Headquarters
New Zealand Defence Force
Defence House
Private Bag 39997
Wellington Mail Centre
Lower Hutt 5045
New Zealand

OIA-2025-5581

17th December 2025

[redacted]@vuw.ac.nz

Dear [redacted]

I refer to your email of 16 October 2025 to the Ministry of Defence, under the Official Information Act 1982 (OIA), requesting information on approved artificial intelligence (AI) tools. Your request was partially transferred to the New Zealand Defence Force (NZDF) for consideration and is addressed below.

1. *A list of all AI tools that are currently approved for use by staff at your agency.*

The only AI software tool approved for use by all NZDF personnel and staff is Copilot Chat, a Microsoft 365 application.

The following AI enabled tools are used in a restricted capacity for research purposes in data processing, sensor processing, and in support of modelling activities:

- ChatGPT
- Dalle-2
- Github Copilot
- Azure Machine learning services
- Azure OpenAI services
- Microsoft Copilot
- Microsoft teams - transcription and summarisation
- AiZynthFinder
- Meta Llama2

Information regarding this research is withheld in accordance with section 6(a) of the OIA.

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

Relevant documentation for Copilot Chat is provided at enclosures one and two. Where indicated, information for internal NZDF purposes is withheld in accordance with section 9(2)(k) of the OIA to avoid the malicious or inappropriate use of that information, such as phishing, scams or unsolicited advertising.

With respect to the research use, only information that is unclassified and non-sensitive, that is publicly available or eligible for public release, is to be used with generative AI tools. Documentation is withheld in accordance with section 6(a) of the OIA.

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.*

There are two paid licenses for ChatGPT.

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

A copy of the high-level risk assessment for Copilot Chat is provided at enclosure three. Where indicated, the names of those who have provided advice are withheld in order to maintain the effective conduct of public affairs in accordance with section 9(2)(g)(i) of the OIA. As noted above, research risk assessments are withheld in accordance with section 6(a) of the OIA.

You have the right, under section 28(3) of the OIA, to ask an Ombudsman to review this response to your request. Information about how to make a complaint is available at www.ombudsman.parliament.nz or freephone 0800 802 602.

Please note that responses to official information requests are proactively released where possible. This response to your request will be published shortly on the NZDF website, with your personal information removed.

Yours sincerely

GA Motley
Brigadier
Chief of Staff HQNZDF

Enclosures:
1. Copilot Chat – Getting started guide
2. Copilot Chat FAQs
3. Copilot high-level risk assessment

TE OPE KĀTUA O AOTEAROA
DEFENCE FORCE

Te Pou Hangarau Matihiko
Information Command

USING MICROSOFT 365 COPILOT CHAT IN THE NZDF

Get Started with Copilot Chat in the Defence Information Environment

What is Copilot Chat?

Microsoft 365 Copilot Chat is an Artificial Intelligence (AI) tool, designed to help personnel work smarter with secure, web-sourced responses

It can help you:

- Answer questions
- Generate content and ideas
- Summarise and rewrite text
- Search for information
- Interact with files you have selected (Word, Excel, PowerPoint, PDFs).

Copilot Chat has undergone a Security Risk Assessment (SRA) and can be used at a RESTRICTED and below level in the Modern Desktop, inside the Defence Information Environment (DIE).

Why Copilot Chat?

With so many AI tools available, it's important to understand why **Copilot Chat** is the preferred option for the NZDF:

✔ Built for Work

Copilot Chat is designed specifically for workplace productivity. It integrates with Microsoft 365, making it easy to use alongside Teams, Outlook, Word, Excel, and more.

🔐 Enterprise-Grade Security

Unlike many public AI tools, Copilot Chat is governed by Microsoft's enterprise security and compliance standards. Your data stays within your organisation's boundaries and is **not used to train the model**.

🔍 Grounded in Trusted Sources

Copilot Chat provides web-grounded answers to ensure responses are secure, relevant, and reliable.

🆓 No Extra Cost to Get Started

Copilot Chat (the free version of Copilot) is available to all users with a Microsoft Entra ID (i.e., those on the Modern Desktop). It's a safe, supported way to explore generative AI without needing to access, use or sign up for third-party services.

Responsible use of Copilot Chat

Responsible use of Copilot Chat

Copilot Chat can only be used in the RESTRICTED and below environment.

