

18 August 2026

[REDACTED]
[REDACTED]

Dear [REDACTED]

Thank you for your request made under the Official Information Act 1982 (OIA), received on 21 July 2026. Following my previous response to your OIA dated 5 February 2025 (available here: <https://www.ird.govt.nz/-/media/project/ir/home/documents/oia-responses/march-2025/inland-revenues-definition-of-maori-data-and-maori-data-sovereignty-or-governance-strategy.pdf?modified=20260115020515>) you have requested any updates to the following seven questions.

1. *How does your organisation define Māori data as opposed to non-Māori data?*
2. *Does your organisation have Māori Data Sovereignty and or a Māori Data Governance policy or strategy? If yes, I request a copy.*
3. *I also request a copy of your organisation Data Governance strategy/policy/policies?*
4. *Has your organisation had with any success or no/limited success, implementation of any Māori Data Sovereignty Principles or Māori Data Governance? If yes, please provide details of the implementation and how you measured its success.*
5. *How many fte are allocated to Māori Data practices in your organisation?*
6. *What country/countries are the majority of your organisation's data stored?*
7. *Which Cloud Provider(s) do you use?*

You also requested the following:

8. *A list of all artificial intelligence systems, machine learning models, and automated decision-making (ADM) tools currently in use, in procurement, or in pilot within the agency, including the name, purpose, vendor/developer, and date of deployment of each.*
9. *A list of all generative AI tools approved for staff use, and copies of any internal policies or guidance governing staff use of generative AI.*
10. *Copies of all Te Tiriti o Waitangi analyses, assessments, or advice prepared in relation to any AI or ADM system in use or in development, including the identity of who prepared them.*
11. *Records of any engagement, consultation, or partnership with iwi, hapū, or Māori organisations regarding the design, procurement, deployment, or governance of AI or ADM systems, including dates, parties, and outcomes.*
12. *Any documents identifying whether Māori data (including data about Māori individuals, whānau, hapū, iwi, or mātauranga Māori) is used to train, fine-tune, test, or operate any AI system, and any governance arrangements applying to that data.*

13. *Copies of any advice, briefings, or policies referencing Māori data sovereignty, Māori Data Governance, or WAI 2522 in relation to the agency's AI or ADM systems.*
14. *Copies of all Privacy Impact Assessments, Algorithmic Impact Assessments, human rights assessments, or equity analyses completed for any AI or ADM system, including any assessment of differential impacts on Māori.*

On 29 July, you clarified your request relating to question 8, and that you were requesting the list of umbrella IT systems which feature automated decision making, and not the individual tools within these systems.

Questions 1, 2, 4, 5, 6 & 7

I have no updates to the information provided in my previous response dated 6 March 2025. Therefore, your request for these updates is refused under section 18(e) of the OIA, as the information requested does not exist.

Question 3

In the time since my previous response to you, Inland Revenue has published a new version of our Data, Information and Knowledge Policy. This document was published on 1 November 2025 and is enclosed as **Appendix A**.

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

As required by section 9(1) of the OIA (in relation to information withheld under section 9(2)(a)), 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.

Questions 8 & 9

The table attached as **Appendix B** details all AI systems, machine learning models, and generative AI approved at Inland Revenue for staff use.

Automated decision-making currently operates in every significant enabling software application in use within Inland Revenue, based on preset business rules. This includes tax and social policy administration software (START), security operations, corporate services systems, enterprise service management and productivity tools. All business rules require human oversight from staff upon creation.

The significant enabling software applications in use at Inland Revenue are ServiceNow, Oracle, START, M365 Applications, Adobe, Snowflake, Genesys, SAP, Verint and Windows.

Question 10

Your request for copies of all Te Tiriti o Waitangi analyses, assessments, or advice prepared in relation to any AI or ADM system in use or in development, including the identity of who prepared them, is refused under section 18(e) of the OIA, as the information requested does not exist.

Although no documentation of the kind described in your request exists, Inland Revenue requires project leads to complete a preliminary risk assessment before progressing any generative or agentic AI-related development, proof of concept, pilot, or deployment activity. The risk assessment process requires project leads to identify initiative-specific risks and controls and to consider issues including the potential transfer of data outside New Zealand and any ethical implications of the proposed AI tool.

Question 11

Your request for records of any engagement, consultation, or partnership with iwi, hapū, or Māori organisations regarding the design, procurement, deployment, or governance of AI or ADM systems, including dates, parties, and outcomes is refused under section 18(e) of the OIA, as the information requested does not exist.

During Inland Revenue's Business Transformation programme (2017-2022), we regularly engaged with Māori stakeholders and iwi to inform them of upcoming changes but not specifically relating to automated decision-making systems within the newly deployed tax and social policy administration software, START.

Governance groups have been set up within Inland Revenue which provide oversight and direction for how we use algorithms and AI. These include the Artificial Intelligence Oversight Group and the Data and Information Group. The members of both 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 perspective into Inland Revenue's work.

Inland Revenue has a strategic advisor (Māori) who provides advice and strategic thinking to the Executive Leadership Team on issues of strategic importance to the Māori Crown relationship. They are a member of Inland Revenue's AI Oversight Group and provide verbal advice within meetings.

Inland Revenue also has an internal team named Te Kāhui Tūhono who provide guidance to contribute to and support the adoption of Māori indigenous approaches that are underpinned by Kaupapa Māori and Te Ao Māori principles and values. A representative of this team is a member of Inland Revenue's Data and Information Group.

Where appropriate, data scientists, AI initiative project teams or governance body members will consult members of Te Kāhui Tūhono and the strategic advisor (Māori) on issues relating to their work at hand and understanding and embedding a Te Ao Māori perspective.

Question 12

Your request for any documents identifying whether Māori data (including data about Māori individuals, whānau, hapū, iwi, or mātauranga Māori) is used to train, fine-tune, test, or operate any AI system, and any governance arrangements applying to that data, is refused under section 18(e) of the OIA, as the information requested does not exist.

Question 13

I am releasing, two documents within scope of your request. Inland Revenue's *Policy for Users of AI at IR* guide is provided to you, enclosed as **Appendix C** and an excerpt from document *Artificial Intelligence Direction and Strategy* enclosed as **Appendix D**.

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

Please note that the remaining contents from Appendix D are outside the scope of your request. This information has not been considered for release and has been withheld as 'not in scope'.

Question 14***Human rights assessments, equity analyses or assessments of differential impacts on Māori***

Inland Revenue has not completed any human rights assessments, equity analyses or assessments of differential impacts on Māori, for any AI or automated decision-making tools. Therefore, your request for a copy of these is refused under section 18(e) of the OIA, as the information requested does not exist.

Algorithmic Impact Assessments

The majority of algorithms that Inland Revenue 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 Algorithm Charter for Aotearoa and a full algorithm impact assessment has not been completed.

Where the Charter does apply to Inland Revenue'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. All activities undertaken as a consequence of these algorithms are subject to human oversight and human decision making. Full testing of all algorithms is undertaken; however, a formal algorithm impact assessment has not been completed for any tool as no tool has met the threshold to do so under the Charter.

Your request for a copy of all algorithmic impact assessments completed for any AI or automated decision-making system is refused under section 18(e) of the OIA, as the information requested does not exist.

Privacy Impact Assessments

There are three documents within scope of your request. I am releasing, the May 2026 document M365 Copilot: Privacy Impact Statement to you in full, enclosed as **Appendix E**.

The remaining two documents are refused under section 18(d) of the OIA, as the information is publicly available:

Item	Date	Document	Website address
1.	18 July 2023	Viva Topics: Privacy Threshold Assessment	2025-11-14-agency-approved-ai-tools-usage-guidelines-licensing-privacy-assessment.pdf
2.	7 February 2023	START Analytics: Privacy Impact Assessment	2025-11-14-agency-approved-ai-tools-usage-guidelines-licensing-privacy-assessment.pdf

Some information has been removed from the documents listed above, on the OIA grounds described in the documents.

Right of review

If you disagree with my decision on your OIA request, you have the right to ask the Ombudsman to investigate and review my decision under section 28(3) of the OIA. 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 ([ird.govt.nz](https://www.ird.govt.nz)) as this information may be of interest to other members of the public. This letter, with your personal details removed, may be published in its entirety. Publishing responses increases the availability of information to the public and is consistent with the OIA's purpose of enabling more effective participation in the making and administration of laws and policies and promoting the accountability of officials.

Thank you again for your request.

