



14 November 2025

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

Dear [REDACTED]

Thank you for your request made under the Official Information Act 1982 (OIA), received on 16 October 2025. You requested the following:

- 1. A list of all AI tools that are currently approved for use by staff at your agency.*
- 2. Any documentation outlining the conditions, guidelines, or policies attached to the approval and use of these tools.*
- 3. For each approved tool that is not free to use, please provide the number of paid licenses or subscriptions the agency currently holds. I confirm that I do not require any commercially sensitive information (e.g. licence costs), merely the number of authorised users.*
- 4. Copies of all completed Cloud Risk Assessments and Privacy Impact Assessments (or equivalent documents) for each of the approved AI tools.*

Items 1 and 3

The table on the following page details all current uses of AI within Inland Revenue (IR), and the number of paid licenses or subscriptions IR currently holds. This may be through a license, a subscription, or consumed as a service.

There are several items within the table which have AI functionality included in existing enterprise tools. In these instances, IR has been using the product for some time, and subsequently, the supplier has added AI features. Examples of this are Snowflake Cortex AI and Microsoft's E5 licence.

There are several acronyms within the table for which I will provide additional detail here.

- **E5 licences:** IR has 5250 Microsoft E5 licences, which includes the standard M365 tools alongside features for risk management and data protection, and Copilot chat.
- **START:** This is IR's tax and social policy administration software supplied by FAST Enterprises. IR has one enterprise license for this product.
- **DIP:** IR's Data Intelligence Platform (DIP) is a SAS Managed Hosted Service, using Snowflake. This is a consumption-based agreement and there are no licences.
- **ServiceNow:** IR has 5301 IT Service Management licences.

Table 1: Uses of AI within Inland Revenue

Name	Description	Product Status	AI Technology	Licenses/subscriptions
ABBYY FineReader 16	Text recognition and document conversion tool, used to convert PDFs into excel.	In production	Optical Character Recognition Machine Learning	45 licences
Ability to pay model	Analyse the customer's circumstances and recommends the best action for that customer	In production	Machine Learning	1 IR enterprise licence for START
AI Futurist	Enables querying and summarisation of content.	In production	Generative	3 licences
Assurity Intelligence	Test scenario generation	Pilot	Generative	6 (consumption-based service)
Copilot Chat (Bing/Browser)	AI-powered chat service	Available	Generative	Part of E5 licence
Coveo	IR's public websites and internal staff intranet and knowledge base use Coveo as a search platform. Coveo uses machine learning to continuously learn and improve from user searches and patterns to inform the best results to display.	In production	Machine Learning	1 IR licence
Dragon Naturally Speaking	Screen reader	In production	Machine Learning	12 licences

Name	Description	Product Status	AI Technology	Licenses/subscriptions
DDoS protections - AWS Shield - Azure DDoS Protection - Cloudflare - Magictransit - F5 BigIP - Oracle DDoS	IR uses a range of tools to prevent our systems and services from unexpected outages due to network attacks.	In production	Machine Learning	1 IR subscription
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 development	Machine Learning	1 consumption-based service (DIP)
Figma	Prototyping software that enables Inland Revenue to develop mock-ups of intended changes to products and services across both e-services and internal/external Inland Revenue websites	In production	Machine Learning	10 paid licences
Financial intelligence network detection	Links, matches and identifies multi-dimensional risks of users via operational and strategic visualisation	In production	Machine Learning	1 consumption-based service (DIP)
Genesys Agent Assist	Creates summaries of conversations with contact centre agents for post-call notes.	In production	Generative	642 licences
GST integrity model	A predictive model to assess the risk of GST returns requesting refunds.	In production	Machine Learning	1 IR enterprise licence for START

Name	Description	Product Status	AI Technology	Licenses/subscriptions
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	1 consumption-based service (DIP)
MarianMT	Translating text for digital forensics	Available	Neural machine translation	Free of charge
Microsoft 365 Copilot	Copilot is integrated into the M365 suite of products and is designed to enhance staff productivity.	Phased roll out	Generative	3000 licences
Microsoft Copilot Studio	Build, customise, and deploy AI-powered copilots and chatbots for business scenarios.	Proof of Concept	AI Agent	Part of M365 Copilot licence
Microsoft Defender	An enterprise-capable host protection solution that is integrated with a range of other Microsoft Apps, observes activity on devices for potential malicious behaviour.	In production	Machine Learning	Part of E5 licence
Microsoft Teams Premium	Provides recaps of Teams meetings.	Pilot	Generative	50 licences
Microsoft power BI	Dashboard/reporting software that connects to multiple Azure services.	In production	Machine Learning	Part of E5 licence
Microsoft Purview	Portfolio of products that span data governance, data security, and risk and compliance solutions.	In development	Machine Learning	Part of E5 licence
Microsoft Teams Voice Isolation	Separate user's voice from other sounds and voices in Microsoft Teams calls and meetings.	In production	Machine Learning	Part of E5 licence
Overdue income tax return RIT prediction	Predicts residual income tax (RIT) on overdue returns.	In production	Machine Learning	1 consumption-based service (DIP)

Name	Description	Product Status	AI Technology	Licenses/subscriptions
Posit Connect	Supports the deployment of AI-powered data science solutions.	Available	Generative	100 licences
Power Automate	Low Code solution that supports automating tasks.	In production	Machine Learning	Part of E5 licence
Propensity to read letter or log-in to myIR	Helps IR to select and use the right channels to communicate with customers.	In production	Machine Learning	1 consumption-based service (DIP)
Qualtrics	Analysis of survey information from customers.	In production	Machine Learning	30 licences
Receipt, invoice, statement and tax/employer return review	Text recognition	In production	Optical Character Recognition Machine Learning	1 IR enterprise licence for START
ServiceNow Accelerators	Enhances IT service management and helpdesk functionality with automated workflows.	Proof of concept	Generative	Part of IR existing ServiceNow contract
SharePoint Advanced Management (SAM)	Analyses site activity and detect unusual patterns, such as potential oversharing or risky access behaviour.	In production	Machine Learning	Part of M365 Copilot licence
SharePoint Sections with Copilot	Support content authors by helping them structure content into meaningful sections.	In production	AI agent	Part of M365 Copilot licence
Snowflake Cortex AI	Allow users to ask business questions in natural language and receive direct answers from structured data.	Proof of concept	Generative	1 consumption-based service (DIP)

Name	Description	Product Status	AI Technology	Licenses/subscriptions
Tableau Desktop	Data analytics and graphing tool used for analysing and visualising performance test results.	In production	Machine Learning Natural Language Processing	5 licences
Viva Topics	Provides people with personalised content feeds and recommendations of communities to join or follow.	In production	Machine Learning	Part of E5 licence
Windows Hello for business	Authentication for IR devices.	In production	Biometrics	5400 devices
Z scaler	Detection and classification of web traffic and websites.	In production	Machine Learning	5400 devices
ZoomText	Screen magnification software for accessibility purposes.	In production	Optical Character Recognition	4 licences

Item 2: Any documentation outlining the conditions, guidelines, or policies attached to the approval and use of these tools.

Inland Revenue AI Policy and Guidelines

Copies of Inland Revenue's *Artificial Intelligence Staff Use Policy* and *Artificial Intelligence Use Case Guidelines* are publicly available on IR's website by searching for [Inland Revenue's use of AI, list of AI technologies and purposes, governance framework and benefits of AI use](#) (pages 23 to 35). Therefore, your request for IR's AI policy and guidelines is refused under section 18(d) of the OIA as the information is publicly available.

Conditions and guidelines attached to the approval and use of specific tools

a) Proofs of Concept

Those tools which are currently in Proof of Concept (PoC) stage (Microsoft Copilot Studio, Snowflake Cortex AI and ServiceNow Accelerators) have conditions on use. We are intentionally ring-fencing the PoC activities for these tools to one or two specific use cases. This focused approach allows us to more effectively measure their impact and evaluate their suitability.

Each tool currently in PoC has defined conditions for use to ensure clarity and consistency during this evaluation phase. As at today's date:

- Microsoft Copilot Studio is approved for a single use case of a Technology Services Agent.
- Snowflake Cortex AI is approved for use in one data domain.
- Seven ServiceNow Accelerators are approved using demo instances with no IR data or configurations.

b) Tools in production/available for use

Genesys Agent Assist: Requires a training module to be completed. As this is available for staff as a guideline, I have enclosed a copy of this.

Item	Date	Document	Decision
2.1	undated	eLearn: Genesys Cloud Conversation Summarisation	Released in full

Copilot Chat (Bing/Browser): This has three conditions of use,

- the staff member's role being approved for using the tool,
- completion of the Copilot fundamentals eLearn, and
- the staff member has agreed to the required Do's and Don'ts.

The following table details these guidelines and conditions and my decision on their release:

Item	Date	Document	Decision
2.2	First published 6 November 2024, last updated 6 October 2025	Internal intranet article showing the roles approved to use the tool	Released in full

Item	Date	Document	Decision
2.3	First published 12 November 2024, last updated 6 November 2025	eLearn Copilot Fundamentals	Refused under section 18(d) of the OIA as the information is publicly available
2.4	First published 24 November 2024, last updated 6 October 2025	Copilot Chat Do's and Don'ts	Released in full

Financial Intelligence Network Detection Engine

One document is held which lays out conditions of use and guidelines for this tool.

This document is withheld under the following grounds:

- i. section 6(c) of the OIA, on the basis that making the information available would prejudice the maintenance of the law, being the Commissioner's statutory responsibilities for the administration of the tax system. Disclosure of the information into the public sphere would enable the recipient of the information to potentially take deliberate steps to evade detection of their unlawful actions.
- ii. section 18(c)(i) of the OIA, on the basis that making the requested information available would be contrary to the provisions of a specified enactment, namely section 18(3) of the Tax Administration Act 1994 (TAA). That is because if the information contained in the conditions and guidelines was made available, this could prejudice the maintenance of the TAA for essentially the same reasons as set out in (i).

Item	Date	Document	Decision
2.5	27 April 2021	Financial Intelligence Network Detection Engine – Technical Documentation	Withheld in full under section 6(c) and 18(c)(i) and of the OIA

GST Integrity Model

One set of guidelines is held for this tool. The information held in these guidelines, along with the conditions set out in the *Memo: Implementing the new GST Integrity Model*, are withheld in full under the following grounds:

- i. section 6(c) of the OIA, on the basis that making the information available would prejudice the maintenance of the law, being the Commissioner's statutory responsibilities for the administration of the tax system. Disclosure of the information into the public sphere would enable non-compliant (or potentially non-compliant) actors to evade their GST obligations.

- ii. section 18(c)(i) of the OIA, on the basis that making the requested information available would be contrary to the provisions of a specified enactment, namely section 18(3) of the TAA. That is because, if the information contained in the conditions and guidelines was made available, this could prejudice the maintenance of the TAA for essentially the same reasons as set out in (i), as there would be ability for a recipient of the information to potentially evade their obligations.

The conditions around the use of this tool sit within the START Analytics Domain. The Integrity Manager Operational Governance Group is responsible for overseeing the execution of our compliance approach as it specifically applies to this tool. The primary purpose of this group is to ensure appropriate business oversight and governance of the use of this tool to validate (where applicable) both the correct filing of and amendments to returns across all tax and social policy products.

Item	Date	Document	Decision
2.6	1 April 2024	Memo: Implementing the new GST Integrity Model	Withheld in full under sections 6(c) and 18(c)(i) and of the OIA
2.7	First published 20 December 2022, last updated 24 June 2025	GST Integrity Review (internal staff guidelines published on intranet)	Withheld in full under sections 6(c) and 18(c)(i) and of the OIA

M365 Copilot

This is currently being rolled out organisation-wide at IR and has four conditions of use:

- a checklist and readiness assessment is completed by a senior leader in each business area,
- the staff member's role is approved for using the tool,
- completion of the Copilot fundamentals eLearn, and
- the staff member has agreed to the required M365 Copilot Do's and Don'ts.

The following documents in scope of your request are enclosed and detail the conditions of use and the guidelines. Additionally, the Copilot fundamentals eLearn referenced in item 2.3, is publicly available.

Item	Date	Document	Decision
2.8	1 October 2025	Extract from minutes of Strategic Investment Board meeting	Released in full
2.9	9 October 2025	Internal intranet page information: M365 Copilot rollout information for senior leaders	Released in full

Item	Date	Document	Decision
2.10	First published 6 October 2025, last updated 30 October 2025	M365 Copilot Dos and don'ts	Released in full

Item 4: Copies of all completed Cloud Risk Assessments and Privacy Impact Assessments (or equivalent documents) for each of the approved AI tools.

The following Privacy Impact Assessments and Privacy Threshold Assessment are in scope of your request and enclosed. Some information has been withheld under section 9(2)(a) of the OIA, to protect the privacy of natural persons.

Item	Date	Document	Decision
4.1	Undated draft	DRAFT: M365 Copilot Privacy Assessment	Released in full
4.2	18 July 2023	Viva Topics: Privacy Threshold Assessment	Partially released with some information withheld under section 9(2)(a)
4.3	7 February 2023	START Analytics: Privacy Impact Assessment	Partially released with some information withheld under section 9(2)(a)
4.4	10 February 2025	Microsoft Teams Voice Isolation: Privacy Impact Assessment	Released in full

All tools with an AI component considered for use by Inland Revenue undergo a risk assessment which includes security and privacy. All risk assessments within the scope of your request are withheld in full under the following grounds:

- i. section 6(c) of the OIA, on the basis that making the information available would prejudice the maintenance of the law, being the Commissioner's statutory responsibilities for the administration of the tax system. Disclosure of the information into the public sphere would enable non-compliant (or potentially non-compliant) actors to evade their tax or social policy obligations or otherwise enable persons to interfere with or cause damage to tax administration systems.

- ii. section 9(2)(b)(ii) of the OIA, on the basis that disclosure of much of the information would be likely unreasonably to prejudice the commercial position of both the parties involved and Inland Revenue. I have taken into account any other considerations which might render it desirable, in the public interest, to make that information available. However, those considerations do not outweigh the good reasons noted above for withholding the information.
- iii. section 18(c)(i) of the OIA, on the basis that making the requested information available would be contrary to the provisions of a specified enactment, namely section 18(3) of the Tax Administration Act 1994. That is because if the information contained in the risk assessments was made available, this could prejudice the maintenance of the TAA for essentially the same reasons as set out in (i), as there would be ability for a recipient of the information to potentially circumvent system security.

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) as this information may be of interest to other members of the public. This letter, with your personal details removed, may be published in its entirety. Publishing responses increases the availability of information to the public and is consistent with the OIA's purpose of enabling more effective participation in the making and administration of laws and policies and promoting the accountability of officials.

Thank you again for your request.

Yours sincerely



Cate Robertson
Enterprise Leader, Strategic Architecture



Genesys Cloud: Conversation summarisation

Welcome to Genesys Cloud: Conversation summarisation course. This course can be completed as either self-paced learning or in a facilitator-led workshop.

The purpose of this course is to learn about Genesys Cloud's conversation summarisation feature.

By the end of this course you will be able to:

- describe what Genesys Cloud's conversation summarisation feature is
- explain your responsibilities when using conversation summarisation
- identify the limitations of conversation summarisation.

If you are not familiar with the Genesys Cloud system and taking notes in START, it is recommended that you complete courses **T1 AAO Genesys Cloud: Inbound and outbound calls** and **T1 AAO Notes on customer accounts** before you complete this course.

When you are ready to start the course click the start button.



What is conversation summarisation?



Using conversation summarisation

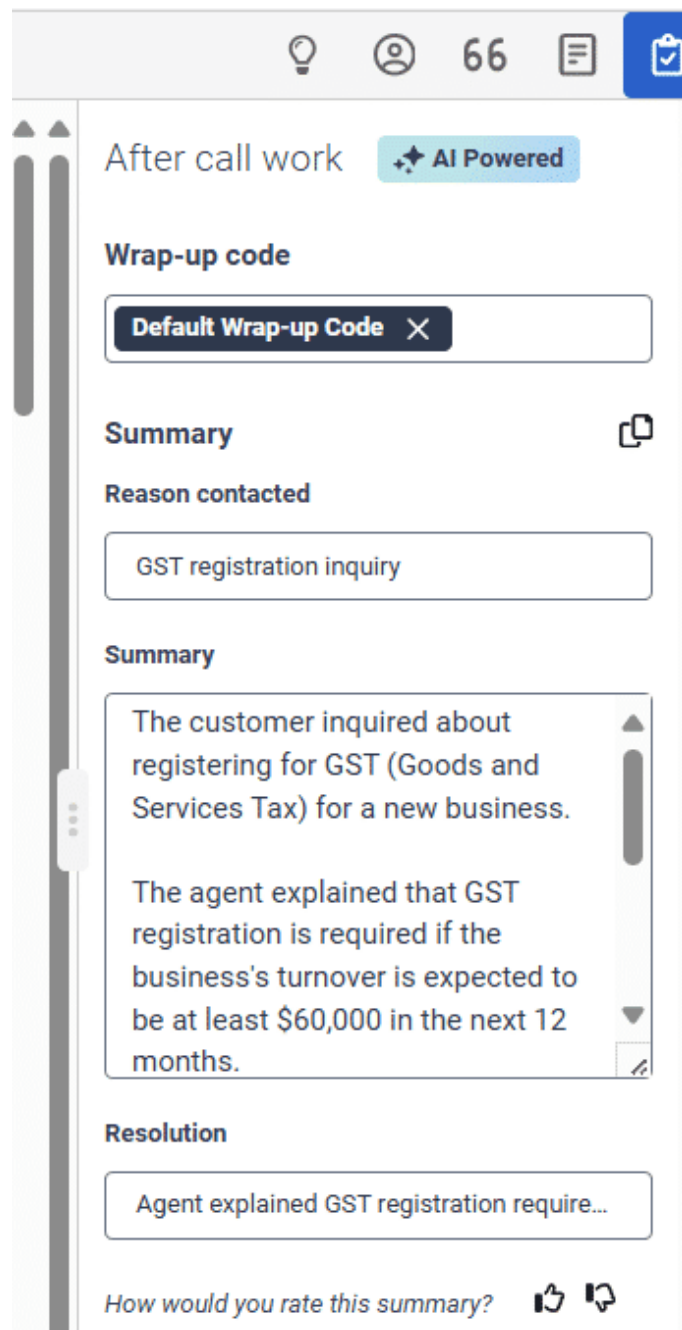
Limitations

Summary of key learnings

What is conversation summarisation?

Overview

Conversation summarisation is a feature within the Genesys Cloud system which uses generative AI (Artificial Intelligence) to **automatically generate summaries** of conversations between customers and agents during a call.