NZDF personnel must use Copilot Chat safely and responsibly. This means you must:

- Continue to follow NZDF policies around information security and acceptable use of digital systems and devices;
- Ensure that you do not upload or enter any data or information above RESTRICTED level into Copilot Chat. If you are in any doubt of the classification level, do not upload the document or information;
- Review any AI-generated (GenAI) content to ensure its accuracy. Results produced by GenAI are not guaranteed to be correct, and must be checked by a human before they are shared or included in documents; and,
- Declare the use of Copilot Chat if you have used it in the creation of NZDF documentation and briefings.
- Please refer to the following documents for more information:
  - DFI 60.20 Acceptable Use of Digital Systems and Devices
  - DFO 51(1) Protective Security
  - CISO Directive 01/2023 Restrictions and Allowances on NZDF Use of Generative Artificial Intelligence

Getting started with Copilot Chat

- Copilot vs Copilot Chat
- Where can you access Copilot Chat?
- How can I use Copilot Chat?
- Getting started
- Things to try in Copilot Chat
- The importance of prompts
- Further resources

Copilot Chat vs Copilot

Copilot Chat is the free version of Copilot, meaning that functionality is limited compared to the 'paid-for' Copilot subscription service, which the NZDF does not have.

Copilot Chat features

Feature	Copilot Chat (Free)
Cost	Free with Microsoft 365 E5 license
Data Access	Web data only
Integration with Office Apps	✗ No integration
Enterprise Data Protection	✓ Yes
AI Agents	✗ No
Image Generation	✓ Yes
File Uploads	✓ Yes (4 images per 24 hours)
Copilot Pages (Loop)	✓ Yes
Content Referencing (e.g., people, emails)	✗ No
Copilot Analytics	✗ No

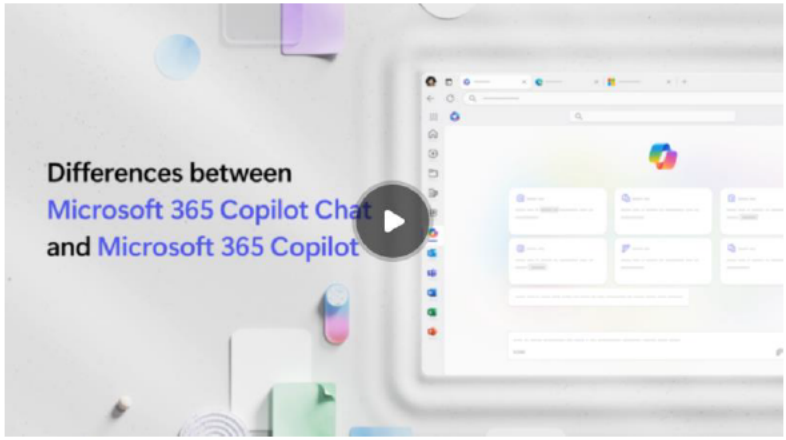

How Copilot Chat works with and without a Microsoft 365 Copilot license

Where can you access Copilot Chat?

We have created a visual guide on how to access Copilot Chat depending on the device, which you can find on the Copilot Chat DDMS page or by clicking here.

Where does Copilot Chat work?	
NZDF desktop computer signed into Modern Desktop	Yes
NZDF managed laptop	Yes
NZDF managed iPad or iPhone	Yes
NZDF managed Android phone	Microsoft Edge (Browser Only)
DIXS	No
DDMS or Sharepoint	No

How can I use Copilot Chat?

	MS Outlook	Word/Excel/PowerPoint/One Note	MS Teams	Microsoft Edge internet browser	File upload	Image generation
NZDF Modern Desktop	× Copilot Chat is not available in Outlook on the Modern Desktop.	**Copilot Chat is not integrated into M365** × No "Copilot" ribbon or task-pane for real-time prompts. × No AI-driven formula suggestions, slide generation, or context-aware drafting directly in the canvas.	✓ The Copilot Chat icon in the compose box lets you open the chat pane. × It does not auto-summarise meeting transcripts, or read your private channels or chat history.	It brings up a browser-side chat ✓ You can research, summarise information, and brainstorm with web content. × It has no access into your internal intranet, DDMS, SharePoint sites, or Teams pages.	✓ You can manually upload files—Word, Excel, PowerPoint, or PDF—for ad-hoc Q&A and summarisation.	✓ Copilot Chat can generate a limited number of images a day per user, based on the prompt you provide.
NZDF managed laptop	× Copilot Chat is not available in Outlook on an NZDF managed laptop.					
NZDF managed iPad or iPhone **(At this point NZDF managed Android devices can only access Copilot Chat through Microsoft Edge internet browser)**	The Copilot Chat icon in the ribbon launches a side-pane chat. ✓ You can ask it to summarise, rewrite, or draft replies based on the currently open email. × It cannot scan your entire mailbox, calendar, or organisational directory.					