Yours sincerely



Jesse Thwaites
Enterprise Leader, Data & Information

Appendix B

Table 1: Artificial Intelligence in use at Inland Revenue

Name	Purpose	Product/ Project Status	Technology	Vendor	Deployment Date
ABBYY FineReader 16	Text recognition and document conversion tool, used to convert PDFs into excel.	In production	Optical Character Recognition, Machine Learning	ABBYY	November 2019
Ability to pay model - Individuals	Analyse the customer's circumstances and recommend the best action for that customer	In production	Machine Learning	FAST Enterprises LLC	1 August 2023
Ability to pay model - Non-Individuals	Analyse the customer's circumstances and recommend the best action for that customer	In production	Machine Learning	FAST Enterprises LLC	30 April 2024
AI Futurist	Enables querying and summarisation of content.	In production	Generative	Futures Platform	25 January 2025
Articulate 360 with AI	E-Learn development software	Pilot	Generative	Articulate Global Inc	September 2025
Assurity Intelligence	Test scenario generation	In production	Generative	Assurity Consulting Limited	November 2024 (proof of concept), April 2025 (pilot), February 2026 (production)
Automate production Contact Centre testing	Provides continuous monitoring of Contact Centre services and faster identification and diagnosis of production issues	In production	Machine Learning	Assurity Consulting Limited	15 April 2024

Bank Account Analysis Tool (BAAT)	Source, Disposition and asset accretion analysis	In production	Machine Learning	Microsoft – Excel (developed in house)	2012
Brandwatch	Media monitoring	In production	Machine Learning, Natural Language Processing	Falcon.io APS	July 2024
Canva AI	Assist in the creation of social media assets such as images, videos and banners	Proof of Concept	Generative	Canva	April 2026
CCH iKnowconnect	Legal and tax research queries and summarisation	Available	Generative	Wolters Kluwer	2 March 2026
Chainalysis Reactor	Blockchain Analysis	In production	Machine Learning	Chainalysis Inc	1 February 2019
Coveo	IR's public websites and internal intranet use Coveo, a machine-learning search platform that continuously improves results based on user behaviour.	In production	Machine Learning	Coveo	January 2021
Dragon Naturally Speaking	Screen reader	Available	Machine Learning	Nuance	2011
DDoS protections - AWS Shield - Cloudflare - F5	IR uses a range of tools to prevent our systems and services from unexpected outages due to network attacks.	In production	Machine Learning	Oracle, AWS, Azure, Cloudflare, F5 Inc	AWS Shield: September 2017 Cloudflare: May 2026 F5: June 2016
Email-user link prediction	Predict the similarity of email addresses and usernames such that IR can identify the probability that they are controlled / used by the same real-world person.	In production	Machine Learning	Snowflake	8 November 2021

Financial intelligence network detection	Links, matches and identifies multi-dimensional risks of users via operational and strategic visualisation	In production	Machine Learning	Snowflake	June 2025
Genesys Agent Assist	Creates summaries of conversations with contact centre agents for post-call notes	In production	Generative	Genesys and One NZ	May 2025
Graph Entity Resolution	Analyses and compares information held by IR to external datasets provided by third parties to determine if records are referencing the same entity.	In production	Machine Learning	Snowflake	December 2024
GST integrity model	A predictive model to assess the risk of GST returns.	In production	Machine Learning	FAST Enterprises LLC	April 2024
Inoreader – RSS feed	Summarisation and automated insight generation from open-source information from selected, publicly available sources.	Pilot	Generative	Innologica	June 2026
MAP	Federated search and trace domestic and offshore corporate entities to understand ownership, control, and cross-border structures	In production	Machine Learning	Snowflake	June 2025
MarianMT	Translating text for digital forensics	Available	Neural machine translation	Meta	20 October 2024
Microsoft 365 Copilot	Chat-style interface to both the web and internal information in Microsoft 365 applications, respecting user access permissions	Phased roll out	Generative	Microsoft	25 September 2025
Microsoft 365 Copilot Agents	Pre-built agents	Phased roll out	Agentic	Microsoft	25 September 2025

Microsoft Copilot Chat (Bing/Browser)	AI-powered chat service	Available	Generative	Microsoft	March 2024
Microsoft Copilot Studio	Create and deploy AI copilots and chatbots for business needs	Proof of Concept	AI Agent	Microsoft	March 2025
Microsoft Defender	Enterprise security solution integrated with Microsoft apps to monitor devices and identities for malicious activity.	In production	Machine Learning	Microsoft	2022
Microsoft Power BI	Data visualisation and analytics platform for building dashboards and reports	In production	Machine Learning	Microsoft	1 October 2018
Microsoft Purview	Integrated solutions for data governance, security, and regulatory compliance within Microsoft products. • Enabling Data Security Posture Management, • Copilot specific controls on content, Insider risk management and compliance manager enabled to help identify interactions in AI.	In development	Machine Learning, automation	Microsoft	November 2025
Microsoft SharePoint Advanced Management	Analyses site activity to detect unusual patterns.	In production	Machine Learning	Microsoft	26 September 2025
Microsoft SharePoint Sections with Copilot	Helps create structured sections for SharePoint pages.	In production	AI agent	Microsoft	October 2025

NET	Search and visualise relationships across Inland Revenue's customer graph to quickly understand how people, entities, and assets are connected	In production	Machine Learning	Snowflake	June 2025
Oracle Cloud	Enterprise Performance Management module. Predictive AI to assist with budgets, forecasting, and planning projections.	Proof of Concept	Machine Learning	Oracle	April 2026
Overdue income tax return RIT prediction	Predicts residual income tax (RIT) on overdue returns.	In production	Machine Learning	Snowflake	2016
Posit Connect	Supports the deployment of AI-powered data science solutions.	Available	Generative	Posit Software	July 2021
Power Automate	Low Code solution that supports automating tasks.	In production	Machine Learning	Microsoft	24 Feb 2022
Propensity to read letter or log-in to myIR	Supports IR in using the right channels to reach customers.	In production	Machine Learning	Snowflake	November 2023
Qualtrics	Analysis of survey information from customers.	In production	Machine Learning	Qualtrics	December 2020
Receipt, invoice, statement and tax/employer return review	Text recognition	In production	Optical Character Recognition Machine Learning	FAST Enterprises LLC	25 October 2005
Redwood Initiative stages 1-9	Transition to Oracle's Redwood UX to provide a modern, consistent experience across cloud applications (HCM, ERP, SCM, CX)	In production	Machine Learning	Oracle	June 2025 to present day (staged deployment)
Snowflake Cortex AI	Enables users to query data in natural language and get direct answers	In production	Generative	Snowflake	6 March 2026
START	Tax and Social Policy administration software	In production	Machine Learning	FAST Enterprises, LLC	2016

Streem	Media monitoring	In production	Machine Learning, Natural Language Processing	Streem NZ Limited	1 June 2025
Tableau Desktop	A data analytics and visualisation tool for analysing and graphing performance test results	In production	Machine Learning Natural Language Processing	Tableau Software	November 2023
TRM	Blockchain analysis	In production	Machine Learning	TRM Labs, Inc.	July 2025
Viva Engage	Delivers personalised content and community recommendations for internal Inland Revenue communities. Addition of AI was included with the deployment of M365 copilot.	In production	Machine Learning	Microsoft	25 September 2025
Westlaw NZ	Legal and tax research queries and summarisation	Available	Generative	Thomson Reuters	2 March 2026
Windows Hello for business	Secure authentication for IR-enabled devices	In production	Biometrics	Microsoft	October 2023
Z scaler	Detection and classification of web traffic and websites.	In production	Machine Learning	Z Scaler	2019
ZoomText	Screen magnification software designed to improve accessibility	Available	Optical Character Recognition	Vispero	2019

Kaupapahere Mōhiohio me te Raraunga

Data, Information and Knowledge Policy

This policy guides the safe access to, collection, storage, use and sharing of data, information and knowledge for staff at Inland Revenue (IR).

This includes access and use through Artificial Intelligence (AI).

Why we have this policy

The purpose of IR's Data, Information and Knowledge Policy is to:

- Detail the obligations for IR staff with access to IR data, information and knowledge.
- Promote a culture of stewardship, and safe and ethical access to, collection, storage, use and sharing of data, information and knowledge.
- Empower our people to pause and reflect on their responsibilities when accessing or handling IR's data, information and knowledge.
- To ensure IR's data, information and knowledge is visible, protected, accurate, properly managed and used in ways that enable IR to achieve Our Strategy.
- Support New Zealanders to have trust and confidence in our use of their data and information.

Context

IR is responsible for complying with legislative obligations, including those in the Inland Revenue Acts, the Public Records Act 2005, the Privacy Act 2020, the Official Information Act 1982, and the Public Service Act 2020.

With AI now widely available across IR, data, information and knowledge are more accessible than ever. It's essential that all IR staff understand and uphold our shared responsibilities when accessing and using this information.

Who this policy applies to

This policy applies to all IR staff, including permanent employees, fixed term employees, inter-agency secondees, consultants and contractors.

Your responsibility

As someone working for IR, you are responsible for understanding and complying with this policy.

