personal information from? If not the person concerned, explain why not and what bullet point exception above is relevant. If the information is already held by IR, explain whether the person initially provided it themselves. Is the project using the information for a directly related purpose to why it was obtained?

If you have identified any risk to complying with principle 2, complete the Table of risks and mitigations.

IPP 3 Tell people what you're going to do with their information

At the point of collection, you must tell people:

- *what information you are collecting*
- *what you're going to do with the information*
- *intended recipients of the information*
- *whether it's voluntary or required by law (and any consequences if they don't provide it)*
- *rights of access to, and correction of, their information.*

You don't have to do this upfront if you believe:

- *non-compliance would not prejudice the interests of the individual*
- *it's necessary to collect information to uphold or enforce the law or protect public revenue*
- *compliance would prejudice the purposes of collection*
- *compliance is not reasonably practicable*

Summary of how IR will comply:

How will IR tell people everything in the bullet point list? Is there a privacy statement or policy that people will be directed to? Is this a one-off exercise or are there on-going implications? If we are not going to tell people what we're doing with their information, which of the exceptions applies?

If you have identified any risk to complying with principle 3, complete the Table of risks and mitigations.

IPP 4 Manner of collection of personal information

Only collect personal information by means that are lawful, fair and does not intrude to an unreasonable extent upon the personal affairs of the individual concerned

Summary of how IR will comply:
How is the personal information to be collected? What method is being used? If a recording device will be used explain why and whether people will be told.

If you have identified any risk to complying with principle 4, complete the Table of risks and mitigations.

IPP 5 Storage and security of personal information

IR must ensure that:

- *there are reasonable security safeguards to protect information against loss, unauthorised access, misuse or disclosure; and*
- *if it is necessary to disclose information outside IR, everything reasonable must be done to prevent unauthorised use or disclosure of the information*

Summary of how IR will comply:
What steps are taken to keep the information safe? Safeguards can be physical or technical. Does the system design enhance privacy and security? Are the security measures commensurate with the sensitivity of the information?
Consider where the information will be stored and controls defining who may access it, is there an audit trail? Will there be external access to a system, has it received approval from CISO? Will staff receive training? Are there mechanisms in place to identify data/security breaches?
Are there documented security procedures for the collection, transmission, storage and disposal of the information?
If third parties are involved, they must sign the IR820 Certificate of Confidentiality.

If you have identified any risk to complying with principle 5, complete the Table of risks and mitigations.

IPP 6 Access to personal information

Where a person's information is held in a way that it can be readily retrieved, they are entitled to:

- *obtain confirmation of whether the information is held; and*
- *have access to their information (subject to the Tax Administration Act or withholding grounds in sections 49-53 of the Privacy Act)*

Summary of how IR will comply:

If an individual asked for access to this information, would it be readily retrievable? Would there be any reason to withhold it from a requester? For instance, disclosure may be refused in some circumstances if doing so would prejudice an investigation, or would breach someone else's privacy.

IR has processes to enable individuals to request access to their personal information. This can be done through myIR and there is also information on the main IR website about making Official Information Act requests that refers to making Privacy Act requests.

This initiative does not change or impact on that process or an individuals' ability to request access to their information. The personal information used will be able to be retrieved and provided to on request [amend this paragraph if not correct].

If you have identified any risk to complying with principle 6, complete the Table of risks and mitigations.

IPP 7 Correction of personal information

Everyone is entitled to:

- *ask that their personal information be corrected; and*
- *if it is not corrected, have a statement attached to the original information saying that correction was sought but not made*

Summary of how IR will comply:

If IR is made aware that incorrect or corrupt information has been obtained, can it be corrected? Is there a process for customers to dispute information used? Are there limitations to IRs ability to correct, for instance character limits in data fields or unable to flag incorrect information?

If you have identified any risk to complying with principle 7, complete the Table of risks and mitigations.

IPP 8 Accuracy of personal information to be checked before use

Before using personal information, reasonable steps should be taken to ensure it is accurate, complete, relevant, up to date, and not misleading.

Summary of how IR will comply:

Explain what steps are taken to ensure the information is accurate before it is used. Has the information been supplied directly by the individual or been checked with the individual? Is the process automated or is human judgment applied? How damaging will it be if information is wrong or misleading? (The more damaging it will be, more extensive steps should be taken to check accuracy).

If you have identified any risk to complying with principle 8, complete the Table of risks and mitigations.

IPP 9 Don't keep personal information for longer than necessary

Personal information must not be kept for longer than needed for the purpose for which we collected it.

Summary of how IR will comply:

How long will IR hold the information? Ask the IKM team if there a requirement under the Public Records Act to keep this information for a specific time. If not, what would be a reasonable time to keep it and how will it be disposed of? If information is shared with a third party, how long will they hold the information for?

If you have identified any risk to complying with principle 9, complete the Table of risks and mitigations.

IPP 10 Limits on use of personal information

Only use personal information for the purpose you got it for, unless:

- *it's used for a directly related purpose*
- *the source of the information is publicly available*
- *it's necessary to use it to detect or investigate an offence or assist court or tribunal proceedings*
- *it's necessary to use it to protect public revenue*

Summary of how IR will comply:

Outline all intended uses of the information and, in particular, if information may be used for another purpose than it was collected. Be clear about the purpose for having the information (review your response to IPP3 and why you said you were collecting the information) – is this what customers will expect or been told?

If you're using information for a different purpose from the one for which it was obtained, how do you justify this?

If you have identified any risk to complying with principle 10, complete the Table of risks and mitigations.

IPP 11 Limits on disclosure of personal information

Only disclose personal information if you've got a good reason such as:

- *the individual authorised you to disclose*
- *disclosure is one of the purposes for collecting the information (and people were told at the point of collection – this links to IPP3)*
- *it's a directly related purpose to why the information was obtained*
- *it's necessary to disclose it to detect or investigate an offence or assist court or tribunal proceedings*
- *it's necessary to disclose it to protect public revenue*

Summary of how IR will comply:

Outline known circumstances when the personal information may be disclosed, who may receive it and for what purpose. Note: this does not include circumstances that are not foreseen at the time of collection. Does the Tax Administration Act permit the information to be disclosed?

If you have identified any risk to complying with principle 11, complete the Table of risks and mitigations.

IPP 12 Disclosure of personal information outside NZ

Only disclose personal information to foreign persons or entities if it:

- *is carrying on business in NZ so subject to the Privacy Act OR*
- *is subject to privacy laws that provide comparable safeguards to those in the Privacy Act, OR*
- *is required to protect the information in a way that provides comparable safeguards (in contract or by agreement)*
- *is a participant in a binding scheme for international disclosures of personal information prescribed in regulations by the NZ Government.*

Summary of how IR will comply:

Is personal information being disclosed overseas? If so, has IR carried out appropriate due diligence checks required under this privacy principle. The OPC has created a decision tree for this principle. Note: It is not considered a disclosure of information to use cloud service providers if they are only holding the information for safe custody or processing.

If you have identified any risk to complying with principle 12, complete the Table of risks and mitigations.

IPP 13 Unique identifiers

Unique identifiers must not be assigned to individuals unless this is necessary to carry out functions efficiently.

If assigned, steps must be taken to ensure that it is assigned only to an individual whose identity is clearly established; and risk of misuse of a unique identifier is minimised (for example, by showing truncated account numbers on receipts or in correspondence).

Summary of how IR will comply:

Is the IRD number being used? If so, is this a new use for that number? Is a new unique identifier being created and assigned to people? If so, explain why this is necessary?

If you have identified any risk to complying with principle 13, complete the Table of risks and mitigations.

4. Ethical Assessment

Using and analysing data can introduce risks around the illegal or unethical use of data. IR must ensure it has ethical data practices and processes to maintain customer trust.

Inappropriate use of data may result in discrimination of subjects either directly or indirectly. Processes and systems may discriminate at the input stage, perhaps because information going into datasets is biased against some individuals or groups, or the system collects information in ways that are more intrusive with respect to some individuals or groups than others (without good cause), and/or at the output stage, when the recommendations of a system or process may discriminate without cause.

Does the initiative use ethical data practices?	**Y/N**	**If yes, explain your response**
Is the proposal likely to result in some members of a group being treated differently to one another?		
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?		
Will the proposal discriminate against some people?		
Can you foresee any harm to individuals in using the data in the way intended?		
Are we identifying and managing bias or discrimination?		
Does the data to be used specifically identify Māori or a Māori collective?		
Have you considered how the proposal contributes to the active protection of Māori interests?		

5. Algorithms and automated decision making

Complete this section only if the initiative will involve algorithms or automated decision making.

The Algorithm Charter for Aotearoa New Zealand is a commitment by government agencies to manage how algorithms will be used to strike the right balance between privacy and transparency, prevent unintended bias and reflect the principles of the Treaty of Waitangi.

IR signed up to the Charter and commits to making an assessment of the impact of decisions informed by our algorithms. We further commit to applying the Algorithm Charter commitments as guided by the identified risk rating.

Delivering clear public benefit	**Y/N**	**If yes, explain your response**
Who will benefit from the development of this system?		
Have any assumptions been made in design or planning?		
Is it likely people will suffer an unintended adverse impact as a result?		
Have associated policies and decisions been evaluated for fairness and potential bias and have a solid grounding in law?		
Ensure data is fit for purpose		
How accurate, precise, consistent, and complete is the data quality? (This may already have been answered at IPP8)		
Are we re-using data that was originally collected for another purpose?		
How are we identifying and managing bias or discrimination? (This may already have been answered in the ethical assessment)		
Transparency		

Are data use and analytical processes well documented, and the decisions they inform described in clear, simple, easy-to-understand language?		
Will decisions be explainable and auditable?		
Are there real-time feedback loops and a 'kill switch' if post-deployment bias is found?		
Have our customers been informed about how the system will work, and how they may seek more information or exercise review rights?		
Understand the limitations		
Do any limitations exist in terms of the collection or use of the data?		
Who is represented in the dataset, and who is not, which might lead to historic or representation bias?		
Human oversight		
Will decisions informed by the algorithms involve human judgement and evaluation?		
Can we test to see how well the algorithm is working compared to human decision-making?		
Will the automated decision-making process be regularly reviewed to make sure it's still fit for purpose?		

Not in scope

Nominate a point of contact for public inquiries about algorithms, and provide a channel for challenging or appealing of decisions informed by algorithms		

6. Risk Assessment

6.1 Table of risks and mitigations

Using the table below, describe the privacy risks you've identified through the privacy assessment and how you propose to mitigate and manage those risks. See the Risk Rating Tool at Appendix 1.

Ref No	Description of risk	Consequences for IR or individuals	Existing controls that help manage risks identified	Suggested action to reduce or mitigate risk
RSK-01	*Example: Information is collected which is not considered necessary for IR to carry out its functions (IPP 1)*			
	Example: There are insufficient controls to prevent unauthorised access to datasets (IPP 5)			
	Example: Information obtained is disclosed to agencies that are not entitled to it (IPP 11)			
	Example: Information gathered is associated with the wrong taxpayer or entity (IPP 7 & IPP 8)			
	Example: Information is kept for longer than is necessary (IPP 9)			

6.2 Summary of risks

Add a narrative summary of the project risk assessment, including an assessment of the severity of any potential impacts (ie. could individuals or IR be harmed if the risk is not mitigated?).

Also think about how the risks may be controlled in future, for example having a governance or working group overseeing the project or arranging for an audit to be conducted.

7. Recommendations

Based on the suggested actions in the risk table, below are summarised recommendations to minimise the impact on privacy. These should be agreed with the senior responsible owner5.

Ref	Recommendation	Agreed Y/N
R-01		

Inland Revenue
Te Tari Taake

[Initiative name]

Privacy Threshold Assessment

Prepared by:

Date:

Supply ID:

About this Document

The purpose of this document is to demonstrate that privacy has been considered in a project or process that involves personal information. The Analysis pulls together relevant information to determine whether a full Privacy Impact Assessment (PIA) should be completed and records IRs decision of why a PIA has not been done. It will answer the following questions:

1. Does this proposal involve a new way of managing personal information?
2. Does the proposal raise a significant privacy risk for IR?
3. Is a full privacy impact assessment required?

Project Summary

1.1 Description

Provide a brief description of the initiative including how many people will be affected or impacted.

1.2 Purpose of the change

Describe the existing process/system and explain the problem you are addressing and the purpose of the change.

1.3 Public benefit

Is there a clear and demonstrable link between the proposed use/reuse of the data and a beneficial outcome? Describe how it will benefit IR or the public or projected benefits to IR or to the individuals affected.

1.4 Privacy Enhancement

Privacy impacts can be positive (privacy-enhancing) or negative (privacy-invasive). Describe how your proposal is going to enhance the safeguarding of customer information eg information that was previously stored in spreadsheets is now in a dedicated repository; the number of transactions in a process have been reduced.

1.5 Personal information to be used

In the table below, describe:

- the personal information that will be collected, used and/or disclosed
- the source of the information
- the purpose of the information for your project.

Note: "Personal information" is any information about an identifiable living person. A person doesn't have to be named for the information to be identifiable.

Type of personal information	Source of information	Purpose of using the information
e.g. full name		

1.6 Governance

Outline who has been engaged to date including sponsor or senior leaders, groups that have been consulted and approvals/endorsement to date.

Name of person or group	Business Unit	Approved, Consulted, Informed etc

2. Privacy assessment

2.1 Areas that are risky for privacy

Some types of projects are commonly known to create privacy risks. If the project involves one or more of these risk areas, it's likely that a full Privacy Impact Assessment (PIA) will be valuable.

Use this checklist to identify and record whether your proposal raises certain privacy risks.

Does the project involve any of the following?	**Y/N**	**If yes, explain your response**
Does the initiative involve a substantial change to an existing policy, process or system?		
Is it linked to a practice or activity that is listed on a risk register?		
Collection	**Y/N**	**If yes, explain your response**
Will IR be collecting personal information that it doesn't currently collect?		
Is collecting this information necessary for IR to carry out its functions?		
Where or who is the information being collected from?		
Storage, security, and retention	**Y/N**	**If yes, explain your response**
Does the initiative change the way personal or sensitive information is stored, secured or managed?		
Where will the information be stored?		
Who will have access to the information?		
How long will the information be retained?		
Does it involve transferring personal information offshore, using a third-party contractor?		
Use, disclosure, and accuracy	**Y/N**	**If yes, explain your response**
Is the information currently held by IR?		

