NEW ZEALAND
FOREIGN AFFAIRS & TRADE
Manatū Aorere

New Zealand Ministry of
Foreign Affairs and Trade
Manatū Aorere

195 Lambton Quay
Private Bag 18−901
Wellington 6160
New Zealand

T +64 4 439 8000
F +64 4 472 9596

OIA 30497

19 November 2025

Requesters name redacted for proactive release

Tēnā koe Requesters name redacted for proactive release

Thank you for your email of 16 October 2025 in which you request the following under the Official Information Act 1982 (OIA):

> 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.*
>
> *If any part of this request is declined, please provide the specific grounds for refusal under the Act. I would prefer to receive this information in a machine-readable electronic format.*

On 14 November 2025, the timeframes for responding to your request were extended by an additional 3 working days due to the consultations necessary to make a decision on your request (section 15A(1)(b) of the OIA refers).

The Ministry of Foreign Affairs and Trade (the Ministry) does not allow the use of AI with classified information. The Ministry does not yet have the supporting documentation to assess AI tools though its formal software approval process. Therefore, part one of your request is refused under section 18(e) of the OIA as the information requested does not exist. Staff may use online AI sites -except those that are blacklisted- with unclassified information, providing that they follow the Ministry's *Guide to using AI services at MFAT* (which has been released to you under part 2 of your request). However, following a risk assessment (which has been released to you under part 2 of your request) the browser-based Microsoft Co-Pilot Chat is the Ministry's preferred AI tool.

As the Ministry does not have any approved AI tool's part three of your request is therefore refused under section 18(e) of the Official Information Act, as the information does not exist.

For a document schedule of the information in scope of part two and part four of your request please refer to appendix one. Some information is withheld under section 6(a) of the OIA, to avoid prejudicing the security or defence of New Zealand or the international relations of the New Zealand Government.

Please note that it is our policy to proactively release our responses to official information requests where possible. Therefore, our response to your request (with your personal information removed) may be published on the Ministry website: www.mfat.govt.nz/en/about-us/contact-us/official-information-act-responses/

If you have any questions about this decision, you can contact us by email at: DM-ESD@mfat.govt.nz. You have the right to seek an investigation and review by the Ombudsman of this decision by contacting www.ombudsman.parliament.nz or freephone 0800 802 602.

Nāku noa, nā

Sarah Corbett
for Secretary of Foreign Affairs and Trade

Appendix 1

Document schedule for information in scope of OIA 30497				
Request		**Documents in scope**	**Decision**	**Document location**
2	*Any documentation outlining the conditions, guidelines, or policies attached to the approval and use of these tools.*	Information policy	Released in full	Page 1-5 of the collated document
		ICT Acceptable Use Policy	Released in full	Page 6-9 of the collated document
		Guide to using AI services at MFAT	Released in full	Page 10-13 of the collated document
4	*Copies of all completed Cloud Risk Assessments and Privacy Impact Assessments (or equivalent documents) for each of the approved AI tools.*	Privacy Impact Assessment (PIA)	This document is the PIA for all Microsoft 365 products. Sections not relating to AI have been removed as "out of scope".	Page 14-40 of the collated document
		Cybersecurity considerations	Partially released under the OIA. Information redacted under the following sections of the OIA: - 6(a). **Please note:** This assessment was completed in April 2025 and reflects the security environment at that time. Some details may no longer be accurate. **Please note:** Operations "Falcon" and "Kiwi" referred to in the Cyber considerations document are fictional and used as examples only.	Page 41-48 of the collated document

Information policy

Policy

Purpose

The purpose of this policy is to ensure that data and information are well-managed and well-used, and can be used to provide a full evidentiary history of Ministry decisions and inform it's operational and policy decision-making.
This policy has several related standards. Please ensure that you read the correct standard in conjunction with this policy: Declassification and release of historical information standard, History and heritage collection development standard, Official Information Act standard, Data content standard, and Digitisation standard.

Application

This policy will apply to all Ministry of Foreign Affairs and Trade staff, contactors or contracted providers (including vendors) who create, manage or share information with the Ministry.
It encompasses data and information in all formats, including intranet, website and social media content. It includes records.
In this policy the term "information" will be used to include data, content and information.

Responsibilities

The following roles have specific responsibilities under this policy; details are outlined later in the policy:

Role	Responsibility
Executive Sponsor, Public Records Act 2005 (Deputy Chief Executive, People and Operations)	• Ensure that the strategy and policy adopted by the Ministry supports information and records management. • Ensure that strategic and operational planning aligns information and records management with the corporate objectives and business activities of the Ministry; and information and records management is integrated into processes, systems and services. • Ensure that the resources needed to support information and records management are known and sought in funding decisions; and that staff with appropriate skills are available to implement information and records management strategies. • Ensure that information and records management is monitored and reviewed to ensure that it is implemented.

Released under the Official Information Act 1982

Chief Data Officer	• Ensure compliance with organisational information policies, legislation and other government requirements. • Oversee the design, implementation and maintenance of systems, tools, policies, processes and practices that operationalise and are compliant with this policy and that meet business requirements. • Establish any required procedures, guidelines or practices needed to ensure the integrity, quality and effective use of Ministry information assets. • Monitor the effectiveness of controls in place to retain integrity, quality and compliance.
Information Stewards (including data and content stewards)	• Define their information asset. • Review and approve security access controls in relation information assets. • Communicate and promote the value of their information asset. • Monitor and improve conformance with information processes, standards and guidelines. • Improve the quality of their information asset. • Manage, resolve or escalate information asset related issues.
Managers	• Ensure employees, contractors and contracted providers understand their information obligations and comply with Ministry information policy, processes, and standards.
All Ministry staff, contractors and contracted providers	• Understand the responsibilities that this policy and its associated standards and guidelines place upon them. • Manage and use information appropriately by looking after it as an asset, storing and handling it in accordance with Ministry processes and using it with integrity. • Create accurate, complete, impartial and accessible records of decisions and actions. • Save records in approved Ministry recordkeeping systems, so that they can be easily found and used by others, now and in the future. • Create good quality information – accurate, factual, non-defamatory and well-described. • Share and reuse information with colleagues at the Ministry of Foreign Affairs and Trade.
Data, Content and Information Management Specialists	• Support the effective management of information assets at both an enterprise and operational level by implementing this policy in support of business information requirements, applying any legislation or

Released under the Official Information Act 1982

	government directive, and adopting and promoting procedures, guidelines and practices, to support organisational information management initiatives.

Context

At the Ministry of Foreign Affairs and Trade, information is one of our greatest assets. This policy sets out the various requirements for how the Ministry shall manage and use information to:

- demonstrate transparency and accountability;
- support better operational and policy decisions; and
- unlock the value of our information.

Principles

The Ministry will manage its information obligations in accordance with the principle of kaitiakitanga. The Ministry shall demonstrate kaitiakitanga by stewarding its information so that:

- users have trust and confidence in it;
- transparency and accountability is demonstrated;
- information is managed and used with integrity and manaakitanga – treating it as an Aotearoa New Zealand asset; and
- the Ministry approach reflects te ao Maori needs and interests in information.

Policy statements

Well-managed

1- The Ministry will manage information well and with integrity throughout its life; and value it as an organisational asset:
    - All Ministry staff, contractors and contracted providers will create reliable and trustworthy records as evidence of transactions, decisions and actions.
    - Ministry staff will describe information assets to recognised standards.
    - The Ministry will provide approved compliant recordkeeping systems and storage locations; and make them secure from unauthorised access, disclosure, alteration, deletion, loss or destruction.
    - Ministry staff will store records in those systems and/or locations.
    - Access to, use of and sharing of information by Ministry staff must be in line with our legal requirements and obligations, including the New Zealand Official Information Act 1982, Privacy Act 2020, Code of Conduct, Contract and Commercial Law Act 2017, Public Records Act 2005 and the Protective Security Requirements.
    - Information will be managed in accordance with international protocols, conventions and obligations.

Released under the Official Information Act 1982

- Information will be managed by the Ministry through its life-cycle including long-term preservation and access, and catering for technological obsolescence.
- Information will be kept for as long as needed for business, legal and accountability requirements. Staff with appropriate delegations will then systematically dispose of information when authorised and legally appropriate to do so, using a managed process. In the case of information, which must be transferred to Archives New Zealand, the Ministry will declassify then release that information.
- Information identified as being of significant value – whether to inform decision-making or of historical and cultural importance – shall be managed by the Ministry in such a way as to support deriving value from that information.

2- All Ministry projects to develop new systems and/or processes will ensure that information and records management is built into those systems and/or processes.

3- The Ministry will manage its information as a taonga and a historical asset, looking after it on behalf of Aotearoa New Zealand in such a way that its value can be harnessed.

4- Information of interest to Maori will be identified, and management of and access to that information will reflect understood Maori needs and interests.

5- If the Ministry identifies gaps and or risks in the flow of information, it will document them, mandate any required improvements, and report on progress.

Subject to effective governance

6- The Ministry will demonstrate leadership in the governance of information, both within the Ministry, and with that NZ Inc information for which it is the custodian.

7- The Ministry shall maintain executive and operational roles which are accountable for records and information management.