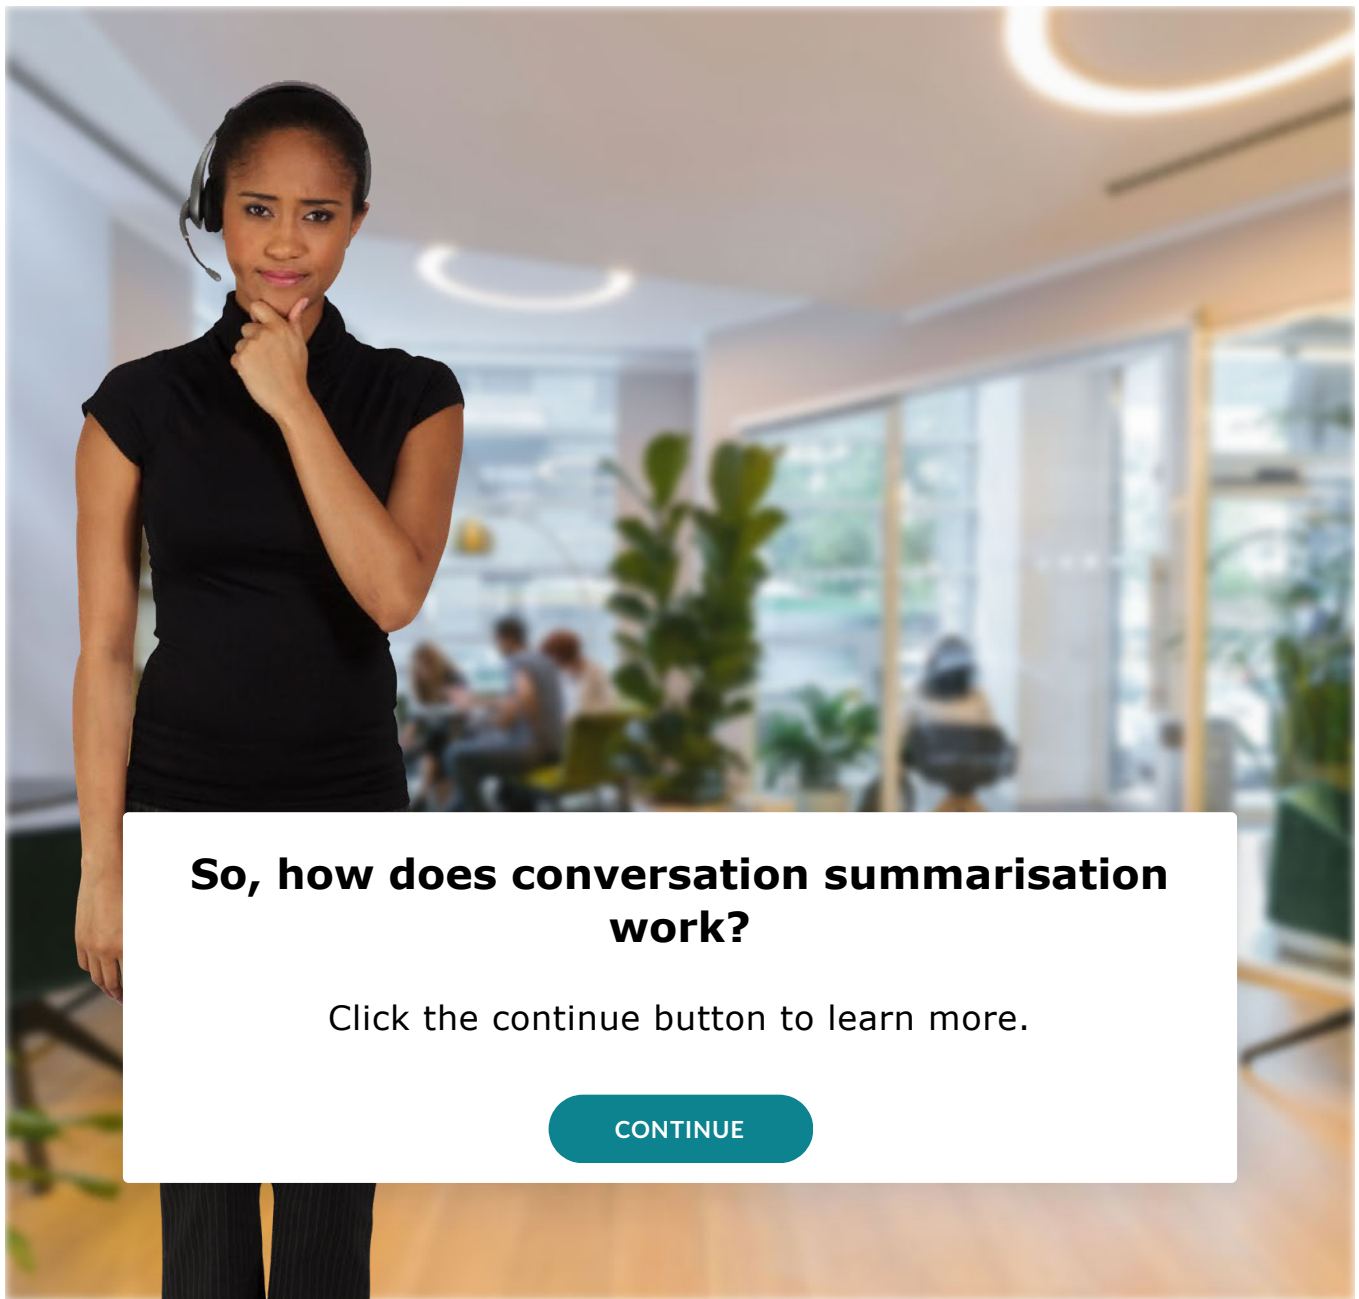
Example of a summary generated by conversation summarisation in Genesys Cloud.



The term '**agent**' is used by conversation summarisation to refer to the staff member involved in the call. For Inland Revenue, this could be Customer Service Officers (CSOs) or

other Inland Revenue staff who handle customer calls using the Genesys Cloud system.

CONTINUE



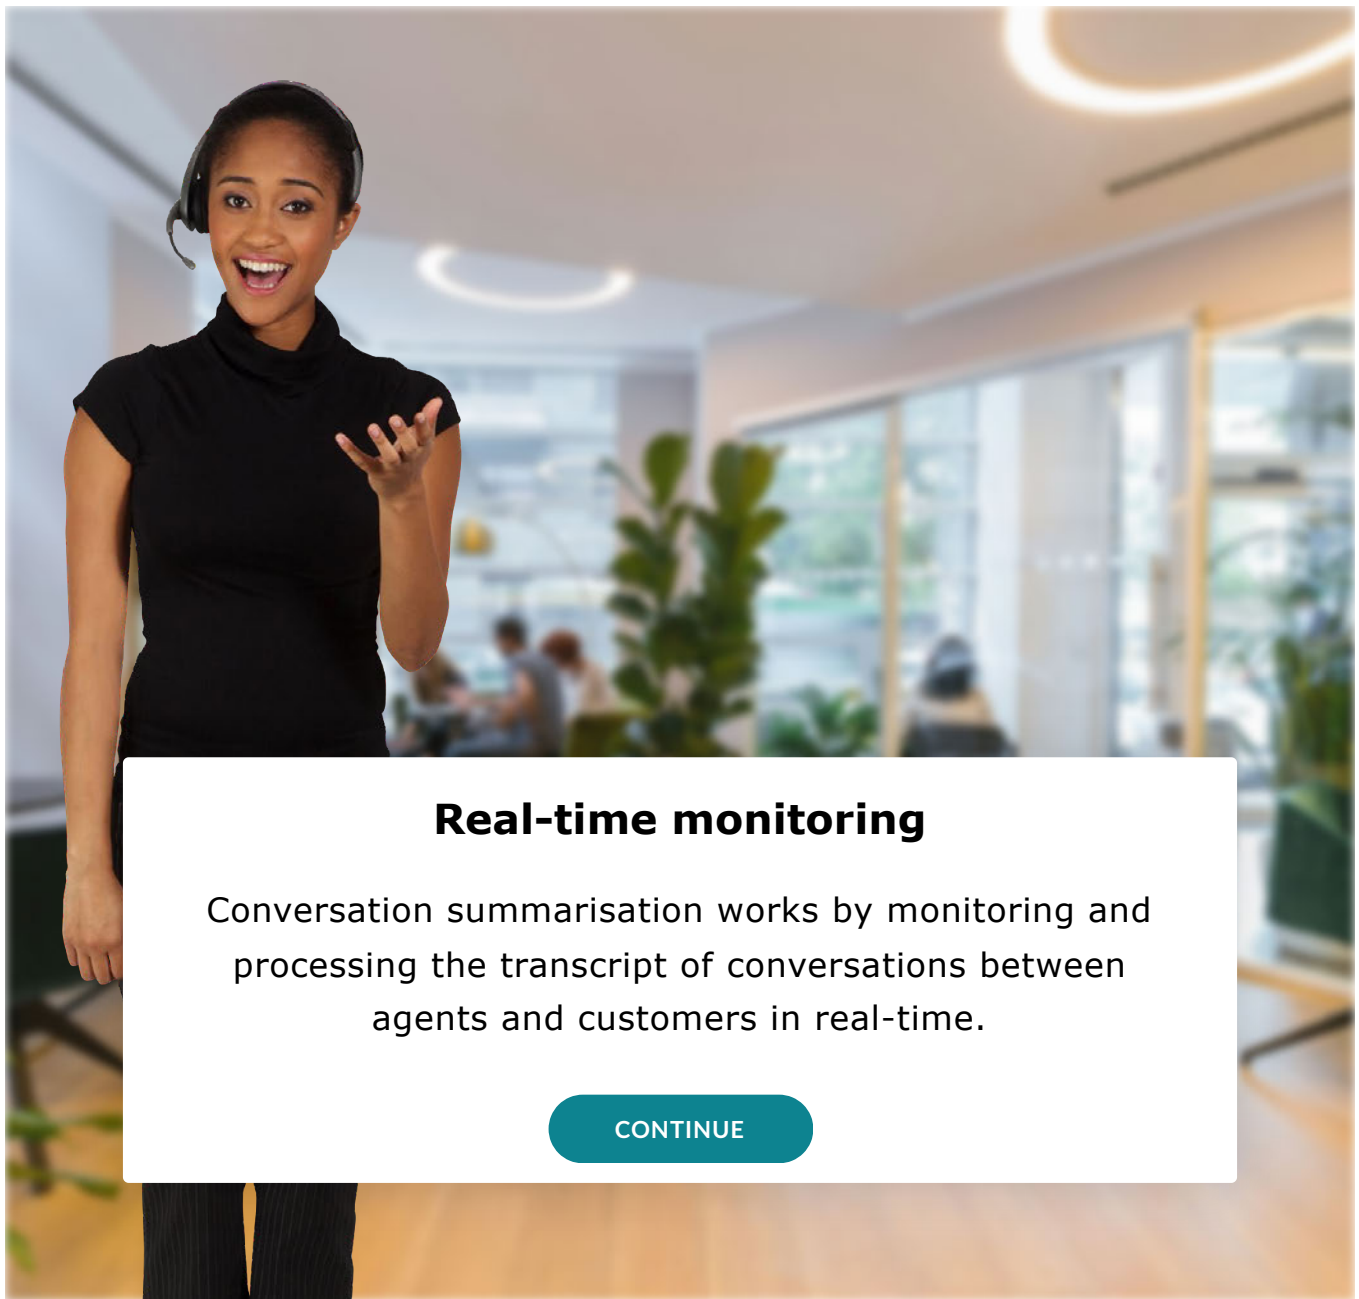
So, how does conversation summarisation work?

Click the continue button to learn more.

CONTINUE

Scene 1 Slide 1

Continue → Next Slide



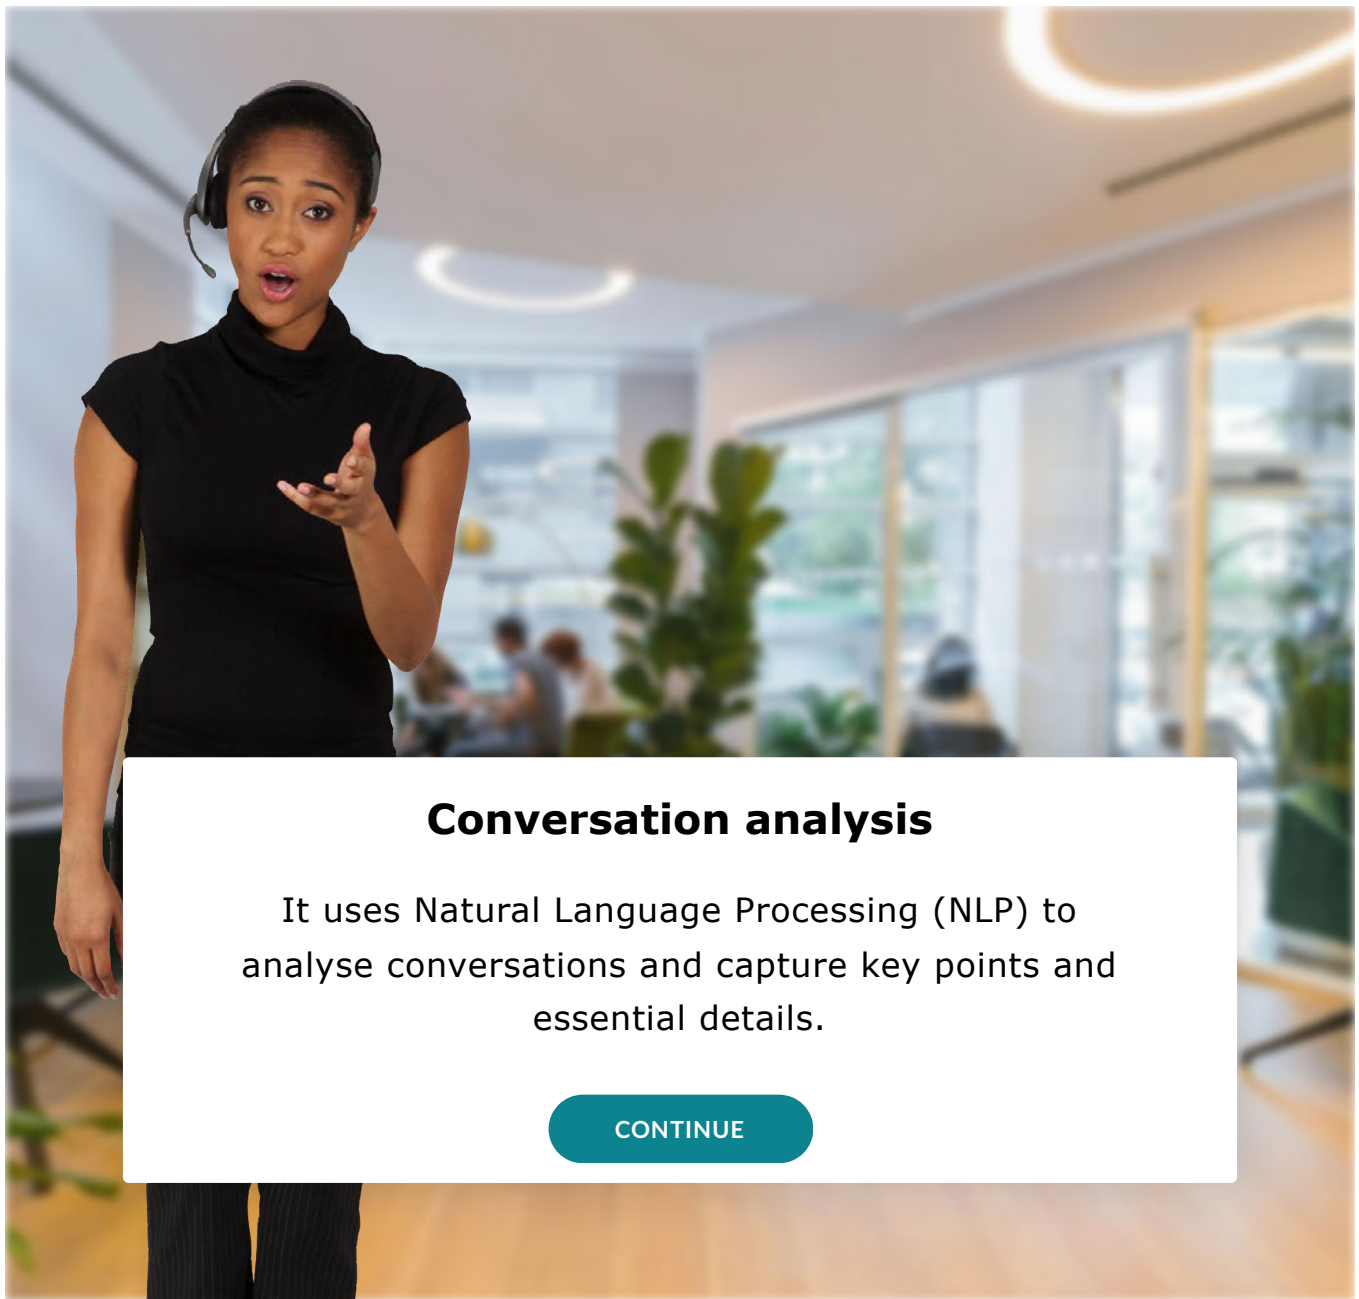
Real-time monitoring

Conversation summarisation works by monitoring and processing the transcript of conversations between agents and customers in real-time.

CONTINUE

Scene 1 Slide 2

Continue → Next Slide



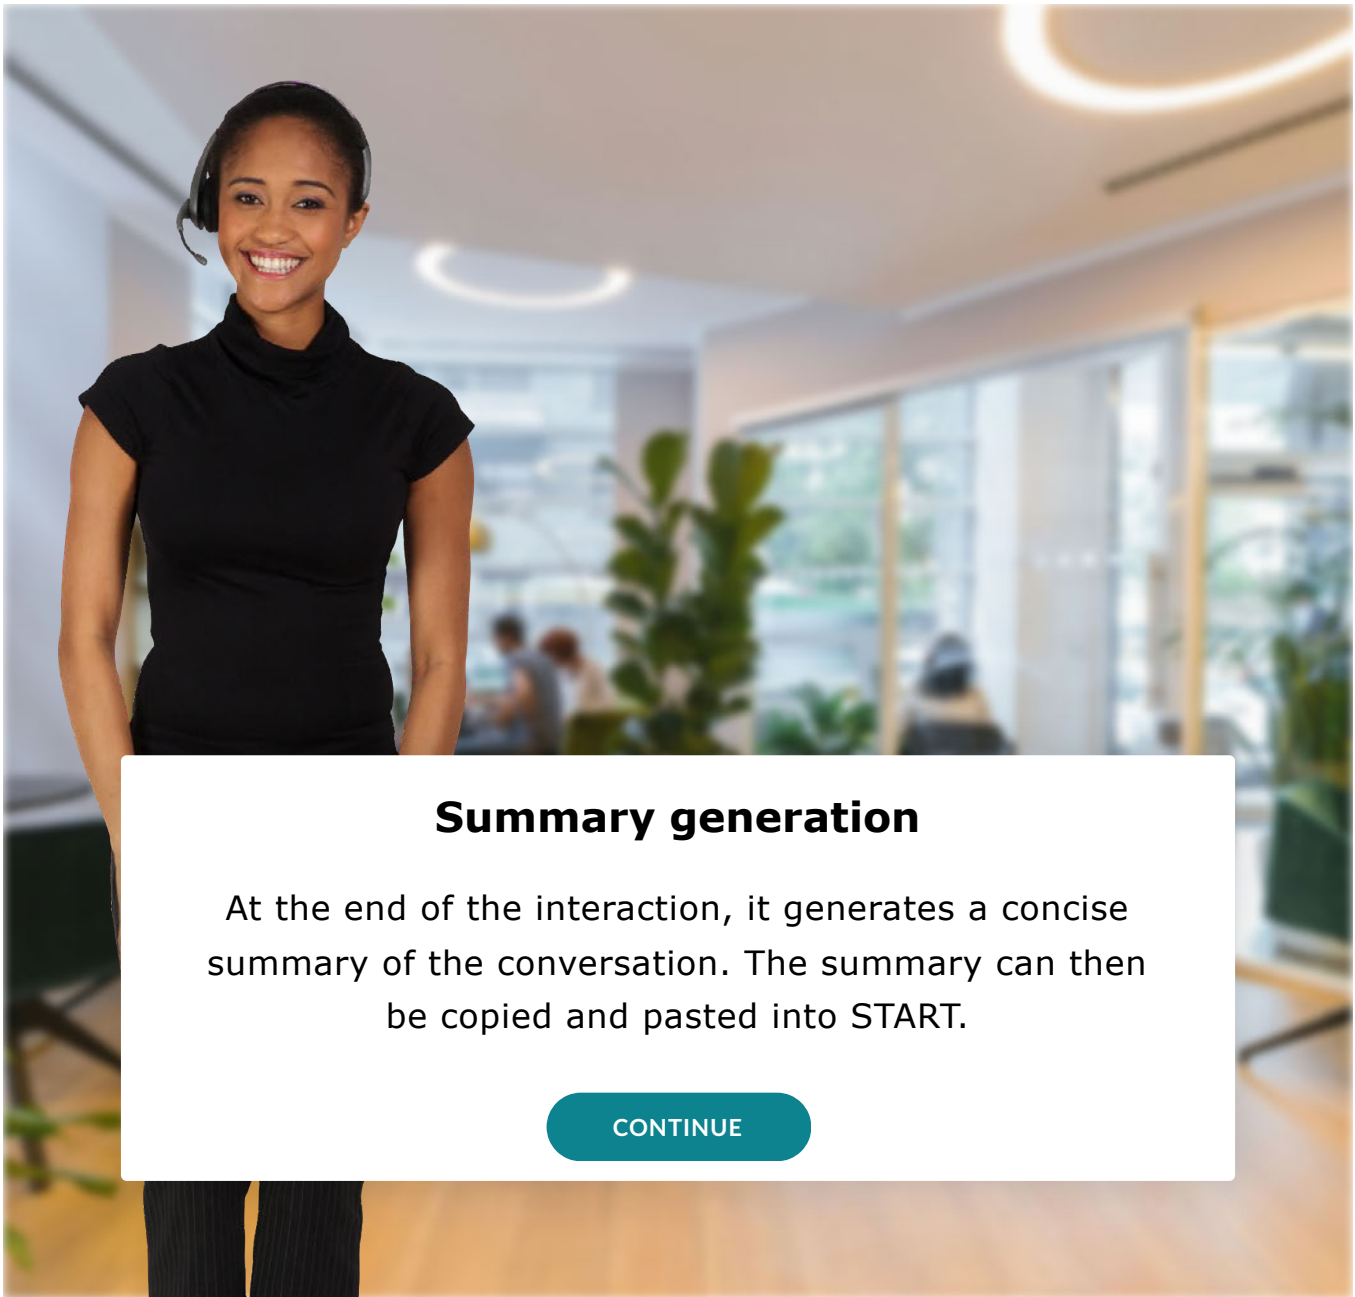
Conversation analysis

It uses Natural Language Processing (NLP) to analyse conversations and capture key points and essential details.