If yes to the above question, for what purpose does IR hold the information?		
Will the initiative use or disclose information for a different purpose to why it was obtained?		
Will IR be sharing personal or taxpayer information with another organisation?		
Describe the data quality – is it accurate, consistent, and complete?		
What processes are in place to ensure and maintain data integrity?		
Access and identification	**Y/N**	**If yes, explain your response**
Will the information be stored on a customer or staff member's record?		
Does the initiative affect how people can access information IR holds about them?		
Does this involve a new way of identifying individuals?		
Other considerations	**Y/N**	**If yes, explain your response**
Is there a way to achieve the purpose of the project using less identifiable data?		
Would people be surprised by this use of their information?		
If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?		
Does the initiative involve tracking or monitoring of movements, behaviour or communications?		

3. Ethical considerations

3.1 Areas that may raise ethical issues

Using and analysing data can introduce risks around the unethical use of data. IR must ensure it has ethical data practices and processes to maintain customer trust.

Does the project use ethical data practices?	**Y/N**	**If yes, explain your response**
Is the proposal likely to result in some members of a group being treated differently to one another?		
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?		
How are we identifying and managing bias or discrimination?		
Can you foresee any harm to individuals in using the data in the way intended?		
Does the data to be used specifically identify Māori or a Māori collective?		
Have you considered how the proposal contributes to the active protection of Māori interests?		
Use of algorithms or AI	**Y/N**	**If yes, explain your response**
If using algorithms or AI is there confidence the output is robust, and assumptions are met?		
Will decisions informed by an algorithm or use of AI involve human review and evaluation?		
Will any automated decision-making process be regularly reviewed to make sure it's still fit for purpose?		

4. Risk assessment

If you answered "Yes" to any of the questions above, use the table below to give a rating – either **Low (L)**, **Medium (M)**, or **High (H)** – for each of the aspects of the project set out in the first column.

For risks that you've identified as Medium or High, indicate (in the right-hand column) how the project plans to lessen the risk (if this is known).

Aspect of the Project	Rating	Describe any risks and how to mitigate them
Level of information handling **L** – Minimal personal information will be handled **M** – A moderate amount of personal information (or information that could become personal information) will be handled **H** – A significant amount of personal information (or information that could become personal information) will be handled	**Low** **Medium** **High**	
Sensitivity of the information **L** – The information will not be sensitive (name, IRD number, or job title) **M** – The information may be considered to be sensitive (contact details, date of birth plus name plus IRD number, financial information, biometric data) **H** – The information will be highly sensitive (health or financial details, information about high profile individuals)	**Low** **Medium** **High**	
Significance of the changes **L** – Only minor change to existing functions/activities **M** – Substantial change to existing functions/activities; or a new initiative	**Low** **Medium** **High**	

H – Major overhaul of existing functions/activities; or a new initiative that's significantly different	
Interaction with others **L** – No interaction with other agencies **M** – Interaction with one or two other agencies **H** – Extensive cross-agency (that is, government) interaction or cross-sectional (non-government and government) interaction	**Low** **Medium** **High**
Public impact **L** – Minimal impact on IR and customers **M** – Likely to have some impact on customers due to changes to the handling of personal information; or changes may raise concern or media attention **H** – High impact on customers and the public, and concerns over aspects of project; widespread media interest likely	**Low** **Medium** **High**

5. Summary of privacy impact

The privacy impact for this project has been assessed as:	Tick
Low – There is little or no personal information involved; or the use of personal information is uncontroversial; or the risk of harm eventuating is negligible; or the change is minor and something that the individuals concerned would expect; or risks are fully mitigated	
Medium – Some personal information is involved, but any risks can be mitigated satisfactorily	
High – Sensitive personal information is involved, and/or several medium to high risks have been identified. **You must complete a full Privacy Impact Assessment**	

Inadequate information – More information and analysis is needed to fully assess the privacy impact of the project.	

6. Reasons for the privacy impact rating

Briefly summarise your reasons for rating the proposal as low, medium, or high.

7. Document sign-off

Position	Name	Business Unit	Sign-off Date
Sponsor or Business Owner			
Privacy Officer		ED&I	
[Add others as appropriate]			

Not in scope

Excerpt from Centre of Enterprise Data and Analytics: Intervention Campaign Process Map

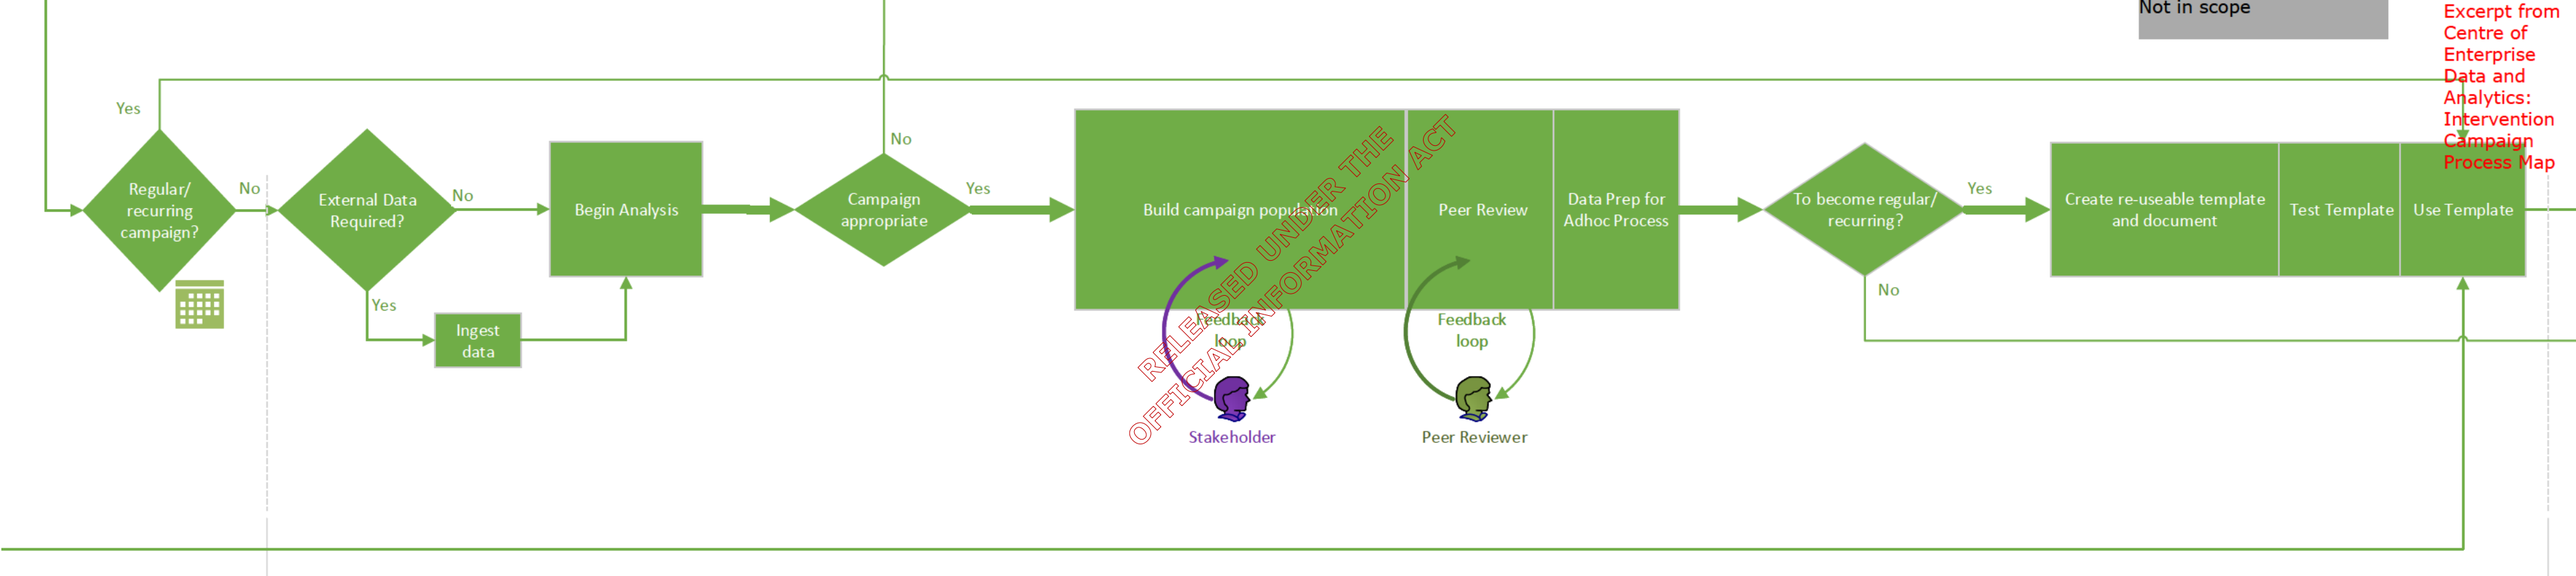

Artificial Intelligence use case guidelines

AI use case guidelines:

AI is an area of huge growth. Organisations across public and private sectors are moving to adopt these solutions and identify the ways that these solutions can support them in achieving better outcomes, and also how these solutions pose risk. IR has developed an AI staff use policy which will be available soon which should be read first by all staff at IR to understand what is permitted, and what restrictions have been put in place.

To provide additional guidance on how staff should use these tools the classification of the information or question space is important to understand. The tables below are designed to provide more clarity around specific use cases. It is impossible to represent all use cases within IR in a single view, however this table will be regularly updated with new use cases as IR evolves its understanding of this landscape.

s18(c)(i)

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 be leverage. This is not an exhaustive list and requests for additional rows and any questions can be made to InformationSecurity@ird.govt.nz. Compare your use case against this table to help verify if its permitted. All rows

	Generative AI /LLM	Integrated AI	Machine Learning	Business Rules

s 18(c)(i)

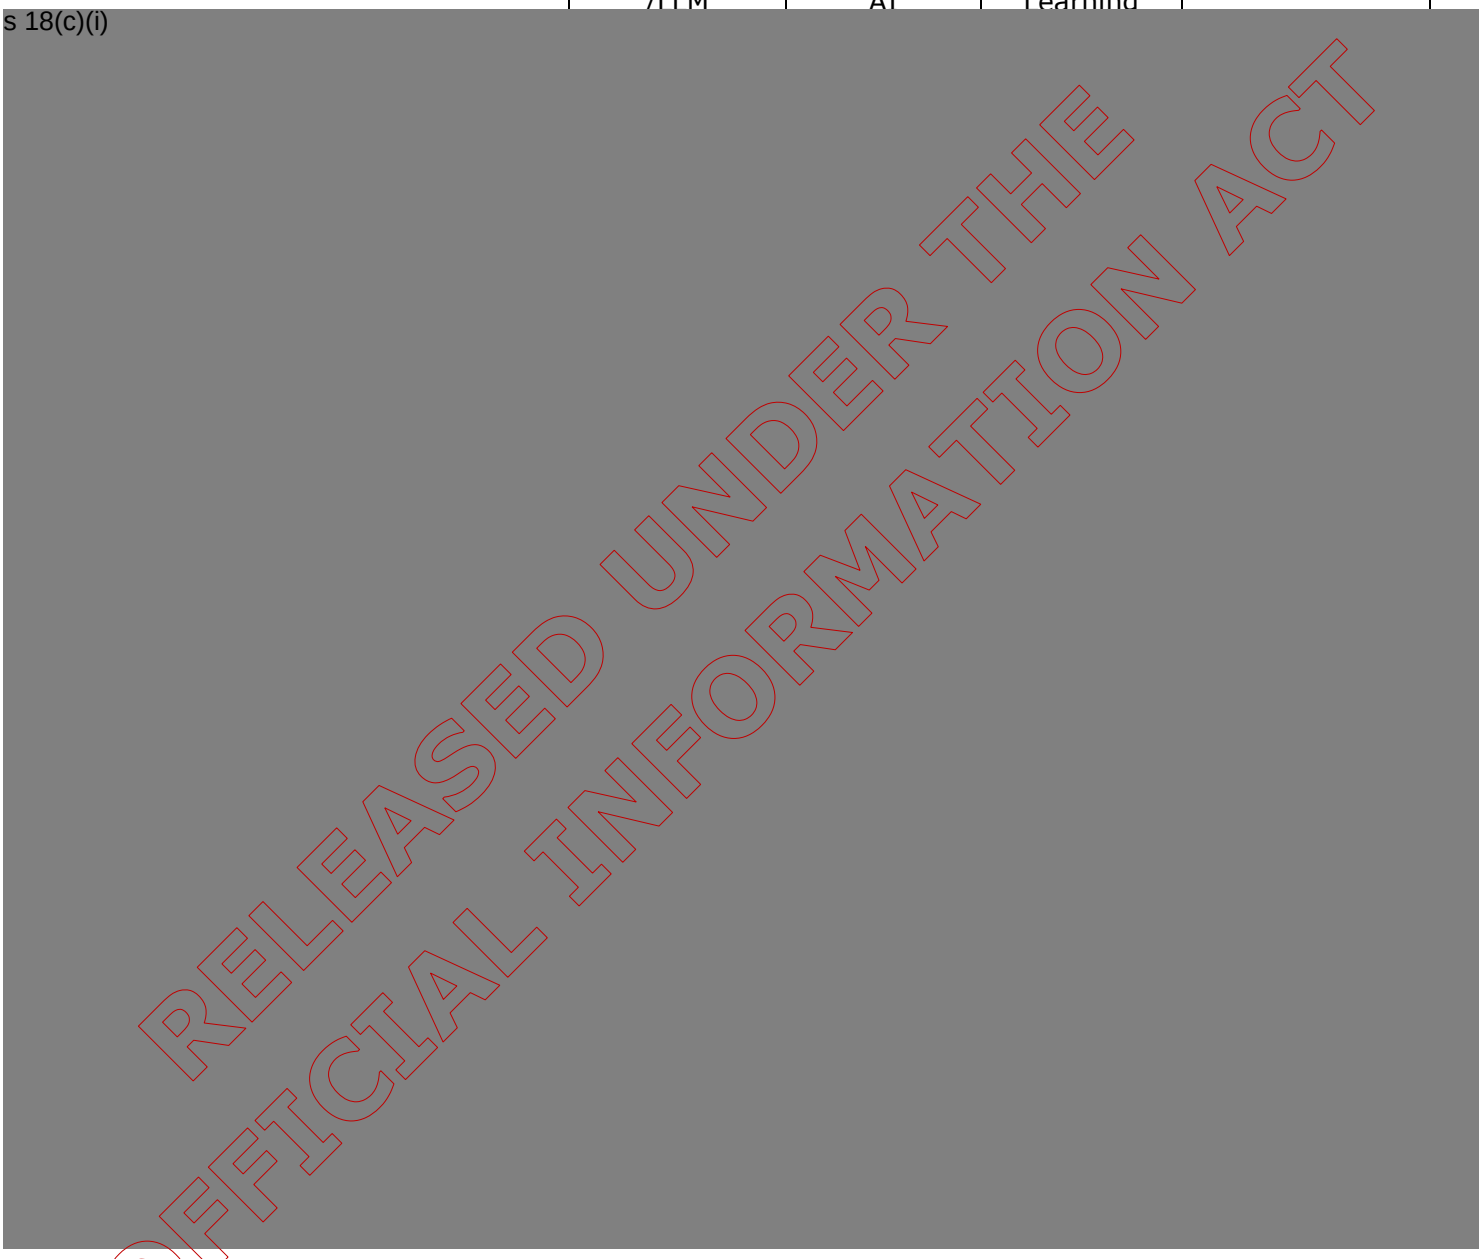

Appendix 1: Definitions

The following definitions apply throughout this document:

Generative AI/ LLM	These are direct interaction models, which describe how the user interacts with the solution. A generative AI/LLM is often used in a question-answer format (such as many of the OpenAI ChatGPT systems, or MS Co-pilot) LLM refers to the concept of a "Large language model", 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 doesn't 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.
Machine learning	ML is a more traditional approach where 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)
Business rules	These are the simplest kinds of computational "intelligence", this includes decision trees (ex. if value X is above 100, then go to this step), simple coding logic such as if/else statements and for/while/until loops. e.g. 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.