8- Te Ministry will ensure that all information assets and products has assigned stewards who proactively manage their assigned Ministry assets.

9- Stewards will be assigned for information created within any functions outsourced by the Ministry, conducted in partnership with the Ministry, and by the Ministry itself.

10- The Ministry will maintain effective information and records management governance frameworks.

11- The Ministry will identify what information is needed by the Ministry in its statutory roles, and the activities which create information to ensure:

- information is only collected for specific policy, operational business or legislative purposes, and
- information supports the purposes for which it was collected.

Unlock and use

Released under the Official Information Act 1982

12-Information will be made available to benefit Aotearoa New Zealand citizens and businesses, iwi, the public sector, counterparts, partners and international institutions and organisations, unless there are reasons to restrict access.

13-Internally, the Ministry will make information open by default unless there is a good reason to restrict access.

14-Ministry staff will be trusted to balance "need to know" with "responsibility to share".

15-The Ministry prefers digital over paper which means we digitise business processes, in line with legal requirements; and will store born-digital information digitally, rather than printing it out and storing it in paper form.

16-The Ministry will enable the fullest appropriate use of our information through:

- better data access and transformation into reports and dashboards to support evidence-based decision making and internal and external reporting,
- adoption of the CC-BY Creative Commons license over our published information, which permits others to distribute, remix and build upon Ministry work, even commercially, as long as they credit the Ministry for the original creation,
- proactively releasing and opening up our information through publishing high-value information to our website and to data.govt.nz,
- use of common data and information metadata schemas, such as JSON, and the data.govt.nz metadata schema,
- publishing high-value corporate information such as policies, guides, forms and people information to our intranet,
- publicly releasing our high value historical information, and creating or facilitating the creation of new knowledge products off that information,
- creating new use and utility from existing information, and
- enabling safe internal and external content collaboration.

Exceptions management and consequences of policy breach or non-compliance

Any deviation from this policy must be approved in advance by the appropriate authority as defined in the Instrument of Delegation; and informed to the Ministry Data and Informaiton Governance Group. If this is not specifically defined in the Instrument of Delegation, then the Deputy Secretary of the relevant division where the deviation is to occur must approve any deviation in advance. The deviation must be notified to the policy owner and sponsor, and the Divisional Manager, Audit and Risk.
Unauthorised breaches of this policy may be viewed seriously by the Ministry as a breach of the Code of Conduct. The Code of Conduct states that "employees are expected to fully comply with Ministry policies in their work". Depending on the circumstances, breach of this policy may result in disciplinary action, up to and including dismissal. Any such breaches should be notified to the policy owner and sponsor and the Divisional Manager, Audit and Risk.

Released under the Official Information Act 1982

ICT acceptable use policy

Policy

Purpose

The purpose of this policy is to set out the acceptable usage of the Ministry's ICT Assets required by all Users.
This policy must be read in conjunction with the related standards.
Application
This policy will apply to:

- All users of Ministry ICT assets, including staff, contractors, consultants, support providers and authorised staff from other agencies.
- Visitors to the Ministry who have authorised use of the Ministry's guest Wi-Fi network (HuiNet or Te HonoNet).
- All Ministry ICT assets.

Responsibilities

The following roles have specific responsibilities under this policy:

Role	Responsbility
Delegated authorities	• As set out in the Instrument of Delegation.
Chief Information Officer (CIO)	• Overall responsibility for the Ministry's information and communications technology (ICT) assets, including setting the Ministry's overall ICT strategic direction.
Chief Information Security Officer (CISO)	• Overall responsibility for the Ministry's cybersecurity and ICT security, including setting the Ministry's cybersecurity and ICT security strategic direction.
Managers	• Monitor compliance with this policy and associated standards, processes and guidelines. • Confirm in the annual Audit and Risk control self-assessment survey that they and their staff understand and comply with their obligations under this policy. • Ensure all ICT security breaches and incidents are immediately reported to the IMD Service Centre and SORD.
Users	• Familiarise themselves with this policy and associated standards, processes and guidelines. • Use all Ministry ICT assets in accordance with this policy and associated standards, processes and guidelines.

Released under the Official Information Act 1982

	• Immediately report any actual or potential ICT security risks, concerns, incidents and breaches, including suspected compromise or inappropriate use of Ministry ICT assets to the IMD Service Centre, SORD, and the user's manager. • Immediately report any damaged, lost, stolen or compromised Ministry ICT assets to the IMD Service Centre, SORD, and the user's manager. • Comply with the Code of Conduct in all use of the Ministry's ICT assets.
Security and Organisational Resilience Divisions (SORD)	• Provide advice and training relating to the security and secure usage of ICT assets.
Information Management Division (IMD)	• Provide secure, performant ICT assets, and training in how to use them appropriately.

Released under the Official Information Act 1982

Context

The Ministry relies on performant, high-quality, accessible and secure information to deliver its business, at pace, around the world. This means that our Information and Communications Technology (ICT) must be characterised by security, availability and integrity to maintain business continuity and the trust of the New Zealand government, the public and international partners.

Principles

The Ministry's ICT assets will be used in accordance with the following principles:

- The Ministry's ICT assets are used in an acceptable, ethical and lawful way.
- The Ministry's cyber security is not jeopardised.
- The integrity of the Ministry's ICT assets and the information stored within is protected against loss, damage, disruption, unauthorised disclosure and misuse.
- Disruptions to Ministry information management services and activities are minimised and adequately recovered from.
- All information transacted through a Ministry system is owned by or under custodianship of the Ministry.
- All access to and use of Ministry systems may be monitored, collected, and provided to senior leaders and relevant third parties, for information assurance purposes, with the approval of the Chief Information Officer (CIO).

Policy statements

ICT Asset security and access

- The security and integrity of the Ministry's ICT assets, data or information must be protected from loss and unauthorised use, modification or release.
- Users will only access the Ministry's ICT assets, data or information when they are authorised and have a business need to do so.
- Users will adhere to the Ministry's standards to ensure the security and integrity of Ministry ICT assets.

ICT asset use

- Ministry ICT assets must not be used for any illegal activity, personal commercial purposes, personal political purposes, misrepresentation or in a way that would contravene the Ministry's Code of Conduct.
- Users will use the Ministry's ICT assets in a way that minimises any adverse impact on the functionality of the Ministry's systems and the access of other users.
- Personal use of the Ministry's ICT assets by authorised users is permitted within reasonable limits, provided it aligns with the Ministry's standards and incurs minimal cost. Incidental personal use is acceptable as long as it does not interfere with official duties or compromise security.

Associated Policies, Standards and Guidelines

- The use of all Ministry ICT assets must comply with the ICT acceptable use standard.
- The use of all Ministry mobile devices and accessories must comply with the Mobile device security standard.
- The use of all Administrator accounts must comply with the ICT administrator standard.
- The disposal of all ICT equipment must comply with the IT equipment disposal standard.
- The use of generative Artificial Intelligence must be informed by the Ministry Guide for using AI Services at MFAT.

Exceptions management and consequences of policy breach or non-compliance

Any deviation from this policy must be approved in advance by the appropriate authority as defined in the Instrument of Delegation. If this is not specifically defined in the Instrument of Delegation, then the Deputy Secretary of the relevant division where the deviation is to occur must approve any deviation in advance. The deviation must be notified to the policy owner and sponsor, and the Divisional Manager, Audit and Risk. Unauthorised breaches of this policy may be viewed seriously by the Ministry as a breach of the Code of Conduct. The Code of Conduct states that "employees are expected to fully comply with Ministry policies in their work". Depending on the circumstances, breach

Released under the Official Information Act 1982

of this policy may result in disciplinary action, up to and including dismissal. Any such breaches should be notified to the policy owner and sponsor and the Divisional Manager, Audit and Risk.

Released under the Official Information Act 1982

Guide: Use of Artificial Intelligence (AI) Services at the Ministry of Foreign Affairs and Trade.

Purpose

The purpose of this guide is to help Ministry employees use Artificial Intelligence (AI) on Ministry devices in a way that is safe, responsible, and protects our information. With the recent launch of a six-month campaign to support using Microsoft CoPilot Chat on Ministry laptops and desktops, and the ongoing consideration of how staff utilise other AI services, this guide is intended to provide practical information for staff to use AI.

This guide should be read in conjunction with these FAQS - Using AI Services at MFAT, which provide more detailed advice on common questions regarding AI use, particularly CoPilot Chat.

Who must follow this guide

- All Ministry employees and contractors working for, or on behalf of MFAT.
- Anyone using AI services provided by external parties for Ministry work.

Definition of GenAI

GenAI refers to artificial intelligence (AI) systems which use data and computing power to enable machines (or software) to perform tasks that typically require human intelligence. In general, GenAI systems work by ingesting and processing large amounts of data about a particular issue or context quickly and analysing it to aid research and understanding. GenAI can create realistic text, images, music, and other media that resemble human creativity.

AI services you can use at MFAT

Microsoft CoPilot Chat: MFAT has rolled out a six-month campaign to introduce Microsoft Copilot Chat, a GenAI service which is available on Ministry laptops and desktops on the Orange network. The purpose of this campaign is to support staff in using this AI service and gather information about the opportunities and advantages Copilot Chat offers.