Our policy principles

1. We all have responsibilities to protect and maintain New Zealand's data, information and knowledge for the benefit of all

We are all Whakahaere Haepapa Pūnaha (stewards) of the data, information and knowledge IR holds. Our responsibilities include:

- complying with all IR's contractual and statutory obligations,
- protecting the integrity of the tax system,
- applying standards and models and using IT systems to ensure consistent quality, supply and lifecycle management, and
- appropriately classifying information and data and handling it accordingly.

We also recognise that other organisations, communities and people, present and future, depend on our data, information and knowledge.

2. We only collect data and information that is required to deliver clear and evidenced business outcomes

- We will only collect data and information when it's necessary, permitted by law and where it relates to, or supports, a function or duty lawfully conferred on the Commissioner.
- To support this principle, the [Information Collection Framework](#) should be considered.

3. We use data, information and knowledge for the lawful purposes that it was collected for

- We only use information for the purposes it was collected for, unless it can be legally used or disclosed for other purposes, such as, linked with other datasets to yield insights.
- We have well-defined governance arrangements and decision-making accountability across data and information collection and use.

4. We only dispose of data and information in accordance with Retention and Disposal Authorities

- We retain data and information in accordance with our legal obligations and ensure they are properly recorded and accessible.
- We must maintain institutional memory - knowledge, records, and insights that support decision-making and continuity over time.
- Information and data are disposed of in accordance with Disposal Authorities detailed in IR's [Records Management Specification](#).

5. We maintain and improve the quality of our data, information and knowledge – as it is integral to delivering current and future IR services

- We prioritise quality maintenance, remediation and resolution across the enterprise as an 'always on' business capability.

- We recognise that quality management is essential for effectively managing strategic assets that support successful initiatives.
- We promote a data-driven culture where data is viewed as a valuable asset, decisions are based on data insights rather than intuition alone, and every employee understands the importance of data quality and its role in generating reliable information and actionable knowledge.

6. We have transparency of data and information-related decisions, processes and controls

Our governance, management and use of data and information is transparent and accessible, with comprehensive oversight and clear accountability. We act ethically and with integrity in delivering our business outcomes.

Transparency and accessibility includes:

- protecting Revenue Information and Sensitive Revenue Information,
- internal visibility of decisions, processes and controls,
- strong stewardship to support 'open government' by making official information available to the public where appropriate,
- cross-Agency collaboration by designing processes and systems that support interoperability and sharing of data, information, knowledge and services across government, and
- complying with all privacy, security, data, information and knowledge protection legislation.

7. Valuing diversity through Manaakitanga

We recognise that different people and communities place different value on data and information. We respect this by upholding the mana of individuals and groups, and by engaging in ways that reflect their needs and expectations. This includes:

- recognising and valuing the knowledge and experience held by our people,
- supporting our people in making appropriate ethical decisions regarding data, information and knowledge, and
- strengthening the Māori-Crown relationship by integrating te Tiriti o Waitangi and te ao Māori perspectives into our work (including mātauranga Māori).

8. Maintaining alignment to Our Strategy and current priorities

The data, information and knowledge we collect, steward and use responsibly will contribute to improving ora for current and future generations.

Stewardship of our data ensures that information and knowledge are treated as valuable assets.

Roles and responsibilities

The broad accountabilities and responsibilities for the stewardship and use of data, information, and knowledge are itemised below.

Role	Responsibility
Executive Leadership team	Supports the Policy Owner with the implementation of this policy by: • leading its implementation in their respective business functions, • ensuring this policy has been implemented within their business unit and is being complied with, and • being accountable for the stewardship and use of data, information, and knowledge within their business unit.
All IR staff, including permanent employees, fixed term employees, inter-agency secondees, consultants and contractors	• use this policy to define and guide their activities in the execution of their work for IR, • understand and comply with this policy, • report any non-compliance with this policy to IR's Information Governance Team.

Glossary

The terms data, information and knowledge are used in various ways depending on perspective.

Term	Meaning and example
Data	Data can be defined in several ways, such as: • most simply as facts used to analyse or plan something, which traditionally has meant facts and figures. • in the wider sense — as data sources and analytical approaches have become richer, data can take the form of numbers, research and analyses of these for greater understanding.
Information	Information is produced when data is combined with context and consequently generates meaning. From a business perspective, information is meaningful. Information is an important asset to IR, as it helps us to make informed decisions and deliver effective and efficient services to our customers. When information is managed well, it allows for IR to operate as an open, accessible and transparent government agency.

Knowledge		Knowledge is the articulated intellectual property of IR that conveys our understanding, experience and accumulated learning on a given topic. Knowledge content enables the application of judgment and decision-making, provision of guidance and completion of tasks and processes. Knowledge content typically has the following characteristics: • it is used in the context of making decisions, completing tasks or informal self-guided learning, • its purpose is to improve people's ability to make decisions and act on them, • it is living information that evolves continuously, • it is reusable and available on-demand, • it is synthesised from a range of sources, and • it is highly contextualized.
Whakahaere Pūnaha	Haepapa	The te reo Māori descriptor used for stewardship. • Whakahaere = Management • Haepapa = Responsible • Pūnaha = System
IR staff		All permanent employees, fixed term employees, inter-agency secondees, consultants and contractors.

Finding out more

IR follows the New Zealand Government's Security Classification System to ensure information is handled appropriately based on its sensitivity. The four main classification levels used are **UNCLASSIFIED**, **IN-CONFIDENCE**, **SENSITIVE** and **RESTRICTED**. Each level requires different handling, marking and sharing protocols. For detailed guidance, refer to the [User Guide on Information Classification](#) or contact the Information Security team at 18(c)(i)

Additional resources include:

- [Inland Revenue Enterprise Risk Management Policy](#)
- [Inland Revenue's Artificial Intelligence Staff Use Policy](#)
- [Inland Revenue's Information Security Policy](#)
- [Inland Revenue's Data Handling Standard](#)
- [Inland Revenue's Information Handling Standard](#)
- [Inland Revenue's Code of Conduct](#)
- [Inland Revenue's Records Management Specification](#)
- [Inland Revenue's Information Collection Framework](#)
- [Te Mātāwai – Information Stewardship](#)

Government directives:

- Aotearoa New Zealand's commitment to the [International Open Data Charter](#)
- [Data Protection and Use Policy \(DPUP\)](#)
- [Government Data Strategy and Roadmap](#)
- [Protective Security Requirements](#)
- [New Zealand Data and Information principles](#)
- [Public service principles - stewardship](#)

Document control	Version 2.0
Policy version application date	1 November 2025
Approved by	David Shanks, Deputy Commissioner, Enterprise & Integrity Services 9(2)(a) 
Review dates	Date reviewed: 31 October 2025 Next review: 1 June 2026
Policy owner	Enterprise Leader, Data, Analytics and Insights
Policy contacts	Domain Lead, Information Governance & Sharing Technical Lead, Information Governance

Ka whakamahi nga kaimahi AI i nga kaupapa here

Policy for Users of AI at IR

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

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 explore and adopt 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 meets our obligations under the Revenue Acts and Privacy Act 2020 and under two key AI governance directives – the [Public Service AI Framework](#) and the [Responsible AI Guidance for the Public Service \(GenAI\)](#). This includes embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of Te Tiriti o Waitangi. A full list of NZ Government authoritative guidance and resources consulted in the creation of this policy can be found in the 'Further Information' section at the end of this document.

In addition to these absolute requirements, we must robustly assess the security, privacy, operational and technical impacts that a new AI solution or use case brings and provide appropriate transparency. Appropriate transparency includes determining what is published on the public facing IR website titled [Our Use of Artificial Intelligence](#) and our response to Official Information Act requests on our use of AI.

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

Who this policy applies to

This policy applies to:

- IR employees and contingent workers responsible for managing third-party vendor relationships
- All contingent workers and non-workers who have access to IR data and systems

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

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 with 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 a 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) 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.

If the requested exemption is for a specific AI application, this needs to be requested through the AI use case submission form.

Our policy

1. Principles

To help guide our thinking and behaviour, IR 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 IR that need to be considered, and appropriate action taken.