Document control	Version 1.1
Approved by	s 9(2)(a)
Review dates	Date approved: 4 October 2023, re-issued 1.1 version 18.12.2023 Next review: Six months from approval date
Guidelines owner	Jay Harris – Chief Information Security Officer
Guidelines contact	InformationSecurity@ird.govt.nz

Customer-facing staff guidance

Makayla Stewart
Change Analyst (L2)

This a quick reference guide for our people who support and work directly with our customers, and applies across all customer groups and types. Note, this information applies to all people at IR.

When and what types of Artificial Intelligence can be used is guided by our AI policy and AI Guidelines.

Tax and Large Language Models (SoloGPT)

With the rise of large language models, companies are now creating apps or chatbots that provide customers with tax advise. An example of this is SoloGPT and TaxGPT.

If a customer mentions or questions about the use of these bots, advise the customer that Inland Revenue's websites are where they need to go for tax information as they are the source of truth or they need to consult with an accountant for tax advice. Large language models can produce incorrect information or be targeted to an overseas tax jurisdiction which could steer customers wrong or lead to poor outcome.

If you receive contact from a customer about one of these apps, tools or services, please continue with your normal process and then email AI@ird.govt.nz with the details of the customer contact.

Do

- Advise customers that official resource we have created (websites, guides, forms etc) and our organisation are the source of truth.
- Advise customers asking about our use of Artificial Intelligence that their best and fastest avenue is to make an Official Information Act request through our online facility. Governance and Ministerial Services see Official In formation Requests for more information. Customers can find out more information and make their requests here: About OIA requests. If the customer is unwilling to do so, please receive the information request and contact Governance and Ministerial Services.
- Only use approved business tools for their intended purpose to complete your work.
- Be aware of the potential for bad actors to use AI systems to create misleading, incorrect, discriminatory or biased information, and to pose a security risk.
- Understand that any information that is produced either in part or whole by an AI system by your use

Don't

- Read any of this material out to customers. If asked for this material, or any information held by IR in a general sense, consult with Governance and Ministerial Services before taking any actions.
- Put customer or IR information into AI systems that have not been explicitly approved for that use, see AI application register for a list of approved tools and their purpose.
- Use publicly available AI services such as Chat GPT or Bing Chat to make responses to customers.

of that system is your responsibility and ultimately, Inland Revenue's.
- Seek guidance from your leader if you are unsure.

Unsure? Get in touch

AI Artificial Intelligence

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 a *user* if you are using or considering the use of AI products and services for an Inland Revenue purpose and are:

- An Inland Revenue employee
- A contingent worker (consultant, contractor 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. 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 perspective 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 unfair 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 through Inland Revenue's AI Oversight Group and AI Working Group (contact InformationSecurity@ird.govt.nz). This ensures adequate testing, review and consideration is applied to all use cases, and prevents effort duplication and risk.

Existing common 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 InformationSecurity@ird.govt.nz.

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 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 specifically checking for 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.
- Make clear through visual, verbal or written indicators/mechanisms the use of AI/LLM and other intelligent solutions in content they produce or contribute to.
- 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.
Information Security Team	• Responds to any report of abuse, misuse, or non-compliance with this policy.
Procurement team	• Ensures appropriate checks are undertaken and understanding is held of vendors where AI is in the services provided. Eg: 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, limit 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.
All IR people, including permanent employees, fixed term employees, agency temps, consultants, and contractors	• 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 / 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. A generative AI/LLM is often used in a question-answer format. Examples: ChatGPT, Bing Discover and MS Co-pilot solutions. These are solutions capable of generating entirely new content and leverage large language models, which more traditional systems do not.
AI integrated solutions	These are not generative, but often focus on adjusting output information based on learned knowledge. This may be very visible (like a chatbot or assistant) or may be hidden from the user, such as a ranking algorithm (used in search engines). Examples: spellcheck, 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: DIP data investigations, Website and Haukainga search, Viva insights.
Business rules	These are the simplest kinds of computational intelligence. This includes decision trees (eg: 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.

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

Matawai page:

- Artificial Intelligence - Home (sharepoint.com)

Additional resources include:

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

Document control	Version 1.1
Approved by	s 9(2)(a)
Review dates	Date approved: 4 October 2023, re-issued 1.1 version 18.12.2023 Next review: Six months from approval date
Policy owner	Jay Harris – Chief Information Security Officer
Policy contact	InformationSecurity@ird.govt.nz

Inland Revenue
Te Tari Taake

Kaupapahere Haumarutanga korero

Information Security Policy

This policy outlines the approach to information security for Inland Revenue. This policy is our overarching information security policy and is supported by our other specific information security policies and standards.

Why we have this policy

IR must be trusted by its customers and partners as an organisation for protecting the information they share with us or that is generated by us.

This information security policy aims to provide the high-level direction for information security at IR. It provides the guidance for how information which IR collects, processes, shares or stores, in whatever form, is kept secure. This will assist us in complying with our obligations under the Tax Administration Act 1994 (TAA). Below are two key extracts from the TAA which set out those obligations:

> *"Every Minister and every officer of any government agency having responsibilities under this Act or any other Act in relation to the collection of taxes and other functions under the Inland Revenue Acts are at all times to use their best endeavours to protect the integrity of the tax system. – Section 6(1)."*

> *"A revenue officer must keep confidential all sensitive revenue information and must not disclose the information unless the disclosure is a permitted disclosure that meets the requirements of sections 18D to 18J – Section 18(1)."*

There are other Legislative, government requirements and industry standards may also govern and control the way we handle and protect information. These are listed in the 'Finding out more' section below.

Who this policy applies to

This policy applies to you if you are:

- An IR employee
- A contractor working for IR or with access to IR systems/information
- A third-party with access to IR systems/information

Third-party vendors and contractors are also responsible for any sub-contractors (or any sub-sub-contractors, etc) which have access to IR assets. Sub-contractors of third-party vendors must also meet all IR requirements and be approved by IR.

It also applies to all information systems, networks and mobile devices that store, transfer, process, connect with or communicate to our applications and databases, along with internet-based services and services that may be located overseas (or in the cloud).

The requirements of this policy apply wherever you are, including whether you are working remotely, travelling overseas or in a business continuity crisis mode (for example, being away from your normal place of work due to a disaster).

In this policy, "information" refers to information in any format that IR creates, receives, uses, holds or maintains.

An "information system" includes the entire architecture of a service or any subcomponent that stores, processes or transfers information.

Your responsibility

As someone working for IR, you are responsible for understanding and following this policy. This means:

Compliance Measurement
The Assurance, Integrity or Cyber Security Operations teams may verify compliance to this policy and in conjunction with the IR Code of Conduct through various methods, including but not limited to, business tool reports, internal and external audits, and feedback to the policy owner.

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

Non-Compliance
If you have been found not to have complied with this policy, then the appropriate action will be taken.

Our policy

1. Principles

To help guide our Information Security thinking and behaviour IR has defined some key security principles. These principles will not only guide our thinking but help to embed our secure by design approach to our information security for both digital and non-digital pathways. Where these principles are not able to be met, it may present additional risks to IR that need to be considered and appropriate action taken.

There are three main information security principles:

- **Confidentiality** – information is accessed to those authorised to have access.
- **Integrity** – Making sure that the information is not modified by unauthorised users, and changes by authorised users are tracked to ensure the integrity is maintained.
- **Availability** – information/systems are accessible when authorised users need it.

There are six supporting principles:

- **Authentication** – we will verify a user's identity to ensure that the person requesting access is authorised.
- **Non-repudiation** –the ability for a system to prove that a specific user sent a message and that it hasn't been modified.
- **Need to know** – a person will be provided with only the information that they need to successfully fulfil their role.

- **Least privilege** – we will only authorise users the privileges needed to undertake their duties.
- **Defence in-depth** – we will use multiple and coordinated security countermeasures to protect the information assets in our organisation.
- **Segregation** – we will ensure our user roles, systems architecture and design incorporate separation and segregation to establish trust zones, define security domains and enforce boundaries.

2. This Policy

It is important that your actions foster a culture of security at IR. You, and any staff you are responsible for (including contractors/non-workers) must be adequately trained by completing the standard onboarding security training, the required annual refresher training and continue to regularly complete security training in our Learning Management Platform.

IR has specific policies and standards which provide further guidance on different areas of Information Security. It is important that these other policies and standards are used as the foundation for anything we do at IR. Our corporate Information Security Intranet provides the collection of guidance documents, Policies and Standards to help you in your day-to-day IR life or help you to comply when undertaking projects.

This Policy sets the overall direction of Information Security for IR and authorises Policies which need to be read in conjunction with this one. For example, our library of Policies and Standards cover, expectations such as:

- Never giving out or share your usernames and passwords.
- Using different passwords for different IR systems and personal accounts.
- Taking care to use the correct address when sending emails.
- Never using a generic account that doesn't identify you or an account that doesn't belong to you, to access IR information.
- Classification of IR information to manage the security and privacy of internal and external customers.
- Ensuring conversations of a sensitive nature are not overheard by others, either at your place of work or in public places.
- Taking care to keep documents and information on portable devices protected using encryption and physical security measures if you take them out of the office.
- Never representing IR using your private email address, or forward work emails or business information to your private email address.
- Accessing and using IR information and systems only for authorised IR business purposes.
- Never disabling, circumventing, or removing the security mechanisms on IR provided devices (i.e. laptops, mobile phones).
- Never forwarding unknown or suspicious email attachments.
- Never clicking on links or attachments in suspicious emails.
- Conducting IR business securely always, even in a crisis.
- Allowing access to all your equipment and files used in your work at IR if requested for security checks.
- Immediately reporting any actual or suspected loss, theft, or improper use of or access to IR information by raising a request in ServiceNow
- Following all documented procedures in response to any ICT security incident.

3. Organisational Accountabilities and Responsibilities

The Commissioner of IR is accountable for information security. The Chief Information Security Officer (CISO) has been delegated authority for generation, maintenance and approval of Policies and Standards enabling this Information Security Policy. The Commissioner, Chief Security Officer and the CISO are responsible for assuring the Executive Leadership Team, the public, and other stakeholders that IR is meeting the Government's information security requirements and is keeping IR's information safe and secure.

Finding out more

IR has a number of other policies, standard and other guidance to support this Information Security Policy. The below provides additional resources:

- Policies and Standards
- Information Security Corporate site

Sign off

Document control	1.2
Approved by	Peter Mersi – Commissioner
Review dates	Date reviewed: August 2023 Next review: August 2025
Policy owner	Peter Mersi – Commissioner
Policy owner signature	s 9(2)(a)
Policy contact	informationsecurity@ird.govt.nz

Version History

Version	Date	Sections amended	Summary of amendment
1.1	21/6/23	All	Rename of policy to Information Security Policy, change of signature to Commissioner
1.2	22/8/23	All	Updated links to new KX page

What is Artificial Intelligence

Makayla Stewart
Change Analyst (L2)

Artificial Intelligence (AI) is a hot topic at the moment. We already use a lot of AI technology in many aspects of our day to day lives - from personalised recommendations on sites like Netflix and Spotify, to using facial recognition to unlock phones & devices instead of a password.

This page will give you a broad overview of the exciting, and emerging, world of AI and the opportunities and risks it presents for Inland Revenue.

What is Artificial Intelligence?

Artificial Intelligence (AI) is the field of computer science that seeks to create systems with the ability to imitate intelligence that is typical of humans, like perceiving, reasoning, planning, decision-making and problem-solving.

Systems are designed so they can learn by example and continuously improve rather than being programmed for specific tasks.

Modern AI systems use algorithms to learn patterns from data. The practice of an AI system processing data is known as training and is key to how AI systems function.

The data used to train an AI system will depend on what it's being designed to do. For example, a self-driving car will be trained on road maps and driving manuals. Whereas a facial recognition system would be trained on different faces and facial features. Training enables AI to complete tasks with a human-level of intelligence at scale.

The quality and amount of data used to train an AI is crucial. The higher quality the data and the more data used, the more reliable and accurate the AI will

become. AI systems that use low quality data or small datasets will have less reliable and accurate results.

Types of Artificial Intelligence

There are many ways to categorise or group different types of AI including by functionality and capability. The following is a traditional perspective on types of AI:

All existing types of AI are known as Artificial Narrow Intelligence or Weak AI. They are built to perform a specific task autonomously and perfectly using human like capabilities.

There are two types of AI known as Artificial General Intelligence and Artificial Super Intelligence, these do not currently exist. Artificial General Intelligence are systems that can learn, perceive, understand and function completely like human beings whereas Artificial Super Intelligence, has similar processing capabilities that humans do, has a greater memory and is faster.

For more information see our related resources above.

What are the opportunities?

AI presents an opportunity for IR to deliver improved services across our organisation. As an organisation we need to be deliberate in our decision making to balance risk, opportunity and return on our investments.

