


# The Australian Privacy Act and Generative AI

A Guide for Customers



# Contents

<b>Executive Summary.....</b>	<b>5</b>
<b>Introduction.....</b>	<b>6</b>
<b>Generative AI in businesses and the public sector .....</b>	<b>6</b>
Our approach to privacy and generative AI.....	6
Purpose and structure of this paper .....	6
<b>Part 1: Responsibly using AI – Microsoft’s AI journey and leveraging our tools and resources .....</b>	<b>8</b>
Responsible AI .....	8
Tools, commitments and resources to assist your AI deployment .....	9
<b>Part 2: The Privacy Act compliance framework in the context of AI .....</b>	<b>10</b>
What is the Privacy Act and who does it apply to?.....	10
Leverage established principles to comply with regulatory frameworks when using AI solutions.....	10
Compliance with the Privacy Act is a shared responsibility .....	11
How does Microsoft support customers with their Privacy Act compliance obligations? .....	11
Protecting the data of our customers – Microsoft’s privacy commitments in the AI era .....	12
Key obligations under the Privacy Act in the context of the procurement and use of generative AI services .....	15
• APPs 1 and 5 (Open and transparent management of personal information and notification of the collection of personal information).....	15
• APP 3 (Collection of solicited information) .....	15
• APP 6 and APP 9 (Use or disclosure of information, including government-related identifiers) ....	15
• APPs 12 and 13 (Access to and correction of personal information).....	16
• APP 11.1 (Security of personal information) .....	18
• APP 11.2 (Destruction of Personal Information) .....	18
• Section 12 of the Government Agencies Code (Privacy Impact Assessments).....	18
• APP 8 (Cross-border disclosure of personal information) .....	19
Our continued compliance with data protection regulation and open dialogue with key regulators across the globe.....	20
<b>Part 3: Copilot for Microsoft 365 .....</b>	<b>21</b>
What is Copilot for Microsoft 365 and how does it work?.....	21
How does Copilot for Microsoft 365 use personal information? .....	23
Security for Copilot for Microsoft 365.....	24
Data residency .....	24

<b>Part 4: Azure OpenAI Service .....</b>	<b>25</b>
What is Azure OpenAI Service and how does it work? .....	25
Preventing abuse and harmful content generation.....	27
How does the Azure OpenAI Service use personal information? .....	28
Security for Azure OpenAI.....	29
Data residency .....	29
<b>Part 5: Conclusion.....</b>	<b>30</b>
<b>Appendix 1: Business opportunities arising from generative AI .....</b>	<b>31</b>
AI transformation opportunities.....	31
General use cases for Copilot for Microsoft 365 .....	31
Department and employee-specific use cases .....	32
Industry-specific use cases.....	32
<b>Appendix 2: Opportunities in the public sector .....</b>	<b>33</b>
<b>Appendix 3: Frequently asked questions (FAQs).....</b>	<b>35</b>
<b>Appendix 3: Additional resources .....</b>	<b>39</b>









# Executive Summary

- The [use cases for generative AI](#) present an exciting opportunity to improve the quality of services and operational efficiency. At Microsoft we want to empower our customers to harness the full potential of new technologies like generative artificial intelligence (generative AI) while complying with their obligations under the Australian Privacy Act 1988 (Cth) (Privacy Act) and the Australian Privacy Principles (APPs).
- Microsoft is committed to ensuring its AI systems are developed responsibly and in a way that is worthy of people's trust. We drive this commitment according to [six key principles](#) that support the objectives of the Privacy Act and the APPs.
- When considering compliance with the Privacy Act in the context of the procurement and use of generative AI services, the APPs apply in the same manner as they do for processing personal information in any other context (e.g. the use of cloud services). So, while AI technology may be new, the principles and accordingly the processes for risk assessment and compliance with the APPs remain the same. Hence, to ensure compliance with the Privacy Act, organisations should be confident to approach Microsoft's AI services in the same way as they have approached procuring and using other cloud services.
- Microsoft's existing privacy commitments including those provided in [Microsoft's Data Protection Addendum](#) extend to our AI commercial products and demonstrate our compliance with the European Union's General Data Protection Regulation (GDPR) and other privacy laws applicable to us providing product and services, including the Privacy Act. Customers can rest assured that the [privacy commitments](#) they have long relied on when using our enterprise cloud products also apply to Copilot for Microsoft 365 and the Azure OpenAI Service. Customers can therefore be confident that their valuable data is safeguarded by industry-leading data governance and privacy practices in the most trusted cloud on the market today.
- There are a number of [key obligations under the Privacy Act](#) that organisations need to consider when procuring generative AI services. In this paper we have included details of these obligations and the associated support and resources Microsoft can offer including in relation to international transfers of personal information, transparency, rights of individuals, technical and organisational security measures, and PIAs. Although it is the responsibility of customers to ensure their own compliance with the Privacy Act, Microsoft's approach to privacy, including compliance with the Privacy Act, supports how you can fulfil your compliance obligations.
- Our customers' data belongs to our customers. Microsoft does not claim ownership of any customer prompts or output content created by Microsoft's generative AI solutions. In addition, no Customer Data (including prompts or output content) is used to train foundation models without customer permission.
- As the regulatory landscape evolves and we innovate to provide new kinds of AI solutions, Microsoft will continue to offer industry-leading tools, resources and support to demonstrate our enduring commitment to meeting the needs and demands of our customers in their AI journey.

# Introduction

## Generative AI in businesses and the public sector

In today's rapidly evolving business landscape, organisations are increasingly pressured to innovate, achieve greater efficiency and to enhance customer experiences. Similarly, there is increasing demand on governments to accelerate their digital and technological capabilities to help deliver on their responsibilities and meet citizen demands while operating with often constrained budgets.

Generative AI can help organisations seek a competitive edge by utilising the potential of generative AI solutions. By automating routine tasks and providing deep analytical insights generative AI solutions can help businesses stay competitive and responsive to market dynamics. There is no doubt AI is poised to shape the future of how organisations operate.

The business value of AI is clear: it helps organisations operate efficiently, perform better, achieve more, and gain the insights required to make better decisions. In addition, investment in AI solutions has been shown to have a positive impact on an organisation's bottom line.

These efficiencies can also be captured by governments by harnessing the potential of digital transformation technologies. Effective public service delivery is both a responsibility and an opportunity for governments. Ultimately it means working efficiently with the resources available, delivering great outcomes for society, while safeguarding people's privacy and wellbeing. Getting this balance "right" requires an understanding of the best tools available, used in a responsible and secure manner, particularly given the need to safeguard personal information.

Generative AI presents an exciting opportunity for both private and public sector organisations. For private sector organisations, generative AI solutions can optimise an organisation at every level and uncover new valuable opportunities within the business. Similarly, for public sector organisations,

if integrated in a considered manner that safeguards personal information, AI has the potential to reduce administrative workload, increase efficiency of services and help public servants make better, faster decisions – all while increasing civil servant satisfaction and much more.

However, delivering this type of impact with AI innovation needs to be balanced by ensuring the organisation selects efficient and trustworthy AI solutions and that these are implemented in a responsible and secure manner, taking into account the need to safeguard personal information.

## Our approach to privacy and generative AI

At Microsoft we want to empower our customers to harness the full potential of new technologies like generative AI while complying with their obligations under the Privacy Act to ensure the privacy and security of their data.

We have a long-standing practice of protecting our customers' information. Our approach to [Responsible AI](#) is built on a foundation of privacy and we remain dedicated to upholding core values of privacy, security and safety in all our generative AI products and solutions. As the use of AI solutions expands, our customers can be confident their valuable data is safeguarded by industry-leading data governance and privacy practices in one of the most trusted clouds on the market today. Customers can rest assured that the privacy commitments they have long relied on when using our enterprise cloud products also apply to our enterprise generative AI solutions that are backed by Microsoft's Data Protection Addendum, including Copilot for Microsoft 365 and Azure OpenAI Service.

## Purpose and structure of this paper

As an industry and thought leader in AI, we have developed this paper to demonstrate how customers subject to the Privacy Act can use Copilot for Microsoft 365 and the Azure OpenAI Service in compliance with their legal obligations and demonstrate our solution's ability to assist customers to comply with the Privacy Act.

This paper is set out as follows:

## Part 1

---

Examines the meaning of responsible AI, the six key principles and approach to responsible AI that guide Microsoft's development of AI products, and demonstrates the tools and resources Microsoft offers to assist your AI deployment.