In compliance with the [Public Service AI Framework](#) and the [Responsible AI Guidance for the Public Service](#) (GenAI), AI tools with approved use cases at IR are required to comply with the following foundational principles:

Foundational Principles for Government	
Governance	Establish and uphold assurance mechanisms to support the delivery of safe and secure AI systems. This includes conducting assessments on risk and impact at regular intervals and disclosing the use of AI within IR to the public
Security	AI applications and tools are to be secure and have access only to the information required to perform functions, evaluating who has access to the tool and the information they have access to, and assess our vendors use of AI. Unauthorised access, misuse or security incidents involving the use of AI tools/solutions will be reported to Information Security
Procurement	Planning for the needs and requirements of IR and not acquiring additional AI tools for the sake of it. 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 AI functionality

Skills and Capabilities	Ensure adequate training, resources and guidance are provided to our staff on the expectations and safe use associated with using approved AI products at IR. This includes showing both the benefits and risks associated with incorrect usage of AI.
Misinformation and Hallucination	Maintain critical thinking skills over all AI outputs and educate IR personnel of the limitations and dangers of relying on accepting AI output as the sole source of truth. This includes providing advice on the best methods of prompting
Accountability and Responsibility	Ensure human oversight is maintained at every step of the process, especially when evaluating or using AI outputs to assist in our work. Human in the loop is enforced at each stage to combine the abilities of AI with the knowledge and the collective experience of our people

And comply with the following principles associated with customer experience with government:

Principles for Government Customer Experience	
Transparency	Clearly explain and document how decisions are informed by AI and where algorithms have generated content in a form digestible to all staff. This applies both internally, promoting transparency within our working teams, and with our customers and partners around how we will utilise these tools
Bias, discrimination, fairness and equity	Strive to create an environment where AI is used in a way that is fair, impartial, and unbiased. This includes checking for biases that may be present in written, oral, or visual formats that may look to discriminate, defame, or otherwise harm our people or the public.
Accessibility	Consistently review the tools we have available to ensure the needs of our people and our customers are addressed, and that can provide a fair and accessible workplace and online space for all.
Privacy	Public trust in how Inland Revenue manages information makes us all information stewards. Our AI systems should respect the law, human rights, diversity and the autonomy of individuals and not result in discrimination against individuals, communities or groups. We can facilitate this by ensuring that data and information is correctly classified, used, and, where required, a Privacy Impact Assessment has been undertaken.
Māori, Pacific Peoples and ethnic communities	We will embed a Te Ao Māori perspective in the deployment and use of AI algorithms consistent with the principles of Te Tiriti o Waitangi. This includes regularly engaging with and encouraging our people to join communities of interest to expand their understanding

2. Adoption of AI solutions

Our environment supports transitioning from the proof of concept/ pilot stage to the deployment and operation stage for trustworthy AI systems. Any new use case will follow approved governance processes to consider experimentation in a controlled environment in which AI systems can be tested, and scaled up, as appropriate, and may require rereview at a later date. Any new proposed use case for an AI, Large Language Model (LLM) tool, AI integrated tool, Agentic tool, or the use of a tool with new or existing AI capabilities

should be escalated through IR's AI governance process including the AI Security Sub Working Group and the AI Oversight Group.

Existing approved workplace technologies that use AI may or may not require new consideration or exemption. Some examples are listed in the [AI Application Register](#).

The form to submit an AI use case can be found [here](#) for tools not currently approved for use at IR. Following the established use case process ensures adequate testing, review and consideration is applied to all use cases and prevents effort duplication and risk. For any questions about submitting a use case, contact 18(c)(i)

Note: Just because an AI tool can be accessed, this **DOES NOT** mean it has been approved for use at IR. While various applications have been blocked for use with IR devices, applications are being consistently reviewed to determine impact to IR and its customers, and blocked applications or tools are subject to change. If you are unsure about the current approval status of a tool or application for use with your business unit or information, contact 18(c)(i)

3. Creation and Use of AI Content

3.1 Users Must:

- 3.1.1. Regularly assess and verify the accuracy of all outputs generated by AI, LLMs, agents (agentic or otherwise), or systems, including but not limited to textual, audio, or visual media.
- 3.1.2. Review AI generated output for potential bias, ethical concerns, and/or unintended messaging.
- 3.1.3. Clearly indicate when AI, LLMs, or other intelligent solutions contribute to content such as visual media, business processes or decisions, especially in customer-facing situations. For example, if a post is created by AI with minimal human oversight for internal use or is intended to be sent to customers, this **MUST** be clearly signposted.
- 3.1.4. Acknowledge that any information produced in part or whole by an AI system will be attributable to the IR staff working with that information, and ultimately IR will be responsible for the quality and outcomes of that information.
- 3.1.5. Apply for an exemption via the formal approval process. This may include approval from the CISO, Privacy, and, where required, Legal and Data if there is a need to use **SENSITIVE Revenue Information (SRI)** or **Personally Identifiable Information (PII)** in an **approved** AI tool. Consider the principles of [IR's Enterprise Risk policy](#) and [framework](#) to ensure a wider risk consideration is given for the use. This includes the potential impact to the security, reputation, or financial standing of IR.
- 3.1.6. Seek the approval of the AI Oversight Group via the AI use case approval process for proposed use cases of unapproved AI/LLM/Agents or other intelligent applications or tools.

3.2. Users **MUST NOT**

- 3.2.1. Use an AI/LLM/Agent or other intelligent application or tool that has not been approved by IRs AI governance process for use.
- 3.2.2. Use an approved tool for unapproved or unintended purposes.
- 3.2.3. Use GenAI to generate code or **MUST** have an exemption filed and approved. This includes using FAST/START code with any AI tools.
- 3.2.4. Use the intellectual property of another organisation or person via an AI tool (including but not limited to Agents and LLMs) without explicit permission from the intellectual owner (e.g. we **MUST NOT** put START data or code into an AI tool – this is the intellectual property of FAST).
- 3.2.5. Use any information classified as **SENSITIVE** or above into an AI tool or application unless this is for use with an application that has been specifically approved for use at this level by the AI Oversight Group. For approved cases see the [AI application register](#) for approvals and their approved use cases.
- 3.2.6. Information classified as **Sensitive Revenue Information** (SRI) or Personally Identifiable Information (PII) **MUST NOT** be used in an approved AI tool without a security exemption approved by the CISO. For approved cases see the [AI application register](#) for approvals and their approved use cases.
- 3.2.7. Intentionally generate any misleading, illegal, discriminatory, defamatory, or otherwise offensive content.
- 3.2.8. Without explicit approval of the tool from the AI Oversight Group, Inland Revenue staff **MUST NOT** process or use in AI/LLM/Agents or other intelligent application or tool, any information that is reasonably capable of being used to identify an Inland Revenue customer or member of Inland Revenue staff.
- 3.2.9. Process or use any of the following classified information types without express approval and consideration from the AI Oversight Group (e.g. through an approved use case submission where it has explicitly stated that the tool will be used with this classification of information.)

IN-CONFIDENCE	COMMERCIAL
IN-CONFIDENCE	CUSTOMER
IN-CONFIDENCE	LEGAL PRIVILEGE

- 3.2.10. Provide automated financial, legal, tax advice, guidance, through AI/LLM/Agents or other intelligent application or tool to another party (either internal or external) without human involvement.
- 3.2.11. Access, use or transfer IR data or information to a personal or non-IR device to input or upload into an unapproved AI application. This includes any AI application that is enabled by default

such as Apple Intelligence, Google search (which uses Gemini), or Google Translate, and does not have an option to opt out.

Note: The above is not a comprehensive list of prohibited uses of AI. If in doubt, email the Information Security mailbox at 18(c)(i)

4. Disposal of Content

The retention and disposal of any AI information or knowledge will be handled in conjunction with IR's current [retention and disposal rules](#) and processes.

Content generated by IR-approved AI use cases may not store data in a non-IR-owned tenancy. Any data or information being stored, processed, or disposed of by a third party **MUST** comply with the clauses outlined in both the [Data Handling Standard](#) and the [Information Handling Standard](#). This includes requiring data to be hosted in a New Zealand/Australia data centre unless provided explicit permission from the AI Oversight Group.

Roles and responsibilities

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

Who	What
Executive Leadership Team	Review IR's approach and use of AI as part of Governance Board meetings
AI Oversight Group (AIOG)	- Provide oversight and direction for AI systems and use at IR, acting as the key governance approvers. - The approval group for proposed AI/LLM use cases or solutions, proposed use of Agentic AI and deployment of AI agents.
AI Security Sub Working Group (AI SSWG)	- Provides initial review of AI use case proposals - Assesses any security, privacy, architecture or legal impacts that may affect IR, including staff and customers.
All employees, non-workers and contingent workers	- 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 IR Information Security Team.

Chief Information Security Office Team (including SecOps, Integrity, and Assurance)	• Responds to any report of abuse, misuse, or non-compliance with this policy. • Monitor user behaviours with AI tooling. • Educate users on the capabilities and acceptable use of emerging technologies • Chairs the Security Sub-Working Group to provide an initial security assessment alongside the Privacy Officer, Corporate Legal and a Security Architect.
Procurement team	• Ensures appropriate checks are undertaken and understanding is held of vendors where AI is in the services provided. E.g. conduct market research on vendors and offerings, covering privacy, security and ethical risks.
Business and Technical owners of AI/LLMs, intelligent solutions and responsible people	• Ensure consistent and fit-for-purpose assessments are conducted of AI/LLMs and AI integrated solutions, including where these capabilities may be integrated into existing products. • Govern and maintain an awareness of the use of AI/LLMs through IR's vendors and partners, and also IR's fourth and fifth parties. • Govern and maintain relevant / useful guidance for staff. • Consider the new high level risks IR may be exposed to and manage the risk as per IR's Enterprise Risk policy and framework. • Consider the carbon and environmental impact of IRs consumption of AI/LLM solutions. • Ensure satisfactory controls are available to govern the use of this information. This includes limiting the collection of IR information by AI tools and products as much as possible, including submitted query and returned answers from the AI vendor.