Copilot Chat is configured with enterprise level data protection controls to safeguard the information put in and received back from being shared or used in undisclosed ways without MFAT permission. **For this reason, Copilot Chat is the Ministry's preferred AI service due to the greater controls it offers.**

While staff are currently able to use other permitted AI services on Ministry devices, Copilot Chat offers greater controls for protecting our information. We therefore strongly encourage staff to use Copilot Chat.

The DeepSeek AI App is not permitted, read about why here.

Released under the Official Information Act 1982

Ask for advice

We recognise that AI technology is rapidly changing and new services will be emerging frequently. If you are unsure of any AI services that you would like to use, seek advice from KIAServices@mfat.govt.nz

What AI can be used for at MFAT

AI can be used for a range of functions. Staff are encouraged to see where AI may be helpful to their work. Possible uses of AI services (including Copilot Chat) include:

- Summarising large quantities of information
- Analysing large information sets
- Reducing time searching for or summarising information on the internet
- Collating information from multiple sources
- Assisting with initial drafting and formatting of documents and emails
- Translation of media and documents
- Production of infographics and generating images
- Brainstorming, and generating or testing new ideas

While AI services can be a useful tool for assisting decision making, people are ultimately accountable for the use of AI-generated outputs. Human oversight over AI functions is necessary, noting that AI can produce errors/hallucinations, draw on datasets that are out of date, and may include bias and a lack of diversity and views.

Permitted Classification Level

Only UNCLASSIFIED information can be used in generative AI services. This means that all prompts fed into, or documents or information uploaded onto an AI service must be UNCLASSIFIED with no endorsements. This can include UNCLASSIFIED information that is not publicly available i.e. UNCLASSIFIED information that is stored on GDM or the MFAT Orange network. Information on classification definitions can be found in the Classification Quick Reference Guide

Personal and private information must not be used in an AI service unless this information is publicly available. Do not input or search for your own or others personal or private information. Guidance on what information is considered publicly available can be found here. Guidance on MFAT's general obligations with regard to personal and private information can be found here.

What you need to know when using AI at MFAT

In addition to the above requirements for classification, personal, and private information, use of AI by MFAT staff should be done in accordance with the following guidelines:

- **Consistent with MFAT's ICT Acceptable Use Policy, staff should only undertake work using their Ministry devices. This includes the use of generative AI for work purposes.** All AI services being used on Ministry devices are monitored as part of the ICT Acceptable Use Policy. Staff should not use AI for work purposes on their personal devices.

Released under the Official Information Act 1982

- **Staff should keep informed of which AI services are not permitted on Ministry devices.** The DeepSeek AI App is not permitted, read about why here.

- **Staff should exercise good judgement when using AI services**, noting that AI can provide incorrect information (hallucinations) and draw on information sources that may be out of date or include bias and a lack of diversity and views. Staff are encouraged to examine the sources that an AI service has drawn on for providing outputs, and to cross-check these where appropriate, particularly where AI is used as a basis for decision making. Staff should ensure that AI services are drawn from sources that are accurate, credible, and reputable.

- **Under the Public Service AI Framework, MFAT must ensure that its use of AI is trustworthy and responsible**. Guidance on what the trustworthy use of AI looks like can be found in the OECD AI Principles. Cabinet approved the OECD AI Principles as the key direction for responsible AI use in New Zealand in July 2024.

- **People are accountable for AI use and its outputs.** While AI can assist with decision making, staff should always maintain human oversight and review of AI outputs.

- **Respect Māori Data sovereignty.** If staff are using AI across Māori Data, respect and uphold the mana and dignity of Māori people, language, images, culture, resources and environments. AI services do not always authentically represent indigenous cultures; this depends on the data sources they use. Staff should always maintain informed oversight to cross check and validate the confidentiality and integrity of Māori information.

- **Staff should note that AI prompts, inputs and responses are public records, subject to the Public Records Act 2005.** All information generated through the use of AI at MFAT, including prompts, inputs and responses, are considered official information. Accordingly, this information is subject to the Official Information Act and if requested, will be reviewed and considered for public release.

- **Report breaches.** Report any privacy, information and security breaches immediately, in accordance with the Ministry's Security, Information and Privacy Policies.

Resources associated with this guide

MFAT AI guidance is aligned with, and underpinned by, the following resources.

Internal resources:

- Information Policy
- MFAT Artificial Intelligence Governance Framework
- MFAT Artificial Intelligence Policy *(currently under development)*
- ICT Acceptable Use Policy
- Privacy Policy
- Security Policy
- Māori Data Governance Strategy

External resources:

Released under the Official Information Act 1982

- Public Service AI Framework Strategy
- OECD AI Principles
- Algorithm charter for Aotearoa New Zealand - data.govt.nz
- Public Records Act 2005.
- Official Information Act 1982

Released under the Official Information Act 1982

Digital Workplace Transformation Programme - Short Form Privacy Impact Assessment (PIA)

This document serves as a platform-level Privacy Impact Assessment (PIA) for the implementation of Microsoft 365 (M365) and Azure. It is intended to be a living document, subject to updates and revisions as advised by the Privacy Advisor. This approach acknowledges that as the program progresses and incrementally rolls out various capabilities, our understanding of the privacy implications may evolve. Consequently, this PIA will be revisited and adjusted in line with these developments to ensure ongoing compliance and to address any emerging privacy concerns effectively. Our commitment is to maintain a high standard of privacy and data protection throughout the lifecycle of the Programme.

Description of your project/initiative

What is the project aiming to do?

The Ministry has made the decision to store its information in the cloud, using Infrastructure (IaaS), Platform (PaaS) and Software as a Service (SaaS) products as part of Microsoft Azure and Microsoft 365 (M365). The Ministry has established the Digital Workplace Transformation Programme to deploy Microsoft 365 and Azure.

M365 Deployment

M365 Application	Description
Out of scope	

Released under the Official Information Act 1982

M365 Application	Description
Out of scope	

Released under the Official Information Act 1982

M365 Application	Description
Out of scope	
CoPilot Chat	Microsoft 365 CoPilot Chat is an AI-powered assistant that helps users interact with the web (and Microsoft 365 apps with the Copilot license) through natural language. It can automate tasks, generate content, provide insights, and answer questions, making it easier to work efficiently and creatively within the Microsoft 365 ecosystem. CoPilot Chat is available to MFAT users already but is only configured to interact with publicly available data on the web. CoPilot Chat interactions are logged in Microsoft Purview for system administrators to monitor. Users are not to enter personal or classified information.

Azure Deployment

The scope of the Azure deployment is limited to implementing the base configuration i.e. setting up the foundation. The Programme does not plan migrate any existing on premises systems to the Azure platform. When migration of existing on premises workloads does happen the Privacy Adviser will be consulted. Due to the breadth of the platform individual services have not been listed.

Released under the Official Information Act 1982

What type of personal information is involved? (e.g email addresses, phone numbers or bank account details)

Due to the diverse range of functions and activities conducted by the Foreign Ministry, a wide variety of personal information may be involved. This may include, but is not limited to:

Contact Information such as email addresses, phone numbers, and physical addresses of individuals, both within and outside the Ministry.

Identification Details such as passport numbers, national identification numbers, or other government-issued identification information.

Financial Information like, bank account details or financial records may be processed, especially in relation to financial transactions or official reimbursements.

Travel Information including flight itineraries, hotel bookings, and visa information for staff or individuals associated with the Ministry's activities.

Diplomatic Correspondence containing personal information, which may include sensitive discussions and negotiations.

Biographical Information in the of background information about individuals, including educational history, employment records, and professional affiliations.

Health and Medical Records when health-related information may be processed, particularly for diplomatic and consular services.

It is important to note that the specific type of personal information involved may vary depending on the nature of the Ministry's activities and functions

Additionally, the Ministry may generate personal information when discussing, communicating about, or forming opinions on matters that pertain to individuals. Lastly, like any employer, the Ministry collects and generates personal information about its employees.

Released under the Official Information Act 1982

How is personal information collected, obtained and used?

This Programme is uplifting the Ministry's technology and is not performing business process redesign (however there will be process changes related to the administration of the new solutions). By introducing M365 platform the Programme will not introduce new collection of personal information from additional sources, nor does it introduce the utilisation of personal information in new ways or intentional disclosure to new agencies. Therefore, any risks presently associated with how the Ministry collects, obtains and uses personal will remain unchanged which is includes the incidental collection of personal information.

Nevertheless, it is important to acknowledge that the Ministry's decision to entrust another entity (in this Microsoft) with the storage and processing of its data/information does require it to relinquish some of control. This decision raises specific risks that need to be considered.

Will personal information be disclosed overseas? If yes, please explain how and why.

In general, the core applications provided by the M365 platform and the Azure platform will be stored in Australia. There are some exceptions to this, as shown in the table below.

M365 Service	Data Residency	Types of Information
Out of scope		

Released under the Official Information Act 1982

M365 Service	Data Residency	Types of Information
Out of scope		

Released under the Official Information Act 1982

M365 Service	Data Residency	Types of Information

Out of scope

CoPilot Chat	•	Users are not to input personal information into CoPilot. Guidance has been drafted for staff with a specific privacy section to ensure compliance.