## Part 2

---

Shifts focus to the structure and requirements of the Privacy Act and how Microsoft can support customers to embrace our AI solutions while continuing to meet their compliance obligations under the Privacy Act.

## Parts 3 and 4

---

Dedicated to an in-depth exploration of Copilot for Microsoft 365 and the Azure OpenAI Service, and how these services can be utilised in compliance with the Privacy Act.

## Part 5

---

Concludes the paper, reflecting on the insights shared and the future trajectory of AI and data protection regulation.

## Appendix 1

---

Showcases some of the exciting opportunities that generative AI presents for businesses across various industries and the public sector.

## Appendix 2

---

Addresses some frequently asked questions (FAQs) that customers have regarding embracing AI in a Privacy Act compliant manner.

## Appendix 3

---

Provides links to additional resources customers can reference to supplement and expand their understanding of the information provided in this paper.



# Part 1:

## Using AI responsibly –

## Microsoft's AI journey and leveraging our tools and resources

### Responsible AI

From streamlining employee tasks to accelerating the delivery of products and services, and improving healthcare, education, public safety and transport, AI has the potential to transform both businesses. However, with this great power comes great responsibility, and it is therefore essential that AI is developed and deployed responsibly. Microsoft has taken a principled role by developing comprehensive AI responsibility policies and tools, grounded in work we have been doing for many years.

The responsible use of AI is, of course, a topic businesses and public sector organisations around the world have actively addressed in recent years. Through leading discussions, developing approaches and strategies, and implementing these in their operations, the responsible use of AI to deliver more productive, efficient, innovative and inclusive products and services within the private and public sector is on the rise.

[Learn more about Governing AI: A Blueprint for the Future](#)

At Microsoft, we are committed to making sure AI systems are developed responsibly and in a way that is worthy of people's trust. We drive this commitment according to **six key principles**, which broadly align with the Australian AI Ethics Principles (that are designed for adoption by both the private and public sectors) and support the objectives of the Privacy Act and the APPs:

- **Fairness:** AI systems should be designed to treat all individuals fairly, without bias or discrimination.
- **Reliability and safety:** AI systems should be reliable and safe, with built-in mechanisms to prevent errors and minimise harm.
- **Accountability:** The creators of AI tools and the developers who leverage them should be accountable for their systems.

- **Privacy and security:** AI systems should respect individuals' privacy and data security.
- **Inclusiveness:** AI systems should be designed to be accessible and usable by everyone, including individuals with disabilities.
- **Transparency:** AI systems should be transparent and explainable, with clear documentation of their functionality and decision-making processes.

These principles can be used by customers to evaluate AI systems and processes in use or under consideration in the context of the Privacy Act, as explored in Part 2 below. Within Microsoft, we have established our Office of Responsible AI, which sets AI governance policies for the entire company, advises our senior leadership team on AI issues and enables engineering and compliance teams across the company to build according to responsible AI principles, all while ensuring that as a corporation we are continuing to examine and improve our ethical stance as new capabilities and challenges arise.

[Learn more about Microsoft's principles and approach to Responsible AI](#)

In May 2024, we published our inaugural [Responsible AI Transparency Report](#), which builds on our internal [Microsoft Responsible AI Standard](#). This report provides insights into how we build applications that use generative AI; make decisions and oversee the deployment of those applications; support our customers as they build their own generative AI applications; and learn, evolve and grow as a responsible AI community.

Organisations should develop and be governed by responsible AI strategies, and these strategies should incorporate principles, practices, tools and governance to enable those across the organisation to assess, adopt and manage their use of AI.



When potential risks are understood and carefully managed, organisations can realise the promise of AI. Forward-looking leaders will ensure that their commitment to responsible AI is not an afterthought but is baked into their organisation's innovation pipeline. This allows businesses to harness the power of AI to improve their products and/or services and allows the public sector to improve the services they provide and benefit society as a whole. You can find [several exciting examples of how to use generative AI in Appendix 1.](#)

## Tools, commitments and resources to assist your AI deployment

To support our customers and empower compliant use of AI in their deployment, Microsoft offers a range of solutions, tooling and resources – from comprehensive [transparency documentation](#) to a suite of tools for data governance, risk, and compliance assessment. Dedicated programs such as our industry-leading [AI Assurance Program and AI Customer Commitments](#) further broaden the support we offer customers in addressing their needs.

Microsoft's [AI Assurance Program](#) helps customers ensure that the AI applications they deploy on our platforms meet the legal and regulatory requirements for responsible AI. The program includes support for regulatory engagement and advocacy, risk framework implementation and the creation of a customer council.

For decades we've defended our customers against intellectual property claims relating to our products. Building on our previous [AI Customer Commitments](#), Microsoft announced our [Customer Copyright Commitment](#), which extends our intellectual property indemnity support to both Copilot for Microsoft 365 and our Azure OpenAI Service. Now, if a third party sues a customer for copyright infringement for using Copilot for Microsoft 365 or the Azure OpenAI Service or for the output they generate, we will defend the customer and pay the amount of any adverse judgements or settlements that result from the lawsuit, as long as the customer has used the guardrails and content filters we have built into our products.

Microsoft has also developed a range of solutions to support our customers with data governance, with Microsoft Purview. You can find further detail on [how Microsoft Purview can support compliance with the Privacy Act in Part 2.](#)



# Part 2:

## The Privacy Act compliance framework in the context of AI

### What is the Privacy Act and who does it apply to?

The Privacy Act and the Australian Privacy Principles (also known as the “APPs”) set the bar in Australia for privacy rights, information security and compliance. At Microsoft, we value privacy as a fundamental right, and we believe our commitment to complying with Privacy Act plays an important role in assisting customers to protect and enable the privacy rights of individuals under the Privacy Act.

Microsoft is committed to complying with the Privacy Act and providing an array of products, features, documentation and resources to support our customers in fulfilling their compliance obligations under the Privacy Act. The Privacy Act applies to “APP Entities”. These are Australian Federal Government agencies and organisations who meet the applicable thresholds (for example, have an annual turnover of A\$3,000,000) and have an ‘Australian link’. This includes organisations that are incorporated in Australia or carry on business in Australia. Most State and Territory public sector agencies are covered by specific State or Territory privacy legislation (which is outside the scope of this paper).

At a high level, the Privacy Act and the APPs impose a set of data protection rules on the processing of personal information with its objectives including:

- Promoting the protection of the privacy of individuals.
- Balancing that protection with the interests of entities carrying out their functions or activities.
- Promoting responsible and transparent handling of personal information by entities.
- Facilitating the free flow of information across national borders while ensuring the privacy of individuals is respected.
- Providing a means for complaints about alleged interferences with privacy.

### Leverage established principles to comply with regulatory frameworks when using AI solutions

When we think about the Privacy Act in the context of leveraging generative AI and taking advantage of the opportunities presented by this technology, the starting point is that the APPs apply in the same manner as they do for processing personal information in any other context, including when using the cloud. So, while the AI technology may be new, the principles and accordingly the processes for risk assessment and compliance with the Privacy Act remain the same.

It is also helpful to recognise that the Privacy Act is technology-agnostic and therefore does not prevent organisations from embracing opportunities to use generative AI.

As such, applying established Privacy Act assessment processes is a great way for organisations to harness the revolutionary potential of AI and deliver great outcomes for society while safeguarding people’s privacy and wellbeing. Microsoft has a long-standing history of collaborating with and assisting organisations in pursuit of their digital transformation priorities while complying with the requirements of the Privacy Act, including in relation to the transition from on-premises to cloud computing. Organisations can approach Microsoft’s generative AI solutions by leveraging the approach they have used in procuring our cloud services.

Cloud computing is essential for accessing the potentially groundbreaking AI technology, and the hyper-scale cloud is, therefore, the foundation for deploying AI. Azure’s enterprise-grade protections that form part of Copilot for Microsoft 365 and the Azure OpenAI Service provide a strong foundation upon which customers can build their data privacy, security and compliance systems to confidently scale AI while managing risk and ensuring compliance with the Privacy Act.



## Compliance with the Privacy Act is a shared responsibility

Compliance with the Privacy Act is a shared responsibility. Microsoft is committed to complying with all laws and regulations applicable to Microsoft and its generative AI tools and services.

As a Microsoft customer, you will need to determine how these tools and services will be used and what personal information will be processed to enable you to ensure you are using such tools in a compliant manner.

To assist you with that, we have designed our generative AI tools and services with privacy and data protection in mind and provide our customers with information, features and contractual commitments to support you in your compliance and accountability obligations under the Privacy Act. The following sections in Part 2 delve into this in more detail and provide you with information to support your assessment of the use of Microsoft's generative AI tools and services in compliance with the Privacy Act.