Terms and Definitions

Term	Definition and Example
Agentic AI	A type of AI that contain complex reasoning skills allowing them to act autonomously, adapt in real time and solve problems based on context and objectives.

	Agent assists are designed to handle day-to-day tasks in collaboration with humans providing instructions.
AI Integrated Solutions	These are solutions that already exist in the IR 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). IR, in many cases, does not have an option to prevent or control the AI component of these solutions, and increasingly will be included in all software tools. System specific guidelines may need to be developed based on the specific tool and use case. Examples: Grammar and spelling checkers, GitHub Co-pilot, Viva answers, Siri, Bing GPT, Spotify and Netflix recommendations.
Artificial Intelligence (AI)	AI refers to the ability for computer systems to perform tasks that typically require human forms of intelligence. This includes tasks such as learning, perceiving, reasoning, planning, which allows machines to complete decision-making tasks and simulate thought adjacent to that of a human.
Business Rules	These are the simplest kinds of computational intelligence. This includes decision trees (e.g.: if value X is above 100, then go to this step), simple coding logic such as if/else statements and for/while/until loops. IR heavily utilizes business rules within START (e.g. calculators, leave and timesheet delegations, customer decision trees, Excel Macros) to help sort and process customer requests.
Generative AI	A generative AI is a direct interaction model which describe how the user interacts with the solution. It is often used in a question-answer format (such as ChatGPT/OpenAI or Microsoft Copilot models). However, GenAI models can generate other forms of media such as images, video, or audio. These are solutions capable of generating entirely new content and may leverage large language models. Examples: ChatGPT, Bing Discover, and Microsoft Copilot solutions.
Intelligent Application or Tool	Intelligent applications are software systems that embed artificial intelligence to augment or automate human workflows, using data-driven models that learn, adapt, make decisions, and act autonomously or semi-autonomously within defined business contexts. Intelligent tools are AI-enabled applications, systems, or components that use

	machine learning, generative AI, or algorithmic models to perform tasks, analyse data, or assist in decision-making, often autonomously.
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.
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 and an LLM. 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: Data Intelligence Platform (DIP) data investigations, Website and Haukāinga search, Viva insights.

Further Information

This policy should be read in conjunction with the following internal references:

- Inland Revenue’s [AI Use case guidelines](#) and the wider supporting AI material on the Te Mātāwai page:
 - [Application register](#)
 - [Artificial Intelligence](#)

IR has several additional relevant policies, standards and resources to support this policy including:

- [Data Information and Knowledge Policy](#)
- [Information Security Policy](#)
- [Risk management policy](#)
- [Use of business tools policy](#)
- [End User Policy](#)
- [Software management standard](#)
- [Acceptable Use Standard](#)
- [Consuming Online Services Standard](#)
- [Data Handling Standard](#)
- [Delivering Online Services Standard](#)

- [Electronic Communications Standard](#)
- [Information Handling Standard](#)
- [Information Classification Guide](#)

IR has drawn on several external resources in the development of this Policy including:

- [NZ Public Service AI Framework](#)
- [Guidance for Safe Use of AI in the Public Sector](#)
- [Paris Call for Trust and Security in Cyberspace: Forging Global Cooperation on AI Risks: Cyber Policy as a Governance Blueprint](#)
- [OECD AI Principles](#)

Document Information

Document control	Version 1.7		
Approved by	9(2)(a)		31-Mar-2026 15:41 NZ
First released	18 December 2023		
Last updated	3 March 2026		
Review dates	Next review: Six months from approval date		
Policy owner	Jay Harris – Chief Information Security Officer		
Policy contact	18(c)(i)		

Version History

Version	Date	Sections amended	Summary of amendment
1.1	18 December 2023	All	Document created and approved
1.2	20 August 2024	All	Amendments to the phrasing of content and the policy template. Added links to relevant sites.

1.3	27 May 2025	All	Updates to the policy to cover data, code bases, and updating to align with government requirements.
1.4	1 September 2025	Users MUST NOT	Removed general reference to IN-CONFIDENCE and added table with “IN-CONFIDENCE” labels that need to be protected from AI view.
1.5	3 October 2025	3.0 Creation and Use of AI Content	Added 3.1.7 and 3.2.6 to clarify the use of SRI and PII in approved tools
1.6	12 January 2026	All	Updates to remove reference to AIWG and update the wording/processes
1.7	3 March 2026	Title	Updated to remove the word ‘staff’

Excerpt from paper to Inland Revenue’s Strategic and Investment Board.

Date: 6 August 2025

Author: Brijesh John (Head of Digital Innovation)

Title of paper: Artificial Intelligence Direction and Strategy

Note: The Enterprise AI principles were approved by the Strategic and Investment Board.

Enterprise AI principles

We have developed eight enterprise principles for AI.

Principle	Rationale	Department impact
1. AI as a system: Design composite solutions combining rules, machine learning, generative AI, and optimisation techniques	Single AI models have limitations; integrated approaches deliver superior outcomes	Enhanced decision quality: Complex tax/policy decisions benefit from multiple AI techniques working together
2. Transparency and contestability: All AI decisions must be explainable, auditable, and challengeable within 48 hours	Public trust requires understanding how AI affects New Zealanders; democratic accountability is essential	Customer trust: Clear explanation of tax assessments, accessible challenge mechanisms protect individual rights
3. Human-AI collaboration: AI augments human expertise without replacing critical judgement in decisions	Maintains stewardship of tax system while amplifying staff capabilities	Staff empowerment: Time savings on routine tasks; humans focus on strategic and high value activities
4. Māori data considerations: AI development must uphold Te Tiriti o Waitangi principles and Māori data governance	Cultural safety for Māori population; constitutional obligations under Te Tiriti	Cultural stewardship: Consider collaborative AI systems design, align to existing and emerging government guidelines
5. Proportional safeguarding: Security measures and privacy controls matched to data sensitivity and risk levels	Protect New Zealanders' data while enabling innovation; comply with Privacy Act 2020	Risk management: For example, Quantum-resistant encryption for sensitive data, tiered access controls. Capability and focus to navigate technology changes and threat landscape
6. Adaptive governance: AI policies and frameworks evolve with technological advances and legislative changes	Rapid AI evolution requires responsive governance; maintain compliance with emerging regulations	Strategic agility: Annual strategy refresh, Parliamentary oversight ensures democratic accountability

Principle	Rationale	Department impact
7. Interoperability first: AI systems use open standards enabling seamless data sharing across government	Maximise innovation investment; avoid vendor lock-in; enable cross-agency collaboration	System efficiency: API-first architecture, federated data sharing with the government system.
8. Beneficial automation: AI deployment must demonstrate measurable value creation for New Zealanders and staff	Align with Public Service AI Framework's human-centric mandate; justify public investment	Value delivery: revenue protection, significant autonomous customer resolution, measurable return on investment

The AI value proposition

For New Zealanders	For Inland Revenue	For New Zealand
Effortless Government Interaction - AI eliminates friction while maintaining human touch	Sustainable Transformation - Service expansion without proportional cost increases	Fiscal Sustainability - Enhanced revenue supporting broader public investment
Proactive Intelligence - Services anticipate and prevent problems before impact	Strategic Workforce Evolution - Staff focus on high-value activity such as policy and customer advocacy	Cross-Agency Synergies - Federated AI multiplying government-wide returns
Cultural Respect - Te Tiriti-aligned design with efficient scalable integration	Operational Resilience - Consistent service during peaks and disruptions	Economic Growth - Reduced compliance burden through automation
Transparent Accountability - Clear AI decisions with efficient challenge processes	Innovation Leadership - Replicable model for government AI adoption	Global Leadership - Responsible AI governance model for international community

M365 Copilot

Privacy Assessment (condensed)

Prepared by: Dawn Swan

Date: December 2025, Appendix 1 added May 2026

Supply ID: N/A

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

From October 2025 Inland Revenue (IR) started rolling out Microsoft 365 Copilot (M365 Copilot), in a staged approach, to staff. This is a generative artificial intelligence (AI) tool integrated into the Microsoft ecosystem.

Users of M365 Copilot can use chat and search functions and the tool has access to existing content within Microsoft applications including emails, documents, calendar events, Teams chats, and files stored in OneDrive and SharePoint. It also has Microsoft's pre-built agents Researcher, Analyst, and Project Manager.