For the purpose of the Privacy Act, this is not considered a disclosure since Ministry is still remain's responsible for the personal information stored in Microsoft's public cloud. In rare circumstances a Microsoft Employee may access personal information in so far as to troubleshoot and provide technical support. In these rare cases, disclosure will be performed in accordance with the Privacy Act, IPP 12. It should be noted that the, the Ministry has enabled a feature called the 'Customer Lockbox' within its M365 tenant, this ensures that Microsoft can't access the Ministry's content to do service operations without the Ministry's explicit approval.

Jurisdictional Risk

Storing personal data on overseas servers, such as in the case of using Microsoft Cloud Services, entails jurisdictional risks. The New Zealand Privacy Act allows for the transfer of personal data overseas for storage and processing, but maintains that the principal agency in New Zealand retains liability and accountability for this data. To address these risks, core app data is being restricted to Australian Data Centres and with some non-core applications to USA data centres. The data stored in Australia has the benefit from

Released under the Official Information Act 1982

the protections contained in the Australian Privacy Act 1988, which is similar to New Zealand's Privacy Act and based on OECD Privacy Guidelines. The data stored in the USA similarly has the broadly equivalent protections to NZ privacy legislation.

The Australian Act imposes obligations on entities through its 'Australian Privacy Principles' (APPs), like mandatory breach notifications and record-keeping for lawful access. However, it lacks EU adequacy status due to certain exemptions, which do not impact the use of Microsoft Cloud Services. Under the Australian Act, third-party providers share equal liability and obligation with the principal agency, meaning they cannot offload legal responsibilities onto customers. The USA also provides broadly equivalent privacy protections.

What measures will be in place to ensure privacy is protected? (e.g are there safeguards in the contract or are you storing it in GDM?)

Data Encryption

Microsoft 365 and the Azure platform uses encryption to protect data both in transit and at rest. This means that data is encrypted when it's transmitted over the internet and also when it's stored on Microsoft's servers.

Identity and Access Management

Microsoft 365 and the Azure platform provides robust identity and access management controls. Multi-factor authentication (MFA) is available to enhance security by requiring users to provide multiple forms of authentication before accessing their accounts.

Compliance and Certification

Microsoft regularly undergoes third-party audits and certifications to ensure compliance.

Privacy Controls

Microsoft provides privacy controls and features that allow administrators to configure and manage privacy settings for users and data. These controls can include policies for data retention, data loss prevention, and more.

Released under the Official Information Act 1982

Privacy by Design

Microsoft incorporates privacy considerations into the design and development of its products and services, following principles of Privacy by Design.

Transparent Data Handling

Microsoft provides transparency about how it handles customer data through its Privacy Statement and Service Trust Portal, which offers information on data processing and security practices.

Data Portability

Microsoft offers tools and mechanisms for customers to export their data from Microsoft 365 services, ensuring that customers can retain control over their information.

Incident Response and Reporting:

Microsoft has incident response plans in place and promptly reports any data breaches or security incidents to affected customers.

Customer Data Location:

Microsoft allows customers to choose the geographic location where their data is stored, helping to meet data sovereignty and compliance requirements.

Note: Some of the safeguards need to enabled through configuration and are not able enabled by default.

Released under the Official Information Act 1982

High level privacy analysis

This section contains a high level privacy analysis, with a focus on considering the project against the information privacy principles. The tables below use the following "assessment" of risk:

None	Low	Moderate	High

See MFAT's risk assessment toolbox to assist with your risk assessment: http://bpportal.orange.mfat.net.nz/AuditRisk/Pages/Risk-management.aspx

Alignment with the Information Privacy Principles (IPP)

Please describe below in the right hand column how you plan on meeting the IPPs. A summary of what the IPPs cover is at the end of the document for ease of reference.

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
Collection Principles 1, 2, 3 and 4				
As described above, the Programme is introducing M365 services, this will entail moving key systems such as the email service and GDM to the Microsoft cloud. Therefore, the Ministry will store and process within the Microsoft Cloud solution the personal information its collects and holds (in these systems), with the exception of any government information security classified as above RESTRICTED. However, it should be noted Programme is not changing the business processes of the Ministry related to the purpose, source, or the manner in which personal information is collected. Therefore, any current risks associated with	**Principle 1** The move to the Microsoft cloud will not change the purpose for which the Ministry is collecting personal information. The Ministry should be only collecting personal information to fulfil its statutory functions and activities, including its function as an employer. Any existing risks such as over collection will remain unchanged. **Principle 2** The move to the Microsoft cloud will not change the sources used by Ministry to collect personal information. The	Low	**Principle 3** 1. The Programme has an extensive change and communication plan. As part of this the Programme will inform users that the M365 platform is hosted in the cloud. 2. The Ministry already states in its published Privacy Statement its uses third party providers to store and process its data. See Clause 'Storage and Security' https://www.mfat.govt.nz/en/privacy/#bookmark0	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
personal information collection will remain including incidental collection.	Ministry must collect personal information directly from the individuals concerned unless one of the IPP2 exceptions applies. Any existing risks such as collecting information from third party sources about an individual etc will remain unchanged. **Principle 3** The Ministry fails to provide its customers or employees with notice about the storage and processing of personal information on an offshore cloud platform. **Principle 4** The move to the Microsoft cloud will not change the manner in which the Ministry is collecting personal information. The Ministry should not be collecting personal information in ways that are unlawful or, in the circumstances, unfair or unreasonably intrusive. Note: Information should be only up to and including Restricted and excluding NZEO endorsed.			

Released under the Official Information Act 1982

High level analysis		Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
	Information storage, security and disposal Principles 5 and 9				
The Ministry is entrusting its data to a third party and this requires it to transfer the data to the cloud and to data centres in other countries. This means that data protection is a key risk for the Ministry – while data is in transit and at rest – and we must ensure that using Microsoft services does not put our data at any more risk of harm.		**Principle 5** 1. The Microsoft Cloud might not meet the Ministry's security requirements.	Low	The risk of MS Cloud falling short of the Ministry's security requirements is mitigated in the following ways: 1. Microsoft's Cloud services offer very comprehensive security features, including strong encryption in transit and at rest (i.e. when data is moving across the internet and when it is stored in Microsoft's servers). 2. As part of Microsofts contractual obligations it promises to provide technical and organisational security measures that comply with relevant ISO standards e.g. ISO 27001, ISO 27002, and ISO 27018. This far exceedes the Ministry's present onpremises capabilities. Refer to Microsoft Products and Services Data Protection Addendum Last updated November 15, 2023	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
	2. There is a risk of deliberate or accidental access or disclosure of personal information stored on the cloud services made possible by a Ministry staff member.Some specific examples include: • Software misconfiguration leading to inadvertent disclosure of personal information. • Accidentally sharing information (e.g., file) with an unintended party.	Medium	1. To prevent accidental misconfiguration the the programme has the taken the following actions: • Production and testing environments separate. With changes being testing in a non-production environment before deployment. • Implemented role based access to control to ensure only authorised staff can make changes. • Following IMD's change management process which 2. To prevent accidental sharing of information by the end user the following actions will be taken: • The M365 services will be closed by and only opened by design • Monitoring of information flows to make sure that only permissible information is exiting the cloud services – through data loss prevention policies and the like.	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
	Principle 9 1. There is a risk of Microsoft or one of its subprocessors retains Ministry personal information for longer than it expected by Ministry or legally required as per the terms of the contract or NZ Law.	Low	• Comprehensive change manage management providing training and guidance. • Reinforcing existing Ministry privacy guidance. This risk mitigated by the following measures: 1. During the term of the Ministry has full control over the retention and disposal of its information stored in Microsoft's public cloud and can delete it directly. 2. As part of Microsoft's contractual obligations, it will retain customer data for 90 days after the termination of a subscription so that the customer may extract it. After 90 days, the data is deleted (unless there is a legal requirement to retain it).	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
	2. The Ministry retains personal information for longer than it is required.	Low	1. The Programme is engaging KIA to perform Information and Data Apprasial's before deploying a new Microsoft cloud services so it can determine the appropriate retention period for the information.	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
Access and correction Principle 6 and 7				
The Ministry must respond to an individual's requests for access or correction of personal information in compliance with principles 6 and 7 and, in particular, within Privacy Act timeframes.	The use of the Microsoft Cloud solution might impact the Ministy's ability to meet its obligations under principles 6 and 7 due to an inability to access the information for reasons outside Ministry's control.	Low	Microsoft is contractually obligated to provide the Ministry the ability to access, extract and delete data stored in each of its online services at all times. Therefore, the Ministry should be able to meet its legal obligations pertaining to Principle 6 and 7 of the Privacy Act. *"At all times during the term of Customer's subscription or the applicable Professional Services engagement, Customer will have the ability to access, extract and delete Customer Data stored in each Online Service and Professional Services Data"*	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
Accuracy of information Principle 8				
The Ministry must take reasonable steps to ensure that personal information is accurate etc before using or disclosing it. However, the Microsoft services are simply containers for the information being produced or collected by the Ministry's business processes. Therefore, the Microsoft services do not materially influence the accuracy of the information held within it.	The move to the Microsoft cloud will not change the accuracy of information being produced or collected by the Ministry's business processes. Any existing risks associated with the accuracy of the information being produced or collected by the Ministry will remain unchanged and continue to persist.	Low	Users must ensure that inputed information is accurate.	Low
Use and disclosure Principles 10 and 11				
Use and Disclosure By permitting other agencies, including Microsoft and its subprocessors, to store Ministry's information, the Ministry could be exposing the data to an increased risk of misuse.	Microsoft or its subprocessors could use personal information about Ministry's customers or employees for other purposes not related to the purposes for which the Ministry collected it (such as for send out travel advisory notices).	Low	Microsoft has made a contractual commitment to only use the Ministry's data to deliver its online services. Furthermore, Microsoft is contractually responsible for ensuring that its subprocessors will not use the Ministry's data for any purpose other than delivering the specific services Microsoft has requested. Lastly, if Microsoft uses the Ministry's data for other purposes, it will be in breach of the Australian Privacy Act. "When providing Products and Services, Microsoft will not use or otherwise process	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
			Customer Data, Professional Services Data, or Personal Data for: (a) user profiling, (b) advertising or similar commercial purposes, or (c) market research aimed at creating new functionalities, services, or products or any other purpose, unless such use or processing is in accordance with Customer's documented instructions." "Microsoft will not disclose or provide access to any Processed Data except: (1) as Customer directs; (2) as described in this DPA; or (3) as required by law. For purposes of this section, "Processed Data" means: (a) Customer Data; (b) Professional Services Data; (c) Personal Data; and (d) any other data processed by Microsoft in connection with the Products and Services that is Customer's confidential information under Customer's agreement. All processing of Processed Data is subject to Microsoft's obligation of confidentiality under Customer's agreement. " "Microsoft is responsible for its Subprocessors' compliance with Microsoft's obligations in this DPA. Microsoft makes available information about Subprocessors on a Microsoft website. When engaging any Subprocessor, Microsoft will ensure via a written contract that the Subprocessor may access and use Customer Data, Professional Services Data, or Personal Data only to deliver the services Microsoft has retained them to provide and is prohibited from using Customer Data, Professional Services Data,	