## How does Microsoft support customers with their Privacy Act compliance obligations?

As more organisations wish to leverage generative AI, many are looking to Microsoft not only as a service provider, but as a trusted partner helping them to meet their compliance obligations under the Privacy Act.

The first step towards compliance is understanding how Microsoft's generative AI services work including how they process personal information. Our comprehensive transparency documentation and information help you understand how our AI tools work and what choices our customers can make to influence system performance and behaviour.

In Part 3 and Part 4 of this paper we provide specific information and links to additional resources you can use to help enhance your understanding of these products and services.

[Jump to Part 3 to find out more about Copilot for Microsoft 365](#)

[Jump to Part 4 to find out more about the Azure OpenAI Service](#)

This knowledge provides the foundation for compliance with a number of key obligations under the Privacy Act. We will explore these [key obligations and the associated support that Microsoft offers customers later in this Part 2](#) but first we will address the seven core privacy commitments Microsoft offers to its customers in the AI era.





## Protecting the data of our customers – Microsoft’s privacy commitments in the AI era

Microsoft’s existing privacy commitments extend to our AI commercial products, as explained in a [blog post from our Chief Privacy Officer Julie Brill](#). You can rest assured that the privacy commitments you have long relied on when using our enterprise cloud products also apply to our enterprise generative AI solutions that are backed by [Microsoft’s Data Protection Addendum](#), including Copilot for Microsoft 365 and Azure OpenAI Service.

The following seven commitments apply to “Customer Data”, which is defined in [Microsoft’s Product Terms](#) as all data, including all text, sound, video or image files and software, that are provided to Microsoft by, or on behalf of, our customers through use of an online service. All inputs (including prompts)<sup>1</sup> and output content<sup>2</sup> are Customer Data. In accordance with [Microsoft’s Data Protection Addendum](#) the customer “retains all right, title and interest in and to Customer Data”.

- 1. We will keep your organisation’s data private.** Your data remains private when using Copilot for Microsoft 365 and Azure OpenAI Service and is governed by our applicable privacy and contractual commitments, including the commitments we make in [Microsoft’s Data Protection Addendum](#) and [Microsoft’s Product Terms](#).
- 2. You are in control of your organisation’s data.** Your data is not used in undisclosed ways or without your permission. You may choose to customise your use of Copilot for Microsoft 365 or Azure OpenAI Service, opting to use your data to fine-tune models for your organisation’s own use. If you do use your organisation’s data to fine-tune, any fine-tuned AI solutions created with your organisation’s data will be available only to you.
- 3. Your access control and enterprise policies are maintained.** To protect privacy within your organisation when using enterprise products with generative AI capabilities, your existing permissions and access controls will continue to apply to ensure your organisation’s data is displayed only to those users to whom you have given appropriate permissions.

1. “Inputs” means all Customer Data that the customer provides, designates, selects or inputs for use by a generative artificial intelligence technology to generate or customise an output including any customer prompts.

2. “Output Content” means any data, text, sound, video, image, code or other content generated by a model in response to Input.

**4. Your organisation's data is not shared.**

Microsoft does not share your data with third parties without your permission. Your data, including the data generated through your organisation's use of Copilot for Microsoft 365 or Azure OpenAI Service – such as prompts and responses – are kept private and are not disclosed to third parties.

**5. Your organisation's data privacy and security are protected by design.** Security and privacy are incorporated into all phases of design and implementation of Copilot for Microsoft 365 and Azure OpenAI Service. As with all our products, we provide a strong privacy and security baseline and make available additional protections that you can choose to enable. As external threats evolve, we will continue to advance our solutions and offerings to ensure world-class privacy and security in Copilot for Microsoft 365 and Azure OpenAI Service, and we will continue to be transparent about our approach.

**6. Your organisation's data is not used to train foundation models.** Microsoft's generative AI solutions, including Copilot for Microsoft 365 and Azure OpenAI Service capabilities, do not use Customer Data to train foundation models without your permission. Your data is never available to OpenAI or used to improve OpenAI models.

**7. Our products and solutions comply with global data protection regulations.** The Microsoft AI products and solutions you deploy are compliant with today's global data protection and privacy regulations. As we continue to navigate the future of AI together, including the implementation of Europe's AI Act and other global laws, organisations can be certain that Microsoft will be transparent about our privacy, safety and security practices. We will comply with global laws that govern AI and back up our promises with clear contractual commitments.

## Microsoft's privacy commitments apply to AI



We will keep your organisation's data private.



Your organisation's data is not shared.



You are in control of your organisation's data.



Your organisation's data security and privacy are protected by design.



Your access control and enterprise policies are maintained.



Your organisation's data is not used to train foundation models.



Our products and solutions comply with global data protection regulations.

You can find additional details about how Microsoft's privacy commitments apply to Azure OpenAI and Copilot for Microsoft 365 [here](#) and the [FAQ: Protecting the Data of our Commercial and Public Sector Customers in the AI Era](#).







## Key obligations under the Privacy Act in the context of the procurement and use of generative AI services

There are several obligations under the Privacy Act, which organisations need to consider when procuring generative AI services. This section considers some of the key obligations and what associated support and resources Microsoft can offer to your organisation to help you comply.

- **APPs 1 and 5 (Open and transparent management of personal information and notification of the collection of personal information)**

APPs 1 and 5 of the Privacy Act require APP Entities to provide individuals with certain key information about how their personal information will be collected, held, used and disclosed. This includes providing a clearly expressed and up-to-date privacy policy and (where reasonable in the circumstances) providing individuals with a collection notice at, before, or as soon as practicable after, the point at which personal information is collected. If you deploy a new technology (such as Copilot for Microsoft 365 or Azure OpenAI Service) and intend to use such technology in a way that is not reflected in your existing privacy policy and collection notices, then you will need to update your privacy policy and collection notices to reflect these new processing activities.

**How we help you comply:** The information set out in this paper and available in our transparency resources noted below is intended to assist your understanding of how Copilot for Microsoft 365 and Azure OpenAI Service process data and the extent to which additional information (if any) needs to be communicated to individuals. Additional product-specific information is available at [Data, Privacy and Security for Azure OpenAI Service](#); [Data, Privacy and Security for Microsoft Copilot for Microsoft 365](#); [Copilot in Dynamics 365 and Power Platform](#); and [FAQs for Copilot data security and privacy for Dynamics 365 and Power Platform](#).

- **APP 3 (Collection of solicited information)**

APP 3 requires APP Entities to only collect personal information that is reasonably necessary for, and directly related to, one or more of their functions or activities. If the information is sensitive information, the individual generally must also consent to its collection. Further, collection must be by lawful and fair means.

**How we help you comply:** We will not contact any individual data subjects to collect personal information from them on your behalf. You are responsible for ensuring that any personal information provided to Microsoft complies with the requirements of APP 3, including ensuring that the personal information was collected lawfully and in accordance with any necessary consents.

- **APP 6 and APP 9 (Use or disclosure of information, including government-related identifiers)**

APP 6 requires APP Entities to use or disclose personal information only for the purpose it was collected or, in certain circumstances, a 'secondary purpose' (i.e. any purpose other than the purpose for which the information was collected). These circumstances in which personal information can be used for a secondary purpose include where:

- The individual has consented to the APP Entity using or disclosing their personal information for a secondary purpose.
- The secondary purpose is within the individual's reasonable expectations and is related to the purpose for which the information was collected (or, in the case of sensitive information, directly related).

More stringent restrictions to government-related identifiers (e.g. an Australian passport number, an Australian State or Territory driver's licence number or Medicare card number). APP 9 prohibits APP Entities from using government-related identifiers unless it is reasonably necessary to verify the identity of the individual or another (limited) exemption applies.

**How we help you comply:** Each customer remains responsible for ensuring its use and disclosure of personal information to Microsoft is consistent with APP 6 and 9. However, as outlined in the “Nature of Data Process; Ownership” and “Disclosure of Process Data” sections of Microsoft’s Data Protection Addendum, Microsoft supports your compliance by contractually committing to:

- Use and process Customer Data only to provide you with Microsoft services and products. Microsoft will not use or otherwise process Customer Data for user profiling, advertising, market research or for any other purpose, unless such use or processing is in accordance with your documented instructions.
- Not disclose Customer Data (including Customer Data that has been processed by Microsoft) to a third party without your permission or as required by law.