When a user interacts with M365 Copilot (e.g., asks it to summarise a document, draft an email or search for content), the system sends the prompt used and relevant contextual data (retrieved via Microsoft Graph) to a large language model (LLM) hosted in Microsoft's secure cloud. The LLM processes the input and returns a response to either the M365 Copilot interface, or within the user's Microsoft app if they send the prompt via Word, Excel Outlook, Teams etc.

M365 Copilot is powered by Microsoft's Azure OpenAI model which is a Microsoft hosted version of GPT that resides within Microsoft's Australasian service boundary. Because it sits within the Microsoft boundary, it is shared across Microsoft customers, but OpenAI has no access to this model. The data is accessed in real time via Microsoft Graph APIs, which serve as the connective layer between M365 and the user's content.

M365 Copilot only accesses information that the user has permission to access. If a user does not have access to personal information, including customer information, as part of their role then it will not be returned in a response.

The project team has undertaken a risk assessment that contains further details and context: [Risk Assessment for Microsoft 365 Copilot.docx](#)

Corporate Legal undertook a legal risk assessment and found the legal risk of IR using M365 Copilot was acceptable.

1.2 Purpose of the change

A pilot of M365 Copilot involved 20 participants from the Enterprise Information & Knowledge Governance group. The pilot focused on M365 Copilot's ability to assist with meeting administration and demonstrated significant productivity gains:

- 89% of users reported daily use, with an average of five hours saved per week.
- 94% found time savings in content creation, and 83% in document retrieval.
- 100% of respondents said M365 made them more productive, and 89% reported improved work quality

Due to the success of that pilot, M365 Copilot was made available to the Policy team and Tax Counsel Office.

On 25 September 2025, IR's executive team approved the organisation-wide rollout of M365 Copilot. It was agreed this would be done in carefully managed waves from October 2025 to mid-2026. Before business areas start using M365 Copilot, senior leaders are asked if their teams are ready by checking data readiness and capacity of staff to undertake learning activities. Senior leaders decide whether to approve or defer M365 Copilot for their teams.

Use of the AI tool is part of a broader digital strategy and aims to help IR people to work smarter, reduce admin, and focus on higher value tasks.

Users are reminded use of M365 Copilot falls within the Use of Business Tools policy and there may be consequences if misused.

1.3 Benefits


With the adoption of M365 Copilot, staff are expected to reduce time spent on several manual and repetitive tasks, including:

- **Manual summarisation of meetings and documents:** M365 Copilot can automatically generate summaries of Teams meetings, emails, and Word documents, reducing the need for staff to manually compile notes or action points.
- **Drafting routine communications:** Staff will spend less time composing standard emails, reports, or internal updates, as the tool can generate first drafts based on context.

- **Searching for information across Microsoft 365:** M365 Copilot's integration with Microsoft Graph enables intelligent search across Outlook, SharePoint, OneDrive, and Teams, reducing time spent locating documents or past conversations.

1.4 How M365 Copilot works

2. M365 Copilot receives an input prompt from a user either directly in the Copilot app, or via Copilot's integration with Microsoft 365 applications such as Word, Excel, PowerPoint, Teams etc.
3. M365 Copilot then pre-processes the input prompt through an approach called grounding; this improves the specificity of the prompt, to get answers that are relevant and actionable to the specific task. The prompt can include text or link to content discovered by M365 Copilot (e.g. SharePoint and OneDrive files, emails, meeting transcripts, chats).
4. M365 Copilot sends this prompt to the LLM for processing. M365 Copilot only accesses data that the individual user has existing access based on their existing Microsoft 365 role-based access controls.
5. M365 Copilot receives a response from the LLM.
6. M365 Copilot takes the response from the LLM and post-processes it. This post-processing includes other grounding calls to Microsoft Graph, responsible AI checks, security, compliance and privacy reviews, and command generation.
7. M365 Copilot returns a response to be reviewed by the user.
8. Content is encrypted both at rest and in transit



1.5 Personal information to be used

M365 Copilot has access to internal information that is contained within Microsoft 365 applications. This includes emails, documents, calendar events, Teams chats, and files stored in OneDrive and SharePoint.

M365 Copilot does not have access to START (IR's main customer database), the Data Intelligence Platform (DIP) or STAX (corporate records repository), but will have access to any extracts of content from these tools that have been stored in Microsoft locations such as SharePoint. Content may include sensitive or regulated content such as personal information or Sensitive Revenue Information (SRI). It is not possible to be specific in this assessment about what personal information may be contained in documents, the table below outlines categories of information that could be returned by the tool.

Type of personal information	Source of information	Purpose of using the information
User identity and profile information, such as name, email address, phone number, job title, department, and organisation	Inland Revenue	To link a licence to a user
General information: including internal communications, meeting notes, documents authored or accessed by IR employees and contractors.	Microsoft 365 apps like Outlook, Word, Excel, and Teams	For one of the permitted use cases
Personally Identifiable Information (PII): There is potential for PII to be accessed during summarisation or drafting tasks, if users have saved PII in documents. This should only be done if there is a business need for this information to be stored in a Microsoft app.	Emails, chats or documents stored in Microsoft 365 apps containing identifiable and contextual information. This could include employment-related information.	For one of the permitted use cases
Sensitive Revenue Information (SRI): Documents or communications may contain sensitive revenue information (tax or social policy information about customers).	Documents such as letters, memos, emails, Excel spreadsheets or notes if stored in a Microsoft 365 apps or used in a prompt. This could include taxpayer information.	For one of the permitted use cases

1.6 Use cases

When access to M365 Copilot was given to IR's frontline staff (Customer and Compliance Services) further guidance was issued around the use of SRI and acceptable use cases. It was recognised that this group frequently accessed and used customer information (the previous business groups that accessed M365 Copilot did not) and frequently record customer-related information in Word, OneNote, Notepad or Shared documents.

Issued in February 2026 the guideline *Using Sensitive Revenue Information with AI products or services* provided clarity on how SRI can and cannot be used with AI products or services when carrying out IR duties. It noted:

Everyone at IR is expected to treat SRI as confidential and only access or use it where:

- *there is a legitimate business need,*
- *the use is permitted under law, and*
- *the use supports carrying out IR duties.*

This obligation applies regardless of role, system, or work location (including working from home or co-located environments).

IR provides many approved tools to support our people in delivering outcomes for New Zealanders. SRI is approved for use with AI products or services when used appropriately and within established controls.

The guideline provides examples of appropriate use cases including:

- Where data is exported from START, AI products and services can assist with data quality checks, analysing patterns in audit outcomes that may indicate emerging compliance risks, analysing a spreadsheet to extract trends of note and insights
- Finding publicly available information about a customer and summarising it,
- Extracting and summarising information from a screenshot or handwritten note provided by a customer.
- Summarising call or case notes or work-in-progress content
- Genesys Cloud creating a first draft of notes for an inbound call, or
- Summarising and formatting detailed notes about actions taken on a customer's account. Final documentation related must be saved to the customer record.
- Searching a shared workspace for documents or folders about a case being worked on.
- Reviewing a web message to enhance clarity and flow of the message.
- Improving the structure and language used for feedback being sent to another IR staff member.

1.7 Governance

There has been wide engagement for AI being introduced in IR, including approval by the Executive Leadership Team and endorsement by both the AI Working Group, AI Oversight Group and Strategic Investment Board.

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?	Y	IR has a Policy for Users of AI at IR and use case guidelines. These set out IR's approach to using AI safely and securely in the workplace and governance groups decide when AI tools will be implemented. Within M365 Copilot users are permitted to use: • information classified up to and including SENSITIVE information • commercially In-Confidence information; and • information that is Sensitive Revenue Information (SRI) under the Tax Administration Act (e.g. reasonably capable of being used to identify a taxpayer, and personal information including information about IR staff).