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
			or Personal Data for any other purpose. Microsoft will ensure that Subprocessors are bound by written agreements that require them to provide at least the level of data protection required of Microsoft by the DPA, including the limitations on disclosure of Processed Data. Microsoft agrees to oversee the Subprocessors to ensure that these contractual obligations are met." Microsoft states in its agreement it will not release any third party unless it is required by law. Upon receipt of any other third-party request for Processed Data, Microsoft will promptly notify Customer unless prohibited by law. Microsoft will reject the request unless required by law to comply. If the request is valid, Microsoft will attempt to redirect the third party to request the data directly from Customer.	

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
Lawful Request Lawful requests for data are those made under legal authority by government, law enforcement, or national security agencies, typically through court orders or legislative provisions. Storing data in an onshore cloud solution could subject it to lawful requests both domestically and internationally through law enforcement cooperation agreements. Australia's legal framework for lawful access by law enforcement and national security agencies is similar to New Zealand's, with due process and oversight requirements.	Microsoft may be required by law to disclose Ministry data to third parties, including law enforcement agencies. While it may not breach the Privacy Act (by virtue of section 10), it may undermine public trust.	Low	Microsoft will only disclose or provide access to any Processed Data as required by law provided that the laws and practices respect the essence of the fundamental rights and freedoms and do not exceed what is necessary and proportionate in a democratic society and, as applicable, to safeguard one of the objectives listed in Article 23(1) of GDPR. Microsoft will not provide any third party: (a) direct, indirect, blanket, or unfettered access to Processed Data; (b) platform encryption keys used to secure Processed Data or the ability to break such encryption; or (c) access to Processed Data if Microsoft is aware that the data is to be used for purposes other than those stated in the third party's request. In support of the above, Microsoft may provide Customer's basic contact information to the third party.	Low

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
However, amendments to Australian law, particularly the Telecommunications Act 1997, aimed to aid law enforcement and intelligence agencies in intercepting encrypted communications. These amendments primarily target communication providers and manufacturers, and less so data storage providers like Microsoft. Therefore, requests under this Act would likely focus on Microsoft's consumer messaging services rather than its enterprise cloud services. Microsoft is transparent about lawful requests, regularly publishing a Law Enforcement Request Report. Recent reports show a low number of such requests in Australia, with no content disclosure in the last five years. There's also a low risk of overseas governments accessing information stored outside their territory, even with laws like the US CLOUD Act, which Microsoft strives to minimize in impact. Overall, the likelihood of data stored in Australian or US cloud services being subject to lawful requests from Australian or US authorities is low. There's no indication that storing data in Australia or another overseas location would inherently increase the risk of lawful requests from those countries.				

Released under the Official Information Act 1982

High level analysis	Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
It should be noted that the Ministry is NOT a pioneer amongst New Zealand government agencies in using Microsoft cloud services domiciled in Australia and the US, for example, NZ Police, IRD, MSD, NZTA, NZDF all use M365 and Azure.				
Overseas disclosure Principle 12				
The Programme intends to provide the Ministry with the ability to connect, communicate and collaborate with the Ministry's foreign partners through MS Teams.	Sensitive personal information could be inadvertently disclosed to foreign partners. Sensitivie personal information may then be captured and subject to foreign states' jurisdiction and beyond MFAT's ability to retain control over that data.	Low	Staff are still subject to our code of conduct and acceptable use policy and should not disclosing Publicsing existing privacy policies and guidelines when launching this feature. Overseas agencies will also be subject to similar code of conduct, acceptable use policies and data security to provide further assurance as to security of data	Low

Released under the Official Information Act 1982

High level analysis		Summary of risks of this project (Set out here is there are any issues with meeting the IPP)	Assessment	Summary of mitigations	Re-Assessment
				and information to the extent lawfully permissible.	
	Unique identifiers Principle 13				
The Programme does not intend to create any new unique identifiers.		N/A	N/A	N/A	N/A

Privacy enhancing measures

No additional privacy enhancing measured have been identified at the time of writing this document. If such measures are identified they will be considered in the future.

Is a full PIA required/desirable?

The implementation of M365 and Azure primarily enhances the Ministry's technological capabilities, rather than altering the way personal data is handled within the Ministry or how business processes are conducted. On this basis we do not recommend a full PIA.

Note: We will engage with CLU and seek legal advice when necessary as the programme progresses.

Approvals

Date this short form PIA was approved: _8/12/2023______________________

Person/Role	Involvement in this PIA	Initial/Date
Ravi Gokal – Snr. Business Analyst	Author	

Released under the Official Information Act 1982

Brett Nisbett – Project Manager	Contributor	
Kerry Oakly – Project Manager	Contributor	
James Gallagher, Associate Counsel – Corporate Legal Unit	Checked legal	29/11
Shanell Christian, Senior Privacy Adviser	Reviewed updated wording	SC 20/06/2024 04/03/2025 Avepoint 13/05/2025 CoPilot
Jeremy Salmond - Privacy Officer	Checked privacy compliance	JS 20/06/2024 06/03/2025 21/5/2025 Copilot

Released under the Official Information Act 1982

Released under the Official Information Act 1982

Summary of Information Privacy Principles

Principles 1, 2, 3 and 4: Collection and Purpose of the information

- The purpose for which the information is being collected is lawful.
- The information is provided by the individual it relates to.
- If the information is not directly from the individual, you can collect if it would not prejudice the individual concerns, the information is publicly available or to collect from the individual would prejudice the purpose for which the information is being collected. Note: If you are not collecting personal information directly from the individual and have any concerns contact CLU.
- The individual providing the personal information needs to know that the information is being collected, the purpose and how long it will be stored.

Principles 5 and 9: Security and Storage of the personal information

- Personal information will be held securely and steps will be taken to ensure that no information is lost or improperly accessed.
- We should not hold personal information for longer than we need.

Principles 6, 7 and 8: Access and correction of the personal information

- The individual should be able to access their personal information.
- We should be able to retrieve information we have collected.
- If we have collected information we should be able to access information so people can correct it if requested.
- We need to ensure, as far as possible, that the information is accurate especially before we are intending to use it.

Principles 10 and 11: Use and Disclosure

- When we use the information it must be in accordance with or closely associated with the purpose it was originally collected.
- We cannot use information collected for one purpose for an entirely different purpose.
- If we hold personal information and it meets an express exception (such as maintenance of the law or to lessen a serious threat) we can use it for a different purpose.
- We cannot disclose personal information we hold to any other person or agency unless we have reasonable grounds to believe that it is in accordance with the purpose it was collected or it meets one of the exceptions. NB if you are considering relying on IPP 10 or 11 come talk to CLU.

Principle 12: Overseas Disclosure

- When we disclose personal information to a foreign person or entity (including to a foreign government) the information must be subject to comparable privacy protections – however this does not apply in emergency situations for example, where there is a serious threat to health.