From a practical perspective, these commitments mean that:

- Copilot for Microsoft 365 and the Azure OpenAI Service are hosted in Microsoft’s Azure environment and are fully controlled by Microsoft. Customer Data is not used by, or disclosed to, OpenAI who developed the foundation LLM upon which Microsoft’s service is based.
- Where content from your Microsoft 365 tenant is used to augment a user’s prompt and enrich the response, no Customer Data is being used to train the foundation LLM owned by OpenAI. In fact, the foundation LLM retains no information about the prompt that was submitted to it, nor any Customer Data that was used to ground it, nor any responses it provided.
- Data, such as prompts and responses, generated through your organisation’s use of Copilot for Microsoft 365 and Azure OpenAI Service, are not shared with third parties (including OpenAI) without your permission.

Microsoft also implements technical and organisational measures (including encryption) to protect Customer Data and data deletion policies to assist in preventing unauthorised disclosure of Customer Data.

#### • APPs 12 and 13 (Access to and correction of personal information)

APP Entities must ensure they are in a position to comply with their obligations to respond to requests from individuals to access their personal information held by the APP Entity and to correct personal information that (having regard to the purpose for which the information is held) is inaccurate, out-of-date, incomplete, irrelevant or misleading. These are rights granted to individuals under APPs 12 and 13 of the Privacy Act.

**How we help you comply:** In the “Data Subjects Rights; Assistance with Requests” section of [Microsoft’s Data Protection Addendum](#), Microsoft commits to make available to customers (in a manner consistent with the functionality of the services) the ability to fulfil requests from individuals exercising their rights of access and correction. If Microsoft receives such a request directly from an individual to access or correct their personal information in situations where it is processing their personal information on behalf of your organisation, it will redirect the individual to submit its request to your organisation instead. You are responsible for responding to any such requests, but Microsoft will comply with reasonable assistance requests in this respect.

Microsoft has developed additional solutions to assist its customers when responding to requests from individuals to exercise their applicable privacy rights, such as Microsoft Purview and Purview eDiscovery. The features of these products empower our customers to proactively govern their AI usage and adhere to evolving regulatory requirements. This can be valuable; for instance, to improve efficiency in responding to and actioning requests in relation to the rights to request access and correction of personal information that apply under APPs 12 and 13 of the Privacy Act.

[Learn more about Microsoft Purview and its features](#) and how these tools can assist you in the deployment of Microsoft’s generative AI solutions.





- **APP 11.1 (Security of personal information)**

APP 11.1 requires APP Entities to take reasonable steps to protect the personal information they hold from misuse, interference and loss and from unauthorised access, modification or disclosure.

**How we help you comply:** In the “Data Security” section of the [Microsoft’s Data Protection Addendum](#), Microsoft contractually commits to implement and maintain appropriate technical and organisational measures to protect Customer Data and Personal Data (which captures personal information as defined under the Privacy Act) against accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, such data transmitted, stored or otherwise processed.

Those technical measures are set forth in Microsoft’s Security Policy and comply with ISO 27001, ISO 27002 and ISO 27018. Microsoft also contractually commits to encrypting Customer Data (including any Personal Data contained therein), in transit (including between Microsoft data centres) and at rest. Appendix A – Security Measures to [Microsoft’s Data Protection Addendum](#) also contains comprehensive commitments from Microsoft regarding the security of Customer Data, including in relation to the Organisation of Information Security, Asset Management, Human Resources Security, Physical and Environmental Security, Communications and Operations Management, Information Security, Incident Management and Business Continuity Management.

The technical, organisational and security measures described above apply to any Customer Data that customers provide or create when using Copilot for Microsoft 365 and Azure OpenAI Service. You can refer to the information set out above to demonstrate the commitment and measures taken by Microsoft to protect Customer Data (including personal information).

[Jump to Part 3 to find out more about Security for Copilot for Microsoft 365.](#)

[Jump to Part 4 to find out more about Security for Azure OpenAI Service.](#)

- **APP 11.2 (Destruction of Personal Information)**

APP 11.2 requires APP Entities to take reasonable steps to destroy or de-identify personal information they hold and no longer need for any permitted purpose (that is, the purpose for which it was collected or a related secondary purpose).

**How we help you comply:** As outlined in the “Data Retention and Deletion” section of [Microsoft’s Data Protection Addendum](#), customers have the ability to access, extract and delete Customer Data stored in our services at all times during the term of the relevant subscription.

If you cancel your Copilot for Microsoft 365 and Azure OpenAI Services or your subscriptions expire, we would delete your Customer Data within 180 days (unless Microsoft is otherwise authorised to retain it under the Data Protection Addendum).

- **Section 12 of the Government Agencies Code (Privacy Impact Assessments)**

Section 12 of the Privacy (Australian Government Agencies – Governance) APP Code 2017 requires Australian Federal Government agencies to conduct a Privacy Impact Assessment (PIA) for all high privacy risk projects. A project may be high privacy risk if the agency reasonably considers that it involves any new or changed ways of handling personal information that are likely to have a significant impact on the privacy of individuals. In Australia, a PIA must (at least):

- Identify the impact the activity or function might have on the privacy of individuals.
- Set out recommendations for managing, minimising or eliminating that impact.

The obligation to undertake a PIA is expected to be made mandatory for all APP Entities as part of upcoming reforms to the Privacy Act. In the interim, organisations are required under APP 1.2 to take reasonable steps to implement practices, procedures and systems that will ensure compliance with the APPs. This will normally include carrying out a PIA on any high privacy risk projects.



When assessing whether or not a PIA is required, APP Entities should take into account the nature, scope, content and purposes of the collection, holding, use and disclosure of personal information. Therefore, whether or not a PIA is required for the use of Copilot for Microsoft 365 and the Azure OpenAI Service will depend on the particular use case and type of personal information you wish to handle using these services.

Even where it is not legally required, a PIA is good practice and can help you work through the specific data protection risks associated with the implementation of Copilot for Microsoft 365 and/or Azure OpenAI Service for a specific use case. Preparing a PIA may also assist you in meeting your obligations under APP 1.2 of the Privacy Act.

**How we help you comply:** The information contained in this paper and the additional resources to which it refers can assist you with completing a PIA. In particular, the information in:

- [Part 3](#) and [Part 4](#) relating to how Copilot for Microsoft 365 and Azure OpenAI Service handles data will assist identify the impact that the activity or function might have on the privacy of individuals.
- The sections on technical and organisational measures for both Copilot for Microsoft 365 and Azure OpenAI Service will assist in developing recommendations for managing, minimising or eliminating that impact. completing the elements described above.

However, when undertaking a PIA, identifying the impact and developing recommendations will depend on the use case and the nature, scope and content of the personal information involved and will need to be undertaken by you.

#### • APP 8 (Cross-border disclosure of personal information)

APP 8 requires APP Entities to take reasonable steps to ensure that a recipient of personal information who is located outside of Australia does not breach the Australian Privacy Principles (other than APP 1).

**How we help you comply:** The Australian Government's [Better Practice Guide, Privacy and Cloud Computing for Australian Government Agencies](#) states:

*"If an agency shares personal information with a contracted cloud service provider, this may be considered a "use" rather than a "disclosure" under the Privacy Act, depending on the degree of control the agency retains over the personal information. An agency that gives up its control over personal information to an outsider is treated as disclosing that information. An agency that maintains control over personal information is treated as using that information".*

When considering compliance with APP 8 in the context of Copilot for Microsoft 365 and Azure OpenAI Service, the same rules apply as with other cloud-based services. Therefore, using these services does not disclose personal information, as Microsoft processes only personal information provided by its customers in accordance with its customer's instructions.

In addition, if you provision your Microsoft 365 tenant in Australia:

- Any stored content of interactions with Copilot for Microsoft 365 and Azure OpenAI Service is stored at rest in Australia.
- Any Customer Data processed by us will be processed primarily in Australia. As outlined in [Microsoft's Data Protection Addendum](#) (that forms part of your agreement with Microsoft), Microsoft will comply with all laws and regulations applicable to its providing the Products and Services, including the Privacy Act.

## Our continued compliance with data protection regulation and open dialogue with key regulators across the globe

As privacy and data protection laws advance, and norms and requirements evolve in Australia and across the globe; you can be certain Microsoft will be transparent about our privacy, safety and security practices. We will comply with privacy and data protection laws applicable to us providing product and services and back up our promises with clear contractual commitments.

Beyond adhering to the Privacy Act and other regulatory requirements applicable to us, Microsoft prioritises an open dialogue with its customers, partners and regulatory authorities to better understand and address evolving privacy and data protection concerns.