	provided they have the right to access the information. This permitted use of information in M365 Copilot is an approved exception to Policy for Users of AI at IR. Apart from this exception the AI staff use policy continues to apply. Targeted training is provided so users understand how to use M365 Copilot effectively and responsibly. Completion of the Copilot Fundamentals course in Ātea and agreement to specific M365 Copilot Do's and Don'ts is a pre-requisite to getting M365 Copilot access. IR also has AI-fluency e-training that staff are encouraged to complete. This has 8 modules including prompt training, trust and privacy, threats and risks of AI. The M365 Do's and Don'ts include: • I understand I can input SRI in Copilot as part of my IR duties • I will talk to my leader if I have any questions about the SRI I am using with Copilot. • I will handle, store and dispose of information and files used or created in Copilot, in line with the existing guidance, go to Records Management for more details. • If I am working with customer information, I will ensure the customer record in START is appropriately updated in line with best practice. See Add a Note, Best Practice Statement and Our Charter for more details. • I will always review the outputs created by M365 Copilot for accuracy, completeness and quality
--	---

		before saving, sharing or acting on them. • To protect information security and privacy, I will think before recording or transcribing a Microsoft Teams meeting and will notify all participants before I do so. If the meeting changes to a different topic, I will consider if the recording should be stopped. Refer to Recording Microsoft Teams meetings. • I will only use M365 Copilot in line with the use of business tools policy. • I understand IR's Code of Conduct applies to my use of M365 Copilot and my misuse of it may breach the Code, with resulting consequences. • I will use Copilot as a support tool to inform my decisions, not as the final decision maker for answers or recommendations. • I will not use SRI for purposes that are outside of my IR duties.
Collection	Y/N	If yes, explain your response
Will IR be collecting personal information that it doesn't currently collect?	N	The information accessed by M365 Copilot does not collect any new personal information. It collates and analyses information that has already been collected by IR and that users have access to.
Is collecting this information necessary for IR to carry out its functions?		M365 Copilot does not collect any new information. The information it may access is stored in relation to IR functions.
Where or who is the information being collected from?		The information accessed and returned in a search by M365 Copilot will be collected from a user's Microsoft applications. A

		user will have stored and collected that information for a lawful purpose. CCS staff are able to use customer information in M365 Copilot if it is appropriate to do so. One of the main use cases permits staff to use Copilot to draft web messages to customers. This may contain information either collected from the customer or associated with the customer. M365 Copilot does not automatically access IR's START, DIP or STAX systems. START and DIP are where the majority of customer information is stored.
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?	Y	Using the AI tool is new but all data remains within the IR-managed Microsoft 365 environment. Prompts entered into M365 Copilot are not used to train the large language model. IR maintains human-in-the-loop oversight, with users responsible for reviewing and validating AI-generated outputs before being used in official communications or decision-making. Before accessing M365 Copilot users must complete the Copilot Fundamentals course, agree to M365 Copilot Do's and Don'ts, clean up digital storage spaces and check what documents are shared with others. When M365 Copilot was rolled out to IR's frontline staff, further guidance was issued around using SRI. This reminded users that everyone at IR is expected to treat SRI as confidential and to only access or use it where: there is a legitimate business need, the use

		is permitted under law, and the use supports carrying out IR duties. It also outlined SRI may be used with an AI product or service when all of the following conditions are met: • Purpose: The use of SRI is directly related to carrying out your IR duties. • Context: The SRI already exists as a normal and expected part of your work (for example, in documents, notes, spreadsheets, analysis outputs, or collaboration spaces you routinely use). • Activity: The tool is assisting you to analyse, summarise, draft, format, or improve content you are already working with. • Control: You are disposing of copies of SRI created when the information is no longer required. Official copies of outputs related to customers must be stored again against the Customer Record. • Tool: the AI product or service has been approved for use at IR. Refer to AI Working Group - AI submissions register - All Items for approved AI products or services.
If IR is sharing or collecting information how it will it be transmitted?		N/A
Where will the information be stored?		If a user creates new information using M365 Copilot this will be saved and stored in IR's usual repositories, most likely Microsoft files or folders. Customer information must be stored on the customer record in START. No data is stored outside IR's Microsoft 365 tenancy.
Who will have access to the information?		M365 Copilot applies existing permissions so will only access information that a user

		is currently able to see and access. If new content is created, the individual user will determine who is able to access the content through the usual permission process. IR's Cyber Security team can access a log that displays prompts and is able to monitor what information is being used in prompts.
Are there controls to protect against unauthorised access, use, modification or disclosure?		There is guidance available for IR staff about how they should store and restrict access to information to prevent unauthorised access or use. IR accepts there is an increased chance of content being surfaced that people were not aware they had access to. Copilot will only allow people to access information they already have permission to see (for more information see Enterprise data protection), as it respects IR's existing content permissions. However, Copilot might surface information that people <i>do</i> have access to but perhaps were not aware of. For example, if a file in a Teams or SharePoint site was unintentionally shared broadly within IR, Copilot could include it in responses for people (since they have access). Guidance on checks people can complete themselves is available on the Getting started with Microsoft 365 Copilot page. Leaders can also request extra support from the IKM team to check whether specific content is secured. There is a process for people to report if they see a file or document they think

		they shouldn't have access to: Reporting accidental access to a document or file .
How long will the information be retained?		Retention and disposal of content is guided by rules set out in IR's Disposal Authorities under the Public Records Act 2005. Documents saved within M365 applications are retained for as long as needed under the Public Records Act. Content is reviewed by users to determine when content is deleted or archived. Copilot prompts and Teams recordings/transcripts have the same retention policy in the backend of M365 so need to have the same retention period applied. Copilot prompts and Teams recordings/transcripts fall under GDA7 class 1.6 "facilitative records" as more substantive records are created after the point the prompt/recording is made (e.g. documents using the Copilot information, or meeting minutes/notes) Class 1.6 has the retention period of "Retain until material is not required". The Enterprise Information and Knowledge team have decided, using a minimum necessary retention by default principle, that prompts and recordings/transcripts will be retained for 6 months. There is the ability to keep recordings for longer if necessary.
Will IR be notified when the information is no longer needed for the specific purpose?		N/A
Will the information be disposed of? If so, explain how		See above for retention and disposal information.

Does it involve transferring personal information offshore, using a third-party contractor?	Y	IR uses Microsoft as a primary vendor to provide cloud infrastructure and applications. IR's Microsoft environment is held within the Australian region, IR tenancy is hosted in Australia and the AI model is hosted in Sydney.
Can the third party provide assurance it is complying with IR expectations and information is handled responsibly? If so, explain how this will be done.	Y	Security credentials and certifications for Microsoft are publicly available Service Trust Portal Home Page (microsoft.com) :
Use, disclosure, and accuracy	Y/N	If yes, explain your response
Is the information currently held by IR?	Y	
If yes to the above question, for what purpose does IR hold the information?		Information stored by IR staff in Microsoft will be relevant to their role and linked to an IR function.
Will the initiative use or disclose information for a different purpose to why it was obtained?	N	The purposes for which information has been collected and stored has not changed nor has the purpose for its use. It is the manner of processing that has changed. Information held within Microsoft applications will include internal communications, meeting content, and documents authored or accessed by IR employees and contractors. The prompts used by staff, and content returned, is likely to be used for a similar purpose to why it was created. The user will only be returned information they have permission to access. Users agree to only use M365 Copilot to carry out IR duties and not to use it for personal use. Any outputs are to go undergo human review prior to being used or disclosed further.

Will IR be sharing personal or taxpayer information with another organisation?	N	Microsoft does not access IR's tenancy where information is stored.
Has due diligence been done on that other organisation (if not a government agency) to ensure it has reasonable controls in place?	Y	
What processes are in place to ensure and maintain data integrity?		Quality of the data varies by business area. Prior to being given access to M365 Copilot and as part of the 'Dos and Don'ts' users agree to check their OneDrive and shared storage locations such as SharePoint, shared workspaces and MS Teams for access permissions and whether they have draft content. It is recommended to delete draft content that has since been finalised so it is not returned in a search. If users are able to access information they shouldn't be able to see, they are expected to report this via a form set up by the IKM team so remediation work can be undertaken.
Access and identification	Y/N	If yes, explain your response
Will the information be stored on a customer or staff member's record?	Y	If M365 Copilot is used to summarise or clarify customer communications (for example web message content) the final version must be saved on the customer record in START.
Does the initiative affect how people can access information IR holds about them?	N	
Does this involve a new way of identifying individuals?	N	
Other considerations	Y/N	If yes, explain your response

Is there a way to achieve the purpose of the project using less identifiable data?	N	
Would people be surprised by this use of their information?	N	Some people may be surprised that IR is permitting customer information to be used in an AI tool. However, there are sufficient security safeguards to protect IR information (for example information does not leave the IR tenancy) and guidance for staff about appropriate use. M365 Copilot does not use any content that is created.
If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?	N	M365 Copilot may access information a customer has freely provided if frontline staff use the tool to summarise correspondence or customer interactions. It is not anticipated that this use will jeopardise taxpayer's providing their information to IR. The tool does not access START which is where the full customer record is maintained. Given IR's operational context, there may be some documents or communications which contain SRI if the user has a legitimate business purpose to access this content.
Does the initiative involve tracking or monitoring of movements, behaviour or communications?	N	

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?	N	
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?	N	
How are we identifying and managing bias or discrimination?		There may be potential for biased or inappropriate content to be generated. Microsoft deploys AI safeguards such as content filtering and prompt grounding and there will be human review of all outputs. Prior to being given access to M365 Copilot and as part of the 'Dos and Don'ts' users agree to review all outputs from Copilot and look out for limitations, bias and inaccuracies. Later in 2026 IR will also roll-out Microsoft Purview. This tool is capable of blocking content from being used in M365 Copilot.
Can you foresee any harm to individuals in using the data in the way intended?	N	
Does the data to be used specifically identify Māori or a Māori collective?	N	
Have you considered how the proposal contributes to the active protection of Māori interests?	N/A	
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?		Users agree to look out for limitations, bias, and inaccuracies in outputs when using M365 Copilot.
Will decisions informed by an algorithm or use of AI involve human review and evaluation?	Y	Users agree to always review the outputs created by M365 Copilot for accuracy and quality before saving, sharing or acting on them. This includes checking citations in all outputs to ensure Copilot is using the correct and most relevant version of

		information that is needed; and the user understands where Copilot sourced the information from (such as a budget secret folder) so they're aware who they're authorised to share the outputs with.
Will any automated decision-making process be regularly reviewed to make sure it's still fit for purpose?	N/A	M365 will not be making automated decisions. Any outputs are reviewed for accuracy and IR is clear with staff that they are responsible for decisions made based on AI-generated content.