Consider how these AI functions could help us in future:

- Pulling together short summaries of information based on a text prompt such as a question or a meeting recording.
- Identifying connections and patterns in data that humans cannot and adjusting actions it takes in real-time.
- Acting as a chatbot that provides real time guidance while an individual completes a task.
- Automating systems and processes for simple and some complex tasks while flagging more complex work for human review.
- Recognising changing trends and behaviours in our society and within specific groups of people.

There is significant opportunity to achieve better outcomes for IR and our customers, and significant and unique risks that AI poses or increases that we will need to mitigate. To understand how we do this, we need an understanding of how the technology behind AI works.

IR already uses and has invested in AI in a number of areas (see AI application register for more information), your welcome to familiarise yourself with the capabilities we have available to us and understand when/how we should use them. This is something we can do now to deliver better outcomes for our customers.

To understand how we are currently using this technology see Using Artificial Intelligence at IR.

Concepts, categories and definitions within Artificial Intelligence

There are key concepts that will be important in helping you navigate the world of AI.

Concepts

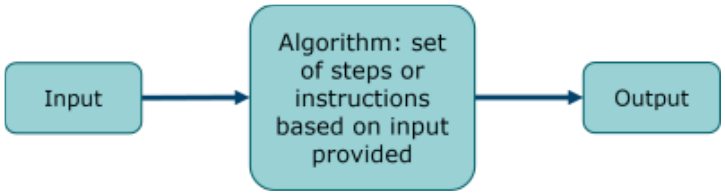

Diagram: How algorithms function at a high level.

An **Algorithm** is a set of steps or instructions for solving a problem or performing a task.

Categories

Business rules are the simplest type of machine 'intelligence', they use simple logic and processing to reach an outcome. Examples of this include decision trees and if/else statements. These machines do not undergo a 'training' process.

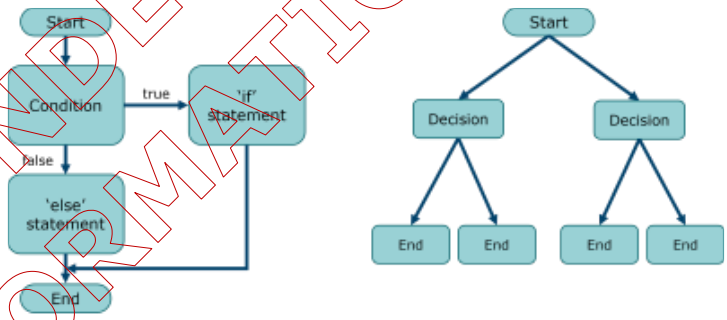

START, Ātea and ServiceNow all utilise business rules to complete some tasks.

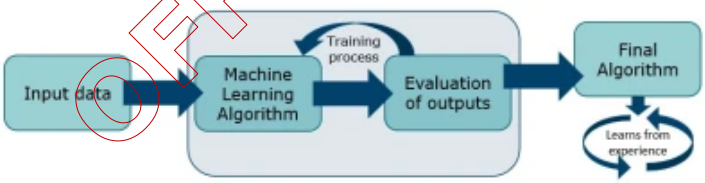

Machine learning (ML) is a branch of AI that enables machines to improve at tasks (e.g. decision making) with experience relying on the data, algorithms and training.

Haukāinga and our external website searches using machine learning to provide the most relevant results to users, see How we use Artificial Intelligence for more information.

AI integrated tools refers to solutions that already exist in our work where a vendor has integrated an AI component into their tool.

In many cases we may not have an option to stop or control an AI component in a particular solution and over time, AI will increasingly be included in all software tools.

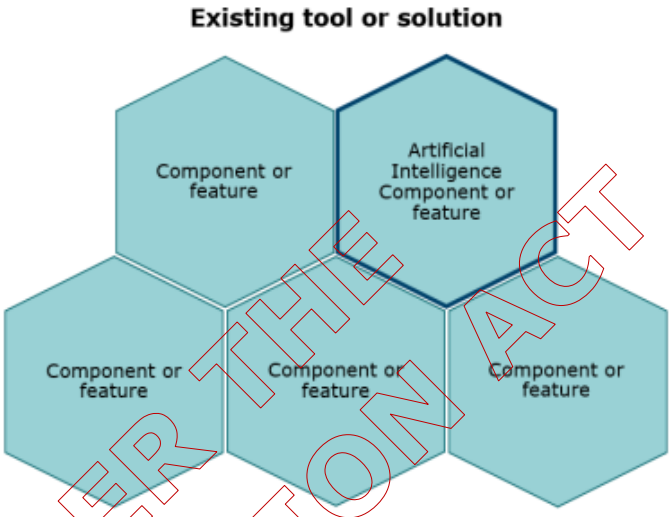

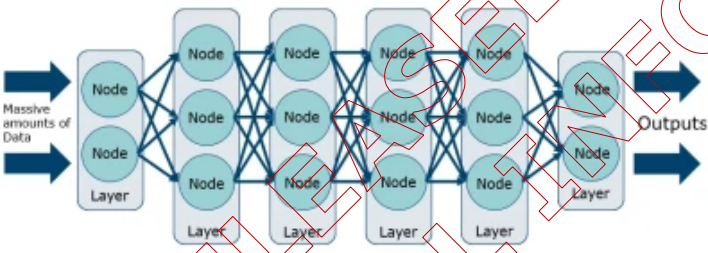

Deep learning (DL) is a type of machine learning that uses neural networks to learn in a way similar to how humans do. It requires massive amounts of data and can learn from its own environment and past mistakes .

Generative artificial intelligence commonly known as GenAI, is a type of AI that can create new content using deep learning.

Gen AI can create a wide range of content including text, audio, videos and computer code.

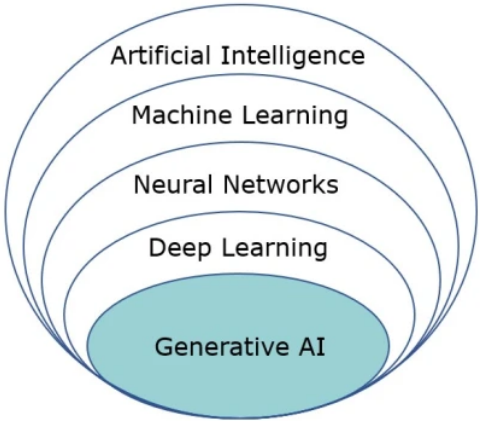

Definitions

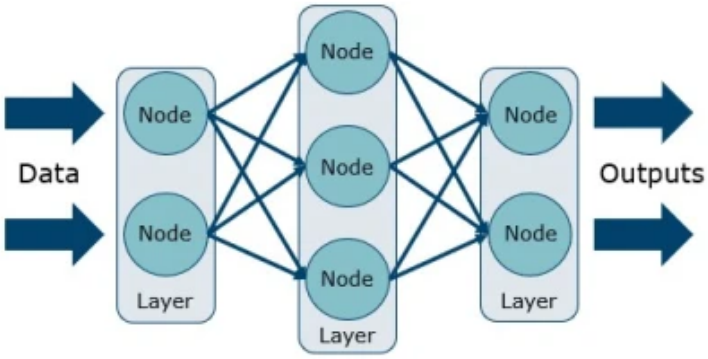

Artificial neural networks (ANN) are a branch of machine learning that enables machines to process data in a similar way to how the neurons in our brains do. Neurons send messages or signals to one another. Neural networks do the same using layers of nodes that send data amongst one-another.

Large language model (LLM) are systems that have processed large amounts of text (usually from the internet) using neural networks. It allows you to input a text prompt and will generate a text response by repeatedly predicting the next word in a sentence or paragraph.

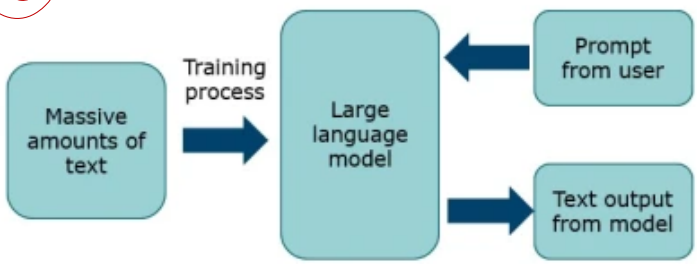

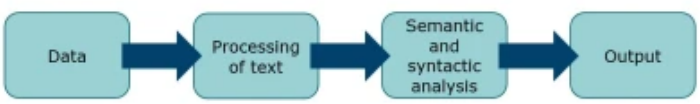

Natural language processing (NLP) aims to give systems the ability to understand and create human speech. It does this through analysing the relationship between words in sentences (syntactic) and the meaning of words together (semantic).

Not in scope

Excerpt from Status Report: Weekly update for the Minister of Revenue

Inland Revenue
Te Tari Taake

Status Report

Weekly update for the Minister of Revenue

Week commencing: Monday 20 July 2020
Date issued: Thursday 16 July 2020

Not in scope

New Topical Issues

New items since the last Status Report (issued on 9 July 2020).

Not in scope

Not in scope

Signing of the Algorithm Charter	As part of ensuring New Zealanders have trust and confidence in the ways agencies use and steward data and information, the Government Chief Data Steward, a function within Stats NZ, has led the development of an Algorithm Charter (in conjunction with agencies) for agency adoption. Algorithms are used to help make good decisions and deliver services that are more effective and efficient. Algorithms can also mitigate the risk that human biases will impact the administration of government services and result in real benefits for everyone. However, the decisions they make or inform can be impacted by bias in the data the algorithm is created from, what is taken into consideration when making a decision, and how well decisions are validated. The Algorithm Charter focuses on demonstrating agency commitment to the management of algorithms where there is a potentially high risk of these biases being included. It also ensures there is sufficient human oversight throughout the development of an algorithm and how it operates to limit biases and reflect what is intended as much as possible. Inland Revenue is already part way through implementing its formalised enterprise data and information governance approach in which the principles of the charter are already established. The Commissioner of Inland Revenue will be signing the Algorithm Charter for Inland Revenue on 17 July 2020.

Data and Information Governance Authority

Inland Revenue
Te Tari Taake

Update on Government Algorithm Charter

Submission details

Title	Update on Government Algorithm Charter
Sponsor	Tina MacLean
Presenter	Doug Lambert
Date	4 August 2020

Recommendations

The sponsor recommends to the Data and Information Governance Authority the following:

1. **NOTE** that Stats NZ's Government Chief Data Steward (GCDS) function has been working since 2018 to raise New Zealander's trust in confidence in Government's use of data.
2. **NOTE** that the Government Chief Data Steward has between August 2019 and July 2020 iterated through several versions of an All of Government Algorithm Charter (the Charter) with iterations increasingly taking onboard feedback from the public, interest groups, and agencies.
3. **NOTE** that the Government Chief Data Steward has published the Charter on 28 July 2020.
4. **NOTE** that the Commissioner of Inland Revenue has adopted the Charter for Inland Revenue.
5. **NOTE** that Inland Revenue will incorporate the Charter into the wider Data and Intelligence work programme to operationalise the Charter commitments.
6. **ENDORSE** the incorporation of the Charter into how Inland Revenue works with algorithms.

Summary

New Zealand is a member nation of the Open Government Partnership and is currently executing Action Plan 3 spanning 2018 to 2020 which has 23 commitments. One of those commitments is the review of Government use of algorithms.

As part of their data system lead role, the Government Chief Data Steward undertook in 2018 a review of operational algorithm use by Government agencies. The aim of the

resulting Algorithm Assessment Report was to let New Zealanders know how the government uses algorithms, and the ways Government makes sure algorithms are being used responsibly and transparently.

Between August 2019 and July 2020 the Government Chief Data Steward have iterated through several versions of an Algorithm Charter. These iterations have involved agency consultation in August and September 2019, public consultation between October and December 2019, and further agency consultation between January and July 2020.

Throughout 2019 and 2020 many people across Inland Revenue engaged on the Charter development and provided feedback which was summarised and sent to Stats NZ. The level of support for an algorithm charter varied. Points raised were focused on these areas:

- Almost all algorithms implemented at Inland Revenue are partial or full automations of business rules defined by legislation and these are already publicly disclosed on the web site.
- Inland Revenue would want to keep some of its algorithms undisclosed to support the integrity of the tax system.
- Determination of likelihood and impact as per the Charter risk matrix are largely subjective and will vary by practitioner and agency.
- Inland Revenue already incorporates most of the Charter mitigations and controls.

The Charter is intended to demonstrate a commitment by government agencies to carefully manage how algorithms will be used to strike the right balance between privacy and transparency, prevent unintended bias and reflect the principles of the Treaty of Waitangi.

Inland Revenue has provided comprehensive feedback on each iteration of the Charter.

The Government Chief Data Steward have settled on a risk assessed approach to determining the degree of scrutiny and safeguards applied to each algorithm an agency operates. Risk is intended to be assessed before any controls and safeguards have been applied to an algorithm.

The Charter commits adopting agencies to:

- Making an assessment of the impact of decisions informed by their algorithms.
- Applying the Algorithm Charter commitments as guided by the identified risk rating.

The Charter commitments are (abridged):

- Maintain transparency by clearly explaining how decisions are informed by algorithms.
- Deliver clear public benefit through Treaty commitments.

- Focus on people.
- Make sure data is fit for purpose.
- Ensure that privacy, ethics and human rights are safeguarded.
- Retain human oversight.

The Charter will be reviewed annually.

In July 2020 the Commissioner of Inland Revenue adopted the Charter for Inland Revenue and Stats NZ published the Algorithm Charter on 28 July 2020.

Founding signatories to the Charter are:

- Te Tari Taake — Inland Revenue Department
- Te Ara Poutama Aotearoa — The Department of Corrections
- Te Tāhuhu o Te Mātauranga — The Ministry of Education
- Te Manatū Mō Te Taiao — The Ministry for the Environment
- The Ministry of Housing and Urban Development
- Te Tāhū o te Ture — The Ministry of Justice
- Toitū Te Whenua — Land Information New Zealand
- Te Puni Kōkiri — The Ministry of Māori Development
- Oranga Tamariki - The Ministry for Children
- The Ministry for Pacific Peoples
- Te Manatū Whakahiato Ora — The Ministry of Social Development
- Te Tatauranga Aotearoa — Statistics New Zealand
- Te Manatū Waka — The Ministry of Transport
- Te Kāhui Whakamana Rua Tekau mā Iwa—Pike River Recovery Agency
- Te Minitatanga mō ngā Wāhine — The Ministry for Women
- Te Hau Tāngata — Social Wellbeing Agency
- Te Ope Kātua o Aotearoa — New Zealand Defence Force
- Te Kaporeihana Āwhina Hunga Whara — Accident Compensation Corporation
- Te Tari Taiwhenua — Department of Internal Affairs
- Te Arawhiti — The Office for Māori Crown Relations
- Waka Kotahi — The New Zealand Transport Agency
- Te Tari Arotake Matauranga — The Education Review Office

There are closely related initiatives underway elsewhere in Government and internationally that New Zealand is involved in:

- In May 2019 New Zealand along with all OECD member nations adopted the OECD AI Principles.
- In November 2019 the Department of Internal Affairs established a partnership project with the World Economic Forum, Reimagining Regulation for the Age of AI, to pilot approaches to artificial intelligence regulation. The first white paper from this project was published June 2020.