We continue to work closely with data protection authorities and privacy regulators around the world to share information about how our AI systems work, thereby fostering an environment of trust and cooperation.





# Part 3:

## Copilot for Microsoft 365

Understanding the potential of generative AI services and how these products and services operate and use personal data is the foundation for customers to determine compliance with a number of obligations under the Privacy Act. This Part 3 provides information and links to various external resources that can help you understand how Copilot for Microsoft 365 operates and provides key information about the product and its features that can be used to assist with completion of a PIA or other data protection assessment/analysis.

### What is Copilot for Microsoft 365 and how does it work?

Copilot for Microsoft 365 is an AI-powered productivity tool that uses "Large Language Models (LLMs)" to work alongside popular Microsoft 365 apps such as Word, Excel, PowerPoint, Outlook, Teams and more. Copilot for Microsoft 365 provides real-time, intelligent assistance that enables users to enhance their creativity, productivity and skills.

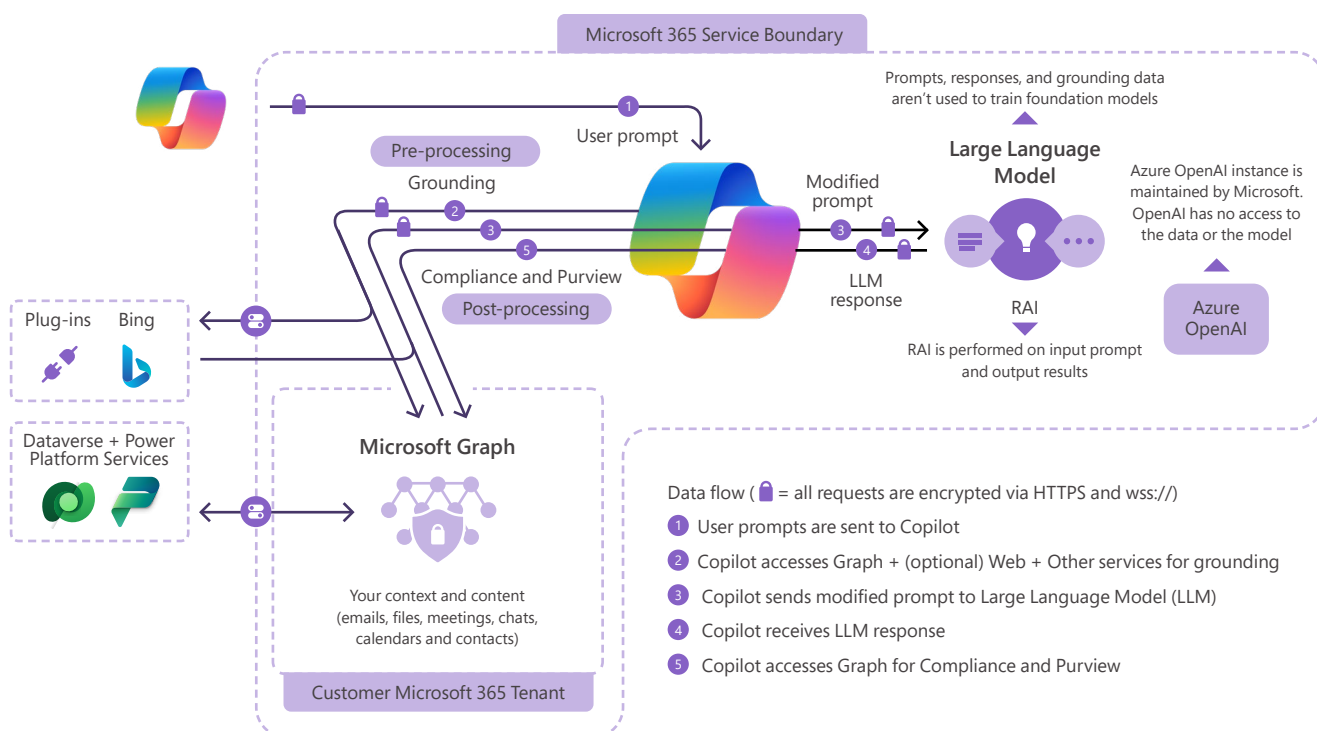
Copilot for Microsoft 365 is built on top of the same cloud infrastructure as its Microsoft 365 applications,

and applies the same principles of confidentiality and privacy to Customer Data that Microsoft has leveraged for years. Copilot for Microsoft 365 adheres to all existing privacy, security and compliance commitments that apply to Microsoft 365 including Microsoft's commitments set out in [Microsoft's Data Protection Addendum](#).

Copilot for Microsoft 365 uses the organisational content in your Microsoft 365 tenant, including users' calendars, emails, chats, documents, meetings, contacts and more only in accordance with existing access permissions. The richness of the Copilot for Microsoft 365 experience depends on the data sources indexed by Microsoft 365. Customers with the most abundant data in Microsoft 365 (Exchange, OneDrive, SharePoint, Teams) will get the best results from Copilot. With access to comprehensive organisational data, Copilot can suggest more relevant and personalised content based on the user's work context and preferences.

The following diagram provides a visual representation of how Microsoft Copilot for Microsoft 365 works.

### Microsoft Copilot for Microsoft 365 architecture



As demonstrated in the diagram above, Copilot responds to prompts from your users. A “prompt” is the term used to describe how you ask Copilot for Microsoft 365 to do something for you – such as creating, summarising, editing or transforming. Think about prompting like having a conversation, using plain but clear language and providing context like you would with an assistant.

When Copilot for Microsoft 365 uses content from the organisation’s Microsoft 365 tenant to enhance the user’s prompt and response, it’s called “grounding”. This differs from training, as no Customer Data is used to train the LLM. The LLM is stateless, meaning it doesn’t retain information about submitted prompts, grounding data, or provided responses.

Copilot for Microsoft 365 leverages an instance of a foundation LLM hosted in Azure OpenAI but does not interact with any services operated by OpenAI (e.g. ChatGPT or the OpenAI API). OpenAI is not a sub-processor to Microsoft, and Customer Data – including the data generated through your organisation’s use of Copilot for Microsoft 365 such as prompts and responses – are not shared with third parties without your permission.

To get the best responses and the most out of Copilot for Microsoft 365, it’s important that you input suitable prompts and avoid certain common pitfalls. Learn more about the skill of prompting: [The art and science of prompting \(the ingredients of a prompt\)](#) and [prompting do’s and don’ts](#).

Copilot for Microsoft 365 is:

- Built on Microsoft’s comprehensive approach to security, compliance and privacy.
- Designed to protect tenant, group and individual data.
- Committed to responsible AI.

Get an inside look at how LLMs work when you use them with your data in Microsoft 365. Learn more about [Copilot for Microsoft 365](#).

Learn about how you can use Copilot in your favourite Microsoft apps by visiting the [Copilot Lab](#).

You can also find more detailed information about Copilot for Microsoft 365 in our [Learn portal](#).

## Copilot and your privacy



### Copilot in Windows

Learn more about how Copilot uses your data to assist you on your Windows device.


[Learn more about your data and privacy](#)



### Copilot Pro (home users)

Learn more about how Copilot uses your data in Microsoft 365 apps at home.

[Read about Microsoft 365 apps and your privacy](#)



### Copilot for Microsoft 365 (IT Pros/admins)

Learn more about how your organisational data is used and protected when using Copilot with Microsoft 365.

[Get details about data, privacy and security](#)



## How does Copilot for Microsoft 365 use personal information?

Copilot for Microsoft 365 provides value by connecting Microsoft's LLMs to your organisational data. Copilot for Microsoft 365 accesses content and context to generate responses anchored in your organisational data, such as user documents, emails, calendar, chats, meetings and contacts. Copilot for Microsoft 365 combines this content with the user's working context, such as the meeting a user is currently attending, email exchanges the user had on a topic or chat conversations the user had in a given period. Copilot for Microsoft 365 uses this combination of content and context to help provide accurate, relevant and contextual responses to the user's prompts.

Copilot for Microsoft 365 can reference web content from the Bing search index to ground user prompts and responses. Based on the user's prompt, Copilot for Microsoft 365 determines whether or not it needs to use Bing to query web content to help provide a relevant response to the user. Controls are available to manage the use of web content for admins.

Abuse monitoring for Copilot for Microsoft 365 occurs in real time without providing Microsoft any standing access to Customer Data, either for human or for automated review. While abuse moderation, which includes human review of content, is available for Azure OpenAI Service, this is not required for Copilot for Microsoft 365. See [Part 4](#) for more information about abuse monitoring in the context of Azure

OpenAI Service (including how customers who meet additional Limited Access eligibility criteria and attest to specific use cases can apply to disable the Azure OpenAI content management features).