CONTINUE

Scene 1 Slide 3

Continue → Next Slide



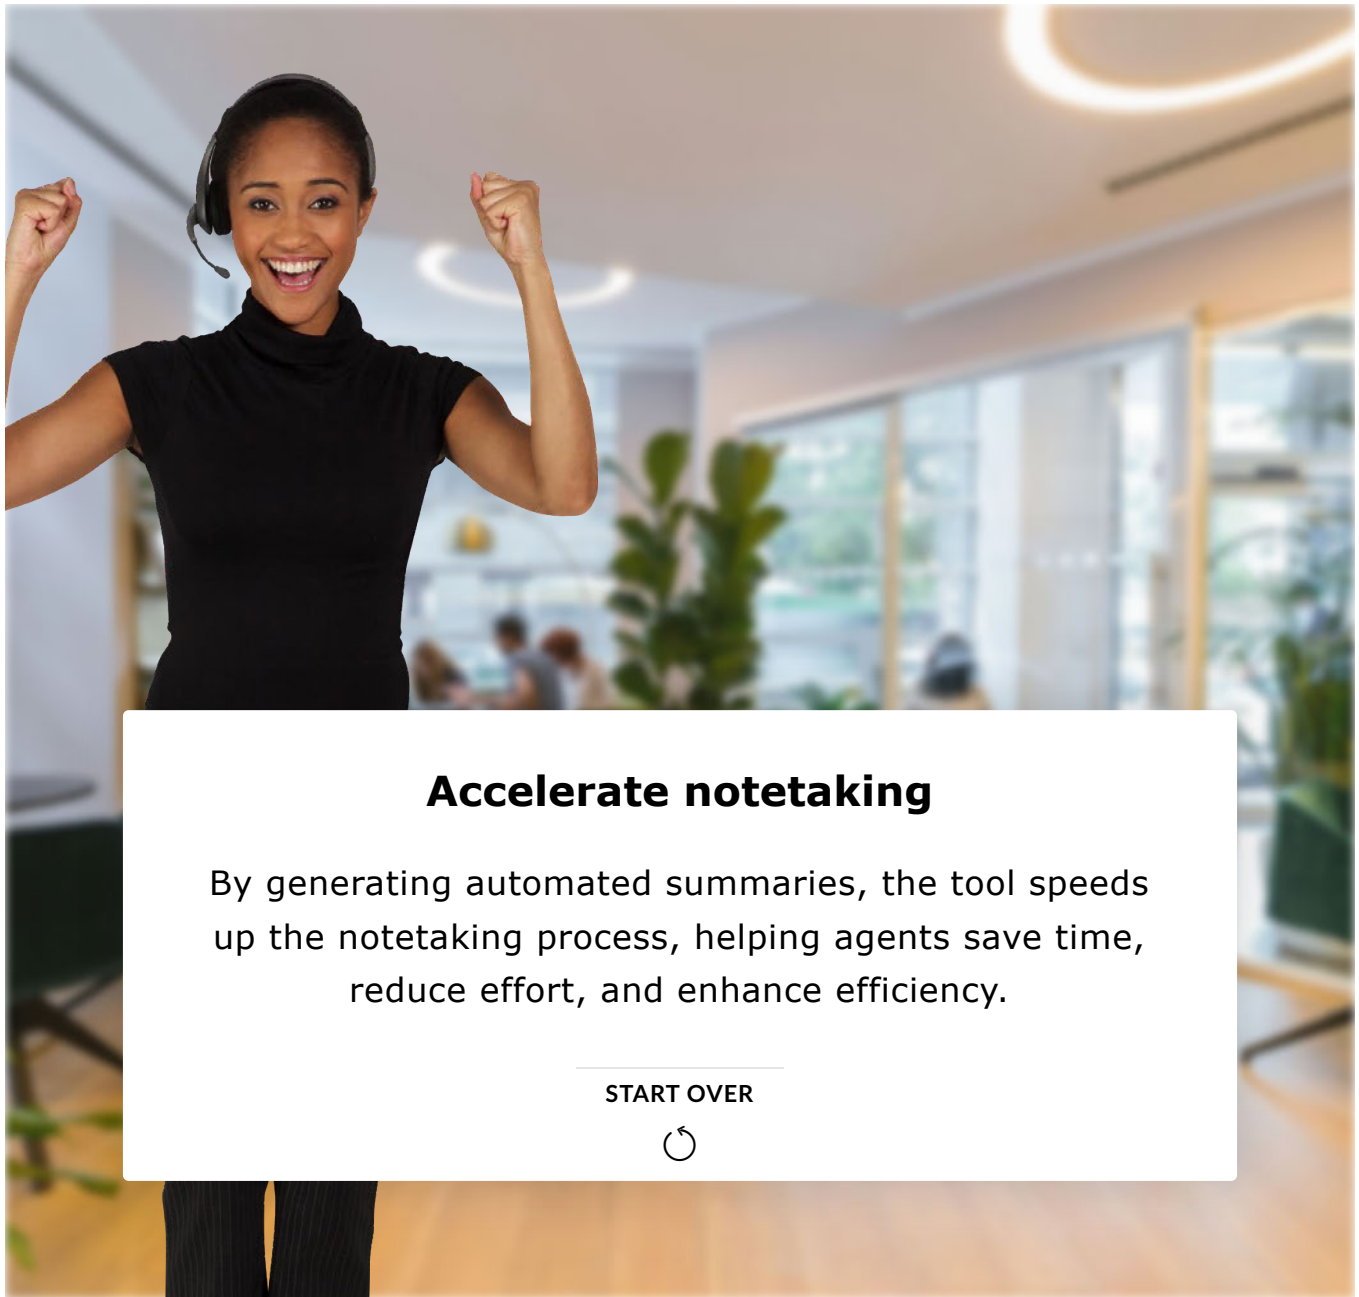
Summary generation

At the end of the interaction, it generates a concise summary of the conversation. The summary can then be copied and pasted into START.

CONTINUE

Scene 1 Slide 4

Continue → Next Slide



Accelerate notetaking

By generating automated summaries, the tool speeds up the notetaking process, helping agents save time, reduce effort, and enhance efficiency.

START OVER



Scene 1 Slide 5

Continue → End of Scenario



Complete the content above before moving on.

Activity

Use what you have learnt so far to answer the question.

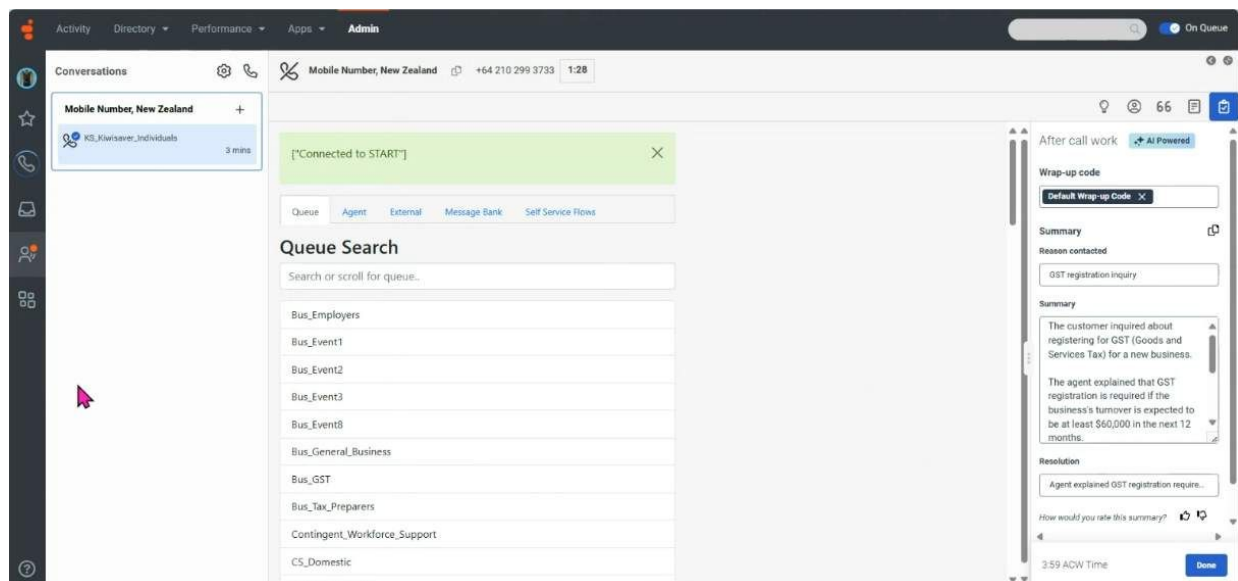
Select which statement below best describes the conversation summarisation feature within Genesys Cloud.

- ☐ It provides detailed transcripts of every customer interaction.
- ☐ It creates automated summaries of conversations between agents and customers.
- ☐ It translates customer conversations into multiple languages in real-time.

SUBMIT

Using conversation summarisation

Conversation summarisation is activated as soon as a call is connected with a customer in Genesys Cloud (either inbound or outbound). When the conversation with the customer ends, conversation summarisation will automatically generate and include a **wrap-up code**, **reason contacted**, **summary**, and **resolution** of the conversation. This information will appear in the After call work box on the right-hand side of the screen as shown in the images below.



Example of a summary generated in the **After call work** box on the right-hand side of the screen.

66

After call work

AI Powered

Wrap-up code

Default Wrap-up Code X

Summary

Reason contacted

GST registration inquiry

Summary

The customer inquired about registering for GST (Goods and Services Tax) for a new business.

The agent explained that GST registration is required if the business's turnover is expected to be at least \$60,000 in the next 12 months.

Resolution

Agent explained GST registration require...

How would you rate this summary?

3:25 ACW Time

Done

Example of a summary generated with a **wrap-up code**, **reason contacted**, **summary**, and **resolution**.

CONTINUE

Must do's when using conversation summarisation

The generated summaries may sometimes contain inaccuracies or miss key points and may not always capture everything we consider important. Therefore, it is important to carefully **review** and **edit** the summaries as needed to ensure that any notes recorded on a customer account accurately reflects your conversation with the customer.

When using conversation summarisation, you must do the following.



Review the summary

Always review the AI-generated summary to ensure it accurately reflects the conversation.



Edit as required

Edit the summary if it contains inaccurate or incomplete information, or if it would not make sense to someone else reading it.



Ensure completeness

Ensure all key points and essential details from the conversation are included in the summary.



What NOT to do

Do not copy and paste the summarisation into Microsoft Copilot as this tool must not be used for any customer-related activities.

For more information, refer to Te Mātāwai page [Microsoft Copilot at IR](#).

CONTINUE

Editing the summary

You can edit the **reason contacted**, **summary**, and **resolution** sections directly in Genesys Cloud. You may also copy them somewhere else (for example notepad) and make your edits there before copying and pasting them into START.

There is only one wrap-up code option, labeled as the **Default Wrap-up Code** and this does not need to be changed.

Copying the summary

There are two options to copy content from Genesys Cloud.

**COPY THE SUMMARY, REASON CONTACTED
AND RESOLUTION**

COPY THE SUMMARY ONLY

You can copy the entire summary, reason contacted, and resolution sections. You should **use this option** when copying and pasting notes into START.

To copy, select the **main** copy icon just above the Reason contacted section as shown in the image below (click to zoom).

After call work

AI Powered

Wrap-up code

Default Wrap-up Code

Copy summary, reason
contacted and resolution

Summary



Reason contacted

Business registration inquiry

Summary

The customer inquired about registering for US taxes, specifically if they need to register if their business income is over \$60,000.

The agent confirmed that registration is required for income above \$60,100.

Resolution

Agent provided the necessary information ...

How would you rate this summary?



How would you rate the voice quality for this

1:42 ACW Time

Done

**COPY THE SUMMARY, REASON CONTACTED
AND RESOLUTION**

COPY THE SUMMARY ONLY

You can copy the summary only.

To copy, hover over the summary box and select the copy icon as shown in the image below (click to zoom).

After call work

AI Powered

Wrap-up code

Default Wrap-up Code X

Summary

Reason contacted

Business registration inquiry

Summary

Copy summary

The customer inquired about registering for US taxes, specifically, if they need to register if their business income is over \$60,000.

The agent confirmed that registration is required for income above \$60,100.

Resolution

Agent provided the necessary information ...

How would you rate this summary?

How would you rate the voice quality for this

2:39 ACW Time

Done

Formatting and adding notes in START

The Te Mātāwai page [Add a note](#) provides information and guidance on when to add notes, what details to include, and how to format notes in START. However, the format of the generated summaries may be different and specific format for notes may not apply when using notes generated by Genesys Cloud.

Below are some key things to consider when editing and formatting the generated summaries and adding notes in START.

- If conversation summarisation has not captured it, add the tax type and period the customer is calling about at the start of your notes. If appropriate, also include the customer's name and role (for example 'INC 31/03/2025 Jim Jones (DIR) called...').
- If the summary includes information you would not normally include (for example validation questions), you can leave them in as long as it does not detract from the note.
- Lodge your notes where you would normally add them. If you wouldn't normally lodge a note, then you don't need to add one just because conversation summarisation provided one.



Refer to Te Mātāwai page [Add a note](#) for more information.

You can also refer to the course **T1 AAO Notes on customer accounts** which covers Inland Revenue's policy and guidelines for writing notes on customer accounts in more detail.

CONTINUE

Activity

Use what you have learnt so far to answer the questions below.

What actions must you take when using the conversations summarisation feature?
Select all that apply.

- ☐ Review the generated summary.
- ☐ Edit the summary as required.
- ☐ Ensure all key points and essential details are included in the summary.

SUBMIT

True or false. You can use Microsoft Copilot to proofread or edit your summary.

- ☐ True.

☐ False.

SUBMIT

Which sections should you copy when copying summaries from Genesys Cloud?

- ☐ The summary, reason contacted, and resolution sections.
- ☐ Only the summary section.

SUBMIT

Other resources

Te Mātāwai:

- [Notes](#)
- [Add a note](#)

- [Microsoft Copilot at IR](#)

Limitations

Conversation summarisation limitations



Conversation summarisation leverages machine learning to effectively analyse and summarise interactions. As with any technology, there are some technical limitations to this feature that you should be aware of. These limitations may change or improve over time as new features and upgrades are made.

Select the start button to learn about some of its limitations.

Missing information or summary

Conversation summarisation may not generate some or all of the summary. This can happen when there is poor audio quality, heavy accents, or the conversation doesn't follow a natural language pattern such as when multiple topics are discussed or when the conversation jumps around from one topic to another.

When a summary can't be created, a message will come up that states:

'AI summary could not be created, write your own instead.'

If this happens, you will need to write your own summary.

Misinterpreted words

Conversation summarisation may misinterpret words which can be caused by varying accents, unclear speech, or poor audio quality.

Limited language

Conversation summarisation is limited to the English language only. When you have Language Assistance or an interpreter on the phone, conversation summarisation will pick up the English conversation between you and the interpreter, but it will not record the foreign language conversation between the interpreter and the customer.

Conversation summarisation may also use American spelling in its summaries (for example 'organ**iz**ation' instead of 'organ**is**ation').

Note: You do not need to change the spelling to New Zealand-English in your notes.

Multiple customers

When multiple customers or entities are discussed within a single call, conversation summarisation cannot distinguish between separate customers/accounts and generates a summary for the entire interaction.

Example:

Lisa receives a call from a tax agent, Mark, who wants to discuss the tax details for three different customers. When the call ends, conversation summarisation generates a single summary for the entire call. Lisa would need to manually edit the summary to separate the information for each customer.

Note: When speaking to a tax agent, conversation summarisation may refer to the tax agent and their client as the 'customer', and the CSO/Inland Revenue staff member as the 'agent'. For clarity, references to the tax agent should be updated to say 'tax agent'.

Role reversals

In some cases, conversation summarisation may reverse the roles in its summaries during outbound calls or callbacks, referring to the customer as the agent and the agent as the customer. This can happen when the customer uses phrases similar to those an agent or Inland Revenue staff member would say, or because the transcript shows the Inland Revenue staff member initiated the call.

Numbers

Numbers are sometimes transcribed into words instead of digits. For example, \$12 would be written as "twelve dollars."

Abbreviations and acronyms

In some cases, only the first letter of abbreviations or acronyms is capitalised. For instance, ACC would be written as "Acc."

Time limit

Conversations summarisation can only generate summaries for conversations within the first 60 minutes. Any part of the conversation exceeding 60 minutes of talk time will not be included in the summarisation. If your conversation with a customer exceeds 60 minutes, make sure to edit the summary to include all relevant points in your notes.

Unnecessary information

Generated summaries may include irrelevant or unnecessary information in its summaries, such as including when a customer was put on hold or validation questions asked.

Summary

In summary, some limitations of conversation summarisation include:

- missing some information or not generating some or all of the summary
- misinterpreting words
- being limited to the English language only
- not distinguishing between multiple customers
- reversing roles
- transcribing numbers as words
- capitalising only the first letter of abbreviations
- summarising only up to 60 minutes of conversation
- including irrelevant information.

These limitations highlight the importance of reviewing and editing the summaries to ensure they accurately reflect the conversation with the customer.

CONTINUE

Genesys Cloud dictionary

Genesys Cloud has a dictionary management tool that allows Inland Revenue to add Inland Revenue specific jargon, terms, and commonly used phrases. This ensures that the AI-generated summaries are more relevant and precise and reflects the unique language and terminology used at Inland Revenue.

We are continuously updating the Genesys dictionary so it can recognise jargon specific to Inland Revenue (for example 'FamilyBoost' or 'KiwiSaver'), and terms unique to New Zealand (for example 'RealMe'). There may be terms that has not yet been added to the dictionary that could result in inaccurate transcripts.



You can provide feedback or suggestions for the Genesys dictionary by submitting a request through the [Support Portal](#) using the '**Request for Services from Real-Time Channel Management Team**' form.

CONTINUE

AI learning

The AI behind conversation summarisation learns from the transcript of the conversations and looks for natural language patterns in the transcripts to generate summaries.

A thumbs up and thumbs down feature is available as shown in the image below which is designed to help the AI understand the quality of the summaries it generates. If the summary is mostly accurate, giving it a thumbs up is recommended, and a thumbs down if the summary is poor or requires significant editing.