Getting started

We recommend you begin by Clicking here to watch the training video delivered by Microsoft

With Copilot Chat you can:

- **Summarise and rewrite**:
  Improve clarity and engagement;
- **Answer questions**:
  Get accurate, web-grounded answers;
- **Create content**:
  Draft documents, presentations, and more;
- **Search for insights**:
  Quickly find relevant data;
- **Upload files**:
  Ask questions about Word, Excel, PowerPoint, and PDF files.

You can also explore these Microsoft modules (opens in ITD)

Microsoft Module: Get started with Microsoft 365 Copilot Chat (24 Min)

Microsoft Module: Explore Microsoft 365 Copilot Chat (1hr 36 Min)

Things to try in Copilot Chat

Copilot Chat can be used in a variety of ways:

- Real-time voice chat
- Dictation and Read Aloud
- Image upload and generation (limits apply)
- Previous chat history
- Copilot Pages

Learn how to use Copilot Chat with a step-by-step program available on Microsoft.com

Click here and sign in with your work account to get started

Want to learn more?

Copilot Skilling

Copilot Learning Hub

Copilot Scenario Library

The importance of prompts

Good prompts are essential for making the most of Copilot Chat.

Try These Prompts:

- "Summarise the latest project report [attach project report]"
- "Rewrite this email to make it more engaging [paste email text]"
- "Generate ideas for our upcoming campaign about [insert topic]"
- "What are the key points from the last team meeting [attach transcript]?"
- "Find recent articles about [insert topic]"
- "Generate an image for my PowerPoint presentation on [insert topic]"
- "Help me write a prompt for…"

Microsoft also has resources available to help you learn to shape effective prompts:

Microsoft Module: Write effective prompts to achieve optimal results (24 mins)

Click here to watch the training video delivered by Microsoft which goes into more detail on how to use prompts effectively.

Further Resources

- Copilot Chat DDMS page
- Click here to watch the training video delivered by Microsoft

Related policy documents:

- DFI 60.20 Acceptable Use of Digital Systems and Devices
- DFO 51(1) Protective Security
- CISO Directive 01/2023 Restrictions and Allowances on NZDF Use of Generative Artificial Intelligence

Questions?

Check out the resources including FAQs on the Copilot Chat DDMS page!

If you have any other questions about the use of Copilot Chat in the NZDF, email s. 9(2)(k)

TE OPE KĀTUA O AOTEAROA DEFENCE FORCE | Te Pou Hangarau Matihiko Information Command

COPILOT CHAT FAQS

Last updated on 07 October 2025.

INFOCOM will continue to update this document. If you have a question you would like to add please email it to s. 9(2)(k)

You can find all of the Copilot Chat resources we've developed on the Copilot Chat DDMS page.

Getting Started with Copilot Chat

What is Copilot Chat?	Copilot Chat is a Generative Artificial Intelligence (GenAI) tool, designed to help personnel work smarter with secure, web-sourced responses. You can use it to: • **Summarise and rewrite documents or information**: Improve clarity and engagement • **Answer questions**: Get web-sourced answers • **Create content**: Draft documents, presentations, and more • **Search for information**: Quickly find relevant data • **Upload files**: Ask questions and interact with files you have selected (Word, Excel, PowerPoint, PDFs).
What is the difference between Copilot and Copilot Chat?	Copilot Chat is the free version of Copilot, meaning that functionality is limited compared to the 'paid-for' Copilot subscription service, which the NZDF does not have. Copilot Chat offers limited integration into M365 apps and is focused on general assistance rather than work specific tasks, meaning that it cannot access DDMS, SharePoint or other NZDF filing systems or intranet sites.
Why Copilot Chat for the NZDF?	Copilot Chat is designed specifically for workplace productivity. It integrates with Microsoft 365, making it easy to use alongside Teams, Outlook, Word, Excel, and more. Copilot Chat, the free version of Copilot (the paid subscription version), is now available to all users via Modern Desktop. It's a safe, supported way to explore generative AI without needing to access, use or sign up for third-party services.