- This might mean we have undertakn due diligence on local laws and assessed them to be comparable or including privacy protections in any contract or information sharing agreement.
- If we are unable to assess the above, then we must expressly bring to the attention of individuals concerned that their information will be disclosed without privacy protections, and obtain their express consent.

Principle 13: Unique Identifiers

- When we collect personal information and we intend to assign a unique identifier (that makes it easier to identify an individual) then we need to meet criteria IPP 12. If we are assigning an identifier come talk to CLU.

Released under the Official Information Act 1982

Addendum 1: AvePoint

AvePoint understands the importance of security and operational risk management and holds the following certifications: IRAP assessed to a PROTECTED level, ISO 27001:2013, ISO 27017:2015, ISO 27701:2019, SOC 2 Type II certified + HITRUST, Cloud Security Alliance (CSA) Security, Trust, Assurance and Risk (STAR) Level 2 as a cloud service provider, ASD Essential 8 Maturity L3, FedRAMP, Microsoft 365 and TISAX. More information can be found AvePoint's Trust Center.

AvePoint has been consistently recognised for its market leading software through various industry accolades and awards. Most recently AvePoint was shortlisted in three categories at the 2023 SaaS Awards, became a finalist in the Global Microsoft Partner of the Year Awards for the Apps and Solutions for MS Teams and was named a finalist for **Australian Cyber Security Awards 2023** for Data/Server Security Company of the Year.

Moreover, AvePoint has a long-standing partnership with Microsoft, being one of Microsoft's very first Microsoft SharePoint ISV when SharePoint (what underpins Microsoft 365) debuted in 2001. Today, AvePoint is the only ISV offering an all-in-one approach to migrating, managing, and protecting customers' Microsoft technology investments. AvePoint is managed by Microsoft both globally and locally in Australia and New Zealand. As a Depth Managed Microsoft Gold Certified Application Development Partner and Gold Certified Collaboration and Content Partner, AvePoint is fortunate to be a part of Microsoft's exclusive Technology Adoption Program (TAP) and Syntex Development program. This status is reserved for the top 1% of Microsoft's partner ecosystem worldwide.

Released under the Official Information Act 1982

NEW ZEALAND
FOREIGN AFFAIRS & TRADE
Manatū Aorere

April 2025

Cyber Security Considerations: Microsoft 365 Copilot chat web-grounded

Prepared by: Principal Adviser Cybersecurity

Released under the Official Information Act 1982

Enabling Microsoft 365 Copilot Chat (Web-Grounded Only) for MFAT

Executive summary:
The deployment of Microsoft 365 Copilot Chat (web-grounded only) presents several risks for the Ministry of Foreign Affairs and Trade (MFAT), primarily concerning data confidentiality, compliance, and the reliability of AI-generated information.

A key risk involves the potential exposure of sensitive information. Although user prompts are protected, the underlying mechanism requires Copilot to query the public Bing search engine. While designed to strip sensitive details, there is a non-zero risk that unique keywords or terms (e.g., confidential project names like "Operation Falcon") could be included in these external web queries, logging their existence outside MFAT's controlled environment. Furthermore, users might inadvertently input restricted text or upload classified documents, potentially breaching handling policies. Even if the input data is protected, the AI's output (e.g., summaries of restricted documents) might lack appropriate sensitivity markings and could be subsequently mishandled or under classified by users.

Compliance risks arise from data sovereignty requirements. AI processing and Bing search queries may occur on global infrastructure outside of New Zealand, potentially violating mandates for restricted data to remain onshore. Enabling the tool might also conflict with existing MFAT policies governing the handling and processing of sensitive information, particularly concerning cloud services and external data interaction.

Further risks relate to decision-making based on AI outputs. Copilot relies solely on public web data, which can be inaccurate, biased, outdated, or deliberately misleading. The AI can hallucinate, presenting misinformation confidently. Its lack of access to internal MFAT context means its responses may be incomplete or inappropriate for specific diplomatic or policy situations. There is a considerable risk that staff may over-rely on the AI's plausible-sounding answers, leading to poor decisions, the propagation of misinformation, or actions based on flawed analysis if outputs are not rigorously verified.

In essence, while the tool offers productivity potential, its use introduces risks of sensitive data leakage via web queries, non-compliance with data sovereignty and handling policies, and flawed decision-making due to reliance on potentially inaccurate or context-deficient web-based AI responses.

Risk Area	**Description of Risk**
Data Confidentiality	Sensitive keywords or project names (e.g., "Operation Falcon") input by users could be unintentionally included in Bing web search queries processed outside MFAT's control and in any country.
Data Handling	AI-generated responses (e.g., summaries of restricted data) may lack appropriate sensitivity labels, leading to potential under classification or mishandling by users downstream.
Compliance	Processing of prompts/queries may occur on servers outside New Zealand, potentially violating data sovereignty requirements for restricted information.
Information Accuracy	AI may provide inaccurate, fabricated (hallucinated), outdated, or biased information based on unreliable public web content.
Decision-Making	Staff may make poor decisions or policy recommendations based on incorrect or misleading AI-generated information derived from the web.
Contextual Relevance	AI responses lack internal MFAT context, potentially providing misleading or incomplete answers for sensitive diplomatic or operational matters.
User Behaviour	Staff may exhibit overreliance on AI outputs, accepting them without necessary critical evaluation or verification due to the AI's confident presentation (automation bias).
Misinformation Amplification	The AI could inadvertently amplify misinformation or biased narratives present on the public internet by incorporating them into its responses.

Released under the Official Information Act 1982

1. Overview of Microsoft 365 Copilot Chat (Web-Grounded Only)

Microsoft 365 Copilot Chat is an **AI-powered chat assistant for enterprise users** that provides conversational answers and content generation using **GPT-4**, with grounding from **public web data** In the **web-grounded only** mode, Copilot Chat draws information **exclusively from the Bing web index** (current internet content) and **does not have direct access to internal Microsoft Graph data** such as emails, documents, or SharePoint files). This is in contrast to the full Microsoft 365 Copilot (the premium licensed product), which can integrate with a user's **meetings, emails, chats, and documents** for context.

In practical terms, Copilot Chat (web-only) functions similarly to an enterprise version of Bing Chat. Users can ask questions in natural language, request summaries or explanations, and the system will retrieve relevant public web results to craft its response. It offers a chat interface accessible via the web (e.g. in Microsoft Edge or the Microsoft 365 app) and is designed for workplace use. Notably, the interface includes productivity features like the ability to **upload files for analysis** and even generate images, within certain limits (Overview of Microsoft 365 Copilot Chat | Microsoft Learn). A **green shield icon** is displayed in the chat UI when using a work account, indicating the session is in a protected enterprise mode (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). This specialized UI and the underlying protections distinguish Copilot Chat from the consumer Bing Chat, making it suitable for organizations such as the Ministry of Foreign Affairs and Trade (MFAT) to consider for internal use.

Because this deployment is *web-grounded only*, it means MFAT would be enabling the base Copilot Chat capabilities **without hooking into any internal MFAT data sources or Microsoft Graph content**. Users cannot retrieve or query organizational files or emails unless they manually provide them to the chat (for example, by copying text or using the file-upload feature) (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn) (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). The absence of Graph integration simplifies the data landscape (all AI responses are based on public information plus any user-supplied content), but it also means **Copilot's answers will not inherently include internal knowledge**. This context sets the stage for evaluating risks and benefits in a government setting.

2. Enterprise Data Protection

Microsoft 365 Copilot Chat includes **Enterprise Data Protection (EDP)** for all user prompts and AI responses when used with an organizational (Entra ID/Azure AD) account (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn). *Enterprise Data Protection* refers to Microsoft's contractual and technical commitments under the standard **Data Protection Addendum (DPA)** and Microsoft's Product Terms for cloud services, which govern customer data handling (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn). In essence, with EDP enabled, any data that MFAT employees input into Copilot Chat or receive from it is treated as **"customer data" owned by the organization**, with Microsoft acting as a data processor bound to strict privacy and security obligations (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn). This has several important implications:

- **No Unauthorized Use of Data:** Microsoft assures that prompts and responses handled under EDP are **not used to train the underlying AI models or improve services outside your tenancy** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). The content of your queries or uploaded documents remains confidential to your organization's session and isn't harvested to refine GPT-4's public knowledge. This is a critical protection for privacy—unlike consumer AI services where user inputs might be retained to improve the AI, Copilot Chat with EDP keeps data isolated.
- **Secure Handling and Storage:** Data exchanged with Copilot Chat is protected by enterprise-grade security measures (encryption in transit and at rest, per Microsoft's compliance standards). In fact, the presence of the green shield in the UI (labelled "Protected") confirms that **Enterprise Data Protection is active**, giving users a visual assurance that their session is covered by these protections (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). For example, if a user uploads a file to Copilot Chat, that file is stored in the user's OneDrive for Business (within MFAT's tenant) rather than on some unmanaged server (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). This means the data remains under MFAT's control and is subject to the same security and compliance policies as other Office 365 content.
- **Auditability and Compliance:** With EDP, **prompts and responses are logged and can be retained** in accordance with organizational retention policies (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). These chat records become auditable, allowing MFAT's compliance or security team to perform eDiscovery or forensic analysis if needed. In other words, Copilot Chat conversations can be treated similarly to emails or messages for compliance purposes. This logging provides accountability and an ability to investigate if a user ever inputs inappropriate data or if the AI returns problematic content. (The specific retention or audit capabilities may depend on licensing and configuration, but Microsoft has aligned Copilot Chat's data retention with the same framework used for Microsoft 365 Copilot (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn).)