How would you rate this summary?



Thumbs up and down feature in conversation summarisation

This feedback helps the AI improve its summarisation capabilities. However, whether or not the summary is edited and where it's edited does not impact the AI's learning model, as it does not use the summaries for learning.



In some cases, customers may ask if AI has been used as part of our interactions. You can explain that we use AI tools to assist with administrative tasks, but our people are always here to accurately answer technical questions and update customer accounts. You can also refer customers to information on our website:

[Our use of Artificial Intelligence \(AI\)](#)

CONTINUE

Activity

Use what you have learnt so far to answer the questions below.

True or false. Conversation summarisation can only generate summaries for conversations within the first 30 minutes.

- ☐ True.
- ☐ False.

SUBMIT

What language/s does conversation summarisation support? Select all that apply.

- ☐ French
- ☐ English
- ☐ German
- ☐ Samoan

SUBMIT

You received a call from a customer to discuss their tax code. During the call, the customer also requested to update their contact details. However, conversation summarisation only summarised the tax code conversation. Should the update to the customer's contact details be included in your notes?

☐ Yes.

☐ No.

SUBMIT

What is Genesys Cloud's dictionary management tool used for?

☐ It is used to add Inland Revenue specific jargon or phrases.

☐ It is used to look look up the definition of uncommon phrases or terms.

SUBMIT

Other resources —

Inland Revenue website

- [Our use of Artificial Intelligence \(AI\)](#).

Summary of key learnings



Conversation summarisation is a feature within the Genesys Cloud system which uses generative AI (Artificial Intelligence) designed to automatically create summaries of conversations between agents and customers.



Generated summaries must be reviewed and edited as required to ensure they accurately reflect the conversation had with the customer.



Conversation summarisation has limitations such as missing some information or not generating some or all of the summary, misinterpreting words, being limited to the English language only, not distinguishing between multiple customers, reversing roles, transcribing numbers as words, capitalising only the first letter of abbreviations, summarising only up to 60 minutes of conversation, or including irrelevant information in its summaries. Therefore, agents must review and edit summaries if it contains inaccurate or incomplete information, or if it would not make sense to someone else reading it.



Genesys Cloud's dictionary management tool allows Inland Revenue to add Inland Revenue specific jargon, terms, and commonly used phrases. This ensures that the AI-generated summaries are more relevant and precise and reflects the unique language and terminology used at Inland Revenue.

The end

That brings you to the end of this course.

Knowledge check

Once you have logged out of this course, return to Ātea and complete the online knowledge check. This has been designed to check your understanding of the information covered in this course. The pass mark is 80%. Please talk to your facilitator or team lead about further coaching and the next steps if you don't pass after your second attempt.

Logging out

To log out of this course click the **save and close** button on the top right of Ātea.

Item 2.2

This page has information about how to get access to Copilot Chat. This is the standard version of Microsoft Copilot, which is available to people in eligible roles.

Note: Starting from October 2025, a more advanced version of Copilot will be rolled out to all IR people in carefully managed waves. To learn more, go to [Microsoft 365 Copilot at IR.](#)

What is Copilot Chat?

Copilot Chat is a generative artificial intelligence (AI) tool that can make your work more productive and enjoyable, and save you a lot of time.

Microsoft offers many different Copilot products. Copilot Chat is a version of Copilot that:

- has a chat-style interface that is accessed via the Edge browser
- does **not** have access to internal IR information, except for information entered by the user.

Can I get access to Copilot Chat? Which roles are eligible?

Copilot Chat is the standard version of Microsoft Copilot. IR people can get access to Copilot Chat if they:

- have completed the [required training](#) and
- are **not** in these roles**:
 - Customer Service Officer (L1 or L2)*
 - Customer Enquiries Assistant
 - Customer Compliance Specialist (L1, L2, or L3)
 - Business Lifecycle Manager
 - Customer Support Administrator
 - Contractor in CCS:I Families
 - Project Team Member (Secondment) in CCS:I Individuals or CCS:I Families
 - Legal Services roles

* **Update March 2025:** People in the below roles can now gain access to Copilot after they have completed the required training. **You must not use Copilot for web messages or any activities related to customers.**

- From 3 February 2025, Community Compliance Officer and Customer Service Officer in Community Compliance
- From 17 March 2025, Customer Service Officer L2

- From 1 August 2025, Customer Enquiries Assistants that are part of a pilot

****Note: Starting from October 2025, a more advanced version of Copilot will be rolled out to all IR people, including customer-facing roles.** To learn more, go to [Microsoft 365 Copilot at IR](#)

Copilot onboarding process at IR

Follow these steps to get started with Copilot at Inland Revenue.

Step 1: Complete required training

To get access to Copilot, you must first complete the 'Copilot fundamentals' course in Ātea and pass the knowledge check at the end of the course. [\[link\]](#)

Note for contractors: if you don't have access to Ātea, you can request it. Please refer to this [FAQ](#) [\[link\]](#).

Step 2: Receive email when your access is ready

Once you've completed the training, our IT team will automatically set up your Copilot access by the end of the next day (if your role is eligible).

When Copilot is ready for you to use, you will be notified by email.

Step 3: Do practice exercises

After you have access to Copilot Chat, complete the practice exercises to learn the key features and understand what the tool can do.

Find out more about Copilot practice exercises [\[link\]](#)

Related reading

Using Copilot [\[link\]](#)

Item 2.4

Copilot Chat Do's and Don'ts

This page has IR's do's and don'ts for using Copilot Chat, which is the standard version of Microsoft Copilot.

Note: If you're looking for the do's and don'ts for the advanced version of Copilot, go to M365 Copilot Do's and Don'ts [\[link\]](#)

Copilot Do's

- I will always **review** the outputs of Copilot for accuracy and quality before sharing or acting upon them.
- I will **explore** how Copilot can support me in my work.
- I will **share** what I learn about Copilot with my team – being open and honest.
- I will **ask** my leader for help or clarification if I need it.
- I will **learn** to use Copilot effectively, understanding that this is a skill that will take time to develop. I will be a **proactive** and **persistent** learner.
- I will **look out** for limitations, bias, and hallucinations when using Copilot.
- I will **consider** if the information I am inputting is allowed (for more detail, see the Don'ts below).
- If I am in doubt over the use of Copilot for a particular purpose I will **check** with my leader.

Copilot Don't's

- I will not accept Copilot answers without applying my own **critical thinking and decision-making** lens.
- I will not input identifiable taxpayer information.
- I will not input sensitive staff information like payroll or personal contact details.
- I will not input data classified at SENSITIVE or above [\[link\]](#)
- I will not input any:
 - Sensitive Revenue Information (such as taxpayer-specific information); or
 - Revenue Information that, if released, could adversely affect the integrity of the tax system or would prejudice the maintenance of the law (such as compliance thresholds or investigative techniques); or
 - Other sensitive information (such as commercially sensitive, employment or personal information).
- I will not use Copilot to translate information into other languages.

Item 2.8

Extract from 1 October 2025 minutes of Strategic Investment Board meeting

Item: M365 Copilot deployment planning

The Board:

- **Agreed** that each business area at tier 3 level will nominate a key contact that has good visibility of their area's needs and can make decisions on behalf of the area, including the most appropriate time to deploy the tool given business needs and priorities.
- **Agreed** that each business area key contact will be required to complete a checklist before access is provided to the tool, confirming matters such as whether training has been undertaken, controls are in place, and readiness assessment has been completed. This checklist will be amended based on feedback and lessons from earlier waves of deployment.

Internal intranet page information: M365 Copilot rollout information for senior leaders

This page is for the Deputy Commissioner or Tier 3 responsible for approving the rollout of M365 Copilot in their business area.

What's happening

On 25 September, our executive team approved the organisation-wide rollout of Microsoft M365 Copilot (M365 Copilot), a powerful Artificial Intelligence (AI) tool. It is part of a broader digital strategy and aims to help IR people to work smarter, reduce admin, and focus higher value tasks. We want to start making the most out of this exciting opportunity.

Before your business area starts using M365 Copilot, we want to make sure you and your teams are ready. We are committed to rolling out M365 Copilot in a controlled way to ensure our people are well supported and risks are appropriately managed. This page has information to help you decide if your business area is ready to get access to M365 Copilot.

Read this page and then complete the Microsoft form at the bottom of this page.

About Microsoft 365 Copilot

M365 Copilot is a more integrated and improved version of Copilot Chat, a tool that many IR people are already using each day.

M365 Copilot is a generative and agentic AI tool that is integrated into the Microsoft 365 applications like Teams, Outlook and Word. This version of Copilot also has access to our internal information that is contained within Microsoft 365. It does not have access to STAX, START or the DIP but does have access to any extracts of content from these tools that have been stored in Microsoft 365 storage locations such as SharePoint.

This tool has been piloted with 140 people in our organisation and is used widely in New Zealand and internationally. It saves time, enhances work quality, and reduces stress and cognitive load.

What you'll need to consider before approving the rollout of M365 Copilot to your team

Are you comfortable with your data readiness?

We will ask you if you understand and are comfortable with the increased chance of content being surfaced that people were not aware they would have access to. Copilot will only allow your people to access information they already have permission to see (for more information

see Enterprise data protection [link]), as it respects our organisation's existing content permissions.

However, Copilot might surface information that people *do* have access to but perhaps were not aware of. For example, if a file in your Teams or SharePoint site was unintentionally shared broadly within IR, Copilot could include it in responses for your people (since they have access).

Consider the Teams and SharePoint sites your people typically use and whether the right people have access to these spaces. Some guidance on checks you can complete yourself is available on the Getting started with Microsoft 365 Copilot page [link].

If you would like extra support from the IKM team to check whether any specific content is secured, you can use this form to request a check: IKM - Request a Permissions Check [link].

Note: Submitting this form is not compulsory and does not stop your area from receiving M365 Copilot if you approve the rollout for your area.

Do your people have capacity to complete mandatory activities to get ready for M365 Copilot?

Before getting access to M365 Copilot, every person must complete our existing Copilot Fundamentals course in Ātea, many people have already completed this course to get access to Copilot Chat. There will however be a new set of 'Dos and Donts' that are specific to M365 Copilot that people will need to agree to. For more information on what we will be asking people to do in preparation for Copilot check out Getting started with Microsoft 365 Copilot.

Do your people have time to complete Copilot learning?

Rolling out M365 Copilot effectively isn't just about our data and technology – our people need to be ready to engage with this change and learning.

When people begin using Copilot, there is a learning period to understand effective prompting techniques and which tasks it can efficiently handle. As our people familiarise themselves with Copilot and develop their prompting skills, some tasks may initially require more time to complete. It is key that people are given time to experiment with Copilot to grow their skills.

Learning opportunities open to you and your people

Adjusting to Copilot takes time, as people learn effective prompting and identify the most valuable use cases. While learning will support this, much of the benefit will come from

ongoing experimentation and exploration. It is important to acknowledge that this will take time.

To provide guidance on how to begin, we will be offering three 45 minute sessions focused on optimising the use of Copilot:

- Introduction to Copilot
- Teams and Copilot Chat
- Documents

These will be run twice weekly for a couple of months along with a weekly optional drop-in session.

We strongly encourage your team members to attend these sessions. There will be a weekly half hour drop-in session run each week where people can ask our trainers questions.

The total time to complete both the required and recommended learning will be **less than 3 hours**.

Consider your area's current workload and whether there are any major events or deadlines that coincide with the suggested timing of your rollout. We can adjust the schedule to accommodate capacity and workload to ensure that your people have the time and space to familiarise themselves with Copilot.

Key contact

This contact role will be the key conduit for information, they will attend briefings and keep you updated on key information and details. You will need to consider now the right person for this role as a person who is knowledgeable about your area and its connection with IR's wider operations, and be capable of making operational decisions where necessary.

Over the coming months we will seek input from the key contact you nominate from your business area to understand what the highest-value use cases are so these examples can be shared with other IR people to help them make the most value from the tool and to assist with benefits measuring and reporting. Purchasing all of our organisation licences is a significant investment, we want to make sure that we are seeing the benefits of this tool. As part of this process, we will also be considering how we measure the benefits from the key use cases

identified for each business area, ensuring that the value and impact of M365 Copilot can be clearly demonstrated and tracked over time.

What might your team use Copilot for?

We will ask you what tasks you think Copilot will be able to help your people with?

How M365 Copilot can help your team

Copilot is designed to help with a wide range of tasks like drafting content, summarising information and answering questions about our organisation. To help you understand the value of M365 Copilot, we encourage you to watch this short video of pilot participants sharing their experiences. Feedback from the pilot has been overwhelmingly positive and we expect similar benefits for people across IR - M365 Copilot – what our people said [\[link\]](#).

Understanding what your team might use it for

If you still aren't sure how this could benefit your team input this prompt into Copilot Chat to get it to help you think about how this tool can benefit your team:

"I want to explore how Microsoft 365 Copilot could be used in my business area. Please ask me questions to understand the nature of the work we do, the types of documents, communications, and processes we use, and any pain points or inefficiencies we face. Once you've gathered enough information, summarise the potential use cases for Copilot that could help improve productivity, decision-making, or collaboration in our area. Keep the number of questions to 3 in total."

Do you agree to share communications with your team to let them know about this change?

As the leader of your business area, it's essential that you visibly support and champion this change to foster engagement and confidence among your team.

Once you have completed the form and approved the inclusion of your team members in the upcoming rollout, you can use the template below to communicate this information.

If you approve of the timing of this rollout: Communications to your team on the rollout – approved [\[link\]](#)

If you do not approve of the timing of this rollout: Communications to your team on the rollout – deferred [\[link\]](#)

Please choose the communication channel that best fits your team's context, such as a Teams post or an email, and adjust the content as needed.

Note: Be sure to include the Experience IT mailbox in your communication so your people know where to direct any questions regarding Copilot.

Complete the rollout approval form

Once you have read the above information, please complete this Microsoft form to confirm you are comfortable proceeding with the rollout to your people: Approving the rollout of M365 Copilot to your teams [\[link\]](#).

M365 Copilot Do's and Don'ts

This page has IR's Do's and Don'ts for using Microsoft 365 Copilot (M365 Copilot).

Before you're given access to the tool, you'll be asked to confirm that you understand and agree to these do's and don'ts.

Note: **These do's and don'ts are for M365 Copilot, which is the advanced version of Copilot.** If you're looking for the do's and don'ts for the standard version of Copilot, go to Copilot Chat Do's and Don'ts [\[link\]](#).

Introduction

Before you start using M365 Copilot, including Microsoft's pre-built agents like Researcher, Analyst, and Project Manager, you'll need to complete the Copilot Fundamentals training and confirm that you understand and agree to the following Do's and Don'ts.

Information that can and can't be used within M365 Copilot

I understand I am permitted to use the following categories of information within M365 Copilot provided I have the right to access the information:

- information classified up to and including SENSITIVE information but not RESTRICTED information or above;
- commercially In-Confidence information; and
- information that is Sensitive Revenue Information under the Tax Administration Act (e.g. information reasonably capable of being used to identify a taxpayer, and personal information including information about Inland Revenue staff).

This permitted use of information in M365 Copilot is an approved exception to our AI staff use policy. This exception to the policy has been approved by Jay Harris (Chief Information Security Officer). Apart from this exception the AI staff use policy continues to apply.

M365 Copilot Do's

- **Before using** M365 Copilot, I will check my OneDrive and shared storage locations (SharePoint shared workspaces, MS Teams) that I am responsible for to ensure the folders are set up with the correct permissions and any obsolete content has been deleted.
- I **understand** that M365 Copilot has powerful search functionality that may give me access to information and folders that I wasn't aware I had access to.
- If I can access information I don't think I should have access to, I will **notify** my people leader and the IKM team to correct the matter. For instructions on how to report this, go to What to do if Copilot surfaces something you shouldn't have access to [\[link\]](#).

- I will always **review** the outputs created by M365 Copilot for accuracy and quality before saving, sharing or acting on them.
- I will **check the citations** in all M365 Copilot outputs to ensure:
 - Copilot is using the correct and most relevant version of information that is needed; and
 - I understand where Copilot sourced the information from (such as a budget secret folder) so I'm aware who I am authorised to share the outputs with.
- Subject to the other dos and don'ts, I will **explore** how M365 Copilot can support me in my work.
- I will **share** what I learn about M365 Copilot with my colleagues – being open and honest. I will act in accordance with my leader's instructions regarding my input on progress, feedback and learnings.
- I will **ask** for help or clarification on using M365 Copilot if I need it.
- I will **learn** to use M365 Copilot effectively, understanding that this is a skill that will take time to develop. I will be a **proactive** and **persistent** learner.
- I will **look out** for limitations, bias, and inaccuracies in outputs when using M365 Copilot.
- To protect information security and privacy, I will **think before recording** a Microsoft Teams meeting and will notify all participants before I do so. If the meeting changes to a sensitive topic, I will consider if the recording should be stopped. Refer to Recording Microsoft Teams meetings [[link](#)]
- **If I am in doubt** over my use of M365 Copilot for a particular purpose, including what information I can use, I will check with my leader.
- I will only use M365 Copilot to carry out my Inland Revenue duties and I will not use it for personal use.
- 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.

M365 Copilot Don'ts

- I will not accept M365 Copilot answers without applying my own **critical thinking and decision-making** lens.
- I will not use Copilot to assist me with formally translating, analysing or producing Te Reo Māori content and instead defer to our internal processes to ensure the resulting work is accurate and culturally appropriate.
- I will not copy and paste taxpayer information from START or Genesys into Microsoft applications. However I understand I am permitted to type such information into M365 Copilot.

M365 Copilot

Privacy Assessment (condensed)

Prepared by:

Date:

Supply ID:

About this Document

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

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

Project Summary

1.1 Description

In November 2024 Inland Revenue (IR) rolled out Microsoft Copilot Chat, in a staged approach, to staff. This is a generative artificial intelligence (AI) tool that has a chat-style interface accessed via the Edge browser and does not have access to internal IR information, except for information entered by the user.

Microsoft 365 Copilot (M365 Copilot) is integrated into the Microsoft ecosystem. It accesses existing content within Microsoft applications including emails, documents, calendar events, Teams chats, and files stored in OneDrive and SharePoint.

When a user interacts with M365 Copilot (e.g., asks it to summarise a document or draft an email), the system sends the prompt 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, which is rendered within the user's Microsoft 365 app (e.g., Word, Outlook, Teams).

The LLM is powered by Azure OpenAI which is a Microsoft hosted version of ChatGPT. This model sits within the Microsoft boundary, meaning that it is shared across Microsoft customers

and 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](#)

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 with licenses 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 Personal information to be used

M365 Copilot has access to internal information that is contained within Microsoft 365. This includes emails, documents, calendar events, Teams chats, and files stored in OneDrive and SharePoint. M365 Copilot does not have access to STAX, START or the DIP but does have access to any extracts of content from these tools that have been stored in Microsoft 365 storage locations such as SharePoint. Some content may include potentially 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 categorises what could possibly be returned by the tool.

Type of personal information	Source of information	Purpose of using the information
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	
Personally Identifiable Information (PII): There is potential for PII to be accessed during summarisation or drafting tasks, especially if users interact with	Emails, chats or documents stored in Microsoft 365 apps containing names, contact details, or other identifiers. This could include employment-related information	
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 if stored in a Microsoft 365 apps	

1.4 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 and AI Oversight Group.

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	<p>IR has an AI staff use policy 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.</p> <p>Within M365 Copilot users are permitted to use:</p> <ul style="list-style-type: none"> • information classified up to and including SENSITIVE information • commercially In-Confidence information; and • information that is Sensitive Revenue Information under the Tax Administration Act (e.g. information reasonably capable of being used to identify a taxpayer, and personal information including information about Inland Revenue staff). <p>provided they have the right to access the information.</p> <p>This permitted use of information in M365 Copilot is an approved exception to IR's <u>AI staff use policy</u>. This exception to the policy was approved by Jay Harris (Chief Information Security Officer). Apart from</p>

this exception the AI staff use policy continues to apply.