Inland Revenue is incorporating the Charter into the wider Data and Intelligence work programme and Data and Information Governance work programme to operationalise the Charter commitments.

Related initiatives

The recommendations originate from or are depended on by these initiatives:

Initiative	Relationship	Contact
Data and Intelligence Programme	This programme will operationalise mechanisms that align with the Charter commitments.	Tina MacLean
Data and Information Governance Programme	This programme will define and establish the governance and capabilities necessary to operationalise mechanisms that align with the Charter commitments.	Doug Lambert

Risks and mitigations

The recommendation(s) address these risks:

Risk	Mitigations and controls to manage risk	Likely consequence if recommendations are adopted
Not adopting the Charter might raise concerns about Inland Revenue's stewardship and use of data for assisted and automated decisions.	Adoption of the Charter.	This sends a clear signal that Inland Revenue is committed to ensuring its stewardship and use of data is beyond reproach.

The recommendation(s) introduce these risks:

Risk	Mitigations and controls to manage risk	Likely consequence if recommendations are adopted
Inland Revenue fails to implement the Charter commitments where appropriate.	Incorporate the establishment and operationalisation of the Charter commitments into aligned work programmes.	Inland Revenue would be exercising its commitments in demonstrable ways and preserving New Zealander's trust and confidence in its use of data.

Supporting material

The Algorithm Charter document spans three pages with the first two covering context, purpose, and the risk assessment foundation. The third page is the Charter proper.

See https://data.govt.nz/assets/data-ethics/algorithm/Algorithm-Charter-2020_Final-English-1.pdf for the published Charter.

JULY 2020

New Zealand Government

Stats NZ
Tatauranga Aotearoa

ALGORITHM CHARTER FOR AOTEAROA NEW ZEALAND

The value of algorithms

Government agencies use data to help inform, improve and deliver the services provided to people in New Zealand every day. Simple algorithms can be used to standardise business processes to ensure scarce resources are distributed equitably. More complex algorithms can be used to distil information from large or complex data sets to support human decision-making and reveal insights that could not easily be revealed by human analysis alone.

These algorithms can be used to help government better understand New Zealand and New Zealanders. This knowledge helps government make good decisions and deliver services that are more effective and efficient. The use of algorithms can mitigate the risk that human biases will enter into the administration of government services and result in real benefits for everyone.

However, the opportunities also bring fresh challenges. For example, human bias could be perpetuated, or even amplified by, algorithms that are not designed and operated in thoughtful ways. Transparency and accountability are critical to ensuring that the public can trust and support the government to use these tools in appropriate ways.

This Charter is a commitment by government agencies to carefully manage how algorithms will be used to strike the right balance between privacy and transparency, prevent unintended bias and reflect the principles of the Treaty of Waitangi.

Definitions

There are a wide range of advanced analytical tools that can fit under the term 'algorithm'. These range from less advanced techniques such as regression models and decision trees, which primarily support predictions and streamline business processes, through to more complex systems, such as neural networks and Bayesian models, which can take on properties of machine learning as they make advanced calculations and predictions.

A good discussion of the different types of predictive algorithms and the challenges of defining these is contained in 'Government Use of Artificial Intelligence in New Zealand' (New Zealand Law Foundation and Otago University, 2019).

The risks and benefits associated with algorithms are largely unrelated to the types of algorithms being used. Very simple algorithms could result in just as much benefit (or harm) as the most complex algorithms depending on the content, focus and intended recipients of the business processes at hand. As a consequence, this Charter does not specify a technical definition of an algorithm. It instead commits signatories to take a particular focus on those algorithms that have a high risk of unintended consequences and/or have a significant impact if things do go wrong, particularly for vulnerable communities.

Review

The Algorithm Charter for Aotearoa New Zealand is an evolving piece of work that needs to respond to emerging technologies and also be fit-for-purpose for government agencies. After twelve months a review of the Algorithm Charter will be conducted, to ensure it is achieving its intended purpose of improving government transparency and accountability without stifling innovation or causing undue compliance burden.

Foundations

The Algorithm Charter is part of a wider ecosystem and works together with existing tools, networks and research, including:

Principles for the Safe and Effective Use of Data and Analytics (Privacy Commissioner and Government Chief Data Steward, 2018)

Government Use of Artificial Intelligence in New Zealand (New Zealand Law Foundation and Otago University, 2019)

Trustworthy AI in Aotearoa – AI Principles (AI Forum New Zealand, 2020)

Open Government Partnership, an international agreement to increase transparency

Data Protection and Use Policy (Social Wellbeing Agency, 2020)

Privacy, Human Rights and Ethics Framework (Ministry of Social Development).

Assessing likelihood and impact

The Algorithm Assessment Report found that advanced analytics and data use are an essential part of delivering public services. Applying the Charter to every business rule and process would be impossible for agencies to comply with and not achieve the intended benefits of the Charter.

However, where algorithms are being employed by government agencies in a way that can significantly impact on the wellbeing of people, or there is a high likelihood many people will suffer an unintended adverse impact, it is appropriate to apply the Charter.

Charter signatories will make an assessment of their algorithm decisions using the risk matrix below. This supports their evaluation, by quantifying the likelihood of an unintended adverse outcome against its relative level of impact to derive an overall level of risk.

The risk rating determines the application of the Charter.

Risk matrix

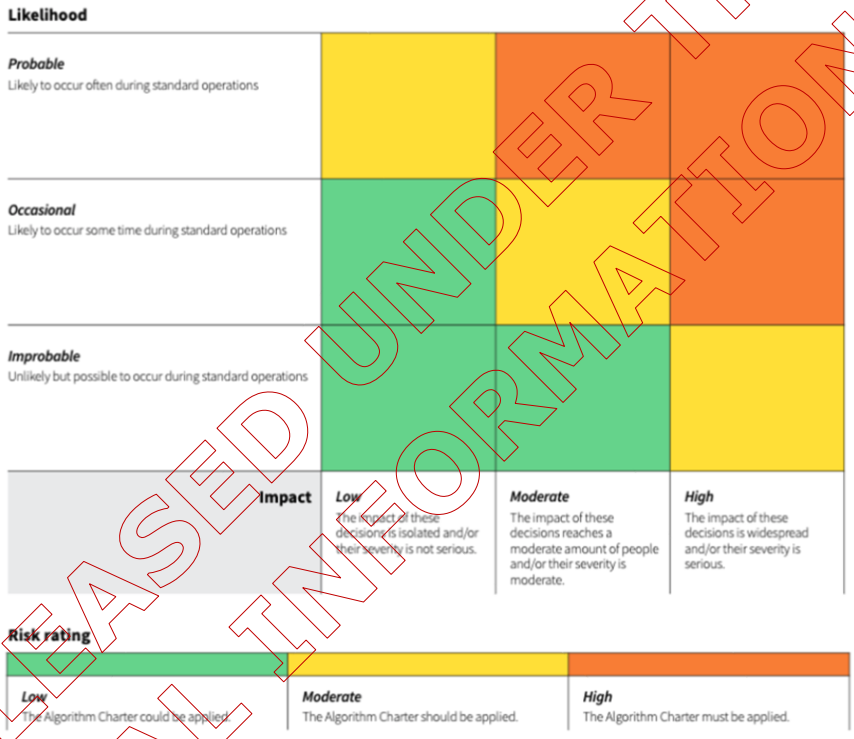

Risk rating

Low	Moderate	High
The Algorithm Charter could be applied.	The Algorithm Charter should be applied.	The Algorithm Charter must be applied.

Application and Commitment

The Charter will apply differently to each signatory. The risk matrix approach means that signatories can focus first on decisions that have a high risk and exclude most of the many business rules that government agencies use every day to give effect to legislative requirements and for business as usual activities. The intention is to focus on those uses of algorithms that have a high or critical risk of unintended harms for New Zealanders. This commitment will be reviewed in twelve months as part of the scope review.

JULY 2020

New Zealand Government

ALGORITHM CHARTER FOR AOTEAROA NEW ZEALAND

This Charter demonstrates a commitment to ensuring New Zealanders have confidence in how government agencies use algorithms. This Charter is one of many ways that government is demonstrating transparency and accountability in the use of data. However, it cannot fully address important considerations, such as Māori Data Sovereignty, as these are complex and require separate consideration.

Commitment:

Our organisation understands that decisions made using algorithms impact people in New Zealand. We commit to making an assessment of the impact of decisions informed by our algorithms. We further commit to applying the Algorithm Charter commitments as guided by the identified risk rating.

Algorithm Charter Commitments:

TRANSPARENCY

Maintain transparency by clearly explaining how decisions are informed by algorithms. This may include:

- Plain English documentation of the algorithm,
- Making information about the data and processes available (unless a lawful restriction prevents this),
- Publishing information about how data are collected, secured and stored.

PARTNERSHIP

- Deliver clear public benefit through Treaty commitments by:
  - Embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi.

PEOPLE

- Focus on people by:
  - Identifying and actively engaging with people, communities and groups who have an interest in algorithms, and consulting with those impacted by their use.

DATA

- Make sure data is fit for purpose by:
  - Understanding its limitations,
  - Identifying and managing bias.

PRIVACY, ETHICS AND HUMAN RIGHTS

- Ensure that privacy, ethics and human rights are safeguarded by:
  - Regularly peer reviewing algorithms to assess for unintended consequences and act on this information.

HUMAN OVERSIGHT

- Retain human oversight by:
  - Nominating a point of contact for public inquiries about algorithms,
  - Providing a channel for challenging or appealing of decisions informed by algorithms,
  - Clearly explaining the role of humans in decisions informed by algorithms.

Signed:

Chief Executive: **Chief Privacy Officer:** **Senior Manager responsible for algorithms:**

Organisation:

Date:

Not in scope

Data and Information Governance Authority

Inland Revenue
Te Tari Taake

Meeting Minutes – 4 August 2020

2.00pm – 3.30pm

Attendees: Mike Cunnington (Chair), Martin Smith, Mary Craig, Cath Atkins, David Carrigan, Dawn Swan, Patrick O'Doherty, Chris Hogg, Carol Feuerriegel, Kirsty Gemmill, Doug Lambert.

Apologies: Tina MacLean.

Secretary: Tanya Williams

Agenda items:

Not in scope

4. Statistics NZ's Algorithm Charter for Aotearoa NZ – presented by Doug Lambert
5. Forward Agenda Review – presented by Mike Cunnington

Note that **Item 2** was not covered in the meeting and members were asked by the Chair to provide feedback directly to the Secretary following the meeting.

Not in scope

4. Statistics NZ's Algorithm Charter for Aotearoa New Zealand

The Data Strategy and Governance Lead covered the intent, development, and adoption of the All of Government Algorithm Charter (the Charter).

The intent of the Charter is to engender the public's trust and confidence in how Government agencies steward and use data algorithms. The Charter takes a risk-based approach to minimising the potential negative impacts to New Zealanders from algorithms.

Potential impacts on IR to implement the Charter highlighted by members were:

- How agencies operationalise the Charter is central to its success.
- It's likely a change in the awareness of principles, processes, and practices will be required along with a strengthening of what is currently in place.
- IR might need to increase its transparency about its diligence in its use of data and algorithms.

Inland Revenue is incorporating the Charter into work underway within the wider Data and Intelligence work programme, and Data and Information Governance work programme, to operationalise the Charter commitments.

Decisions: DIGA endorsed incorporation of the Charter into how Inland Revenue works with data and algorithms.

Not in scope

Not in scope

Excerpt from Status Report: Weekly update for the Minister of Revenue

Inland Revenue
Te Tari Taake

Status Report

Weekly update for the Minister of Revenue

Week commencing: Monday 10 August 2020
Date issued: Thursday 6 August 2020

Not in scope

Not in scope

Media Coverage	
	Not in scope

Not in scope

Not in scope

Reseller news says twenty-one government agencies have signed up to the Algorithm Charter for Aotearoa New Zealand, including Inland Revenue. It's a new set of standards introduced by the government, and a world first, to guide the use of algorithms by public agencies and to give New Zealanders confidence that data is being used safely and effectively across government.

Not in scope

Not in scope

Not in scope

Report to Minister of Revenue: Inland Revenue's contribution to advancing Government data outcomes

Inland Revenue
Te Tari Taake

Inland Revenue report: Inland Revenue's contribution to advancing Government data outcomes

Date:	23 August 2021	**Priority:**	Medium
Security level:	In confidence	**Report number:**	IR2021/334

Action sought

	Action sought	Deadline
Minister of Revenue	**Note** the contents of this report **Refer** a copy of this report to members of the Cabinet Government Administration and Expenditure Review Committee	None None

Contact for telephone discussion (if required)

Name	Position	Telephone
Mike Cunnington	Deputy Commissioner, Information & Intelligence Services	s9(2)(a)

Not in scope

23 August 2021

Minister of Revenue

Inland Revenue's contribution to advancing Government data outcomes

Executive summary

1. Inland Revenue is a very active participant in a number of cross-agency initiatives related to data management and governance. We have adopted the algorithm charter and mandated data standards, and have a very constructive working relationship with Stats NZ as the functional lead for data. Through this work, we are supporting and enabling the New Zealand economy and public service to become increasingly digital.

2. Ensuring that Inland Revenue has appropriate safeguards and protections in place is a critical component of the work we do. Customers entrust us with sensitive information about their financial and personal circumstances, and we take our obligation to protect their information and keep it confidential very seriously. Policies and procedures are in place to ensure that we meet our legislative obligations and retain our customers' trust.

3. We share information with a wide range of other agencies, to help smooth customers' experiences when dealing with government and deliver better outcomes for them.

Recommended action

4. I recommend that you:

a) **Note** the contents of this report.