What do I need to know before I start using Copilot Chat?	**NZDF personnel must use Copilot Chat safely and responsibly.** This means you must:
	• continue to follow NZDF policies around information security and acceptable use of digital systems and devices.
	• ensure that you do not upload or enter any data or information above RESTRICTED level into Copilot Chat. If you are in any doubt of the classification level, do not upload the document or information,
	• review any generated AI (GenAI) content to ensure accuracy. Results produced by GenAI are not guaranteed to be correct, and must be checked by a human before they are shared or included in documents; and,
	• declare the use of Copilot Chat if you have used it in the creation of NZDF documentation and briefings.
	• Please refer to the following documents for more information:
	  • DFI 60.20 Acceptable Use of Digital Systems and Devices
	  • DFO 51(1) Protective Security
	  • CISO Directive 01/2023 Restrictions and Allowances on NZDF Use of Generative Artificial Intelligence
	
	Further details including a user guide can be found on the Copilot Chat DDMS page.

Do I need to do anything to get Copilot Chat?	No, there is no action you need to take other than ensuring your managed devices - phones and tablets - are given the chance to update. Copilot Chat will automatically become available to all users on the Modern Desktop, as well as across all devices - phones, tablets, laptops - as updates are pushed out over the last week of September 2025.
Where do I start?	INFOCOM developed a range of resources, available on the Copilot Chat DDMS page, to help you get started. We suggest you start by: • Bookmarking the Copilot Chat DDMS page, so you can refer back to it as needed; • Click here to watch the training video delivered by Microsoft (works in Modern Desktop) • Referring to the Copilot Chat User Guide if you need extra support; • Checking the FAQs if you have any questions. If the resources above do not provide the information you need, you can email s. 9(2)(k). We will continue to add resources to the DDMS page as they are developed.
Does Copilot Chat only work in the Modern Desktop?	Yes. Copilot Chat does not work in our legacy desktop DIXS.
Can I use other GenAI tools (such as ChatGPT, Gemini etc.) for NZDF business?	All NZDF personnel, contractors, and external vendors must not use other GenAI tools for NZDF official business without prior approval from Chief Information Security Officer (CISO) and endorsement from Chief Data Officer (CDO) through the GenAI Request for Information process (AI RFI form and instructions).

Information and Data Security

Is Copilot Chat safe to use in the Defence Information Environment (DIE)?	Yes. Copilot Chat went through a full risk assessment process, controls were already in place or implemented as part of the process, with the risks fully understood and agreed upon. The speed with which we were able to roll out Copilot Chat was ONLY possible because Copilot Chat inherited the controls from M365. We would not have been able to do this as fast with any other GenAI implementation within the DIE. In INFOCOM, Defence Information Security (DIS) and Data Analytics and Information Management (DAIM) have worked closely to enable this. All prompts and responses entered in Copilot Chat are monitored and logged for security and compliance.
What is Enterprise-Grade Security?	Unlike many publicly-accessible GenAI tools, Copilot Chat is governed by Microsoft's enterprise security and compliance standards. This means data and information stays within an organisation's boundaries and is not used to train public models.
How does Copilot Chat keep NZDF data and information safe and secure?	Copilot Chat is configured with enterprise level data protection controls to safeguard the information put in and received back, preventing it from being shared or used in undisclosed ways without the NZDF's permission. Copilot Chat strips any sensitive or identifying data and information from queries input by users, and ensures responses are secure. Results produced by GenAI are not guaranteed to be correct, and must be checked by a human before they are shared or included in documents.
Can I use Copilot Chat in Secret and above environments?	No. Copilot Chat can only be used in NZDF's RESTRICTED and below Defence Information Environment (DIE), with data and information classified up to RESTRICTED or SENSITIVE. Findings can be subsequently used/transferred up to Secret and above environments, but must adhere to the acceptable use guidelines and requirements above.
Can I upload documents or enter information protectively marked IN-CONFIDENCE, SENSITIVE or RESTRICTED, that also carries an endorsement marking? ***(Endorsement markings set requirements for special handling of the information and/or it's dissemination)***	Yes, provided the endorsement marking can be used within NZDF's RESTRICTED and below Defence Information Environment. Examples of what you can enter or upload include, but are not limited to, information marked STAFF, MEDICAL, COMMERCIAL and BUDGET, for information classified IN-CONFIDENCE or SENSITIVE. Most of the endorsement markings in use within New Zealand government agencies are listed and explained on the PSR website (link opens in ITD). Defence Information Security can provide additional guidance on security classifications and endorsement markings if required. Note: If unsure, it is your responsibility to verify and apply the most appropriate security classification, endorsement markings, and any additional handling caveats.