Targeted training is being provided to help 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.

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 is already held by IR. It does not collect any new information.

Is collecting this information necessary for IR to carry out its functions?

N/A

Where or who is the information being collected from?

The information will be collected from a user's Microsoft applications.

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?

N

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.

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.

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.

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.

We accept there is an increased chance of content being surfaced that people were not aware they would have 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 *do* 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.

<p>How long will the information be retained?</p>		<p>The IKM team published guidance on how to report if people see a file or document they think they shouldn't have access to: Reporting accidental access to a document or file.</p>
		<p>Any documents saved within M365 applications are subject to IR's Retention & Disposal Authority under the Public Records Act.</p> <p>By default M365 Copilot retains prompts for 93 days. As at 23 October 2025 the IKM team is looking into whether we can keep prompts for as short a time as possible (TBC).</p>
<p>Will IR be notified when the information is no longer needed for the specific purpose?</p>		<p>N/A</p>
<p>Will the information be disposed of? If so, explain how</p>		<p>N/A</p>
<p>Does it involve transferring personal information offshore, using a third-party contractor?</p>	<p>Y</p>	<p>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.</p>
<p>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.</p>		<p>Security credentials and certifications for Microsoft are publicly available Service Trust Portal Home Page (microsoft.com):</p>
<p>Use, disclosure, and accuracy</p>	<p>Y/N</p>	<p>If yes, explain your response</p>
<p>Is the information currently held by IR?</p>	<p>Y</p>	
<p>If yes to the above question, for what purpose does IR hold the information?</p>		<p>Information stored by IR staff in Microsoft will be relevant to their role and linked to an IR function.</p>

Will the initiative use or disclose information for a different purpose to why it was obtained?	N	<p>Information held within the 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.</p> <p>Users agree to only use M365 Copilot to carry out IR duties and not to use it for personal use.</p>
Will IR be sharing personal or taxpayer information with another organisation?	N	
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?		<p>Quality of the data varies by business area.</p> <p>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.</p> <p>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.</p>
Access and identification	Y/N	If yes, explain your response

Will the information be stored on a customer or staff member's record?	N	
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	
If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?		It's not expected that M365 Copilot will access information a customer has freely provided. The tool does not access START or DIP. Given IR's operational context, there may be some documents or communications which contain sensitive revenue information if the user has a legitimate business purpose to access this content and it is saved on a Microsoft application.
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
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	
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?		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?

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	<p>Given the types of use cases M365 Copilot will be used for (surfacing and summarising documents and team meeting admin for instance), it is not anticipated a significant amount of personal information will be used but this is possible.</p> <p>Sensitive Revenue Information has been approved to be input into M365 Copilot as an exception to IR's AI staff use policy.</p>
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	Medium	<p>If a document has previously been shared with a user, and their permission has not been revoked, the document may be visible to the user. Due to a lack of security on a document, or it being shared with all of IR instead of specific people, it's also possible that 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</p>

<p>details, information about high profile individuals)</p> <p>Significance of the changes</p> <p>L – Only minor change to existing functions/activities</p> <p>M – Substantial change to existing functions/activities; or a new initiative</p> <p>H – Major overhaul of existing functions/activities; or a new initiative that's significantly different</p> <p>Interaction with others</p> <p>L – No interaction with other agencies</p> <p>M – Interaction with one or two other agencies</p> <p>H – Extensive cross-agency (that is, government) interaction or cross-sectional (non-government and government) interaction</p> <p>Public impact</p> <p>L – Minimal impact on IR and customers</p> <p>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</p> <p>H – High impact on customers and the public, and concerns over aspects of project; widespread media interest likely</p>	<p>Medium</p> <p>information about how staff can report accidental access to a document or file.</p> <p>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.</p> <p>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.</p>
<p>Medium</p> <p>Microsoft has provided training to key IR staff and supports further training for users.</p>	<p>Medium</p> <p>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 (exceptions being ACC and MBIE) 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.</p> <p>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.</p>

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

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

7. Recommendations

Based on the above assessment, below are summarised recommendations to minimise the impact on privacy. These should be agreed with the senior responsible owner.

Ref	Recommendation	Agreed Y/N
R-01		

8. Document sign-off

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

DRAFT

Viva Topics Pilot / Rollout Privacy Threshold Assessment

Prepared by: Michael Wright

Date: 18 July 2023

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?
-

1. Project summary: Viva Topics Pilot / Rollout

1.1 Description of the project

Viva Topics is a knowledge management product from Microsoft. Part of the broader Viva suite of applications, it integrates into our Microsoft 365 (M365) environment.

Microsoft describe Viva Topics as ‘...applying AI to empower people with knowledge and expertise in the apps they use every day, and to connect, manage, and protect content across systems and teams.’

Viva Topics analyses our content and automatically identifies the topics that are relevant to us as an organisation. It then compiles articles known as ‘topic pages’ that bring together key information on each topic, including:

- Alternative names or terms for the topic (including acronyms and initialisms)
- A brief description of the topic
- A list of people in the organisation who potentially hold knowledge on the topic
- A list of resources (sites, pages and documents) relating to the topic
- Linked to other related topics

The AI efforts can be supplemented with human curation over the information produces including, potentially, opening the information to the whole organisation for ‘crowdsourcing’.

Finally, Viva Topics can present this topic information in context across M365 applications. Examples of how it does this:

- Automatic 'highlights' of topics mentioned in SharePoint pages.
- Users can add topics to Teams chat messages or Yammer posts.
- Relevant topics appear pinned to the top of SharePoint search.

Inland Revenue are piloting Viva Topics with the intent of informing a business case for the adoption and rollout of the tool across Inland Revenue.

Note that Viva Topics also serves as an enabler for Answers in Microsoft Viva, a new capability within Viva Engage (previously known as Yammer) that supports knowledge sharing in the form of questions and answers across the organisation. IR will also be piloting Answers in Microsoft Viva in 2023, with a view to including that application in the same business case. This will be covered in a separate Privacy Threshold Assessment.

1.2 Purpose of the change

Viva Topics has potential to enhance IR's knowledge management capability and practices in a number of key areas:

- **Tacit knowledge / people expertise:** Managing and making available information on who holds particular subject matter knowledge across IR.
- **Reducing fragmentation:** Bringing together related resources and knowledge from various teams and spaces throughout the organisation.
- **Collaboration:** Improving visibility of who is doing what, and who knows what, to facilitate better collaboration between different parts of IR.
- **Knowledge discovery:** Making knowledge readily available to people in context, to reduce discovery time and improve productivity.

Improved access to knowledge for our people will improve our ability to make decisions, provide advice and deliver change. By extension, this will lead to improved services and outcomes for our customers and for the country.

1.3 Privacy Enhancement

We expect Viva Topics to be largely neutral from a privacy perspective. The information it aims to synthesise and present is generally already available across the whole organisation, albeit in a different form. However, it does present some privacy risks that need to be managed.

Viva Topics respects the underlying security settings and access controls in place for the source material it uses; it will not present any topic information to a user where they do not already have access to the source material. This includes the very existence of a topic, if revealing the topic name would provide restricted information (for example, the existence of a particular Policy initiative).

Note that some topic information (including title, alternative terms, and description) will become available to the whole of IR if a human curator chooses to confirm and publish an AI-suggested topic, or to create a new topic from scratch.

The information currently in scope for AI topic discovery is currently limited to IR's SharePoint environment (including Teams sites). Managing privacy risks can be achieved through a combination of:

- Security access controls on sites that hold sensitive information (existing practice)
- Applying correct security classifications to our documents (existing practice)
- Excluding specific high-risk workspaces from topic discovery (an IKM responsibility).

- Governance over who takes up the 'knowledge manager' role¹ for Viva Topics, with oversight over topics and the ability to approve/publish or reject suggested topics.

Currently we are excluding a small number of PARS and Budget workspaces from AI discovery. This means that information from these spaces will not be used to identify topics or topic-related resources. In future we expect to exclude any other sites that contain large amounts of customer information or information rated higher than 'IN CONFIDENCE'.

With appropriate controls in place, the only personal information expected to be held in Viva Topics is the information about human expertise – i.e. the links between IR people and topics. This information may be displayed on the topic pages and on people's M365 profile cards.

Broadly speaking, this information is already 'public' in the sense that it is drawn from document metadata and similar information that is already available to readers. However, Viva Topics greatly amplifies the visibility of this information and presents it in new contexts. Note as well that parts of IR already maintain a number of shared skill / knowledge registers for various purposes.

Individuals have the ability to opt out of being listed against a topic. They will be prompted to confirm or remove their own connections when visiting the Viva Topics home site, or potentially via other means such as their daily briefing email.

1.4 Personal information that the project will involve

Type of personal Information	Source of Information	Purpose of information for the project
Knowledge / expertise (at a high level)	Documents and articles across M365 SharePoint (including Teams) IR people.	Facilitating collaboration and knowledge sharing by improving visibility of who holds knowledge on specific topics.

2. Privacy assessment

2.1 Areas that are risky for privacy

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
Information management generally			

¹ Microsoft's term. In IR we typically refer to them as Curators or Facilitators.

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
<p>A substantial change to an existing policy, process or system that involves personal information</p> <p>Example: New legislation or policy that makes it compulsory to collect or disclose information</p>	✓		A new knowledge management system (software) that provides a new layer to our information architecture.
<p>Any practice or activity that is listed on a risk register or your specific business unit's risk register</p> <p>Note: Check your business unit's risk register and with Risk Services</p>		✓	Broadly speaking, enterprise risk 5 is applicable to any new information system such as Viva Topics. IKM owns sub-control 276 <i>Information and Content Management Strategy and Practice</i> which applies here. We are already applying existing best practices to the information in Viva Topics.
Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response

Collection

<p>A new collection of personal information</p> <p>Example: Collecting data not collected before, collecting information about a customer's location</p>	✓		Information about people's knowledge areas will be determined from document metadata. IR people can also manually add people to topics (themselves and others).
<p>Collecting information which is not necessary for IR to carry out its functions</p> <p>Example: Information is not relevant to tax administration</p>		✓	Not all topics will be tax-related, however all should be relevant to running the tax system in some way (including running IR itself).
<p>A new way of collecting personal information</p> <p>Example: Collecting information through a new app or using facial recognition</p>	✓		Use of an AI product plus potential crowdsourcing to identify and record people's subject knowledge areas.
<p>Collecting information from someone other than the individual themselves</p>	✓		Information can also be provided by other IR people.

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
Example: Contacting a person's employer to obtain information			
Storage, security and retention			
A change in the way personal information is stored or secured Example: Storing information in the Cloud (if so also contact ICT for advice)	✓		Information about our people's knowledge areas will be stored in a new way that is much more visible to the whole organisation.
A change to how sensitive information is managed Example: Moving financial records to a new database		✓	
Transferring personal information offshore or using a third-party contractor Example: Outsourcing the payroll function or storing information in the Cloud		✓	Viva Topics stores data in the same locations as IR's wider M365 environment.
Making IR information available to another agency for it to use and retain. Example: If sharing taxpayer information with another agency ensure it is not kept for longer than necessary		✓	
A decision to keep personal information for longer than IR has previously Example: Changing IT backups to be kept for 10 years when previously only stored for 7		✓	Retention of source material used by Viva Topics will not change. How long we keep topics themselves (including linked people expertise) has yet to be determined.
Use or disclosure			
A new use or disclosure of personal information that is already held Example: Sharing information with other agencies in a new way		✓	
Sharing or matching personal information held by different		✓	

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
organisations or currently held in different datasets <i>Example: Combining information with other information held on public registers, or an information matching or sharing agreement</i>			
Linking or matching personal information across government <i>Example: Is there an existing data sharing agreement? Will a new arrangement be necessary?</i>		✓	
Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
Individuals' access to their information			
A change in policy that affects how people can access information that IR holds about them <i>Example: A system that does not allow the ready access of information</i>		✓	No change in policy, however Viva Topics increases people's ability to see certain information held about them by IR and provides them the ability to update that information.
Identifying individuals			
Establishing a new way of identifying individuals <i>Example: A unique identifier, a biometric, or an online identity system</i>		✓	
A new way of linking individuals or entities in a database		✓	
Other considerations	Yes (tick)	No (tick)	If yes, explain your response
Is there a way to achieve the purpose of the project using less identifiable data?		✓	
Would people be surprised by this use of their information?		✓	Pilot users have responded well to the people expertise aspects of

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
			Viva Topics, noting it as the greatest potential value-add this tool offers. No material concerns about privacy have been raised.
If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?		✓	
Are there processes in place to ensure and maintain data integrity?	✓		Existing content access permissions are respected by Viva Topics. IKM are excluding more sensitive information by source. We have a manual curation process in place. People can edit their own topic information. DLP is also in the early stages of being implemented in IR's M365 tenancy.

Does the project involve any of the following?	Yes (tick)	No (tick)	If yes, explain your response
--	------------	-----------	-------------------------------

New intrusions on individuals' property, person or activities

Introducing a new system for searching individuals' property or premises		✓	
Surveillance, tracking or monitoring of movements, behaviour or communications Example: <i>Installing a new CCTV system or GPS in vehicles</i>		✓	
Changes to premises that will involve private spaces where clients or customers may disclose personal information Example: <i>Co-location or changing the location of a reception desk, where people may discuss personal details</i>		✓	

3. Ethics and social licence

3.1 Areas that may raise ethical issues

Does the project use ethical data practices?	Yes (tick)	No (tick)	If yes, explain your response
Will the proposal discriminate against some people?		✓	
Is the proposal likely to result in some members of a group being treated differently to one another?	✓		Potential for some people to be 'favoured' for inclusion on topics as they work in M365 more than others. People associated with topics may experience an increase in demand for their time and knowledge.
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?		✓	
Can you foresee any harm to individuals in using the data in the way intended?		✓	
Have you considered how the proposal contributes to the active protection of Māori interests?	✓		We have engaged with the Domain Principal, Māori Crown Relations, regarding the potential use of Viva Topics to protect te reo and te ao Māori concepts and encourage consistent and proper use across IR. We anticipate that topics relating to Māori terms and concepts will be curated and managed by key members of IR's Māori Network.
If using algorithms is there confidence the algorithm is robust and assumptions are met?	✓		Viva Topics does not involve any material decisions that directly affect people. We are confident the AI behaviour is fit for purpose and subject to human curation and note that it is continually improving.

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).

If you answered “No” to all the questions in 2.1 above, move on to section 3 below.

Aspect of the Project	Rating (L, M or H)	Describe any medium and high 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	L	The only personal information gathered will be high-level association between individual IR people and the topics they may be knowledgeable on.
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, salary information, biometric data) H – The information will be highly sensitive (health or financial details, information about high profile individuals)	L	<p>The personal information gathered and presented in Viva Topics is generally open (to IR people) anyway, as it relates to the work people do for IR.</p> <p>More sensitive personal information should not be involved, as it will be protected by a combination of access controls, security classification, and exclusion of sources from AI discovery.</p>

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	L	Viva Topics is not intended to change any functions or activities. It merely aims to improve productivity and the flow of knowledge across IR.
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	L	The information in Viva Topics is intended for internal IR use only.
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	L	From a privacy perspective, Viva Topics will not affect customers at all.

5. Summary of privacy impact

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

Reduced risk – The project will lessen existing privacy risks	
Inadequate information – More information and analysis is needed to fully assess the privacy impact of the project.	

5.1 Reasons for the privacy impact rating

Viva Topics is an organisational knowledge management product. It aims to improve the utility of what the organisation already knows, rather than providing 'new' information per se. Generally, it is simply making it easier for people to find and access resources they already have access to.

Subject to the controls described above, the only situation where it will handle any personal information is in relation to the topics on which an individual might hold expertise. Broadly speaking this is information already available to the organisation, albeit not in such an accessible form. Exposing this information in this way has been perceived as a logical (and positive) step by pilot participants. There is little risk of any individual being harmed by this.

6. Document Sign off

Position	Name	Sign off Date
Project Manager	Amie Butler, Domain Lead, Information & knowledge Management	18 July 2023
Privacy Officer	Dawn Swan, Privacy Officer	19 July 2023
Aidan Roberts	Information Security Officer, CISO Office	15 June 2023

7. Appendices

7.1 Topic page (example)

Initiative

Te Pou o te Tangata

Alternate names: Organisational behaviours, How we do things at IR, Te Pou O Te Tangata, Te Pou o Te Tangata

Te Pou o te Tangata describes how we do things at IR and will start to replace our existing cultural anchors, values and behaviours from March 2023.

We set out to create a single set of meaningful descriptors that better reflect who we are as an organisation and ensure we can continue to deliver ongoing benefits following our 6-year transformation.

They should connect us more closely to our vision of becoming a world class revenue organisation and reflect the importance of Māori Crown relations.

Te Pou o te Tangata – How we do things at IR reflects te ao Māori concepts which will become embedded in the way we work IR.

This will be supported by new learning for IR people around the descriptors, a transformed organisational induction and learning to support leadership capability.

Policy

Strategy

Project

Sources

Confirmed people

S 9(2) (a)

Amy Hartnell

Domain Specialist (L2)

Key contact for Te Pou o te Tangata.

S 9(2) (a)

Emily Hookham

Contractor

Key contact for Te Pou o te Tangata.