Released under the Official Information Act 1982

It's important to note the scope of Enterprise Data Protection **covers the user's prompts (questions/inputs) and the AI's responses** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). **However, any web search queries generated by Copilot Chat to fetch information from Bing are handled a bit differently**. When Copilot needs to look up information, it formulates a search query derived from the user's prompt and sends that to the Bing search service (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn) (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). Those Bing search interactions are considered an "optional connected experience," meaning they fall under the Microsoft Services Agreement and Privacy Statement (with Microsoft acting as an independent data controller for the search data) (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). Microsoft's enterprise DPA does **not** directly apply to the Bing service calls. The good news is Microsoft has engineered the system to **exclude sensitive details from these Bing queries** – for example, it will not include the full prompt or any complete document text in the search; it only sends relevant keywords and omits user identifiers (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). This design helps ensure that even the portion of the workflow that isn't under the DPA has a minimal privacy exposure. Nonetheless, the fact that Bing searches are outside the enterprise contractual umbrella is something to be aware of in terms of data protection (we will discuss the risk implications of this in Section 4).

In summary, **Enterprise Data Protection greatly mitigates many data privacy concerns**. It effectively extends the same protections MFAT expects from Office 365 (Office documents, Exchange email, Teams chats, etc.) to this new AI chat service (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn). This means that *if* Copilot Chat is enabled for the ministry, all usage by our staff would occur in a controlled, compliant environment rather than the open internet. The data stays within our tenant's sphere of control, we have transparency via logs, and Microsoft contractually commits to handling it with care. This EDP framework is a strong argument in favour of the service's security – it's **specifically built to let organizations leverage AI with their data protected** (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn).

4. Risk of Inputting Restricted Data

A key question for MFAT is whether Microsoft 365 Copilot Chat (web-grounded version) is **secure enough to handle "Restricted" level data**, this requires evaluating confidentiality, compliance, and potential leakage paths.

Protection Strengths: In principle, Copilot Chat with EDP provides a **secure environment comparable to other M365 services**. Since MFAT already entrusts Restricted data to Office 365 (e.g. storing such documents in OneDrive/SharePoint, or sending via Exchange email), then Copilot Chat operates under the same umbrella of security controls. The data a user inputs is encrypted, stays within Microsoft's cloud, and is not visible to other customers. Microsoft's DPA coverage means the content is handled as confidential data (Microsoft term) processed on MFAT's behalf (Frequently asked questions about Microsoft 365 Copilot Chat | Microsoft Learn). Additionally, any prompts or files the user provides are not retained by the AI beyond the session and are **not used to improve the model** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn), eliminating the worry that our sensitive information could seep into a public AI model. For example, if an employee at MFAT were to input a confidential briefing into Copilot Chat, Microsoft contractually commits that this briefing will not be mined to train GPT-4 or be viewed by Microsoft's researchers – it remains our private data (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). This is a crucial distinction from consumer AI services. (Notably, there have been real incidents underlining the importance of this protection: **Samsung employees accidentally leaked trade secrets via ChatGPT** – a consumer service – which retained their prompts for training (Samsung workers made a major error by using ChatGPT | TechRadar). Copilot's EDP would prevent such an outcome because those prompts would stay within MFAT's control, not become OpenAI's data.)

Despite these protections, there are **residual risks and important caveats** when it comes to RESTRICTED data:

- **Exposure via Web Searches:** As mentioned earlier, Copilot Chat relies on Bing web searches to provide up-to-date information. When a user's prompt contains or references Restricted data, the AI might generate a search query to Bing that inadvertently includes some portion of that prompt. Microsoft has implemented safeguards to strip out full sentences or obviously sensitive sequences from these queries (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). Only a few keywords deemed necessary are sent out. For example, if a user asked *"Summarize the implications of Project Falcon (a confidential MFAT initiative) on regional trade"*, the Copilot might extract generic terms like "regional trade implications" for Bing, rather than the specific project name. However, this mechanism isn't foolproof. There is a **risk that unique or sensitive terms could be leaked through search queries** if the AI deems them relevant to find information. Those queries go to the public Bing service, which is **outside of our enterprise control and not covered by the DPA** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). In a worst-case scenario, if a

Released under the Official Information Act 1982

confidential codeword or detail did get included in a Bing query, it could be logged in Microsoft's consumer systems. (Microsoft assures that such queries are not shared with advertisers and are not used to personalize anything (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn), but they *are* stored as any search might be, under consumer privacy rules.) For Restricted data, even a small metadata leak can be concerning. **Example:** If MFAT had a confidential project nicknamed "Operation Kiwi", and an employee unknowingly triggers Copilot to search that term, Bing might not find anything (since it's internal) – but the occurrence of that query exists outside our secure boundary. This risk is low-probability due to the filtering (the AI would likely exclude words it can't contextualize publicly), but it's not zero. **Bottom line**: Users must be cautious about entering unique sensitive terms or full texts; even though the system tries to protect them, they are safest kept offline or summarized abstractly when using the AI.

- **Compliance and Policy Alignment:** We need to consider MFAT's own security policies regarding Restricted data. Some agencies mandate that certain levels of data **not be processed by any cloud service** that isn't explicitly accredited for that level. Microsoft 365 may be authorized for Protected/Restricted data in general (assuming MFAT already uses O365 for sensitive information), but the introduction of an AI service that sends snippets to an external search could require a policy review. If our policy says "Restricted data must not be exposed to the internet," one could argue that **using Copilot Chat on such data is not fully compliant** due to the Bing component. However, if our stance is that Microsoft's cloud (including Bing as a sub-service) is within acceptable risk parameters (given contracts and technical controls), then it might be allowed. It's a nuanced point. Many organizations err on the side of caution: for instance, the University of Colorado explicitly advises that **"Highly confidential data should NOT be entered into Copilot Chat."** (Best practices for using Copilot for the web securely | University of Colorado). This aligns with general cybersecurity best practice: even with enterprise protections, **treat generative AI outputs as potentially public** and avoid inputting your crown jewels. For MFAT, "Restricted" likely qualifies as data that should not be casually input. So while the system is secure, **the safest approach is to *not* feed Restricted (or above) information into Copilot Chat unless absolutely necessary and approved** (Best practices for using Copilot for the web securely | University of Colorado).
- **Data Residency and Sovereignty:** Another consideration for government data is where it is processed. Copilot Chat will route the AI processing to the nearest available data centre, but **it can use other regions if needed to meet capacity** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). Unlike some Office 365 services that guarantee local data residency, the nature of the AI service means your prompt might be handled by servers in another country (e.g., outside New Zealand) if the load is high. For EU customers, Microsoft keeps AI processing in the EU by special arrangements (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn), but for the rest of the world, data could travel to the EU or US. If MFAT has strict requirements that certain sensitive data must remain within a certain jurisdiction, this dynamic could pose a compliance issue. Web search queries, in particular, are **not compliant with the EU Data Boundary** (they go to global Bing systems) (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). For us, this means Restricted data could, in processing, transiently live on infrastructure abroad. All data is encrypted and handled under Microsoft's security, but government stakeholders might need to be comfortable with that trade-off.

Assessment: Microsoft 365 Copilot Chat (web-grounded) is **significantly more secure for sensitive data than consumer AI services**, thanks to EDP and architecture choices. It is likely **sufficient for *some* level of sensitive data** especially if MFAT's cloud risk appetite includes Office 365 usage for such data. However, for Restricted data (e.g. anything that could seriously damage national interests if leaked), the **residual risk may be unacceptable**. The primary worry is the uncontrolled Bing search aspect and the general principle of not exposing critical secrets to an AI. We should assume that anything entered s6(a)

In line with best practices, we recommend **users not to input restricted content into Copilot Chat** unless it's been sanitized. Instead, they can use the AI for more generic tasks or public-source information. This approach lets us reap the benefits (productivity and insights) while containing the risk to an acceptable level.

5. File Attachment Risks

Copilot Chat's interface includes an **"attach" or "+ Add content" button**, allowing users to upload documents directly into the chat conversation for analysis. This is a powerful feature – for example, an employee can feed a policy document or a report to the AI and ask for a summary or specific insights. However, it introduces specific security risks, especially in the context of Restricted data:

- **Document Ingestion Mechanics:** When a user uploads a file in Copilot Chat, **the file is stored in the user's OneDrive for Business under our MFAT tenant** as part of the enterprise data protection measures (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). The AI service then accesses the file from OneDrive to read its content for that session. Importantly, Copilot Chat will **not automatically send the entire file to Bing or outside**; as noted earlier, it may

Released under the Official Information Act 1982

extract key terms if it needs to do a web search, but **the full text stays within the protected environment** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). For instance, if a user attaches a confidential PDF, Copilot will process its text internally (within the Microsoft 365 cloud and the Azure OpenAI model) and might pull out a few topic words to search the web for related info (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). The **service explicitly does not include the whole file in any outbound query** (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). After the user's chat session ends, the document's content is no longer retained in the AI's working memory (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). These design choices mitigate some risk – the file isn't sprayed all over the internet – but we must scrutinize what could go wrong.