How does Copilot Chat work?

Does it use data from inside or outside of the NZDF to generate responses to questions?	Both. Copilot Chat generates answers to user prompts by pulling data from within and/or outside of NZDF's systems.

	1. Questions or prompts specific to the NZDF may be answered using Copilot Chat's own learned data, or by using its limited (controlled) view of Microsoft 365 organisational data (e.g., 'Who is Jane Doe's manager?). 2. When a question or prompt causes Copilot Chat to seek data from outside of NZDF's systems (e.g., 'What is the weather like in Wellington today?'), it does so by "grounding" prompts (stripping them of sensitive and identifying information) and forms Bing search queries using only these keywords.
What does 'prompt grounding' mean?	Prompt grounding means that Copilot Chat strips any sensitive or identifying data and information from queries input by users, and ensures that responses are secure, accurate, and contextually appropriate. It means that any data leaving our internal environment does not include: • The users entire prompt (key words are used only); • Files or content uploaded from OneDrive; • Web pages or PDFs summarised by Copilot Chat; • User or tenant identifiers (organisational information such as names, org structure etc.)
Does Copilot Chat interact with DDMS, Sharepoint or other files or intranet sites such as ILP?	No. Copilot Chat does not reach into official NZDF data or document stores (e.g. DDMS or ILP), outside of whatever document or information a user has uploaded themselves, or a user's Entra ID (organisational information such as a user's team name or manager, for example).

Using Copilot Chat in the NZDF

Can I use Copilot Chat to translate content into different languages, including te reo Māori?	We recommend using it as an initial guide only, and anything that is created or translated must be validated by a suitably qualified individual, especially if it is being used for official use.
As the NZDF is working to reduce energy use and implement sustainable practices, is the NZDF factoring in the high energy use of GenAI when introducing tools like Copilot Chat?	This is something that the NZDF are aware of when considering the use of Copilot Chat and future AI tools. More information will be shared on this in the future.

Having trouble accessing Copilot Chat?

If you are having trouble accessing Copilot Chat, it may be for one of the following reasons:

No access to Modern Desktop	If you do not have access to the Modern Desktop, then you will not have access to Copilot Chat.
Older devices running Windows 10 325	If you have an older device that is running Windows 10 325 then you will not be able to access Copilot Chat through that device. You will be able to access it through the Modern Desktop if you have it.
Trying to access it through DIXS	Copilot Chat is not available on our legacy desktop DIXS.

High-level Risk Assessment:
M365 Copilot Chat (Free Version)

Reviewer(s)	s. 9(2)(g)(i) **DAIM-D (D&A Team)** DIS Cyber DIS Advisory
Date of Review	**May 28, 2025**

A. Background

Microsoft Copilot (free version) is a generative AI assistance tool built on Microsoft's secure enterprise infrastructure. This assessment evaluates the risks associated with deployment and use of Copilot Chat by members of the NZDF and contractors for NZDF purposes. This assessment explicitly focuses on the free version of Copilot that predominantly uses publicly accessible internet content and Microsoft-hosted Large Language Models (LLMs), while providing limited and controlled interactions with organisational data through user-uploaded documents stored securely in OneDrive for Business.

B. Immediate Security Concerns and Clarifications

1. Model perimeter

i. **Note:** The model perimeter is cited as excluding all internal organisational data sources (e.g., Outlook, Teams, SharePoint, OneDrive, and local drives). The free version of Copilot explicitly does not integrate with Microsoft Graph[1] and ensures that **prompts and responses remain within the M365 tenant**, avoiding external visibility or storage (see point (3) below).

2. Aggregation and use of prompts

i. **Note:** According to the official Microsoft documentation on Copilot Chat's enterprise privacy and protections, prompts and model responses used within Copilot Chat are **neither visible nor accessible to Microsoft or third parties**.