Noted

b) **Refer** a copy of this report to members of the Cabinet Government Administration and Expenditure Review Committee for their information.

Referred

Mike Cunnington
Deputy Commissioner, Information & Intelligence Services
23 August 2021

Hon David Parker
Minister of Revenue
/ /2021

Background

5. This report provides information about Inland Revenue's contribution to the development of data governance and management frameworks and standards across the public sector, and our data governance and management practices. It has been prepared in response to the questions asked when the Commissioner and Deputy Commissioner, Transformation appeared before the Cabinet Government Administration and Expenditure Review Committee on 8 July 2021 to discuss the Cabinet paper *Inland Revenue's July 2021 Transformation Update*.

How Inland Revenue manages and uses data and information

6. Inland Revenue has significant data and information collection powers and many information sharing arrangements in place. To ensure we act with integrity and engender trust and confidence amongst New Zealanders, we have established a data and information governance programme. The programme is helping to ensure we have transparent and accountable governance arrangements that provide oversight of our management and use of data, so that we act knowingly, responsibly, and with integrity.

7. We have defined a set of capabilities and established a governance framework (as shown in the diagram below) to:

- Support the delivery of our strategic objectives.
- Ensure information and tools are in place for our people to access guidance and make the decisions they should.
- Provide guidance and decisions where it is most appropriate for a governance group to do so.

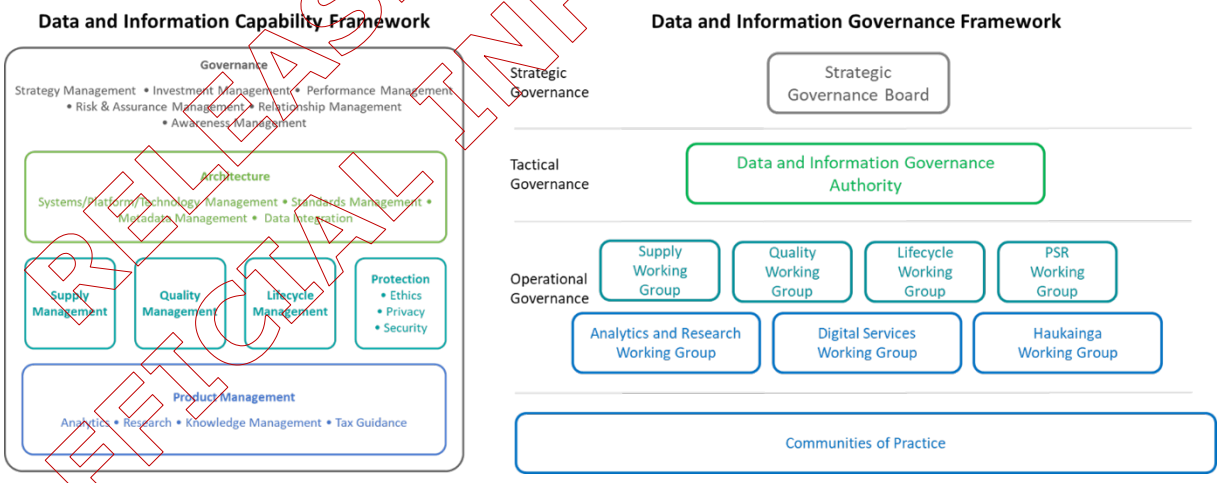

8. The programme is not yet complete and continues to focus on incorporating increased due diligence, including ethical considerations, into our collection, use, and stewardship of data and information. The work we do with other agencies is a key input into the programme.

Algorithm charter

9. Inland Revenue is a signatory to the algorithm charter for Aotearoa New Zealand, and contributed significantly to its development. The Government Chief Data Steward (GCDS), the functional lead role for data delegated to the Chief Executive Stats NZ, has acknowledged our commitment to ensuring the algorithm charter is of a high quality and

readily implementable. Stats NZ has favourably recognised our plan for implementing the algorithm charter.

10. To ensure Inland Revenue meets the commitments in the charter, we have developed a plan setting out what we need to do to improve our governance and management of data. We have also incorporated the commitments of the algorithm charter into our data and information policy, which sets the foundations for our data governance.

11. We have already begun putting the foundations in place and have established a data governance framework as noted above, along with policies and standards for managing data ethically. These are consistent with the commitments in the algorithm charter. Our next area of focus will be working with functional areas, such as our analytical teams, to ensure that they understand and adopt the agreed frameworks, policies and standards.

12. It is important to note that the charter does not define what an algorithm is. Each agency is responsible for identifying the algorithms that will have the highest risk of unintended consequences and/or significant impacts if things do go wrong. We are currently refreshing our approach to cataloguing algorithms and analytical models as part of the data and information governance programme noted above.

13. We are sharing our experiences and what we are putting in place with others. We have discussed the work we are undertaking regarding data governance and management with many agencies including Stats NZ, the Treasury, the Social Wellbeing Agency, the Ministry of Business, Innovation and Employment (MBIE), ACC, and the Ministry of Social Development (MSD). Earlier in 2021, we met with the representatives from the Chilean Government regarding our implementation of the Charter, along with MSD and MBIE.

Mandated data standards

14. Inland Revenue has committed to adopting the data standards mandated by the GCDS. The intent of the current standards is to ensure that agencies exchange key data in the same format, for example all agencies using the same format for dates.

15. We have contributed significantly to the development of the three currently mandated standards for an individual's date of birth, name and street address. These standards are being incorporated into the guidance we give to the agencies that provide us with data to analyse.

16. We are working closely with Stats NZ on the development of further data standards, and have invited Stats NZ to be members of our data governance groups. Together, we are developing a plan to further improve the way our two agencies work together.

17. Importantly, the outcomes and lessons learned from this plan will be shared with other agencies, to contribute to the broader advancement of data management in the wider public sector.

Security and privacy

18. Everyone who works for Inland Revenue has a legislative obligation to protect the integrity of the revenue system and keep sensitive revenue information confidential. It is an obligation we take very seriously.

19. We have strict rules and restrictions on integrity matters such as staff accessing customer information. Routine monitoring is in place to check for potential wrong-doing and ensure that incidents are investigated properly and fairly.

Privacy

20. We routinely conduct privacy impact assessments whenever we are considering the use of new data and information, and when we use existing data in new ways. The assessment sets out what data we are using and why, how we will store it, who will have access to it, and the controls we have in place.

21. Each year, we complete an assessment of privacy capability using the Government Chief Privacy Officer's (GCPO) privacy assessment framework. For the year ending 30 June 2021, we met four of the five core expectations, and all but one of the other 11 elements. We were assessed as having strong maturity in culture, breach and incident management, and policies and process.

22. We report to the Office of the Privacy Commissioner each year on the information we share under information matching and approved information-sharing agreements (AISAs). We have AISAs with the Ministry of Social Development (MSD), the Police, the Department of Internal Affairs (DIA), and the NZ Gang Intelligence Centre. We also liaise with the Office of the Privacy Commissioner on any new policy or process that may impact on privacy.

23. All privacy breaches are reported. For the year ending 30 June 2021, 124 breaches were reported. The number of breaches has been stable. The majority of breaches are minimal or minor, using the GCPO breach-reporting matrix, as only a small number of people were impacted and there was little or no indication of systemic issues.

Security

24. Threats are occurring more frequently and are increasingly sophisticated. Tax agencies are seen as an attractive target, given the data we hold and our large volumes and value of financial transactions. We have plans in place should a threat materialise, so we are well prepared.

25. Although we have very good systems and processes in place to prevent privacy and security breaches, it is not possible to fully mitigate against the possibility of human error.

Assurance

26. Our internal assurance team have a risk-based assurance programme that includes assessing whether our cybersecurity practices and information matching agreements comply with the rules set by the Office of the Privacy Commissioner. The plan is updated every year and endorsed by the Risk and Assurance Committee, which includes external members and provides independent advice to the Commissioner. This year our focus areas are the AISAs we have with MSD, Police and DIA and the information matching we undertake with Customs.

27. The audit by Archives New Zealand (Archives NZ) of public office record-keeping aims to provide a point-in-time view of core information management practices, identifying strengths, and where there might be opportunities for improvement. The audit is based on Archives NZ's Information Management Maturity Assessment framework, which consists of eight categories including governance, self-monitoring, capability, creation, management, storage and access.

28. We are scheduled to be part of Archives NZ's 2021–22 audit programme. We have yet to receive formal notification but, to prepare for the audit, we recently conducted our annual Archives NZ self-assessment survey. This will be used as input into the audit.

29. The audit process involves four stages: pre-audit, in-audit (onsite), post-audit and follow-up. Our people will be involved in the pre-audit for onsite activities. The onsite audit will include interviews and focus groups with a range of our people, including me as the Executive Sponsor, information management specialist staff and technology staff.

Methods of sharing information

Who we share information with

30. We exchange information with a wide range of agencies including MSD, MBIE, Stats NZ, Customs, DIA, the Police, and the Treasury. These arrangements are well documented and have a clearly defined purpose.

How information is shared

31. Information is exchanged in a variety of ways, from storage devices, such as iron keys, to application programming interfaces (APIs). Given the number and diversity of agencies we share information with, we accommodate what others can work with. This means that we do use channels and technology that we would prefer to move on from; however, our approach is to minimise the burden on the agencies we exchange information with.

32. APIs, or gateway services, enable direct machine-to-machine interactions between systems. They are a fast, high-volume, secure channel we are increasingly making available to both public and private sector organisations. APIs enable us to define what data we will share and then make it available for others to access when they need to. For example, to support the administration of the wage subsidy, introduced in response to COVID-19, we enabled MSD staff to access the information they needed directly from our systems, with the appropriate security permissions to verify applications, rather than having to contact us.

33. The IRD number validation API is being used by banks, KiwiSaver providers, tax agents and payroll providers to ensure that their clients are correctly identified, thus improving data quality and reducing error rates. This API is also available to other agencies to use, such as MSD, when they are ready to adopt it.

34. Protecting information is complex, due to the possibility of human error and the many different mechanisms for sharing that are available. All information-sharing activities undergo formal review, testing and acceptance.

Digital identity framework

35. Inland Revenue is actively engaged in advancing the development of the digital identity trust framework being led by DIA. We have considered what our participation could be in terms of digital identity, and have provided example cases to DIA. We are strong advocates for the digital economy, and of business and support initiatives that will advance the digital economy.

Working with Stats NZ

36. Inland Revenue participates in many forums led by Stats NZ or the GCDS, including:

- Digital Government Leadership Group (DGLG)
- Information Group
- Strategic Advice on Integrated Data (SAID), and
- IDI Investment Advisory Group.

37. Inland Revenue has contributed to many initiatives led by Stats NZ or the GCDS, including:

- Algorithm Charter for Aotearoa New Zealand
- Data Investment Plan
- Māori Data Governance
- Statistics Bill, and
- New Zealand Government Data Strategy and Roadmap.

Continuing to develop our data management and use practices

38. We are continuing to implement our data and information governance programme. Two immediate areas of focus are improving how we record and report collection and sharing of data, and building the capabilities of our people.

Recording and reporting

39. All the data and information we collect and share, both inside Inland Revenue and with third parties such as other agencies, will be subject to governance.

- When we propose to collect and/or share data and information, we will notify all potentially involved parties.
- A comprehensive due diligence checklist will guide the involved parties through the process. The checklist requires an ethics assessment to be conducted at the same time a privacy impact assessment is considered.
- The information collected will be recorded in a data and information supply register. Our people will be able to access the register at any time to know what is being collected and shared, who to contact, with links to all relevant material (for example the privacy impact and ethics assessments).
- All collection and sharing of information will require the endorsement of the Data and Information Supply Working Group.
- The Data and Information Supply Working Group and involved parties will be notified of milestones and events such as proposal, implementation, and eventual decommissioning.
- Information will be recorded and reported to support performance measurement, risk management, and assurance activities.

Building the capabilities of our people

40. The data and information environment Inland Revenue operates within is complex, rapidly changing, and can be challenging for our people. Legislation is clear about our data and information collection and disclosure powers, and clear how data and information held by the Commissioner will be handled and safeguarded. However, legislation is not exhaustive and situations arise where interpretation of legislation and individual judgement will vary. In these situations our people, no matter their role, need understanding and guidance to decide what course of action is right.

41. Policies, strategies, principles, frameworks, and standards are necessary, but alone they are insufficient. We are developing straightforward guidance for our people, along with supporting tools to support our people to do what's right with data and information. We are:

- Producing practical guidance that brings together a range of material including our Code of Conduct and our data and information policy to make clear to our people what "doing the right thing" looks like when working with data and information.

- Working with our people who manage and use data and information so they understand how to work with the governance groups.
- Building self-service tools so our people know what data and information is available to them and for what purposes, and what decisions have been made about data and information by governance groups.
- Simplifying engagement with governance groups so guidance and decisions can be more readily accessed.
- Establishing working groups focused on developing areas of practice that span Inland Revenue, for example: the Analytics and Research Working Group, and the Digital Service Working Group.

42. We will continue to share our experiences and what we are putting in place with other agencies and will continue to work closely with Stats NZ to contribute to advancing wider data outcomes.

Update on the Government Data Strategy and Roadmap

Update to Data and Information Governance Authority
10 May 2022

Doug Lambert, Data Strategy and Governance Lead

Not in scope

Update on the Government Data Strategy and Roadmap

The purpose of this discussion

- This discussion is about the **why** there is a Government Data Strategy and Roadmap, **what** that Strategy and Roadmap is, **how** that Roadmap is being delivered, and **who** is involved.
- This discussion is a pre-cursor to Stats NZ representatives and the DIGA discussing the Government Data Strategy and Roadmap further at the next DIGA meeting

Summary

- The first Government Data Strategy was published more than three years ago. The Government Chief Data Steward (filled by the Chief Executive of Stats NZ) system functional lead role had only recently been established. That strategy was aspirational and did not include a roadmap.
- During 2021, Stats NZ facilitated a number of hui to crowd-source an updated strategy and a roadmap. The resultant Government Data Strategy Roadmap is ambitious and with significant dependency on Māori and iwi to collaborate in a co-design manner, and for Stats NZ to be supported by several agencies, including Inland Revenue.
- Unfortunately, COVID impacted the availability of many people, including Māori and iwi representatives, and there have been resourcing challenges.
- Toward the end of Year One of the Roadmap it is being assertively reviewed by the cross-agency Information Group to produce a more reliably deliverable tranche of change.