Microsoft will collect and store data about user interactions with Copilot for Microsoft 365. This will include the user's prompt, how Copilot responded, and the information used to ground Copilot's response ("Content Interactions"). Customer admins can view, manage and search your organisation's Content Interactions. It may be necessary to update your privacy notices for your organisation's users to ensure it appropriately captures any processing of personal information by admins in this context. See [Part 2 for further details of the transparency obligations under the Privacy Act](#).

It is important for Microsoft that our customers' data belongs to our customers. Microsoft does not claim ownership of the content created by Copilot for Microsoft 365. All Content Interactions including user prompts and any output data/content qualifies as "Customer Data" in our [Product Terms](#) and [Microsoft's Data Protection Addendum](#).

All Customer Data processed by Copilot for Microsoft 365 is processed and stored in alignment with contractual commitments with your organisation's other content in Microsoft 365.

Copilot for Microsoft 365 does not use Customer Data to train foundation models without our customers' permission.



## Security for Copilot for Microsoft 365

As noted in Part 2, the Privacy Act requires APP Entities to protect information from misuse, interference and loss and from unauthorised access, modification or disclosure.

The same security and compliance terms apply, by default, to Copilot for Microsoft 365 as apply already for your organisation's use of Microsoft 365. Copilot for Microsoft 365 is hosted in Azure infrastructure and protected by some of the most comprehensive enterprise compliance and security controls in the industry. Copilot for Microsoft 365 was built to take advantage of the security and compliance features that are already well-established in Microsoft's hyperscale cloud. This includes prioritisation of reliability, redundancy, availability and scalability, all of which are designed into our cloud services by default.

Copilot for Microsoft 365 also respects each user's access permissions to any content it retrieves. This is important because Copilot for Microsoft 365 will generate responses based only on information the particular user has permission to access.

Microsoft already implements multiple forms of protection to help prevent customers from compromising Microsoft 365 services and applications or gaining unauthorised access to other tenants or the Microsoft 365 system itself. Below are a few examples of those forms of protection:

- Logical isolation of Customer Data within each tenant for Microsoft 365 services is achieved through Microsoft Entra authorisation and role-based access control. Learn more about [Microsoft 365 isolation controls](#).
- Microsoft uses rigorous physical security, background screening and a multi-layered encryption strategy to protect the confidentiality and integrity of Customer Data.
- Microsoft 365 uses service-side technologies that encrypt Customer Data both at rest and in transit, including BitLocker, per-file encryption, Transport Layer Security (TLS), and Internet Protocol Security (IPSec). Learn more about encryption in Microsoft 365; see [Encryption in the Microsoft Cloud](#).

- Your control over your organisation's data is reinforced by Microsoft's commitment to comply with applicable privacy laws including the GDPR and privacy standards, such as ISO/IEC 27018, the world's first international code of practice for cloud privacy. This, in turn, facilitates your compliance with the Privacy Act.
- For content accessed through Copilot for Microsoft 365 plug-ins, encryption can exclude programmatic access, thus limiting the plug-in from accessing the content. Learn more by reading [Configure usage rights for Azure Information Protection](#).
- As generative AI systems are also software systems, all elements of our Security Development Lifecycle (SDL) apply, from threat modelling to static analysis, secure build and operations to use of strong cryptography, identity standards and more.
- We've also added new steps to our Security Development Lifecycle to prepare for AI threat vectors, including updating the Threat Modelling SDL requirement to account for AI and machine-learning-specific threats. We put our AI products through AI red teaming to look for vulnerabilities and ensure we have proper mitigation strategies in place.

[Learn more about Data, Privacy and Security for Copilot for Microsoft 365.](#)

## Data residency

When you store data generated by Copilot for Microsoft 365 in Microsoft 365 products that already have data residency commitments under the [Product Terms](#), then the applicable commitments will be upheld.

Copilot for Microsoft 365 has been added as a covered workload in the data residency commitments in the [Microsoft Product Terms](#). Microsoft [Advanced Data Residency \(ADR\)](#) and [Multi-Geo Capabilities](#) also include data residency commitments for Copilot for Microsoft 365 customers. Practically, this means that when an organisation provisions their Microsoft 365 tenant in Australia, any stored content of interactions with Copilot for Microsoft 365 is stored at rest in Australia.

# Part 4:

## Azure OpenAI Service

Understanding how generative AI products and services operate and use personal information is the foundation for compliance with a number of obligations under the Privacy Act. This Part 4 provides information and links to various external resources that can help you understand how Azure OpenAI Service operates with key information about the service and its features you can use to assist with completion of a PIA or other data protection assessment/analysis.

### What is Azure OpenAI Service and how does it work?

Azure OpenAI Service is a cloud-based platform that enables customers to build and deploy their own generative AI applications leveraging the power of AI models. Azure OpenAI Service provides customers with access to a set of LLMs for the development of generative AI experiences.

From generating realistic images and videos to enhancing customer experiences, generative AI has proven to be a versatile tool across various industries. The models underlying Azure OpenAI Service can be easily adapted to your specific task including: content design, creation and generation; summarisation; semantic search; natural language to code translation; accelerated automation; personalised marketing; chatbots and virtual assistants; product and service innovation; language translation and natural language processing; fraud detection and cybersecurity; predictive analytics and forecasting; creative writing; and medical research and diagnosis.

Azure OpenAI Service is fully controlled by Microsoft. Microsoft hosts the OpenAI/Chat GPT models in Microsoft's Azure environment and the service does not interact with any services operated by OpenAI (e.g. ChatGPT or the OpenAI API).

OpenAI/ChatGPT owns and trains the foundation LLMs that Microsoft uses, and Microsoft has a licence to offer services that rely on these foundation LLMs.

OpenAI/ChatGPT is not a sub-processor to Microsoft and Customer Data – including the data generated through your organisation's use of Azure OpenAI Service (such as prompts and responses) – are kept private and not shared with third parties without

your permission. [Learn more about the underlying LLMs that power the Azure OpenAI Service.](#)

Azure OpenAI Service can be used in the following ways:

- **Prompt engineering:** Prompt engineering is a technique that involves designing prompts for LLMs. Prompts are submitted by the user and content is generated by the service, via the completions, chat completions, images and embeddings operations. This process improves the accuracy and relevance of responses, optimising the performance of the model.

[Learn more about prompt engineering.](#)

- **Azure OpenAI On Your Data:** When using the "on your data" feature, the service retrieves relevant data from a configured Customer Data store and augments the prompt to produce responses that are grounded with your data.

Azure OpenAI "on your data" enables you to run supported LLMs on your organisation's data without needing to train or fine-tune models. Running models on Customer Data enables you to analyse your data with greater accuracy and speed. By doing so, you can unlock valuable insights that can help you make better decisions, identify trends and patterns and optimise your operations.

One of the key benefits of Azure OpenAI "on your data" is its ability to tailor the content of conversational AI. The model within Azure OpenAI Service has access to and can reference specific sources to support responses. Answers are based not only on its pre-trained knowledge but also on the latest information available in the designated data source. This grounding data also helps the model to avoid generating responses based on outdated or incorrect information.

[Learn more about Azure OpenAI On Your Data.](#)

- **Azure OpenAI fine-tuning.** You can provide your own training data consisting of prompt-completion pairs for the purposes of fine-tuning an OpenAI model. This process fine-tunes an existing LLM using example data. Fine-tuning in this instance refers to the process of retraining pre-trained models on specific datasets, typically to improve model performance on specific tasks or introduce information that wasn't well represented when the base model was trained originally. The outcome is a new "customised" LLM that has been optimised for the customer using the provided examples.

Training data and fine-tuned models:

- Are available exclusively for use by your organisation.
- Are stored within the same region as the Azure OpenAI resource.
- Can be deleted by the customer at any time.


When you upload customised data to fine-tune the results of the LLM, both the Customer Data and the results of the fine-tuned model are maintained in a protected area of the cloud, stored in your tenant – accessible only by your organisation and separated by robust controls to prevent any other access. The Customer Data and results can additionally be encrypted by either Microsoft-managed or customer-managed encryption keys in a Bring Your Own Key format, if a customer so chooses.

In most instances, Microsoft can support and troubleshoot any problems with the service without needing access to any Customer Data (such as the data that was uploaded for fine-tuning). In the rare cases where access to Customer Data is required, whether it be in response to a customer-initiated support ticket or a problem identified by Microsoft, you can assert control over access to that data by using Customer Lockbox for Microsoft Azure. Customer Lockbox gives customers the ability to approve or reject any access request to their Customer Data.