- When users interact with Copilot Chat, their prompts and responses are logged and stored within the M365 boundary for internal auditing (see 4.iii below). Organisational (NZDF) M365 administrators (i.e., admins within the service boundary) can view anonymised and deanonymised usage metrics (e.g., total active users, average daily active users, and activity per application).
- Microsoft staff only access aggregate and anonymised service-level metrics for operational purposes and assure that they do not have access to organisational data as per their Copilot Chat Privacy and Protections and Enterprise Data Protection (EDP) policies.

[1] Microsoft Graph provides Application Programming Interfaces (APIs) to enable access to data and services within an organisation's Tenant (e.g., an API to query data from SharePoint).

*(Refer to: *"M365 Copilot Privacy and Protections"* Microsoft Documentation).*

ii. **Note:** Microsoft explicitly states that prompts and responses **are not stored or used to train foundation models**, including those maintained by Microsoft or OpenAI. Further, when enabled, **web search queries are anonymised and stripped of user and tenant identifiers** prior to external transmission (e.g., Bing queried) via a process referred to as **"(web/prompt) grounding"** *(Step 2, Figure A).*

- That is, when a user submits a prompt, both the query (prompt) and response remain within the service boundary. The response a user receives is formulated using publicly available data and the LLMs themselves. There is no reaching into NZDF data/documents, so classified queries would be answered in terms of what is publicly available or presumed by the model. Prompt grounding, which Microsoft implements by default, means any data leaving the service boundary does NOT include:
    a. The users entire prompt (key words are used only);
    b. Files uploaded from Onedrive;
    c. Web pages or PDFs summarised by Copilot Chat;
    d. User or tenant identifiers.

iii. Figure A: M365 Copilot Chat Architecture Diagram *("M365 Copilot Privacy and Protections").*

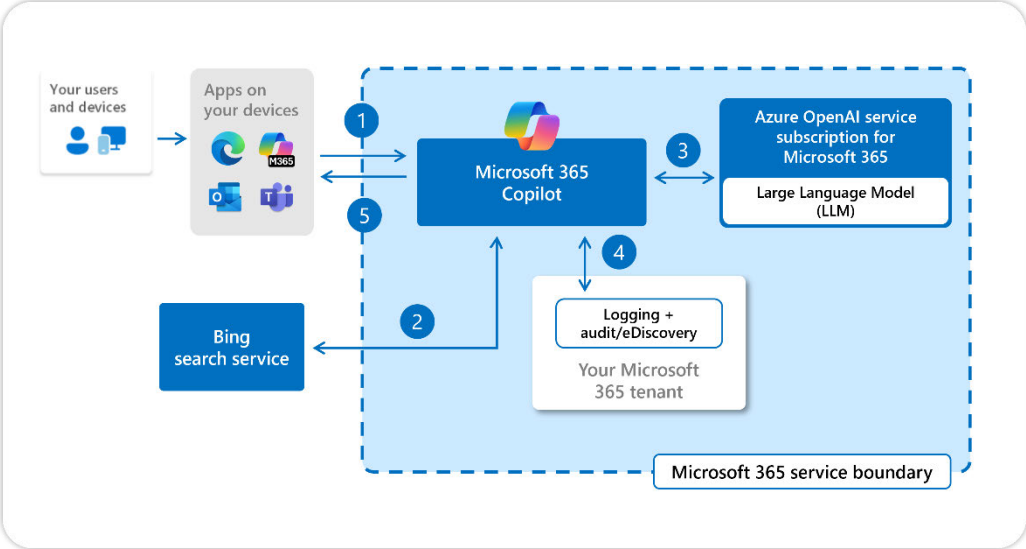

3. Data access within DIE boundaries

i. **Note:** The free version of Copilot Chat does not integrate with Microsoft Graph and thus **cannot autonomously access or interact** with any documents stored within the Defence Information Environment (DIE). That is, Copilot Chat responses consist of publicly available web content and LLM-generated responses, bar user-uploaded data/documents (see point (4) below).