Background on Government data strategies to date

The following is largely an excerpt from online material introducing the current Government Data Strategy and Roadmap:

The first Data Strategy was published more than three years ago. The ambitions of that strategy were:

- making the right data available
- building capability and good practice
- growing effective partnerships
- implementing open and transparent practices.

However, the data landscape rapidly changed. Since the first Data Strategy and Roadmap was published, agencies have been working on a 10-year Data Investment Plan. Stats NZ and the Data Iwi Leaders Group have signed the Mana Ōrite Relationship Agreement to realise iwi data aspirations. A range of initiatives have been developed to increase trust in the data system, including the Data Protection and Use Policy and the Algorithm Charter for Aotearoa NZ. In addition, the COVID-19 pandemic has completely changed the context in which data is collected, managed, and used.

Despite the progress made since 2018, there were some fundamental system issues yet to be resolved:

- data about and for some important topics and communities does not exist
- settings to realise the rights and interests of Māori and iwi do not exist
- many agencies lack capability to take advantage of the power of data
- it is difficult to retrieve and re-use data across the system.

Throughout the process of refreshing the strategy, key stakeholders have participated in workshops and interviews. These stakeholders came from central government, local government, private businesses and non-government organisations (NGO's). They have tested the Government Strategy and Roadmap's direction to ensure it is future-focused, robust, practical, and able to be adapted to different data sources and uses.

The Strategy

The Government Data Strategy and Roadmap will be brought to life through principles adapted from the Data Protection and Use Policy, the Mana Ōrite relationship agreement between Stats NZ and the Data Iwi Leaders Group, and the Public Service Act.

Manaakitanga: Respect and uphold the mana and dignity of the people, whānau, communities, and groups who share their data and information.

Mana Whakahaere: Empower people by giving them a choice and enabling their access to, and the use of, their data and information.

Kaitiakitanga: A shared culture of respect, guardianship, care, and protection for data as a strategic and valued resource, recognising that for some Māori, Māori data is a taonga and iwi-Māori are kaitiaki over their taonga.

Rangatiratanga: Leadership that focuses on common purpose whilst also respecting the autonomy and independence of individuals, groups and agencies.

Whanaungatanga: Strong transparent relationships through respect, integrity, empathy, and commitment to the kaupapa.

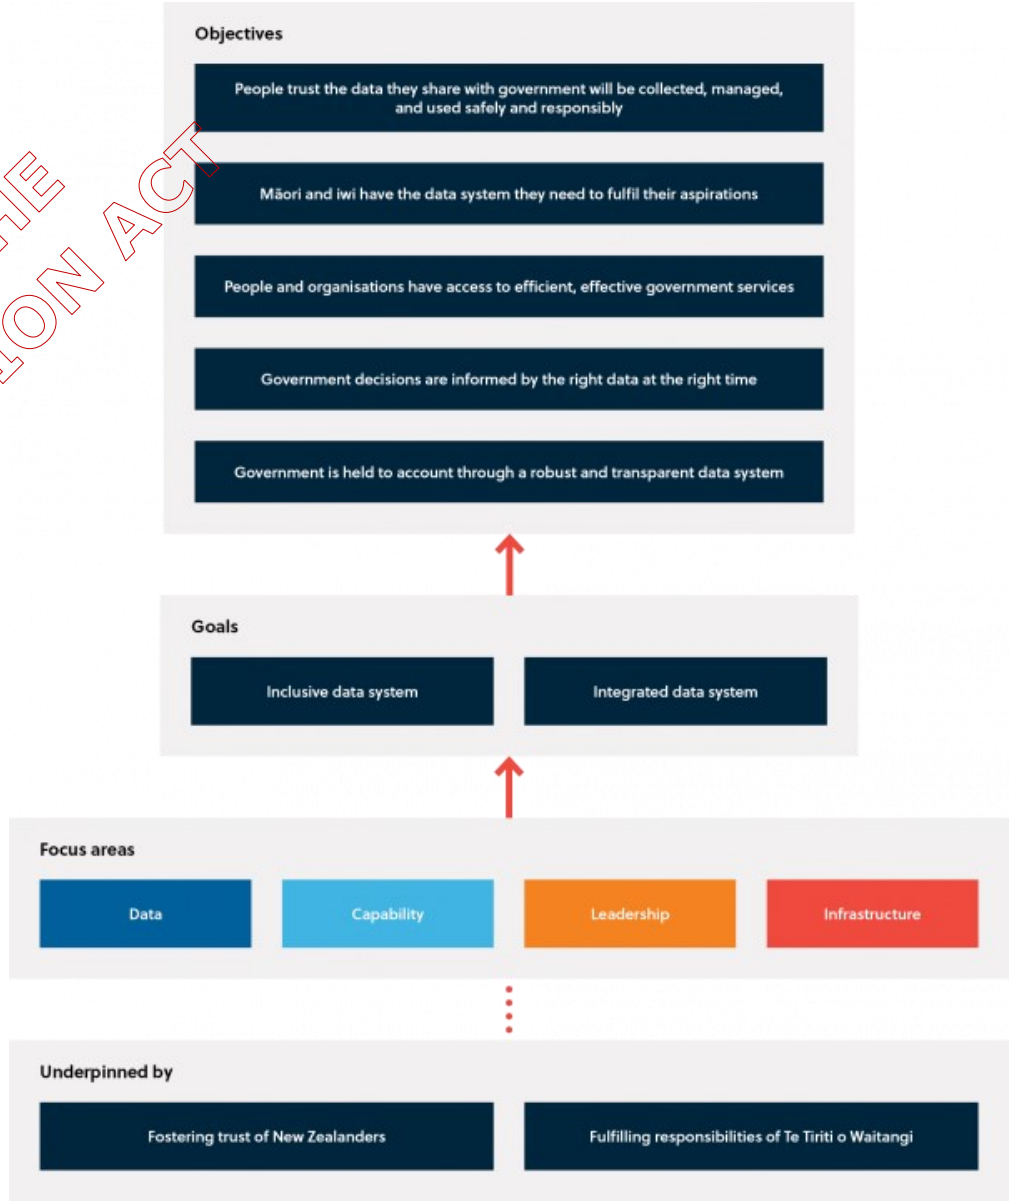

Not in scope

The Roadmap

The Roadmap is ambitious and has significant dependency on collaboration with Māori and iwi. Risks have materialised as issues and toward the end of Year One the cross-agency Information Group is re-shaping the Roadmap. IR is contributing to the highlighted initiatives in varying degrees.

	Year One (2021-2022): Foundations	Year Two (2022-2023): Investment	Year Three onwards (2023-2024): Maturity
Data	Data Investment Plan	Update the Data Investment Plan	Establish quality stamps on government data releases
	Open Data Charter Implementation Plan	Published data quality framework	Establish an easy-to-use mechanism for anyone to identify a data need
	Initial delivery of iwi affiliation data	Māori data stocktake	
Capability	Grow data capability through partnership with iwi and Māori	Develop Analytics and Research in Government (ARG) Graduate Programme pilot	Roll out initiatives based on the data capability framework
	Review Algorithm Charter and offer support to the system	Create a micro-credential for data ethics	Develop a data profession involving Communities of Practice, training, evaluation, and monitoring
	Provide implementation support for using Ngā Tikanga Paihere	Ensure consistent job families and descriptions are used across the system	
	Design and deliver a data system maturity assessment	Review data capability framework	
	Implement Government Chief Data Steward agency partnering model		
	Continue Analytics and Research in Government (ARG) Intern Programme		
Leadership	Revise system governance including advisory gaps	Implement changes aligned to the new Data and Statistics Bill	Implement the system settings for trust and data driven technologies
	Introduce Data and Statistics Act	Co-develop guidance for Māori Data Strategy with Māori	
	Implement new suite of trust proposals in support of the Digital Strategy	Establish inclusive data working group	
	Develop guidance for agencies on Te Ao Māori perspectives on cloud storage	Develop principles and protocols on the production and management of data system assets	
	Finalise and implement Māori Data Governance Model	Deliver next suite of mandated standards	
	Establish an updated approach to mandating and managing standards.		
	Develop Health of the Data System report		
	Develop the Trust Framework for Digitial Identity Services		
Infrastructure	Develop strategic response and seek investment for future of integrated data	Establish the Integrated Data Infrastructure (IDI) Commons	Scope system-wide data integration requirements
	Develop system architecture	Establish consistent approach to release protocols across the system	
	Confirm feasability of a joint property data source	Begin implementation of integrated data strategic response	
	Development of a joint data and analytics platform	Review and agree on consistent data sharing principles and protocols	
	Review Data Lab access requirements	Implement against agreed system architecture	
	Further develop and maintain innovative Māori data platforms		
Key changes	Develop a system-wide plan for government investment in data content, capability, and infrastructure	Use the Data Investment Plan to inform The Budget and agency planning	Create a public mechanism to identify data needs and embed a quality stamp which is consistently used by agencies
	Develop a framework to benchmark agency maturity and develop Government Chief Data Steward (GCDS) services to lift agency capability	Align agency capability to maturity assessment findings and build capability using Government Chief Data Steward (GCDS) services	Build data capability to extend into collection, stewardship, and use of the data cycle
	Review system settings for trust and data driven technologies	Complete the next suite of mandated standards	Implement system settings for trust and data driven technologies
	Develop Health of the Data System report	Align system architecture change across agencies	
	Establish the system architecture required for a connected and inclusive data system	Begin implementation of integrated data initiatives should funding be secured	

Email from Craig Jones (StatsNZ) to Mike Cunnington (Deputy Commissioner) and Tina MacLean (Intelligence Leader – Data)

From: Craig Jones <s 9(2)(a) >
Sent: Thursday, July 28, 2022 1:33 PM
To: s 9(2)(a) ; s 9(2)(a) ; s 9(2)(a) s 9(2)(a) ; Mike Cunnington <s 9(2)(a) >; Steve Murray <s 9(2)(a) >; Tina MacLean <s 9(2)(a) >; s 9(2)(a) ; s 9(2)(a) ; s 9(2)(a) ; Simon Ross <s 9(2)(a) >; s 9(2)(a) ; s 9(2)(a) ; s 9(2)(a)
Cc: Wendy Hamilton <s 9(2)(a) >
Subject: Algorithm Charter review

External Email CAUTION: Please take **CARE** when opening any links or attachments.

Kia ora koutou

As I mentioned at the Analytics DCE/Dep Sec meeting on Wednesday, the independent review of the first year of operation of the Algorithm Charter for Aotearoa New Zealand, is being sent to the office of the Minister of Statistics this week.

Following this, the review will be published on *data.govt.nz* next week. I have attached a copy of the review to this email, for easy reference.

As discussed on Wednesday, I agree with the findings and recommendations in the review. It found that there is strong ('almost universal') support for the Charter across government agencies and subject matter experts. It also found that agencies require additional guidance and assistance if we are to realise the required shifts in the ethical use of algorithms necessary to fully and successfully implement the Charter.

I look forward to developing, with you, actions in response to the recommendations. That includes working on building data ethics capability across the public sector and understanding any regulatory gaps that need to be filled in order to build a trustworthy data ecosystem. It also includes working across agencies to reduce any duplication and make best use of our collaborative resources, and an increased focus on the importance of Māori data governance and engagement with Treaty partners. Alongside this we have the potential to increase the visibility and public awareness of the use of algorithms and the steps we are taking to provide strong oversight and appropriate use, and to encourage public participation in the discussion.

We'll talk more at our regular meetings, and we'll share our plans for the approach we hope to take.

Please do not hesitate to contact me if you have any questions. Thank you, and your teams for your support to this mahi – it really is appreciated.

Ngā mihi nui
Craig

Dr Craig Jones
Kaimātai Tauanga Kāwanatanga Tuarua | Deputy Government Statistician
Pouārahi Tuarua - Kaiārahi Pūnaha Raraunga | Deputy Chief Executive - Data System Leadership
Tatauranga Aotearoa | Stats NZ
DDI +s 9(2)(a) | **M** +s 9(2)(a) | stats.govt.nz

About Aotearoa, for Aotearoa

Data that improves lives today and for generations to come

Email from Mary Craig to IR staff

From: Mary Craig
To: All IR Sites
Subject: Use of Artificial Intelligence (AI) products at IR
Date: Tuesday, 21 March 2023 1:20:25 pm
Attachments: image001.jpg
image002.png
image003.png

[UNCLASSIFIED]

Kia ora

Artificial Intelligence (AI) products or services like Chat GPT have recently become more accessible and have been a hot topic in the news and across social media.

IR is currently considering how we should integrate and/or use these products or services.

While IR works through our use of these solutions, please be aware that these products and services are not currently approved for business use. As a result, they should not be accessed from IR devices and workplaces. As always, we need to keep our confidentiality obligations under Tax Administration Act, Code of Conduct and other policies in mind. Our Security team will be keeping an eye out for any use of these products and services from IR devices and workplaces.

If you have any concerns, believe you have a valid business need to use these products or services please contact InformationSecurity@ird.govt.nz.

Ngā mihi

s 9(2)(a)

Mary Craig (she/her)
Deputy Commissioner – Enterprise Design & Integrity
Inland Revenue | Freyberg Building | 20 Aitken Street | PO Box 2198 | Wellington 6140
Ext: s 9(2)(a) | ☎ s 9(2)(a) | ☎ s 9(2)(a) | ✉ s 9(2)(a)

s 9(2)(a)

s 9(2)(a)

Mike Cunnington (he/him)
Deputy Commissioner - Enterprise Services

Not in scope

Excerpt from Artificial Intelligence Oversight Group: Minutes

AI Oversight Group

Inland Revenue
Te Tari Taake

Meeting Minutes for 7th August 2023, 1.30 pm – 2.30 pm.

Voting members	Brijesh John (Chair), Anil Srinivasa, Cate Robertson, Daniel Blank, Jay Harris, Jesse Thwaites, Malcolm Breadmore, Patrick O'Doherty, Tina MacLean
Standing members	Graham Poppelwell, Conrad Bace, Vanessa Johnson, Virginia Flaus
Attendees	Tanya Williams
Apologies	Underlined above

1. Welcome and Karakia

The Chair opened the meeting with a karakia.

2. Approval of previous meeting's minutes

The Oversight Group **approved** the minutes from the 17 July 2023 meeting.

3. All of Government developments / IR response

Presenter: Brijesh John, Domain Lead, Technology Architecture

Discussed was the usefulness of the Generative Artificial Intelligence guidance, issued for use across the NZ Public Service, for the Oversight Group's responsibilities. The members considered that the guidance is good, common sense and there were no surprises within the guidance. IR should look at adopting the guidance alongside other government departments.