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	Medium	Copilot does not collect new personal information but processes existing data in new ways (e.g., summarisation, drafting). Given the types of use cases M365 Copilot will be used for (surfacing and summarising documents and team meeting admin for instance), a moderate amount of personal information may be used depending on the user's role. M365 Copilot can access and process personal information if such data is stored in Microsoft 365 (emails, documents, Teams, SharePoint, OneDrive). Mitigations/controls • Copilot only accesses information users already have permission to see. • Approved exception allowing SRI use with Copilot, supported by guidance.

		• Mandatory training and agreement to Copilot Do's and Don'ts. • Human review of all outputs before use or disclosure.
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)	Medium	Information available to some users may contain information that is considered sensitive. For example, frontline staff may input a taxpayer's identifying information and financial details into Copilot for it to be analysed (but only if they have a legitimate business need to do so). Users must follow the Use of Business Tools policy and ensure that taxpayer information is saved and used appropriately and only used for a business purpose connected to their role. If a document has previously been shared with a user, and their permission has not been revoked, the document may still be visible. Due to a lack of security on a document it's also possible users will come across information that is not intended to be seen by them. All users of M365 Copilot are reminded to check documents that have been shared to determine if they should still be accessible, and to review security on OneDrive documents. The IKM team has published information about how staff can report accidental access to a document or file. Mitigations/controls • Use of Business Tools policy and role based access controls apply. • Users required to review and clean up OneDrive, SharePoint and Teams permissions. • Guidance and process for reporting accidental access to documents or files. • IKM support available to review and remediate access issues

Significance of the changes L – Only minor change to existing functions/activities M – Substantial change to existing functions/activities; or a new initiative H – Major overhaul of existing functions/activities; or a new initiative that's significantly different	Medium	Introducing M365 Copilot is a change for IR. It has rolled out Copilot Chat but M365 Copilot will access and surface more content to users. Pilots were undertaken to measure its effectiveness, and the executive leadership team approved a carefully managed roll-out complete with training, resources and reporting accidental access to documents. Mitigations/controls • Pilot testing completed before wider rollout. • Executive Leadership Team approval. • Staged rollout with training, guidance, and readiness checks for business areas
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	Medium	Microsoft has provided training to key IR staff and supports further training for users. IR's Microsoft 365 tenancy and the AI model are hosted in Australia (Sydney). Data is processed offshore by Microsoft. Mitigations/controls • Microsoft is IR's primary cloud vendor. • Microsoft security credentials and certifications are publicly available via the Service Trust Portal. • Data remains within IR's Microsoft 365 tenancy
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	Medium	IR's use of M365 Copilot will have minimal impact on customers but may raise concern or media attention. AI tools, in particular M365 Copilot, have not been rolled out across many other government agencies and using AI may still be considered risky for some use cases. IR has not encouraged staff to use AI to communicate with customers and all outputs must be reviewed. Microsoft Purview, when rolled out, can restrict what information is

		permitted in prompts and will allow more monitoring. Use of M365 Copilot has been through IR's AI governance groups (Working Group and Oversight) and been approved by the executive leadership team with mitigations in place. Mitigations/controls • Use approved through IR AI governance groups and Executive Leadership Team. • Staff are not encouraged to use Copilot to communicate directly with customers. • All outputs must be reviewed by humans before use. • Future tools (e.g. Microsoft Purview) can further restrict prompt content and monitoring
--	--	---

5. Summary of privacy impact

The privacy impact for this project has been assessed as:	Select
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	X
High – Sensitive personal information is involved, and/or several medium to high risks have been identified. <u>You must complete a full Privacy Impact Assessment</u>	
Inadequate information – More information and analysis is needed to fully assess the privacy impact of the project.	

6. Reasons for the privacy impact rating

Medium privacy impact - Some personal information and SRI is involved, but risks are considered satisfactorily mitigated through governance, technical controls, and staff obligations.

Some personal information is involved, but risks can be mitigated satisfactorily with the controls and recommendations:

- Access permissions.
- Mandatory training and agreement to Copilot Do's and Don'ts before access is granted.
- Human review of all outputs to ensure accuracy before use or disclosure.
- Regular review of access permissions and reporting of accidental access.
- M365 does not integrate with IR's customer record systems.

If there are changes to the way that M365 Copilot impacts privacy, IR needs to ensure that controls are in place to mitigate any risk that arises.

7. Document sign-off

Position	Name	Business Unit	Sign-off Date
Sponsor or Business Owner	Tim Crook	Enterprise & Integrity Services	27 February 2026
Privacy Officer	Dawn Swan	Enterprise & Integrity Services	2 March 2026
Chief Information Security Office	Yolanda Wilke	Enterprise & Integrity Services	15 April 2026
Change Manager	Makayla Stewart-Wright	Enterprise & Integrity Services	24 February 2026
Initiative Delivery Lead	Craig Massie	Enterprise & Integrity Services	24 February 2026

Review date: May 2027

Appendix 1

When Microsoft 365 Copilot was approved to be deployed across the IR workforce the underlying model, Microsoft Azure AI (Open AI), was run within the Microsoft Cloud trust boundary in Australian datacentres.

In May 2026 Microsoft offered the Claude LLM model as an option in M365 Copilot. Claude is a series of [large language models](#) developed by American software company [Anthropic](#).

When Claude is enabled it will become an option for users to select from the dropdown in M365 Copilot – users can either select Claude or leave it on Open AI.

The key difference with the Claude option is that processing sits outside the AU Microsoft boundary and data is hosted in AWS in the US. So, if IR uses Claude, IR data will be processed in the US. Identity and authorisation, grounding, conversation history and audit processes do not change.

If Claude is used the path of an IR user prompt is:

Step	Description
1	IRD user authored prompt is captured in a Copilot surface (Word, Excel, PowerPoint, Researcher, Teams, Copilot chat, or Copilot Studio agent) running under the user's Microsoft 365 / Entra ID identity in the Australia regional Microsoft 365 tenant.
2	The Copilot orchestration service performs grounding (Microsoft Graph retrieval, file context, web search if enabled, plug-ins / connectors). Grounding sources are accessed under the user's permissions; this stage continues to occur inside the Microsoft Cloud trust boundary.
3	If the orchestrator selects (or the user has selected) Claude, the assembled prompt + grounded context is forwarded by Microsoft to the Anthropic inference endpoint that Microsoft has contracted via AWS. This crosses the Microsoft / AWS boundary to a US-based AWS region.
4	Anthropic processes the data on Anthropic-managed infrastructure inside AWS. As per the Microsoft contract, and Anthropic's no-training commitment, the prompt and response are not used to train Anthropic foundation models.
5	The Claude completion is returned to the Microsoft Copilot orchestrator, post-processed (responsible-AI checks, citations, formatting) and returned to the user inside the Microsoft 365 surface. Conversation state, audit and Copilot interaction history continue to be stored within the Microsoft 365 tenant per existing M365 Copilot data handling.

Privacy risk assessment of enabling Claude

- Microsoft has onboarded Anthropic as a Microsoft sub-processor effective January 2026, which brings Claude under the Microsoft Product Terms including compliance with data protection. Anthropic **operates with Microsoft oversight through contractual safeguards** which provides comparable safeguards to the Privacy Act and complies with IPP 5 and 12. See <https://learn.microsoft.com/en-us/microsoft-365/copilot/connect-to-ai-subprocessor>
- Under the Microsoft agreements customer data is not used to train Anthropic models.
- Anthropic temporarily retains what is necessary to perform the service and data that flows to Anthropic does not persist in the LLM.
- Independent of the Microsoft sub-processor relationship, Anthropic maintains its own externally audited compliance program.

Given the controls in place, enabling Claude in M365 Copilot does not alter the privacy risk of using the tool or have any significant privacy impact.