4. User-uploaded data & documents

i. **Note:** Users can upload data/documents in the free version of M365 Copilot Chat, provided they have access to OneDrive for Business, **which NZDF possesses**. However, **uploaded files are securely stored within the user's OneDrive for Business account inside of a 'Microsoft Copilot Chat Files' folder**, which ensures that any uploaded content remains within the enterprise's M365 tenant and is governed by existing data protection policies. Therefore, the risk of data exposure or unauthorised access remains minimal, provided OneDrive governance and security controls are sufficiently managed.

ii. **Note:** Copilot Pages is accessible to users with a Microsoft Entra ID (formerly Azure active dir.) and either SharePoint or OneDrive storage and ***does not*** require a paid M365 Copilot license. Any content created by Copilot Pages is stored in a user-owned SharePoint Embedded container (one per user). Copilot pages reside in the ecosystem as *.loop* files. These can extend user permissions to others via file sharing settings. However, **conditional access can be enabled to fully block users from opening *.loop* files**.

**(Refer to: "Copilot pages for IT Admins – Sep 2024 update", M365 Copilot Blog).*

iii. **Note:** Out of scope for this assessment, but it is possible to explore E5 license options for auditable Copilot trails consisting of M365 tools for auditing, retention, and discovery (Purview → Exchange Online → SharePoint → OneDrive → SharePoint Embedded [Container], *Figure B*). However, the **free version of Copilot Chat still includes some of the foundational capabilities of the Copilot Control System**.

- *[Microsoft] Purview audit logs can be used to identify how, when, and where Copilot interactions occurred, which items were accessed, and the sensitivity labels of those items. Purview eDiscovery can search for keywords in Copilot prompts and responses to identify "inappropriate, malicious, or risky behaviour". Purview Communication Compliance can detect and alert "…", like personal data or confidential information. Purview Retention Policies can maintain a copy of deleted Copilot conversations. (Paraphrased from MS Documentation).*

**(Refer to: "Where Copilot usage data is stored and how you can audit it", Microsoft Documentation).*

iv. ***Figure B: M365 Copilot usage data and tools for auditing*** *("Where Copilot usage data is stored and how you can audit it").*

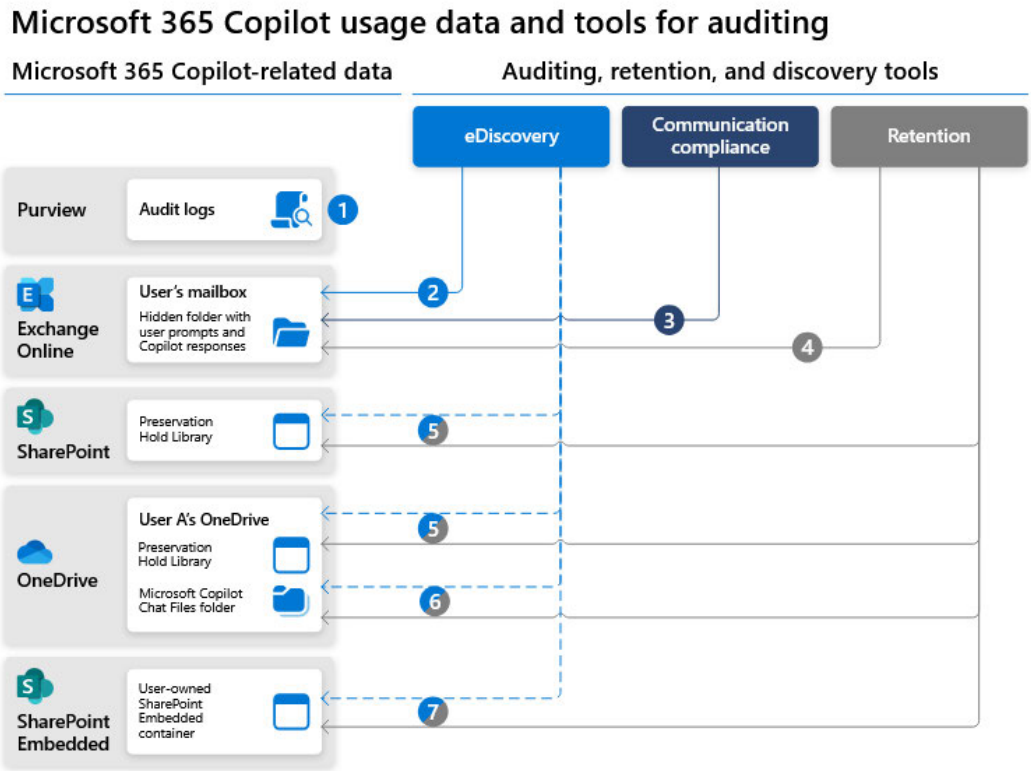

5. Data sovereignty and regulatory compliance

i. **Note:** Microsoft defines their role as a **data processor and not a controller** under the Data Protection Addendum (DPA), thus NZDF retains full ownership and control over NZDF data *("M365 Copilot Privacy and Protections", Microsoft Documentation).*