[Learn more about Azure OpenAI fine-tuning.](#)

Whether content is used to ground prompts using the "on your own data" feature or whether content is used to build a fine-tuning model, the Customer Data is not being used to train the foundation LLM. In fact, the LLM is stateless, meaning it retains no information about the prompt that was submitted to it, nor any Customer Data that was used to ground it, nor any responses it provided. The LLM is not trained and does not learn at any point during this process; it is exactly the same foundation model even after millions of prompts are run through it.

You can find detailed information about Azure OpenAI Services through the [Azure OpenAI Service – Documentation, quickstarts and API reference guides.](#)







## Preventing abuse and harmful content generation

To reduce the risk of harmful use of Azure OpenAI Service, both content filtering and abuse monitoring features are included.

Content filtering is the process by which responses are synchronously examined by automated means to determine if they should be filtered before being returned to a user. This examination happens without the need to store any data and with no human review of the prompts (i.e. the text provided by users as requests) or the responses (i.e. the data delivered back to the user.)

[Learn more about content filtering.](#)

Abuse monitoring is conducted by a separate process. This data may be accessed only by authorised Microsoft personnel to assist with debugging, and protect against abuse or misuse of the system. The human reviewers are authorised Microsoft employees who access the data via pointwise queries using request IDs, Secure Access Workstations (SAWs) and Just-In-Time (JIT) request approval granted by team managers.

[Learn more about abuse monitoring.](#)

This human review may create a challenge for customers, who need to strike a balance between the safety of the system and the risks of external access – even under controlled conditions. To accommodate that balance, Microsoft offers limited access features that allow for approved customer-use cases to opt out of these human review and data-logging processes.

Some customers may want to use Azure OpenAI Service for a use case that involves the processing of sensitive, highly confidential or legally regulated input data but where the likelihood of harmful outputs and/or misuse is low. These customers may conclude that they do not want or do not have the right to permit Microsoft to process such data for abuse detection, as described above, due to their internal policies or applicable law. To address these concerns, Microsoft allows customers who meet additional Limited Access eligibility criteria and attest to specific use cases to apply to disable the Azure OpenAI content management features by completing [this form](#).

If Microsoft approves a customer's request to disable abuse monitoring, then Microsoft does not store any prompts and completions associated with the approved Azure subscription for which abuse monitoring is configured off. In this case, because no prompts and completions are stored at rest in the service results store, the human review process is not possible and is not performed.

## How does the Azure OpenAI Service use personal information?

The diagram below illustrates how your organisation's data is processed by Azure OpenAI Service. This diagram covers three different types of processing:

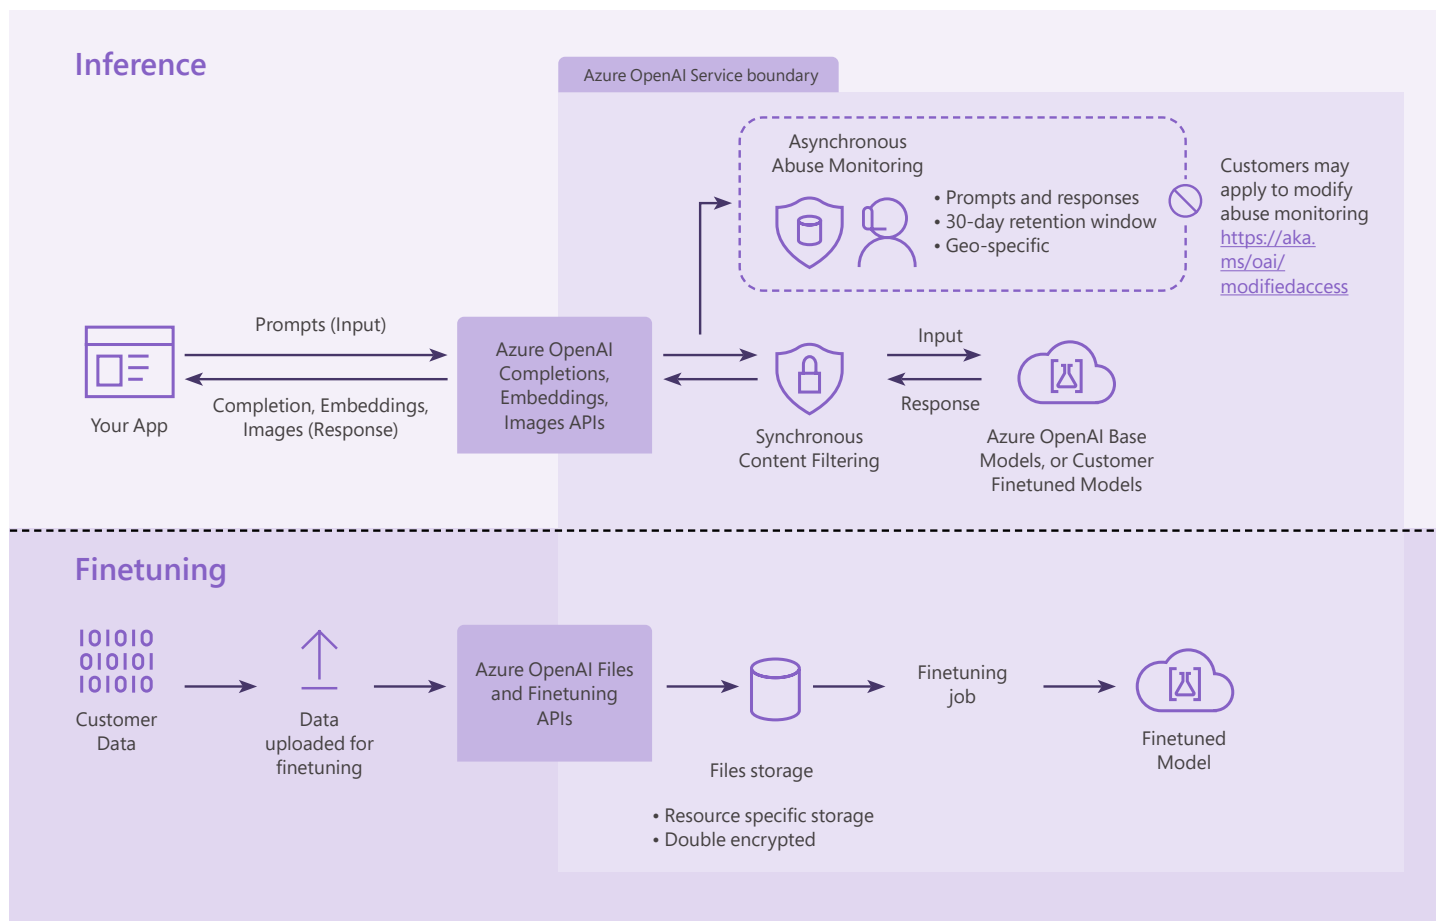
1. How Azure OpenAI Service processes your prompts to **generate content** (including when additional data from a connected data source is added to a prompt using Azure OpenAI "On Your Data").
2. How Azure OpenAI Service **creates a fine-tuned (customised) model** with your training data.
3. How Azure OpenAI Service and Microsoft personnel **analyse** prompts, completions and images for harmful content and for patterns suggesting the use of the service in a manner that violates the Code of Conduct or other applicable product terms.

Customer prompts (inputs) and completions (outputs), embeddings and training data are:

- NOT available to other customers.
- NOT available to OpenAI.
- NOT used to train foundation models without the customer's permission.
- NOT used to improve any Microsoft or third-party products or services.
- NOT used for automatically improving Azure OpenAI models for your use in your resource. (The models are stateless unless you explicitly fine-tune models with your training data.)

Customer fine-tuned Azure OpenAI models are available exclusively for your organisation's use.

## Azure OpenAI | Data flows for inference and training





## Security for Azure OpenAI

As noted in [Part 2 of this paper](#), the Privacy Act requires APP Entities to protect personal information from misuse, interference and loss and from unauthorised access, modification or disclosure.

Security is built-in throughout the development life cycle of all of our enterprise services (including those that include generative AI technology), from inception to deployment.

Azure OpenAI Service is hosted in Azure infrastructure and protected by some of the most comprehensive enterprise compliance and security controls in the industry. These services were built to take advantage of the security and compliance features that are already well-established in Microsoft's hyperscale cloud. This includes prioritisation of reliability, redundancy, availability and scalability, all of which are designed into our cloud services by default.