- **No Automatic Classification Checks:** Because Copilot Chat treats file uploads as user-provided input, s6(a)

  This gap means a **user could inadvertently violate data handling rules** by uploading a document that shouldn't leave a certain system. s6(a)

- **Potential Data Leakage via Responses:** Once the AI has the document, it will incorporate that content into its answer to the user's query. There is a risk that the **AI's response might contain extensive excerpts or details from the document**. If the user then copies that response and shares it (or if the chat transcript is visible to others over the shoulder), sensitive info could leak. For example, if asked to summarize a Restricted briefing, the summary itself is RESTRICTED. Copilot Chat will treat the conversation as protected (EDP applies to the response as well), but outside of the tool, the user must still handle that output carefully. This is more a user handling risk than a platform security flaw, but it's part of the overall risk picture: **the AI might make it easier to redistribute the content of a sensitive file, intentionally or not**.

Guidance: The presence of the file upload feature means we must put **clear guidelines around its use**. The risk of a user uploading a Restricted document without clearance is real. If we enable Copilot Chat, we should strongly consider **instructing staff *not* to upload documents containing Restricted or above information** unless there is a vetted business need and they have permission. If possible, we could look into disabling the file upload capability administratively for most users – although Microsoft's current admin controls primarily mention turning agents on/off, it's unclear if file upload can be turned off independently. If it cannot be disabled, then training and trust become the mitigations.

In conclusion, **the file attachment capability is a double-edged sword**: it boosts productivity (which is a major benefit of Copilot) but can lead to inadvertent data exposure if used recklessly. From a risk perspective, this is one of the more *sensitive aspects of enabling Copilot Chat*. We should treat it with careful policy – likely forbidding its use with classified data and emphasizing that any uploaded content must be something the user is allowed to store in OneDrive and allowed to have the AI process. If these conditions are respected, the risk remains manageable. If not, this feature could become a channel for breaching internal data handling rules.

6. Decision-Making Risks of Web-Grounded AI Responses

When Copilot Chat is used in a government context, a critical risk category is the **accuracy, reliability, and trustworthiness of the AI's outputs**, especially since they are based solely on web information. Relying on AI-generated answers for decision-making can be risky for several reasons:

- **Potential for Misinformation (Hallucinations):** Generative AI models like GPT-4 are known to occasionally produce incorrect or fabricated information that sounds plausible – a phenomenon often referred to as *hallucination*. Even though Copilot Chat grounds its answers with web search results, it may still **generate and confidently present erroneous or false content** if the web data is incomplete or if the model fills gaps with its own assumptions (Latest NIST Guidance Identifies Generative AI Risks and Corresponding Mitigation Strategies | Davis Wright Tremaine). NIST's guidance on AI risk management specifically flags this risk (terming it "confabulation") and notes that the danger is amplified when **users believe and act on the incorrect output** (Latest NIST Guidance Identifies Generative AI Risks and Corresponding Mitigation Strategies | Davis Wright Tremaine). In a ministry setting, such an error could lead to an embarrassing situation at best (e.g., quoting a wrong statistic in a briefing) or a harmful decision at worst (e.g., misinforming policy because the AI provided a misleading interpretation of events).
- **Web Content Quality and Bias:** Copilot Chat's knowledge is only as good as the public information available. The internet can contain outdated data, rumours, biased opinions, or deliberate disinformation. The AI might pick up a news article or blog post that is **not authoritative or is one-sided** and synthesize an answer from it. Unlike a human analyst, the AI may not fully weigh the

Released under the Official Information Act 1982

credibility of sources – it attempts to provide a balanced answer, if possible, but it has no innate understanding of truth. For example, asking a sensitive geopolitical question might yield an answer that aggregates what various media outlets say. Those outlets might have speculative or biased takes. There is a **risk that the AI could present a controversial or unofficial perspective as if it were factual** unless the user specifically checks the sources. Government users must therefore be cautious: any factual statement from Copilot Chat should ideally be **verified against official sources or trusted data** before being used in official work.

- **Lack of Internal Context:** Because we are using the web-only mode, the AI lacks any internal MFAT context or data to cross-reference. It won't know MFAT's policies, internal assessments, or classified intelligence that might drastically alter the interpretation of a public fact. This means its answers might be **missing key context**. For instance, the AI might summarize a foreign country's public stance on a treaty from news sources, s6(a)
  If an employee relied only on the AI's summary, they could draw wrong conclusions. In short, the AI provides a generic perspective and cannot substitute for internal expertise or context that our staff hold.
- **Overreliance and Cognitive Bias:** One subtle risk highlighted in AI ethics research is that users may **over-trust AI outputs due to their fluent and confident presentation** (Microsoft Copilot Security Concerns Explained). The AI will often answer in a very assured tone ("Based on the information, the answer is X…"), which can lend it undue credibility. Studies have shown people might trust an AI's answer **even when they have evidence it might be wrong** (Microsoft Copilot Security Concerns Explained). In a fast-paced work environment, an employee might be tempted to accept Copilot's response at face value, especially if it sounds convincing and saves time. This is dangerous in a government context because policies or external communications based on incorrect information could have serious repercussions. We must guard against this *automation bias*. The **OWASP foundation lists overreliance on AI as a key human-factor risk** in using generative AI (Microsoft Copilot Security Concerns Explained). In practice, this means MFAT staff should be trained to use Copilot Chat as a helpful assistant, **not an oracle**. The output should be double-checked, and users should maintain healthy scepticism.
- **Misinformation Amplification:** Consider scenarios where misinformation exists on the web (which is common in international affairs – think of doctored statistics or false claims circulating in social media). Copilot Chat might inadvertently **amplify such misinformation** by including it in an answer (especially if multiple sources on the web repeat the same false info). For example, if an inaccurate figure about trade volumes is widely reported, the AI could present that figure unless it finds a correction. In a government context, propagating misinformation can erode trust and lead to poor decisions ([PDF] Artificial Intelligence Risk Management Framework: Generative ...). The risk is not that Copilot intends to mislead, but that it **lacks the judgment to exclude unreliable information** consistently.
- **No Guaranteed Source Citation for Facts:** While Copilot Chat does attempt to show which search queries it ran and sometimes provides citation links (it can display the queries and sites it used to compose the answer (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn)), these are not as transparent as one might want. The enterprise chat experience shows a *linked citation section* with the **exact web search queries** used (Microsoft 365 Copilot Chat Privacy and Protections | Microsoft Learn). This helps the user see what was searched (and by extension, the user can infer which sources might have been consulted), but it may not directly list, in the answer, the specific websites or a bibliography. This is slightly different from consumer Bing Chat, which often footnotes specific sentences with source URLs. The onus is still on the **user to click those queries, review the actual web sources, and confirm the information**. If users skip this step, they might take the AI's synthesized answer as vetted truth when it's not. Given time pressures, there is a real risk of that happening.
- **Misleading Conclusions and Analysis:** Beyond factual accuracy, there's the risk of the AI drawing **incorrect conclusions or analogies**. If asked for an analysis or recommendation, Copilot will do its best to "think" through the prompt with the data it has. But it does not reason like a human with domain expertise; it patterns its answer on learned data. It might make an analytical leap that a trained policy officer would avoid. For example, "What is the likely outcome of X conflict given Y?" – the AI might provide a speculative outcome that sounds logical but is not grounded in any official assessment. If an inexperienced user took that and ran with it, it could misdirect our strategic thinking. In essence, **the AI is not a subject-matter expert, but it can sound like one**, and that is inherently risky if we let it drive decisions.

Given these factors, the **risk of misinformation and misleading outputs is one of the highest concerns** in using Copilot Chat in a mission-critical environment. The consequences of acting on a false piece of information can be severe for a foreign affairs agency. As a mitigation, it's imperative to use Copilot Chat as a *support tool* rather than an authoritative source. Users should use it to gather quick insights, draft initial versions of content, or discover publicly available information – but then they must apply **critical thinking and verification**. This aligns with general cybersecurity best practices for AI: incorporate a human review step for important outputs and maintain oversight (Latest NIST Guidance Identifies Generative AI Risks and

Released under the Official Information Act 1982

Corresponding Mitigation Strategies | Davis Wright Tremaine). If we enable Copilot Chat, we will need to **train staff in AI literacy** – understanding that the AI can err and that they are responsible for the final judgement.

In summary, while Copilot Chat can improve efficiency (finding information or summarizing data in seconds), MFAT must weigh that against the **risk of error**. Appropriate use policies (e.g., "verify any AI-provided fact from a second source before including in reports") and perhaps limiting usage to less critical queries can manage this risk. The tool should ideally augment human analysis, not replace it. **Misinformation risk is real but manageable with vigilant user behaviour**; ignoring it could lead to flawed decisions or dissemination of false information.

Released under the Official Information Act 1982