ii. **Note:** The free version of Copilot Chat is built on the same secure foundation as M365 and adheres to Microsoft's enterprise-grade security, compliance, and privacy standards, including **ISO/IEC 42001:2023 certification** (first international standard for AI management systems), **GDPR compliance** and alignment with Microsoft's Responsible AI principles, and **Data residency and encryption** within Microsoft's trusted cloud infrastructure *("ISO/IEC 42001:2023 Artificial intelligence management system", Microsoft Documentation).*

C. Aggregate (High-level) Risk Assessment

1. Risk Identification

i. **Sources:** Public web content, hosted LLMs.

ii. **Events:** Data breaches, unauthorised access, or inadvertent disclosure of sensitive information.

iii. **Vulnerabilities:** Reliance on external, publicly accessible data, potential misuse by end users, and unclear boundaries for anonymisation (grounding).

iv. **Timeframe:** Immediate/ongoing (continuous use). Continued monitoring required.

v. ***Table A: Copilot Chat Aggregate (High-level) Risk Assessment***

Probability (Likelihood)	***Impact (Severity)***	***Exposure***	***Velocity (Speed of Onset)***
Low → Mod	**High**	**Low**	**High**
Microsoft is a trusted partner with stringent security controls, but low residual risks remain (external data sources).	Potential for severe impacts **if data sovereignty, access, privacy, or governance** were compromised.	Minimal exposure to organisational data; moderate exposure through reliance on public content.	Speed of onset is immediate if the breach occurred as a result of real-time service interactions.

2. Risk Mitigations

i. **Preventative Measures:**

a. Data security/loss prevention;

b. Data encryption and residency (i.e., data remains within the boundaries of the NZDF tenant and trusted cloud infrastructure – *true by default **and** avoid autonomous tenant-specific data integrations [e.g., MS Graph] without due investigation*);

c. Compare against existing NZDF configured security controls for data loss prevention of OneDrive;

d. Anonymisation and constraints exercised when copilot accesses external applications or search engines;

e. Appropriate use and change management plans;

f. Policy, training and awareness around responsible use and best-practice prompt engineering for copilot and validation/checking outputs from copilot will be required to support personnel to avoid 'garbage in garbage out' scenarios.

ii. **Detective Measures:** Prompt/response visibility controls and stringent auditing *(– limited free license capabilities with the ability to mature over time).*

iii. **Response Measures:** Copilot control system and Enterprise Data Protection (EDP) *(– limited free license capabilities with the ability to mature over time).* Implement more robust governance and administrative controls *(– would ordinarily possess alignment with AI/RAI Governance and Use policies).*

iv. **Note: Low remaining residual risk**, particularly with respect to anonymisation process and reliance on public content.

D. Final Review and Recommendations

The reviewers of this risk assessment conclude that **Microsoft Copilot Chat (free version) presents a low to moderate risk profile within the DIE and NZDF tenant boundary *and* that more rigorous and clear governance, ownership, and mitigation strategies should be in place and validated as soon as possible to monitor and prevent the risk profile from escalating.**

Further, it would be prudent to distribute internally authored general usage guidelines for Copilot Chat and update existing NZDF guidance regarding GenAI. Specifically, it is recommended to provide a programme of training and education to ensure secure and responsible use of Copilot Chat in the immediate future.

Copilot Chat presents itself as an especially valuable and secure tool within the wider M365 ecosystem, but business ownership, active governance, continuous monitoring, and user compliance training are necessary to ensure the safe uptake and assimilation of the product on an ongoing basis.

Note: The limitations of this review include the following…

i. This paper was written largely by consulting information that is publicly available and published online by Microsoft.

ii. The NZDF tenant may have different configurations than an 'ideal' tenant. As such, NZDF should seek to validate the native auditing, compliance, and security controls before considering further action.

Acceptance of Notes, Review, and Recommendations

CDO	**Date:**
CISO	**Date:**