Specific areas within the guidance identified as needing expansion for IR's needs were:
Prevent AI from being used as shadow IT. For instance, what of low code tools such as PowerApps, what would be our position? As PowerApps is not fully in a managed environment and has a dispersed user group within IR.

The Oversight Group discussed that for any AI technology being investigated for IR's use, there should be a solid use case for progressing an AI tool review, so that IR understands the reasoning and consequences of what we are doing and the wider downstream system implications. This need doesn't differ from other information technologies. The use case perspective can then be built upon with the controls IR goes on to establish, such as testing and checking for bias. Because of the wave of technology out there already and continuing to evolve, strengthening our controls is one of the best things we can do in the AI technology space.

A further consideration not defined in the released guideline is how we ensure that we understand AI and how it fits into our wider technology – the interaction of the AI that IR implements

interacting with each other and with other IR system technologies, and the potential to produce adverse results. Again controls, especially security controls, over AI will be key.

The guideline doesn't provide the next level of detail for capability identification and uplift. There is potential here to build upon or augment our capabilities through automation to help with our capacity efficiencies and in a reduction in human error.

On Ethics concerning AI decisions, IR needs to evolve its operating environment in order to be better across and more cognisant of the biases that can occur with data analytics. The conversations that currently occur are in pockets, so we need to be broadening the spaces where this awareness of and conversations about data ethics and biases do occur in IR. While high level information will be published, a data ethics approach and a data and ethics standard, the really important controls that sit below the standard (i.e. high level controls) nothing has been developed for those layers yet.

The Chair requested members to discuss the following points via email or Teams chat:

- Should our AI policy fully focus on generative AI?
- What level of involvement is needed with the contact points mentioned in the guidance?
- With the document not having the detail we would be looking for to post on our channels, how closely should we follow this guidance in relation to the use of AI for people information? What further messaging should be there on classifications of sensitive and in confidence?

The Oversight Group:

- **Discussed** the applicability of the Generative Artificial Intelligence guidance in an Inland Revenue context.
- **To provide feedback** to the questions above posed by the Chair.

Not in scope

Using Artificial Intelligence at IR

Makayla Stewart
Change Analyst (L2)

Using the definitions and concepts outlined in What is Artificial Intelligence, lets look at how we use AI, what the benefits are and what's coming up.

Its important that we, as an organisation be deliberate in the decisions about which systems and tools we use. Only approved business tools can and should be used for their intended purpose.

On 21 March 2023, our Deputy Commissioner of Enterprise Design & Integrity, Mary Craig, sent out an email to all of IR about Use of Artificial Intelligence (AI) products at IR which outlined that tools such as Chat GPT have not been approved for use.

Since March we have continued to develop guidance for using Artificial Intelligence at IR and are now in a position to provide you with more information about how to treat this technology. It is important that we understand:

- Which tools have been approved for business use
- What the purpose these tools have been approved for
- When these tools should or shouldn't be used

- What data is appropriate for us to input into these tools.

To ensure we approach the use of AI at IR both safely, securely and in line with the requirements of us under the Tax Administration Act 1994, use case guidelines and a staff use policy have been developed to support you.

Guidelines for using Artificial Intelligence

You can find the full version of the AI use case guidelines here AI use case guidelines

Our use case guidelines have been developed using:

- IRs classification labels as the basis (see Information classification and handling for more on classification).
- The definitions for: Business rules, Machine learning, AI integrated solutions and Generative AI / Large Language Model (see Concepts within AI here What is Artificial Intelligence for more information).

Only business tools that have been approved can be used on work devices. Full details of approved applications that have an AI component can be found here Artificial Intelligence Application Register.

The use of Generative AI solutions such as Chat-GPT is still under assessment and IR staff are asked to continue to wait for additional guidance. Some pilot options have been assessed and are progressing; we hope to share more on these in the near future.

If you are unsure whether the actions you are intending to take are within the guidelines, check with your leader or contact AI@ird.govt.nz.

Policy for using Artificial Intelligence

You can find the full version of the AI staff use policy here Artificial Intelligence (AI) staff use policy

This policy applies to all of IR and sets out our approach to safely and securely look at how we can use AI to deliver more effective and efficient services. Our policy has been developed using:

- The definitions for: Business rules, Machine learning, AI integrated solutions and Generative AI / Large Language Model (see Concepts within AI here What is Artificial Intelligence for more information).
- Our obligations under the Inland Revenue Acts, Privacy Act, Te Tiriti o Waitangi, and The Algorithm Charter for Aotearoa New Zealand

As IR people we are responsible for adhering to this policy and reporting any abuse, misuse or non-compliance with the policy to our Information Security Team.

Have an idea for how we should use AI?

Check out how to Get involved with the future of using AI at IR.

Related Resources

AI use case guidelines

Artificial Intelligence Staff Use Policy

Inland Revenue
Te Tari Taake

Publication of IR's transparency obligations under Algorithm Charter for Aotearoa NZ

Artificial Intelligence Working Group

31 January 2024

Graham Poppelwell

Purpose

The purpose of this paper is to present:

- The key insights of the December 2021 review of the Algorithm Charter for Aotearoa NZ (specifically transparency and engagement)
- A brief overview of IR's transparency obligations under the Algorithm Charter for Aotearoa NZ
- Draft content to be published on IR's external website

Not in scope

Recommendations

It is recommended the AI Working Group:

- **Discuss** the transparency obligations of IR under the Algorithm Charter of Aotearoa NZ
- **Review and provide feedback** to proposed external website content

Background

Governance review of IR's Algorithm Charter ('the Charter') commitments

The Charter was independently reviewed in December 2021 (one year after public release). The review report includes an implementation plan (in 3 phases) to assist agencies in meeting Charter expectations. Details of the review and report can be found here.

The Cross Agency Community of Practice recently held discussions, asking how agencies were progressing with phase 2 - Transparency and Engagement. As part of this phase, agencies should consider improved transparency of the Charter, and practices that give effect to Charter commitments. Most agencies have no published view, including IR.

The Charter Implementation plan – Phase 2: Transparency and Engagement

IR is listed as a signatory on the Charter's website and the Charter was mentioned in the 2023 IR Annual Report (see next slide). There was also a 2018 story published Inland Revenue's Tony Morris – let's bust the myth about robot tax investigators (ird.govt.nz)

Draft website content, aligning to IR's Governance programme for Artificial Intelligence, is being prepared for publishing.

Annual report 2023 extract (page 196)

Algorithm Charter for Aotearoa New Zealand

Te Tari Taake, Inland Revenue adopted the Algorithm Charter (the charter) for Aotearoa New Zealand[1] in July 2020. The charter includes 6 commitments intended to shape government agency behaviours in ways that minimise risk to people and maximise the benefits algorithms can provide.

Te Tari Taake, Inland Revenue has a governance programme over our data and information, which includes how we work with the charter commitments and assessing the risk when we use algorithms. Legislation provides us and our customers with specific rights regarding information. We have legal, ethical and business responsibilities to ensure that the information we are a steward of is protected, maintained and developed for the benefit of more than just ourselves.

Background - Phase 2: Transparency & Engagement

From the implementation plan:

Phase 2 – Transparency and engagement

Phase 2 will focus on authentic, high-quality engagement to provide greater visibility and transparency.

Key areas of work in phase 2 may include:

- marketing and communicating to ensure improved transparency of the charter and the agency practices that give effect to the charter commitments
- providing practical advice on engaging with iwi and Māori on data and algorithms to implement the partnership commitment in the charter
- investigating novel forms of citizen participation.

From digital.govt.nz summary:

Key Take aways on managing the risks of GenAI to the public service:

- Be open and transparent in terms of what GenAI is being used for and why. Ensure processes are in place to respond to citizen requests to access and correct information.
- Citizens are concerned about ethical use of GenAI, and the public has expectations about how GenAI is being used, particularly for the public service.
- The public service is held to a higher standard; so, consider your social licence and how to assure transparency, accountability, and fairness in how your agency is using and applying GenAI — whether directly or as part of a wider technology solution.

Background: Algorithm Charter Committments

The charter contains six core commitments:

1. **Transparency** – Maintaining transparency by clearly explaining how decisions are informed by algorithms
2. **Partnership** – Embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi
3. **People** – Identifying and actively engaging with people, communities and groups who have an interest in algorithms, and consulting with those impacted by their use
4. **Data** – Making sure data is fit for purpose by understanding its limitations and identifying and managing bias
5. **Privacy, ethics and human rights** – Ensuring that privacy, ethics and human rights are safeguarded by regularly peer reviewing algorithms to assess for unintended consequences and acting on this information
6. **Human oversight** – Retaining human oversight by:
   - Nominating a point of contact for public inquiries about algorithms
   - Providing a channel for challenging or appealing of decisions informed by algorithms
   - Clearly explaining the role of humans in decisions informed by algorithms.

Algorithm Charter Commitments to transparency:

Maintain transparency by clearly explaining how decisions are informed by algorithms.

This may include:

- plain English documentation of the algorithm
- making information about the data and processes available (unless a lawful restriction prevents this)
- publishing information about how data are collected, secured and stored.

Not in scope

Background: December 2021 Review

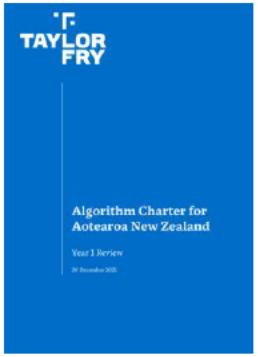

Transparency – findings:

- Algorithm documentation is often too complex for the layperson. This means transparency requires that a range of products are released – including 'plain English' documentation. Converting documentation into 'plain English' can be challenging.
- Some agencies have conflicts because being transparent could indirectly affect the operation and effectiveness of the algorithms e.g. algorithms that support identifying criminal activity.

Responses by other agencies to the review

- Some agencies have lists of algorithms published online. In most cases, this is a new practice and would not have happened without the influence of the Charter.
- Some agencies have performed stocktakes and risk assessments and published the output.
- Some agencies have plans for further publishing of information in the future.
- Some agencies publish code online on GitHub.
- Some agencies have published methodological reports online in publicly available academic forums – these explicitly referred to the principles of the charter.
- Commissioning the work - Responding to the Algorithm Assessment Report - data.govt.nz

Examples of what's published by other agencies

justice.govt.nz

Algorithm use in the Ministry of Justice

Te Tāhū o te Ture - the Ministry of Justice is committed to transparency and accountability in our use of operational algorithms as set out in the Algorithm Charter. Our Ministry currently uses three operational algorithms within the Collection Services Business Unit[1]:

- Task Generator
- Automatic Attachment Orders
- Algorithm to identify new contact details for debtors.

None of these algorithms incorporate machine learning capability.

[1] *This unit is responsible for the collections of court-imposed fines and reparation, infringements issued by third parties such as New Zealand Police and local authorities, legal aid debt and civil enforcement.*

Te Whatu Ora
Health New Zealand

Actions to be taken by Te Whatu Ora

There are no algorithms in use at Te Whatu Ora at the moment, there are significant opportunities in health care in the safe usage of algorithms, machine learning and artificial intelligence.

As kaitiaki of the health system, we're focussed on ensuring that high quality, effective and safe tools are used across the system to enable quality health care. Beyond the review of our use of algorithms, this guide is a significant contribution to enable safe algorithm use across the health sector into the future.

As part of our ongoing work, we'll embed the principles of the Algorithm Charter into the development and maintenance of any algorithms we produce. We expect to review our processes every 12 months, and update our reporting accordingly.

MINISTRY OF BUSINESS, INNOVATION & EMPLOYMENT
HĪKINA WHAKATUTUKI

Algorithms at MBIE

As part of our commitment to the charter, we are implementing an MBIE Algorithm Use Policy. The purpose of this policy is to enable accountability for decisions at MBIE, increase transparency about our algorithm use, and strengthen support systems for our staff working with algorithms.

Algorithms are used at MBIE for one or more of the following reasons:

- **Being more efficient**
  Computers can process large quantities of data quickly, much faster than our staff can manually. We may use algorithms to keep up with large application numbers so our staff can deal with more complex work and respond to the public quicker. For example, the Immigration New Zealand Advance Passenger Processing (APP) which performs validation matching and screens for risks at the border.
- **Limiting human bias from a process**
  As humans, we bring our own subconscious biases into all decision we make. Through automation we can limit and reduce the impact of these biases and provide more consistent decisions and responses. For example, algorithms that assist with recruitment, finance and procurement allocation.
- **Predicting future behaviour**
  Our predictions for the future are never guaranteed, but are still useful tools for planning, weighing up different options and taking proactive action if risky situations arise. Our algorithms help us understand what our future may look like so that we can stay prepared. For example, using predictive analytics to estimate business health.

We will apply the charter where the use of algorithms can significantly impact the wellbeing of people, or there is a high likelihood many people will suffer an unintended adverse impact. We will assess our algorithms for risk and apply the charter accordingly.

Data Science Review Board

We have established a Data Science Review Board to provide MBIE with strategic and practical direction, guidance and leadership for matters relating to data science and algorithm governance.

The Board is made up of MBIE staff and external members, who provide advice on algorithm use and development, and ensure that algorithms have undergone a robust review by quality experts and adhere to accepted standards.

Not in scope

To discuss: suggested content for IR's external website

Inland Revenue's algorithm governance follows All of Government directives.

- IR signed the **Algorithm Charter** for Aotearoa NZ in July 2020.
- IR aligns to the guidance provided on algorithms by the System Leaders Guidance and Office of the Privacy Commissioner. We are also participating in the Department of Internal Affairs led All of Government Programme of Work.
- IR is taking a cautious approach to algorithm use, considering impacts on our customers and Aotearoa in line with the Charter commitments and Government system led guidelines.
- IR has established governance bodies, including an Artificial Intelligence Oversight Group and Artificial Intelligence Working Group with representatives across the business, to provide oversight and direction for algorithm use at IR.

If you have any questions regarding IR's use of algorithms, please direct these to our Official information requests.

The above is minimum information, other suggested wordings are:

- It is a priority for IR to ensure that algorithms are 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.

- IR has published an AI Staff Use Policy and Artificial Intelligence use case guidelines. This policy sets out IR's approach to safely and securely look to use algorithms in the workplace, to help make good decisions and deliver services that are more effective and efficient.

Proposed publishing landing page – in the collection of pages that is About us (ird.govt.nz)

➢ An MS Word version of the drafted wording is here if you wish to provide written feedback to this content.

Inland Revenue
Te Tari Taake