T

Te Pou o te Tangata

Initiative team mailbox

Suggested people

S 9(2)

Sam Evenson

Contributed to resources

S 9(2)

Lynne Barks

Contributed to resources

S 9(2)

Nikita Meehan

Contributed to resources

S 9(2)

Stephanie Alderson

Contributed to 3 resources

S 9(2)

Bronwyn Yates

Contributed to resources

S 9(2)

Jarrod Rendle

Contributed to resources

S 9(2)

TeeJay Bannister

Contributed to resources

S 9(2)

Hannah Swasbrook

Contributed to [Te Pou o te Tangata - where to from here?.as...](#)

See more











Pinned files and pages

	Name	Language	Modified

HAUKĀINGA SEARCH

by

	About Te Pou o te Tangata CO-IP-organisational-behaviours > Site/Pages	English	10 October 20...	Nikola Meehan	Michael Wright
	Te Pou o te Tangata Yammer www.yammer.com/main/any/any?...com/group/any/fcltheZS6kdyuVw...				Michael Wright
	Our Te Pou o te Tangata behaviours CO-IP-organisational-behaviours > Site/Pages	English	4 April	Nikola Meehan	Michael Wright

	Name	Language	Modified	Modified by	Activity
	Discussion guide - Te Pou o te Tangata CO-IP-organisational-behaviours > Shared Documents	English	2 February	Nikola Meehan	Top c. mentione
	Te Pou o te Tangata virtual open home now available News > Site/Pages	English	25 May	s 9(2)(a)	Top c. mentione
	New: Te Pou o te Tangata learning News > Site/Pages	English	4 days ago	s 9(2)(a)	Top c. mentione
	Te Pou o te Tangata open homes are underway News > Site/Pages	English	27 April	s 9(2)(a)	Top c. mentione
	Te Pou o te Tangata - draft descriptors ready for your feedback News > Site/Pages	English	20 October 20...	s 9(2)(a)	Top c. mentione
	Watch: An inside perspective of Te Pou o te Tangata News > Site/Pages	English	30 November ...	s 9(2)(a)	Top c. mentione
	Te Pou o te Tangata – What we heard and next steps News > Site/Pages	English	14 November ...	s 9(2)(a)	Top c. mentione
	Our Te Pou o te Tangata behaviours CO-IP-organisational-behaviours > Site/Pages	English	24 April	Nikola Meehan	Top c. mentione
	Te Pou o te Tangata - where to from here? News > Site/Pages	English	11 September ...	Hannah Swaisb...	Top c. mentione

See more ▾

Pinned sites



Te Pou o te Tangata - Organisational ...
Pinned by you
23 February

Suggested sites



Kawepūrongo - News
Suggested based on resources



Learning Strategy and Organisational...
Suggested based on resources



SharePoint Site Resources
Suggested based on resources



To mātou Tari - Our organisation
Suggested based on resources



Recruitment and Onboarding
Suggested based on resources



Māori Network
Suggested based on resources



Te Pou Herenga - People Leaders Sp...
Suggested based on resources

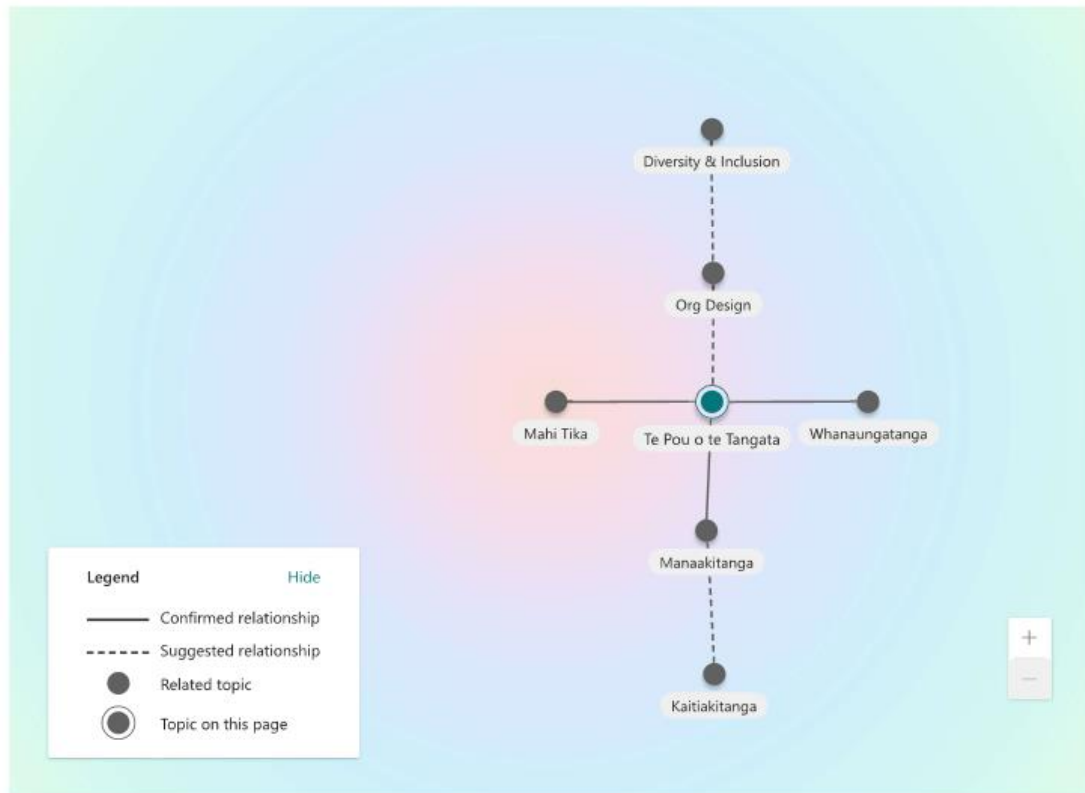


SW-Information & Knowledge Mana...
Suggested based on resources



See more ▾

Related topics



7.2 Topic highlight - SharePoint

OUR ORGANISATION

New: Te Pou o te Tangata learning

New learning on our [Te Pou o te Tangata](#) behaviours is now available to help you better understand Te Pou o te Tangata and what it means for you.

Te Pou o te Tangata project team understands there's no 'one size fits all' when it comes to learning. Our new approach means there's learning to suit different needs and roles, and to meet you wherever you are in your Te Pou o te Tangata journey.

7.3 Topic card - SharePoint

our Te Pou o te Tangata behaviours is now available

ata a

ata p
w ap
u are

6 le
rt e-l
talki

o te
veryo
o be

ss t

be f


set c
! Thi

own Te Pou o te Tangata journey.

Strategy

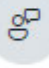
Viva Topics

Te Pou o te Tangata

 View details


Alternate names: Organisational behaviou...

Te Pou o te Tangata describes how we do things at IR and will start to replace our existing cultural anchors, values and behaviours from March 2023. We set out ...
More

 Improve AI suggestions ×
Do the topic details for Te Pou o te Tangata match the context it was highlighted in?

People (10+) >

s 9(2)
(a)

Amy Hartnell 
Domain Specialist (L2)
Key contact for Te Pou o te Tangata.

7.4 Topic tag and card – Teams

12:54

Let's talk about #Te Pou o te Tangata at today's meeting.



Strategy

Viva Topics

Te Pou o te Tangata

View details

Alternative names: Organisational behaviou...

Te Pou o te Tangata describes how we do things at IR and will start to replace our existing cultural anchors, values and behaviours from March 2023. We set out t...

[More](#)

People (10+) >

s 9(2)
(a) Amy Hartnell
Domain Specialist (L2)
Key contact for Te Pou o te Tangata.

s 9(2)
(a) Emily Hookham
Contractor
Key contact for Te Pou o te Tangata.

7.5 Topic result (pinned) in SharePoint search

The screenshot shows a SharePoint search interface. At the top, a search bar contains the text 'Te Pou o te Tangata'. Below the search bar, there are tabs for 'All', 'Files', 'Sites', 'People', 'News', 'Images', and 'Power BI'. Under the 'All' tab, there are filters for 'File type' and 'Last modified'. The main content area displays a pinned result for 'Te Pou o te Tangata' under the 'Strategy' category. This result includes a description, alternate names, and sections for 'People (10+)', 'Resources (10+)', and 'Related topics (7)'. Below the pinned result, there are three additional search results: 'New: Te Pou o te Tangata learning', 'Te Pou o te Tangata - Organisational behaviours', and 'Te Pou o te Tangata learning'.

Te Pou o te Tangata

Alternate names: **Te Pou o Te Tangata**, Organisational behaviours, How we do things at IR

Te Pou o te Tangata describes how we do things at IR and will start to replace our existing cultural anchors, values and behaviours from March 2023. We set out to create a single set of meaningful descriptors that better reflect who we are as an organisation and ensure we can continue to deliver...

People (10+)

s 9(2) (a) Amy Hartnell
Key contact for Te Pou o te Tangata.

s 9(2) (a) Emily Hookham
Key contact for Te Pou o te Tangata.

Resources (10+)

About Te Pou o te Tangata
Nikita Meehan modified on Oct...

Te Pou o te Tangata - Orga...
Site

Related topics (7)

Whanaungatanga

Whanake

Org Design

Was this answer useful?

New: Te Pou o te Tangata learning ...
<https://irnz.sharepoint.com/sites/News/SitePages/New--Te-Pou-o-te-Tangata-learning.aspx>
Stephanie Alderson published 3 minutes ago
...**Te Pou o te Tangata** behaviours is now available to help you better understand **Te Pou o te**...
and what it means for you. **Te Pou o te Tangata** project team understands there's no 'one... yo...

Te Pou o te Tangata - Organisational behaviours
<https://irnz.sharepoint.com/sites/CORP-organisational-behaviours>
...**Te Pou o te Tangata** Corporate Space. Here you will find information on the initiative, **Te... o te**
Tangata behaviours, tools and resources, and what's coming up in the learning space...**Te Pou** ...

Te Pou o te Tangata learning ...
Te Pou o te Tangata - Organisational behaviours
Nikita Meehan modified 4 days ago
...**Te Pou o te Tangata** behaviours. There's no right or wrong order to complete the learning...
we suggest you start with 'Our journey to **Te Po.....Te Pou o te Tangata** behaviours. There's no...

Inland Revenue

Privacy Impact Assessment

Use of START Analytics Manager

Senior Responsible Owner: Tony Morris

Prepared by: s 9(2)(a) Kristina Elfstrom, Al
Warren

Date: 07/02/2023

In Confidence

About this Document

The purpose of this document is to fully consider all privacy risks this proposal raises and how they will be mitigated. You will have completed the Brief Privacy Analysis and information from that document can be included. The PIA will answer the following questions:

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

Document version

Version	Date	Section	Page	Description	By
001	02/09/2022			First Draft	Al Warren
002	08/09/2022			Review	Kristina Elfstrom
003	15/09/2022			Review	s 9(2)(a)
004	27/10/2022	1.1,1.2	4,6	Update project details based on feedback	Kristina Elfstrom
005	17/12/2022	1.9	4, minor updates through out.	Added additional bullet point (1.9) and other minor Updates based on additional feedback	Kristina Elfstrom
006	16/02/2024	1.1,2.1	6,8	Addendum added for the model to be used for non-individual customers. Minor update to list of data points.	Kristina Elfstrom

Document contributors

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

Name	Role	Business Area
s 9(2)(a)	Data Strategy and Governance Lead	Enterprise Information & Knowledge
Dawn Swan	Privacy Officer	Enterprise Design & Integrity
Miriana Stanley	Information Specialist	Information Sharing
Graham Poppelwell	Domain Lead	Information Sharing
s 9(2)(a)	Sponsor	Payments & Assessments Network

Document Sign off

Business unit	Name	Sign off date
Senior Responsible Owner	Tony Morris	16 February 2023
ED&I, Privacy Officer	Dawn Swan	09 February 2023 19 February 2024 for minor amendments in version 006.

Contents

Document contributors	2
1 Introduction	5
1.1 Project summary	5
1.2 Scope of the PIA	6
1.3 Glossary of terms	7
2 Personal information	8
2.1 Information to be used in the project	8
2.2 Information flows.....	9
3 Privacy assessment	11
3.1 Privacy principles and response	11
4 Ethics assessment.....	19
5 Algorithm assessment	20
6 Risk assessment	23
6.1 Table of risks and mitigations	23
6.2 Summary of risks.....	25
7 Recommendations.....	25
Appendix 1 Enterprise risk rating tools	27
Consequence matrix	28
Risk rating matrix.....	32

1 Introduction

1.1 Project summary

1. Inland Revenue intends to begin using START's native Analytics capabilities. We will commence with building 'models' in START's Analytics Manager which will initially be used to aid IR with managing a customer's compliance needs. The information provided by the Analytics Manager is intended to be used alongside other tools in START such as Discovery Manager and Decision Support.
2. The Analytical models will use data readily available within START to assess characteristics about a customer. This project will focus on building Analytical models to assist with optimising how we segment our compliance work. The output provided by the model's analysis will aim to categorise customers so that they receive interventions best suited to their circumstances. The model will use only the information that is fed into it by a person, it will not use anything without being directed to.
3. It is intended this will improve on our current method of treating customers that fall into default with IR, which is typically a 'one size fits all' approach regardless of a customer's position. Additional to this, the size of our customer base means we cannot offer a one-on-one approach to every customer, it is intended that the model's analysis will allow us to better focus our limited resources on customers that will benefit the most from it.
4. Discovery Manager and Decision support (we already make use of these tools in different contexts) are additional tools and will be used to combine the data from analytical models as well as other relevant customer information to support the appropriate compliance pathways for customers. Whether that be direct one-on-one interaction or an automated action such as a follow up notification.
5. It is intended that this will be an iterative project and the outcomes of the proposed pilot may be used to further refine model results if it is deemed necessary to do so. This is to ensure we have the flexibility to make improvements to the model based on feedback provided throughout the pilot.
6. Using the information available in START already will give IR the opportunity to create tailored approaches to compliance that will better serve our customers and allow us to efficiently use our resources.
7. A developer will program START Analytics Manager and/or the Decision Support tool to collect information on certain customer characteristics e.g., income, location, payment history, etc. This information is the same as what a user would collect in a one-to-one interaction with a customer.
8. This project looks to collect the information at a greater scale to both reduce the effort needed by staff when deciding how to resolve customers compliance needs, and to aid in segmenting populations. This will help us ensure our customers receive the intervention type most suitable to their needs.
9. The model output and the data used within it will go through preliminary testing in STARTs Testing environment before progressing to a pilot stage. This preliminary testing has now been completed and has passed.

*Addendum (added February 2024)

The initial version of this document does not specify, but it should be noted that the original version if the analytical model applied solely to Individual customers, the model is now being expanded out to cover non-individual customers. The data used for these entities is covered by the list of data points under section 2.1 however this extension to the model uses no individual customer information.

1.2 Scope of the PIA

10. *Analytics Manager, Discovery Manager and Decision Support Manager* are tools available for use within START. These tools are maintained by the CSI Team.
11. The information that is collected will be stored and accessible within START, and can also be retrieved from its corresponding database; all those with access to the secure database will be able to view and retrieve the stored information.
12. This Analytics model will create a 'Model Score' which is an aggregated view of a customer's characteristics and circumstances that have been fed into it. This score will be viewable in START and is also able to be viewed in the related database (similar to point 9 mentioned above).
13. It is important to note that no 'new' information is proposed to be collected at this stage for the purpose of this project; the project will be utilising existing data within our START system.
14. Information and results produced by models created within START will be used for developing insights to further improve the process of modelling as it develops.
15. While the information provided by the model may be deemed useful there will always be a human element to this review/improvement process which will allow us to ensure only appropriate information/categorisation is utilised.
16. Information collected by the Analytics Manager as well as any output created is subject to the same storage/disposal rules as all data held within START.
17. Results produced by these analytical models will be used to more accurately segment our customer population to help determine what intervention is most appropriate for certain customer groups.
18. Initially, this model will be used primarily for compliance activities, however, this could expand into other areas of customer segmentation as areas for improvement are identified.
19. The initial Pilot will focus on customers who have a payment agreement with IR and have defaulted on their agreement. The involvement of the model in this initial pilot will be to quantify a customer's compliance position and will consider their financial situation (how much income they receive, what kinds of income, if a customer has been affected by COVID and to identify patterns in this income) as well as aspects from their compliance history, such as payment history and timeliness. Knowing this will assist in segmenting the customers who:

- May benefit from a 'nudge' to catch up on payments and/or are likely to self-correct,
- are likely to have debt that is uncollectable, or
- would benefit from a more one-on-one interaction to resolve their debt situation.

20. One of the drivers for the work that Analytics Manager will produce is the Payments & Assessments Network in CCS.

21. Marketing may also make use of the information produced by any analytical model to improve targeting of 'one-to-many' interventions.

22. Out of scope includes any changes to START beyond the building of the facility in Analytics Manager, Discovery Manager and Decision Support Manager.

23. Also out of scope: this is about understanding customer data in such a way that it signals specific interventions; in no way will we be expecting to use the analysis to *automatically* change/adjust/amend customer data.

1.3 Glossary of terms

Term	Meaning
Analytics Manager	START toolset that gathers and combines customer data-points at scale to produce insight into populations as an aid to tailored and targeted intervention.
Decision Support Manager	START toolset that gathers and combines customer data-points for presenting to IR staff when interacting or intervening with a customer, enabling greater efficiency and the completion of a 'whole-of-job' approach.
QA	Quality Assessment – a checking approach designed to determine that the quality (accuracy level) of data is above a specified threshold.

2 Personal information

2.1 Information to be used in the project

Identify and describe the personal information to be used in the project, i.e. full name, phone number, postal/email address, IRD number, employer's name, financial details, family relationships. Is this information already held by IR or is new personal information being collected? What is the nature of the information and the source and who will have access to the information?

'Personal information' is any information that is capable of identifying a living human being. It doesn't have to be sensitive or negative information just so long as someone could be identified by it.

However, the level of sensitivity, and impact on individuals, will influence the privacy risks in managing the information.

24. We plan to explore the use of any or all of the following (in no particular order):

- customer age/age of business
- employment status and history
- financial information including income level and income pattern (rises and falls in income)
- bank details – interest/bank number history
- investment income and Kiwisaver contributions
- income tax return information
- account registration, account status and associated return information (GST, EMP etc)
- filing and payment history
- previous interventions a customer may have had with IR
- insolvency history
- association with other tax types
- industry
- geography
- history of COVID products received
- fraud indicator
- audit history
- integrity history (return filing accuracy)
- segment
- tax agent status
- debt level
- links to associated entities
- instalment arrangement history and status
- deduction history (S157 deduction notices)
- general collection history, collection age, prior collection records and actions taken
- prior debt write-offs/remissions
- customers family information that we hold (working for families/child support data)
- contact history (with IR) across any channel
- MyIR use

25. The information listed above is available in START. As further information becomes available, it may be added into analytical models if it is deemed appropriate and will improve the accuracy of its scoring.

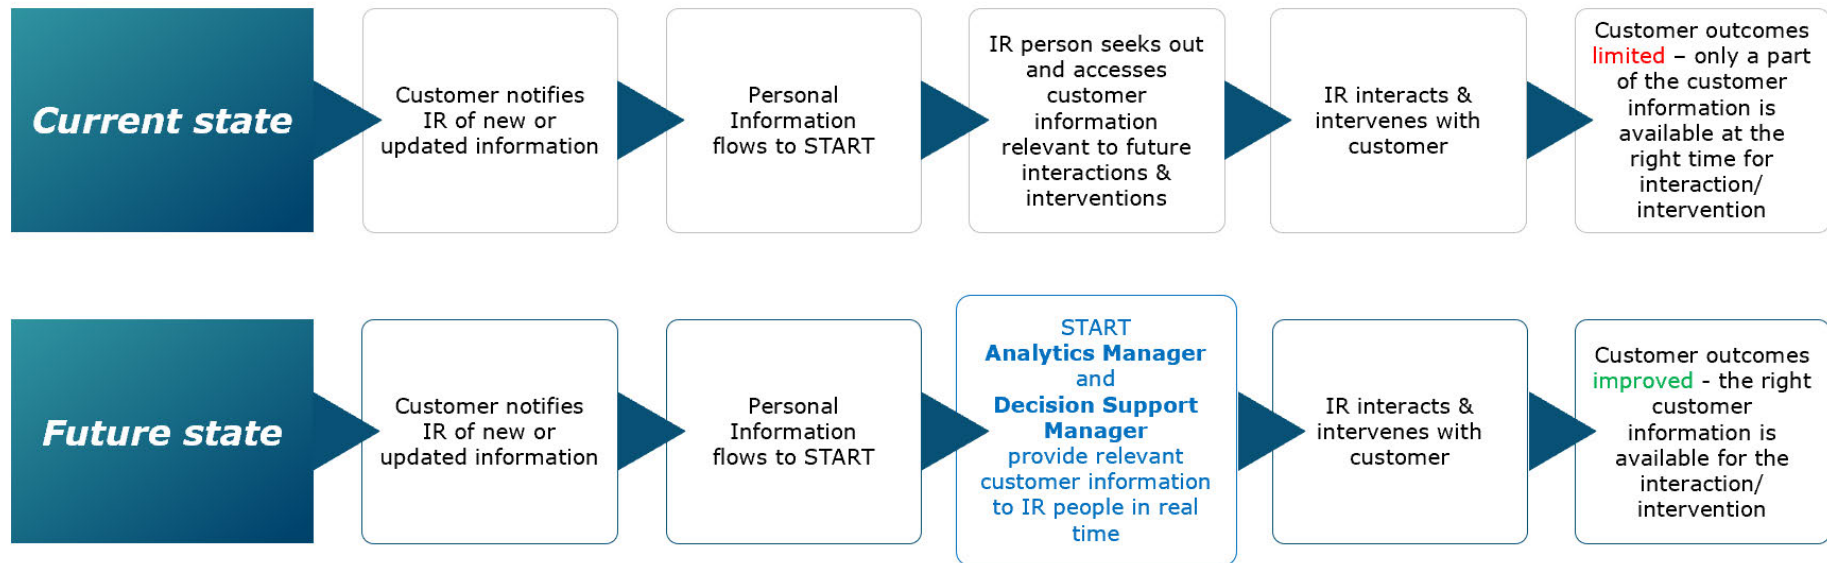
- 26. Those with access to it will remain the same as our BAU approach, i.e., there will be no access level changes based on this pilot.
- 27. The data collated in the analytics model will be accessible in the START system by the CSI team, who manage these tools. However, there are associated data tables that are accessible in START's database which is securely accessible by all data analysts.
- 28. These data tables are also frequently shared to "Snowflake" or the DIP to allow further data teams' access. This is the standard secure process for data access at IR currently and this project has made no additional request to provide or withhold data.
- 29. The insight generated will be worked on across IR including CCS and Marketing.

2.2 Information flows

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

Flow of personal Information

Analytics Manager and Decision Support Manager improve the flow of the right personal information for IR interaction/intervention with customers.



3 Privacy assessment

3.1 Privacy principles and response

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

Description of the privacy principles
<p>Information Privacy Principle (IPP) 1 Purpose of collection</p> <p><i>Only collect personal information if you really need it:</i></p> <ul style="list-style-type: none">• collection must be for a lawful purpose connected with a function or activity of IR; and• collection must be necessary for that purpose
<p>Summary of how IR will comply</p> <p><i>Explain what personal information is being collected and why IR needs it. How will it enable IR to operate? Are you only collecting what you actually need?</i></p> <p>30. All the information collected in START is for a lawful purpose related to tax or social policy administration. The information which this project will utilise is detailed in section 2.1</p> <p>31. It is intended that the model be robustly tested by both technical and business specialists, if at any point a piece of collected information is not deemed useful to the models' outcome, this data point will be removed from the Analytics Manager.</p> <p>32. This initiative doesn't propose to collect additional information from the customer, rather it places the information available to us already within a global customer context, enabling us to see a fuller picture of a customer in one glance. This enables IR to compare customer attributes with similar (or dissimilar) customers.</p> <p>33. This combining of customer personal information is done in pursuit of a lawful purpose: seeking positive customer outcomes, and ensuring IR operates as efficiently as possible.</p> <p>If you have identified any risk to complying with principle 1, complete the Table of risks and mitigations.</p> <p>34. No risks identified above.</p>
<p>Information Privacy Principle (IPP) 2 Source of personal information</p> <p><i>Get it directly from the person concerned wherever possible.</i></p> <p><i>You can collect from another source if:</i></p> <ul style="list-style-type: none">• the information is publicly available• the individual concerned has authorised collection from someone else• non-compliance would not prejudice the interests of the individual concerned• it's necessary to collect the information to uphold or enforce the law or protect public revenue• compliance would prejudice the purposes of the collection

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

- *compliance is not reasonably practicable*

Summary of how IR will comply

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

35. The information used by Analytics manager is information contained in START.

36. This is already deemed to be information collected directly from the person, or lawfully collected elsewhere, such as from the customer's employer, their bank or information provided by other government agencies. IR is complying with this privacy principle.

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

37. No risks identified above.

Information Privacy Principle (IPP) 3 Tell people what you're going to do with their information

At the point of collection, you must tell people:

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

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

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

Summary of how IR will comply

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

38. Inland Revenue collects personal information for lawful purposes connected with our functions and activities of administering the revenue laws. People are advised what their information will be used for when they fill out a form or register for services. All forms contain a privacy statement that links to the publicly available Privacy Policy.

39. Customers are advised that any information provided will be used by IR for tax and social policy administration.

40. The purpose of this work is to make information that we already have, more visible and useful. However, when we collect the information, we don't expressly indicate that the individual details may form part of a wider picture, but this would not be surprising given IR's functions.

41. Customers have not been informed how their information is used in Analytics Manager as this is a new tool and does not prejudice the interests of the individual. The pilot is simply collating available information in a new way.

42. To inform the public about the planned use of the information they provide, we will ensure our algorithms are transparent and open to scrutiny.

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

Information Privacy Principle (IPP) 4 Manner of collection of personal information

Do not collect information in a way that is unfair, unlawful or unreasonably intrusive

Summary of how IR will comply

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

43. The purpose of this work is to make use of information that we already have so IR can make informed decision about what our customers may need. The way the information is collected will not change and is lawful.

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

44. No risks identified above.

Information Privacy Principle (IPP) 5 Storage and security of personal information

IR must ensure that:

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

Summary of how IR will comply

What steps are taken to keep the information safe? Safeguards can be physical or technical. Does the system design enhance privacy and security? Are the security measures commensurate with the sensitivity of the information?

Consider where the information will be stored and controls defining who may access it, is there an audit trail? Will there be external access to a system, has it received approval from IT&C? Will staff receive training? Are there mechanisms in place to identify data/security breaches?

Are there documented security procedures for the collection, transmission, storage and disposal of the information?

If third parties are involved they must sign a Certificate of Confidentiality and abide by IR's Code of Conduct.

45. The pilot does not affect current security settings and all IR security safeguards around customer data will remain the same. This includes role-based access in START and event logging.

46. Combinations of data will be treated in the same manner as the raw individual data with appropriate limitations on access and use.

47. Where reporting is required on combinations of data points, we will continue our current process where individual customer information will not be identified.

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

48. No risks identified above.

Information Privacy Principle (IPP) 6 Access to personal information

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

- *obtain confirmation of whether the information is held; and*
- *have access to their information (subject to withholding grounds contained in the Privacy Act)*

Note: IR has business processes to deal with access requests. If the project will not affect that, say so.

Summary of how IR will comply

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

49. IR has processes to enable individuals to request access to their personal information and there is information on the website.

50. This initiative does not change or impact on that process or an individual's ability to request access to their information. The personal information used will be able to be retrieved and provided to customers upon request.

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

51. No risks identified above.

Information Privacy Principle (IPP) 7 Correction of personal information

Everyone is entitled to:

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

Summary of how IR will comply

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

52. IR has processes to enable customers to request their personal information be corrected.

53. This initiative does not change or impact on that process or a customer's ability to request correction of their information.

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

54. No risks identified above.

Information Privacy Principle (IPP) 8 Accuracy of personal information to be checked before use

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

Summary of how IR will comply

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

55. We want to responsibly check customer information is accurate before use. This initiative does not change or impact on that process or an individuals' ability to request correction of their information.

56. As part of the algorithm transparency process, we will ensure people make decisions about use of algorithmically defined groupings of customers rather than automatically running interventions.

57. Part of that decision-making process will incorporate a data accuracy quality assessment (QA) of the combination of customer data, ensuring both the raw data and the insight based on the algorithm stands up to scrutiny.

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

Information Privacy Principle (IPP) 9 Don't keep personal information for longer than necessary

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

Ask the Information & Knowledge Management team if there is a requirement under the Public Records Act to keep the information for a specific period

Summary of how IR will comply

How long will IR hold the information? Is there a requirement under the Public Records Act to keep this information for a specific time? If not, what would be a reasonable length of time to keep it and how will you ensure it is disposed of?

If information is being shared with a third party, how long will they hold the information for?

58. Inland Revenue's retention of information is governed by a Retention and Disposal Schedule under the Public Records Act 2005. The obligations under the Public Records Act override this privacy principle.
59. Data used in the analytics model will be held in line with all data held within START. There is no intention to share information produced by our analytical models to third parties.

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

Information Privacy Principle (IPP) 10 Limits on use of personal information

Only use personal information for the purpose you got it for

Exceptions include:

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

Summary of how IR will comply

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

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

60. We are using the data for the purpose that it was obtained.
61. We intend to use the combination of data points in pursuit of positive customer outcomes and ensuring IR operates at optimal efficiency; no different to the original purpose of collecting data.
62. Any additional uses of the data beyond their original intent will still be aligned to protecting the integrity of the tax system, protecting public revenue, and pursuing positive customer outcomes.
63. To be certain of this, use of the data from the algorithm (i.e., intervention design and delivery) will not be *automatically* delivered by the algorithm, but rather will require human oversight.

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

Information Privacy Principle (IPP) 11 Limits on disclosure of personal information

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

- *the individual authorised you to disclose*
- *disclosure is one of the purposes for collecting the information (and people were told at the point of collection – this links to IPP3)*
- *it's a directly related purpose to why the information was obtained*

- *it's necessary to disclose it in order to detect or investigate an offence or assist court or tribunal proceedings*
- *it's necessary to disclose it in order to protect public revenue*

Summary of how IR will comply

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

64. This initiative will not change our current policy or approach to the disclosure of personal information. Information generated will be disclosed to the individual concerned.
65. Human checks will be a part of any initiative to ensure appropriate quality control of combined data.

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

66. No risks identified above.

Information Privacy Principle (IPP) 12 Cross border disclosures

Only disclose personal information to foreign persons or entities if it's reasonably believed it:

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

Summary of how IR will comply

Is any personal information being disclosed to a person or entity overseas? If so, how is IR able to demonstrate that it has carried out the necessary due diligence checks required under this privacy principle.

67. This initiative will not change our current policy or approach to the disclosure of personal information across borders. The data will not be disclosed to a person or entity outside New Zealand.

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

68. No risks identified above.

Information Privacy Principle (IPP) 13 Unique identifiers

Unique identifiers must not be assigned to individuals unless this is necessary for the organisations concerned to carry out its functions efficiently

Summary of how IR will comply

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

69. No new unique identifiers will be used.

If you have identified any risk to complying with principle 13, complete the [Table of risks and mitigations](#)

70. No risks identified above.

4 Ethics assessment

i. Areas that may raise ethical issues

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

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

Does the project use ethical data practices?	Yes (tick)	No (tick)	If yes, explain your response
Will the proposal discriminate against some people?	✓		<p>The initiative aims to better segment customer groups, this will allow IR to intervene more efficiently and effectively.</p> <p>While it is probable that an analytics model may discriminate against certain customer groups, it is not intended to be for a negative purpose but more so to ensure that customers receive a more appropriate interaction from Inland Revenue.</p>
Is the proposal likely to result in some members of a group being treated differently to one another?	✓		<p>Part of the problem this initiative hopes to solve is that our current interactions and interventions are too linear and homogenous, leading to unequal and unfair outcomes. We are deliberately hoping to treat different customers differently, but in a transparent way, targeting improved outcomes for all.</p>
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?	✓		<p>Our aim is to positively impact these groups through being able to identify them faster and to be able to tailor and target our interventions.</p>
Can you foresee any harm to individuals in using the data in the way intended?	✓		<p>While all reasonable steps will be taken to robustly test Model outcomes, there is always a risk of a customer being assessed incorrectly.</p> <p>The intended approach is to continually optimise/review models based on user feedback to minimise this risk.</p>

Have you considered how the proposal contributes to the active protection of Māori interests?	✓	This initiative will support our ability to better identify Māori needs and intervene more effectively. We have also considered concepts of trust and data autonomy, elements critical to Te Ao Māori. This initial pilot is not designed specifically with Māori needs in mind.
---	---	--

If you have identified any risk with ethical data practice associated with this proposal, complete the [Table of risks and mitigations](#)

5 Algorithm assessment

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

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

Algorithm principles	Explain your response
<p>Deliver clear public benefit:</p> <p>Who will benefit from the development of this system?</p> <p>Have any assumptions been made in design or planning?</p> <p>Is it likely people will suffer an unintended adverse impact as a result?</p> <p>Have associated policies and decisions been evaluated for fairness and potential bias and have a solid grounding in law?</p>	<p>The beneficiaries of the use of this algorithm are the New Zealand Govt (through more efficient operation of IR) and New Zealanders who will experience more efficacious interactions/interventions with IR.</p> <p>The assumptions are that:</p> <ul style="list-style-type: none"> IR's (raw) is up to date and maintained by the customer.; decisions about how/when to intervene will still be made by trained people; combinations of data points can more usefully and fairly lead to positive outcomes than the raw points on their own; New Zealanders (including Govt stakeholders) are generally comfortable with the use of algorithms, if the right level of risk mitigation is applied. <p>It is possible, but the likelihood diminishes significantly with the risk mitigation approaches described in this paper.</p> <p>The planned use does not fall outside of any current policy or decision-making process, all of which we expect will have been evaluated for fairness, potential bias and be grounded in law. Any move to operate</p>

	under new policy or decision-making processes would require just such an evaluative exercise.
<p>Ensure data is fit for purpose:</p> <p>How accurate, precise, consistent, and complete is the data quality? (This may already have been answered at IPP8)</p> <p>Are we re-using data that was originally collected for another purpose?</p> <p>How are we identifying and managing bias or discrimination?</p>	<p>Data quality is of paramount concern. We recognise that this area can lead to potentially negative outcomes, so our risk mitigation approach requires thorough and transparent QA.</p> <p>All use of data will be in the pursuit of protection of the integrity of the tax system and public revenue. Where there is data collected for explicit purposes, it would either not form a part of this pilot, or we would ensure there was visibility and a means of removing the data where customers felt that was appropriate.</p> <p>We are attempting to eliminate the bias and discrimination that currently occurs when we apply a one-size-fits all approach to all customers. To prevent new forms of bias/discrimination from occurring, we will have both QA checks of data and human decision-making processes between algorithm activity and intervention with customers.</p>
<p>Transparency:</p> <p>Are data use and analytical processes well documented, and the decisions they inform described in clear, simple, easy-to-understand language?</p> <p>Will decisions be explainable and auditable?</p> <p>Are there real-time feedback loops and a 'kill switch' if post-deployment bias is found?</p> <p>Have our customers been informed about how the system will work, and how they may seek more information or exercise review rights?</p>	<p>Our pilot is designed to learn how to do this well. We aim to be able to answer simply 'yes' to this question at the end of the pilot.</p> <p>Again, we haven't built this yet, but knowing we need to do this, it will be part of the experience and outputs of the pilot.</p> <p>The intention of the initial pilot is to incorporate user feedback to improve the Analytical model. There is a standard process to follow for any issues or concerns found within START users can utilise once the pilot stage has completed. At any point the scoring information produced can be removed from production use if it is deemed appropriate to do so.</p> <p>No, but a part of the pilot will be to test how this operates. Given we'll be beginning with an algorithm that mimics what individual IR officers do right now (albeit at a much smaller scale), there may be minimal perceptible difference for customers experiencing the intervention. But we are interested in transparency, so we will be expecting to make available to the public the ways in which the algorithm works.</p>
<p>Understand the limitations:</p> <p>Do any limitations exist in terms of the collection or use of the data?</p>	<p>Our data is primarily maintained by our customers, this can lead to information being incorrect if this has not been maintained.</p> <p>The system itself is very capable of drawing together the data elements into a whole.</p>

Who is represented in the dataset, and who is not, which might lead to historic or representation bias?	The dataset includes all IR customers; however, we have caveats around accuracy of the data across all points, and missing data points which can lead to bias. For example, we have only recently added a 'Māori business indicator' to the system enabling us to consider Māori interests when intervening.
<p>Retain human oversight:</p> <p>Will decisions informed by the algorithms involve human judgement and evaluation?</p> <p>Can we test to see how well the algorithm is working compared to human decision-making?</p> <p>Will the automated decision-making process be regularly reviewed to make sure it's still fit for purpose?</p> <p>Nominate a point of contact for public inquiries about algorithms, and provide a channel for challenging or appealing of decisions informed by algorithms</p>	<p>Yes.</p> <p>Yes, we will be running a process to deliver statistically significant trial comparisons (e.g., randomised controlled trial (RCT) methodology).</p> <p>As this is a pilot, it will be thoroughly reviewed. We anticipate that any future ongoing algorithm use would require regular review periods.</p> <p>Our Data Strategy area, where we consider ongoing data use and governance thereof. A feedback channel into this area would be very useful. Additionally, Compliance Strategy and Innovation (CSI) will be responsible for building and delivering ongoing algorithms; some form of algorithm governance structure aligned to the Data Strategy group could exist.</p>

If you have identified any risk with using the algorithm, complete the [Table of risks and mitigations](#)

6 Risk assessment

6.1 Table of risks and mitigations

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

Ref No	Description of risk	Consequences for IR or individuals	Existing controls that help manage risks identified	Residual current risk	Suggested action to reduce or mitigate risk	Residual risk remaining
RSK-01	<i>Customers are unaware of how we'll use their collected information.</i>	Breach of Information Privacy Principle 3: "Tell people what you're going to do with their information."	We're not gathering new information to inform the algorithms and the customer data is already used by IR officers in the way the algorithm will do at scale.	High Because of the public perception of the use of algorithms.	Adherence to the Algorithm charter ensures we'll be transparent in how customer data is used in algorithms.	Low
RSK-02	<i>Data inaccuracy leads to interventions that deliver negative customer outcomes.</i>	Breach of Information Privacy Principle 8: "Accuracy of personal information to be checked before use"	Our current systems and checks to understand and improve the accuracy of customer data will help to avoid many issues.	High Because we have a lot of data, and often there is very little requirement for either customers or IR to improve the accuracy.	We'll build a QA process that examines both the quality of the raw data entering the algorithm and the algorithm outputs. Agreed and visible thresholds of quality standards will ensure interventions are not able to be delivered with lower quality data.	Medium

Ref No	Description of risk	Consequences for IR or individuals	Existing controls that help manage risks identified	Residual current risk	Suggested action to reduce or mitigate risk	Residual risk remaining
RSK-03	<i>Algorithm use means IR holds onto some insight data longer than what would be required for the original raw data.</i>	Breach of Information Privacy Principle 9: "Don't keep personal information for longer than necessary."	IR already has good rules around retention of data and manages this well.	Low	A process to be built that ensures algorithmically generated insight is not held longer than the raw data that informed it.	Low
RSK-04	<i>Algorithm use leads to interventions that demonstrate discrimination or bias that negatively impacts customers.</i>	Breach of Information Privacy Principle 10: "Limits on use of personal information."	Our current approach is 'one-size-fits-all' and has the potential for bias or discrimination. This is mitigated to a degree by having people at the decision making heart of interventions.	High	We will build a process for use of algorithms that incorporates a requirement for person-centred decision-making between the algorithm and the delivery of the intervention.	medium
RSK-05	<i>Public are unable to enquire about or feedback on the use of algorithms for their data</i>	A breach of the requirement in the Algorithm charter to "Nominate a point of contact for public inquiries about algorithms, and provide a channel for challenging or appealing of decisions informed by algorithms"	Our customers are accustomed to talking to us and questioning us, however this process is new to them.	Medium	We recommend a feedback loop be built back into IR (perhaps at a data governance layer and/or through to the CSI team who build algorithms) ensuring that customers can question and appeal the use of their data.	Low

6.2 Summary of risks

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

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

This initiative has these identified risks:

- Data inaccuracy leading to poor customer outcomes.
- Combinations of customer data points may lead to the potential for unintended and negative discrimination and bias.
- Customers don't know how their data is used, or that we're not keeping it for longer than necessary and they're unaware of a feedback loop to be able to talk with us.

Risks will be mitigated as described above with quality assurance, human-centric decision making and algorithm visibility and feedback processes.

The privacy risk for this initiative is rated as Low. The information being used has been collected by IR for lawful purposes and is only a new use of the information in terms of technology being used.

7 Recommendations

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

Ref	Recommendation	Agreed Y/N
R-01	To run a pilot that will: <ul style="list-style-type: none">• Use START Analytics Manager to collect and assess customer characteristics with the intention of scoring a customer's financial health• To use the information produced to segment and intervene with customers who have failed to adhere to their payment agreements.	
R-02	To build a quality assessment (QA) approach to both the raw data used in the pilot, as well as the algorithmic insight generated by Analytics Manager.	
R-03	To build a person-centric decision-making process on how and when IR will intervene with customers, based on algorithmic output from Analytics Manager	
R-04	To build a feedback approach to our algorithmic use so that customers can see what we're using data for and can question us.	

Ref	Recommendation	Agreed Y/N

Appendix 1 Enterprise risk rating tools

The following tools (Likelihood Matrix, Consequence Matrix and Risk Rating Matrix) are designed to assess the rating of a risk. Using these requires an element of judgement, as a risk may meet multiple levels of consequence or likelihood criteria.

Likelihood matrix

	Description		Influence	Control effectiveness	Probability
Almost Certain	The risk can be expected to occur	OR a very poor state of understanding exists about the risk event	A very poor ability to modify or influence the likelihood of the risk occurring	Almost all of the controls are ineffective, partially effective or not in place	Greater than 95% chance of occurring
Likely	The risk will quite commonly occur	OR a poor state of understanding exists about the risk event	A poor ability to modify or influence the likelihood of the risk occurring	The majority of controls are ineffective, partially effective or not in place	60% - 90% chance of occurring
Possible	The risk may occur occasionally	OR a moderate state of understanding exists about the risk event	A moderate ability to modify or influence the likelihood of the risk occurring	Some controls are effective, and some are ineffective, partially effective or not in place	30% - 60% chance of occurring
Unlikely	The risk could infrequently occur	OR a good state of understanding exists about the risk event	A good ability to modify or influence the likelihood of the risk occurring	The majority of controls are effective. Some are ineffective, partially effective or not in place	5% - 30% chance of occurring
Rare	The risk may occur in exceptional circumstances	OR a very good state of understanding exists about the risk event	A very good ability to modify or influence the likelihood of the risk occurring	Almost all controls are effective	Less than 5% chance of occurring

Consequence matrix

Consequence	Definition	Customer outcomes	Health, Safety & Wellbeing	Reputation	Legislative Compliance	Performance	Initiative delivery consequences		
							Schedule	Delivery Cost	Scope/ Benefit
Severe	The risk event is widespread or significant, impacting IR's ability to achieve it's strategic objectives	Widespread or significant impacts to customers, impacting their ability to receive entitlements or meet obligations Widespread customer complaints >50% increase in compliance costs	Fatality or fatalities to an employee or a visitor attributable to Inland Revenue's actions or inactions	<p>A systemic or perceived failure (either internally or via a third party) that could undermine the reputation of and trust and confidence in IR's revenue collection and social support payment systems.</p> <p>Sustained and ongoing negative multi-channel media coverage</p>	<p>Royal Commission of Inquiry in relation to breach of Compliance</p> <p>Breach of sensitive or highly sensitive information, with serious potential or actual harm to customers. It is likely that more than one type of harm has occurred, and that harm is likely to be ongoing.</p> <p>Information (regardless of its classification) is shared inappropriately or modified or deleted either maliciously or non maliciously, by IR or third parties that presents a widespread or significant impact</p>	<p>Significant impact on ability to meet outcome-level performance measures (e.g. payments made on time, refunds to customers on time)</p> <p>At risk of breaching appropriations</p> <p>More than 5% under or overspend of internal budget</p> <p>Significant impact to revenue base over time</p>	Continuing delays that significantly impact on infrastructure availability, other projects or multiple business areas	<p>>25% variance against year-end cost</p> <p>>50% variance or >\$1m in total project implementation (excl contingency)</p>	<p>>30% or >\$1m benefits not achieved</p> <p>>50% inability to meet critical success factors</p> <p>Significant interruption of >1 day that affects our ability to collect revenue or disburse entitlements</p>

Major	The risk event is major or considerable, impacting IR's ability to achieve some of it's strategic objectives	<p>Major or considerable impacts to customers, impacting their ability to receive entitlements or meet obligations for a short period of time</p> <p>High level of complaints in a single area</p> <p>>25% increase in compliance costs</p>	<p>Injury or illness with permanent or long term severe disabling effects and irreversible health damage. Immediate extensive emergency, medical assistance and hospitalisation</p> <p>Severe psychological trauma requiring long term counselling support. Recovery time of 3 months or longer</p>	<p>A systemic or perceived failure (either internally or via a third party) that may cause long term loss of trust and confidence in IR that could impact service delivery.</p> <p>Ongoing negative multi-channel media coverage</p>	<p>Ministerial Inquiry commissions in relation to breach of compliance</p> <p>Ministerial questions in Parliament in relation to breach of compliance</p> <p>Breach of sensitive or highly sensitive information, with serious potential or actual harm to customers.</p> <p>Information (regardless of its classification) is shared inappropriately or modified or deleted either maliciously or non maliciously, by IR or third parties that presents a major or considerable impact.</p>	<p>Major impact on ability to meet outcome-level performance measures (e.g. payments made on time, refunds to customers on time)</p> <p>3-5% under or overspend of internal budget</p> <p>Major impact to revenue base over time</p>	<p>>35% delay in critical path</p> <p>Major delay that seriously impacts on infrastructure availability, other projects or some business areas</p>	<p>16-25% variance in year-end cost</p> <p>36-50% variance in total project implementation (excl contingency)</p> <p>Partial requested funding provided by will need descopeing</p>	<p>16-30% of benefits not achieved</p> <p>36-50% inability to meet critical success factors</p> <p>Outage between 4-8 hours that affects our ability to collect revenue or disburse entitlements</p>
-------	--	--	---	--	--	--	---	---	--

Moderate	The risk event has a moderate impact, and has some impact on IR's ability to achieve some strategic objectives	Moderate impact for customers, with negligible impact on their ability to receive entitlements or meet obligations Sustained increase in complaints >10% increase in compliance costs	Illness or injury with temporary significant or severe disabling effects. Advanced first aid or attention by medical services or hospital Psychological trauma requiring long term counselling support. Recovery time of 1-3 months	Customers may stop using, or be reluctant to use, a service or delivery channel. The incident may get negative media attention.	Breach of Act or regulation with material effect Either the information is not sensitive/highly sensitive and the potential or actual harm to the customers is more than minor, or the information is sensitive/highly sensitive and the potential or actual harm to the customers is minor. Information (regardless of its classification) is shared inappropriately or modified or deleted either maliciously or non maliciously, by IR or third parties that presents a moderate impact	Impact on an outcome-level performance measure (e.g. (e.g. payments made on time, refunds to customers on time) 1-3% impact on internal budget, can be managed in-year Moderate impact to revenue base over time	21-35% delay in critical path Moderate, or numerous minor delays that impact on other projects or other business areas	10-15% variance against year-end cost 21-35% variance in total project implementation (excl contingency) Some funding cuts but still able to deliver within scope	6-15% of benefits not achieved 21-35% inability to meet critical success factors Moderate inability to collect revenue or disburse social policy entitlements
Minor	The risk event has a minor impact, and has some impact on IR's ability to achieve a strategic objective	Minor impact for customers, with no impact to being able to receive entitlements or meet obligations Short term increase in complaints	Injury or illness with nonsevere or temporary disabling effects requiring first aid and/or referral to	The risk event requires additional communication with stakeholders to address temporary or short-term impact to IR's reputation	Information request from the ombudsman in relation to breach of compliance Breach of Act or regulation with legal rebuke	Impact on multiple performance measures (impacts, services or organisational health) 0.5-1% impact on internal budget, can be managed in-year	5-20% delay in critical path Minor delays that have some impact on a project or business area	2-9% variance against year-end cost 5-20% variance in total project implementation	<5% of benefits not achieved 5-20% inability to meet critical success factors

		>5% increase in compliance costs	a medical professional Psychological trauma requiring some counselling support. Recovery time of up to one month		Minor potential or actual harm to customers. Information (regardless of its classification) is shared inappropriately or modified or deleted either maliciously or non maliciously, by IR or third parties that presents a minor impact.	Minor impact to revenue base over time		(excl contingency)	Minor inability to collect revenue or disburse social policy entitlements
Minimal	The risk event has minimal impact, and has very little impact on IR's ability to achieve a strategic objective	Minimal to no impact on customers Small increase in complaints for a short time <5% increase in compliance costs	Nil or minor injury or illness that does not require medical attention and no recovery time is needed.	The risk event has no real effect on stakeholder perceptions of IR's revenue collection and social support payment systems	Official information request in relation to breach of compliance Little or no potential or actual harm to customers. Information (regardless of its classification) is shared inappropriately or modified or deleted either maliciously or non maliciously, by IR or third parties that presents a minimal impact.	Impact on performance measures (impacts, services or organisational health) Minimal to no impact on internal budgets Minimal to no impact on revenue base over time	<5% delays in critical path	<2% variance against year-end cost <5% variance in total project implementation (excl contingency)	>95% of the critical success factors can be achieved

Risk rating matrix

The matrix below takes the ratings from likelihood and consequence and gives them an overall risk rating.

		CONSEQUENCE							
LIKELIHOOD		Minimal	Minor	Moderate	Major	Severe		Risk rating	Definition
	Almost Certain	Almost Certain / Minimal	Almost Certain / Minor	Almost Certain / Moderate	Almost Certain / Major	Almost Certain / Severe		Extreme	The risk poses a significant threat to being able to achieve IR's strategic objectives
	Likely	Likely / Minimal	Likely / Minor	Likely / Moderate	Likely / Major	Likely / Severe		Very High	The risk poses a major threat to being able to achieve IR's strategic objectives
	Possible	Possible / Minimal	Possible / Minor	Possible / Moderate	Possible / Major	Possible / Severe		High	The risk poses a high threat to being able to achieve one or more of IR's strategic objectives
	Unlikely	Unlikely / Minimal	Unlikely / Minor	Unlikely / Moderate	Unlikely / Major	Unlikely / Severe		Medium	The risk poses a medium level of threat to being able to achieve one or more of IR's strategic objectives
	Rare	Rare / Minimal	Rare / Minor	Rare / Moderate	Rare / Major	Rare / Severe		Low	The risk poses a low level of threat to being able to achieve one or more of IR's strategic objectives

Item 4.4



Enable Voice isolation in Microsoft Teams

Privacy Assessment (condensed)

Prepared by: Phyllida Crawford

Date: 10 February 2025

Supply ID:

About this Document

The purpose of this document is to demonstrate that privacy has been considered in a project or process that involves personal information. The Analysis pulls together relevant information to determine whether a full Privacy Impact Assessment (PIA) should be completed and records IR's 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

Microsoft has added a voice isolation feature to Microsoft Teams that can identify a person's voice and filter out background noise if they've pre-enrolled a voice profile. IR's IT team is planning to make the voice isolation feature available to all IR people as an optional feature.

The voice profile can also be used to identify speakers in meeting transcripts when they use an 'intelligent speaker'. This may be deployed to IR office meeting rooms in the future.

To enrol a voice profile, the user must open Teams settings, and open Recognition in the menu. The user will then need to read some sample text while their voice is recorded in Teams.

When a user enrolls, Microsoft Teams captures their voice to create a unique profile associated with their identity. The profile consists of a set of biometric features that can be used to

enhance the user's voice in Teams meetings and calls and recognise them in meeting rooms with their voice. The biometric features are the unique pitches and sounds from the recording which are then stored in a biometric profile or "voice signature" that the Microsoft Teams tool uses.

For more detail on what these features will look like for IR people, see:

[Voice isolation in Microsoft Teams calls and meetings - Microsoft Support](#)

[Use Microsoft Teams Intelligent Speakers to identify in-room participants in a meeting transcription - Microsoft Support](#)

1.2 Purpose of the change

The purpose of the change is to help address background noise issues that IR people are experiencing, and the potential privacy issues arising from this.

In our open-plan work areas, it can be challenging to focus and communicate effectively on Microsoft Teams. This is because our microphones are so sensitive they pick up background noise such as customer calls, office conversations between staff, or neighbours mowing the lawn. This noise is distracting and makes it hard to hear our colleagues and customers during meetings and calls. There is also potential for a privacy breach when a Teams caller overhears confidential information being discussed in the office.

1.3 Public benefit

The change will primarily benefit IR and the people we communicate with using Microsoft Teams. IR communications will be clearer, and confidentiality will be better protected. If the speaker recognition feature becomes available, it will improve the quality and clarity of meeting transcripts by attributing comments to specific individuals instead of the room.

1.4 Privacy Enhancement

This change is privacy-enhancing as it will reduce the risk of confidential information being overheard.

1.5 Personal information to be used

In the table below, describe:

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

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

Type of personal information	Source of information	Purpose of using the information
Voice recording of the IR person reading out a standard piece of text. The unique pitches and sounds from the recording are stored in a biometric profile or "voice signature" that the Microsoft Teams tool uses.	If the person has chosen to enrol their voice, they will record their voice profile using Microsoft Teams. The person can delete their voice profile anytime.	To isolate the voice and filter out background noise so the person's voice can be heard more clearly on a Teams call or meeting. Future: Attribute comments to the specific individual in meeting transcripts

Note: There is also a feature to record your face profile in Microsoft Teams, which we intend to keep turned off at this stage. The face profile can be used to automatically recognise you, display your name during meetings, and add your name to the participant list as you walk into a meeting room with supported Teams equipment. If we decide to turn on this feature at a later date, we will do a separate privacy assessment.

1.6 Governance

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

Name of person or group	Business Unit	Approved, Consulted, Informed etc
AI Working Group (AIWG) Aidan Roberts Brijesh John Graham Poppelwell	Cross-functional working group	We submitted an AI use case request and were advised that AIWG review not required. Not a use of AI that has much risk or complexity associated with it. We already use this type of tech day to day.

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?

N

Is it linked to a practice or activity that is listed on a risk register?

N

Collection	Y/N	If yes, explain your response
------------	-----	-------------------------------

Will IR be collecting personal information that it doesn't currently collect?

Y

IR does not currently collect voice profiles of IR people.

Is collecting this information necessary for IR to carry out its functions?

N

We intend to make the feature optional for IR people to use.

Where or who is the information being collected from?

The voice profile is collected from the individual only if they choose to register it.

Storage, security, and retention	Y/N	If yes, explain your response
----------------------------------	-----	-------------------------------

Note: responses in this section are based on this Microsoft documentation:

[Overview of voice and face enrollment - Microsoft Teams | Microsoft Learn](#)

Does the initiative change the way personal or sensitive information is stored, secured or managed?

N

At a high-level it will use the same methods as other Microsoft cloud services we already consume. Specific details will apply as noted below.

Same region as Teams: Voice data is stored in the same region as our Microsoft Teams data.

Encryption: The voice data for users is encrypted at rest and in transit and is

		protected by Microsoft's security and privacy policies and practices.
Where will the information be stored?		<p>Compliance store: Voice data is stored in the IR tenancy, within the Microsoft Office 365 trusted compliance store.</p> <p>Local copy: When a user uses voice isolation on their device, a local copy of the voice signature is stored encrypted. This signature expires after 14 days and will be replaced with a new download.</p>
Who will have access to the information?		<p>Individual user: The individual user has full control over their profile and can export (save to their device) or delete their own voice profile data.</p> <p>Admins: IT admins do not have access to export the voice data.</p> <p>Microsoft: Microsoft does not access or share the voice data of users with any third parties, unless required by law or with the user's consent.</p>
How long will the information be retained?		<p>One year retention: The data is retained for one year. User's data will be deleted if it isn't used for one year.</p> <p>User can un-enrol anytime: If the user unsubscribes from this feature in the Teams app, the data will be immediately deleted. If a profile hasn't been used for a year it will be automatically removed.</p> <p>Teams account deleted: The voice profile is deleted within 90 days if a Teams account is deleted.</p> <p>Local copy: The local copy of the voice signature is stored encrypted and expires after 14 days. If the device is lost or stolen, our IT team can perform a device wipe that will remove all local data from</p>

		the device, provided it is connected to and accessible via the internet. https://learn.microsoft.com/en-us/mem/intune/remote-actions/devices-wipe#wipe
Does it involve transferring personal information offshore, using a third-party contractor?	Y	Microsoft is a vendor that currently supplies information technology products to IR including Microsoft Teams. The voice data is stored in the same region as our Teams data, which is usually within the jurisdictions of New Zealand and Australia.
Use, disclosure, and accuracy	Y/N	If yes, explain your response
Is the information currently held by IR?	N	
If yes to the above question, for what purpose does IR hold the information?	N/A	
Will the initiative use or disclose information for a different purpose to why it was obtained?	N	
Will IR be sharing personal or taxpayer information with another organisation?	Y	As noted above, the voice profiles of IR staff, if they choose to enrol, will be stored in the Microsoft cloud service.
Describe the data quality – is it accurate, consistent, and complete?	Y	The Teams app prompts the user to read out a standard piece of text that is recorded to produce the profile. This takes around 30 seconds.
What processes are in place to ensure and maintain data integrity?		As noted above, the app has a standardised voice enrolment process.
Access and identification	Y/N	If yes, explain your response
Will the information be stored on a customer or staff member's record?	Y	The voice profile will be associated with the individual's work account that they use to access Microsoft Teams.

Does the initiative affect how people can access information IR holds about them?	N	The voice profile settings will appear as a new item under the existing settings menu in the Teams app.
Does this involve a new way of identifying individuals?	Y	The voice profile is a new way of identifying individuals in Microsoft Teams calls and meetings. It is optional for staff to enrol.
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	We intend to communicate this change to IR people, making it clear it is optional.
If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?	N	
Does the initiative involve tracking or monitoring of movements, behaviour or communications?	Y	At a future date the voice profile might be used to attribute comments to specific individuals when a meeting room speaker is used. Currently comments are only attributed to the individual if they are using the input audio on their assigned work device.

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	We will offer this as an optional feature that individuals can choose to use

Will the proposal have an impact on vulnerable people or those identified as disadvantaged?	N	We will offer the feature to all IR staff. We do not foresee any specific impact on vulnerable or disadvantaged people.
How are we identifying and managing bias or discrimination?		We will offer the feature to all IR staff. We do not foresee any bias or discrimination arising from it.
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	
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?	Y	We plan to test the feature with a small group to verify that it works well before we start the company-wide roll out.
Will decisions informed by an algorithm or use of AI involve human review and evaluation?	N	Not applicable.
Will any automated decision-making process be regularly reviewed to make sure it's still fit for purpose?	N	Not applicable.

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	Low	Export of the voice profile data is controlled by the individual. Our IT admins will not have access to download or export the voice profile information.

<p>M – A moderate amount of personal information (or information that could become personal information) will be handled</p> <p>H – A significant amount of personal information (or information that could become personal information) will be handled</p>		
<p>Sensitivity of the information</p> <p>L – The information will not be sensitive (name, IRD number, or job title)</p> <p>M – The information may be considered to be sensitive (contact details, date of birth plus name plus IRD number, financial information, biometric data)</p> <p>H – The information will be highly sensitive (health or financial details, information about high profile individuals)</p>	Medium	<p>The voice profile is a type of biometric data. Refer to the Storage, security, and retention section above for how the risks will be mitigated.</p>
<p>Significance of the changes</p> <p>L – Only minor change to existing functions/activities</p> <p>M – Substantial change to existing functions/activities; or a new initiative</p> <p>H – Major overhaul of existing functions/activities; or a new initiative that's significantly different</p>	Low	<p>IR people already use the Teams app extensively. This will be a new, optional feature in the app.</p>
<p>Interaction with others</p> <p>L – No interaction with other agencies</p> <p>M – Interaction with one or two other agencies</p> <p>H – Extensive cross-agency (that is, government) interaction or</p>	Low	<p>IR voice profile data will not be shared with any other agencies.</p>

cross-sectional (non-government and government) interaction		
Public impact L – Minimal impact on IR and customers M – Likely to have some impact on customers due to changes to the handling of personal information; or changes may raise concern or media attention H – High impact on customers and the public, and concerns over aspects of project; widespread media interest likely	Low	This change only affects Microsoft Teams. It doesn't affect the Genesys system, which is the primary system for voice communications with end customers.

5. Summary of privacy impact

The privacy impact for this project has been assessed as:	Tick
Low – There is little or no personal information involved; or the use of personal information is uncontroversial; or the risk of harm eventuating is negligible; or the change is minor and something that the individuals concerned would expect; or risks are fully mitigated	
Medium – Some personal information is involved, but any risks can be mitigated satisfactorily	<input checked="" type="checkbox"/>
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

We expect this use of personal information will be uncontroversial as it will be optional, it provides a direct benefit to the individual, and people will be able to un-enrol their voice profile at any time. Overall the change is privacy-enhancing, and the risk of a privacy breach is no higher than for existing products we use in the Microsoft 365 suite including Microsoft Teams.

7. Recommendations

Based on the above assessment, below are summarised recommendations to minimise the impact on privacy. These should be agreed with the senior responsible owner.

Ref	Recommendation	Agreed Y/N
R-01	Communicate this change to business users before starting company-wide rollout, specifically that it is optional and people can unenrol anytime	

8. Document sign-off

Position	Name	Business Unit	Sign-off Date
Sponsor or Business Owner	Tim Crook	Technology Experience	21/02/2025
Privacy Officer	Dawn Swan	ED&I	3/03/2025