As generative AI systems are also software systems, all elements of our Security Development Lifecycle (SDL) apply: from threat modelling to static analysis, secure build and operations, use of strong cryptography, identity standards and more.

We've also added new steps to our Security Development Lifecycle to prepare for AI threat vectors, including updating the Threat Modelling SDL requirement to account for AI and machine-learning-specific threats. We put our AI products through AI red teaming to look for vulnerabilities and confirm we have proper mitigation strategies in place.

[Learn more about data, privacy and security for Azure OpenAI Service](#)

## Data residency

In relation to:

- **Azure OpenAI On Your Data feature:** Any data sources you provide to ground the generated results remain stored in the data source and location you designate. No data is copied into Azure OpenAI service.
- **Training data and fine-tuned (customised) LLMs:** These are stored within the same region as Azure OpenAI resource in the customer's Azure tenant.
- **Abuse monitoring for customers who use Azure OpenAI service:** The data store where prompts and completions are stored is logically separated by customer resource (and each request includes the resource ID of the customer's Azure OpenAI resource). A separate data store is located in each region in which Azure OpenAI service is available, and a customer's prompts and generated content are stored in the Azure region where the customer's Azure OpenAI service resource is deployed, within the Azure OpenAI service boundary.



# Part 5:

## Conclusion

Microsoft runs on trust. We are committed to security, privacy and compliance across everything we do, and our approach to generative AI is no different. As an industry leader in the provision of generative AI solutions we are trusted by customers across the world and adhere to the strictest privacy and security standards in the industry. We provide superior products and services to our customers, thereby facilitating continued progress towards their digital transformation goals.

Furthermore, we have been intentional about signalling to our customers our willingness and commitment to get our data protection and privacy settings right to ensure compliance with the Privacy Act. We demonstrate this commitment through our contracts, extensive technical documentation (providing details about our data processes and

activities), and the implementation of technical and organisational safeguards to mitigate residual privacy and security risks. This is backed by consistent engagement with regulatory and industry stakeholders with whom we partner on our journey towards responsibility, accountability and integrity in the delivery of generative AI solutions at scale.

As the regulatory landscape evolves and we innovate to provide new kinds of AI solutions, we are keenly aware that organisations will continue to look to us to help decipher and operationalise the requirements of new and existing data protection frameworks. Microsoft will continue to offer industry-leading tools, transparency resources and support and we look forward to the opportunity to demonstrate our enduring commitment to meeting the needs and demands of our customers in their AI journey.



# Appendix 1:

## Business opportunities arising from generative AI

The availability of generative AI solutions has served as an accelerator to the consideration of generative AI use cases. This Appendix sets out several relevant areas of impact for consideration by businesses.

### AI transformation opportunities

The integration of generative AI into business operations is driven by several key opportunities:

- **Enrich employee experiences:** Automating routine and time-consuming tasks frees up human resources to focus on more strategic initiatives. AI-driven processes reduce human error and increase the precision of outputs, from financial forecasting to legal compliance checks.
- **Improve relevant customer engagement:** By providing personalised experiences and rapid responses to customer inquiries, AI can help improve overall customer satisfaction and loyalty.
- **Reshape business processes:** As businesses grow, AI can easily scale to handle increased data and transaction volumes, ensuring consistent performance without proportional increases in operational costs.
- **Bend the curve on innovation:** AI can be leveraged to facilitate the exploration of new business models and services by identifying trends, predicting market movements and customising offerings.

This introduction sets the stage for a detailed exploration of Generative AI's specific applications within different industries, demonstrating how its capabilities are not just theoretical but have practical and transformative impacts on business operations.

### General use cases for Copilot for Microsoft 365

Copilot for Microsoft 365 is designed to enhance operational efficiencies and decision-making across

a wide range of industries. This section outlines the most popular and universally applicable use cases for Copilot for Microsoft 365, demonstrating its flexibility and the value it adds to any business operation.

- **Automated customer support:** Use advanced virtual assistants and chatbots that manage customer inquiries, provide real-time support and resolve issues autonomously. This reduces response times, increases customer satisfaction and decreases the operational costs associated with maintaining large customer service teams.
- **Document automation and management:** Create, format and manage documents. Copilot for Microsoft 365 can generate reports, draft correspondences and prepare presentations based on user inputs. This enhances productivity and ensures consistency across all business communications, allowing staff to focus on more strategic tasks.
- **Data analysis and insights generation:** Analyse large datasets to identify trends, perform predictive analytics and generate actionable insights, which are crucial for decision-making. This helps businesses make informed decisions based on data-driven insights, optimising operations and improving strategic planning.
- **Workflow and process automation:** Automate repetitive and time-consuming tasks such as data entry, scheduling and process tracking, integrating seamlessly with existing systems to streamline workflows. This increases operational efficiency, reduces human error and frees up employees to focus on higher-value activities.
- **Personalised content and recommendations:** Copilot for Microsoft 365 tailors content and recommendations to individual users based on their behaviours, preferences and past interactions, commonly used in sectors like e-commerce, media and content delivery. This enhances user engagement and satisfaction, leading to increased loyalty and revenue from personalised experiences.

## Department and employee-specific use cases

We have developed the Microsoft Copilot Scenario Library to provide guidance by department and employee-specific scenarios to get inspired, empower your workforce and realise value from your Copilot for Microsoft 365 investment. Find more examples by department and role at the following links:

[Use Cases for Finance Department](#)

[Use Cases for Human Resources Department](#)

[Use Cases for Information Technology](#)

[Use Cases for Marketing Department](#)

[Use Cases for Sales](#)

## Industry-specific use cases

This section explores the specific applications of Copilot for Microsoft 365 in three critical industries: legal, banking and healthcare. By highlighting targeted use cases, we demonstrate Copilot's effectiveness in addressing industry-specific challenges and enhancing core operations.

### 1. Legal industry use cases

- **Contract review and analysis:** Automates the review process by comparing contract clauses against legal standards and previous contracts. This increases efficiency, reduces human error and ensures compliance with legal standards.
- **Litigation support:** Assists in organising and analysing vast amounts of case-related data to support litigation processes. This saves time and enhances the preparation and presentation of legal arguments.

- **Compliance monitoring:** Scans continuously for changes in legislation to help firms remain compliant with all relevant laws. This reduces the risk of legal penalties and enhances the firm's reputation for diligence.

### 2. Banking industry use cases

- **Fraud detection:** Utilises AI to monitor transactions in real time and identify patterns indicative of fraudulent activity. This minimises financial losses and protects customer trust.
- **Risk assessment:** Analyses customer data to predict and mitigate potential risks in lending and investments. This enhances the bank's ability to manage and mitigate risk effectively.
- **Regulatory compliance tracking:** Keeps track of all regulatory requirements and ensures the bank complies with financial regulations. This avoids legal penalties and maintains operational integrity.

### 3. Healthcare industry use cases

- **Patient data management:** Manages and secures vast amounts of patient data, facilitating easy access for healthcare providers. This improves the efficiency and confidentiality of patient care.
- **Diagnostic assistance:** Provides support in diagnosing diseases by analysing patient data and medical imagery. This enhances the accuracy of diagnoses and the effectiveness of treatment plans.
- **Remote patient monitoring:** Monitors patients remotely using data from wearable devices, providing real-time health updates to providers. This reduces hospital readmissions and allows for proactive healthcare management.



# Appendix 2:

## Opportunities in the public sector

The availability of generative AI solutions has also served as an accelerator to the consideration of public sector generative AI use cases. This Appendix sets out several relevant areas of impact for consideration by public sector organisations.

- **Citizen services:** Generative AI can help governments and public sector organisations provide enhanced service experiences that make government more accessible and less time-consuming by acting as an “Information Assistant” – answering frequently asked questions, recommending services based on inputs, and even handling simple transactions


Many governments have already experimented with chatbots to help answer simple questions about COVID-19 vaccinations, provide support during tax time and offer answers to common inquiries. Generative AI helps chatbots handle more open domain questions over more sophisticated and complex materials, including rapid responses to a broader range of questions at any time from anywhere, increasing accessibility for citizens while simultaneously increasing government efficiency and reducing administrative burden.

Citizens can even provide a narrative of their current circumstances and discover service options they previously did not know existed. These tools also free up public sector workers to focus on strategic projects instead of being tied down to mundane, repetitive functions such as responding to common questions.

- **Internal efficiency:** Government can be complex even for government employees! Providing public sector workers with the capacity to intuitively search and interact via chat with intranets and public sector materials in an automated fashion eases onboarding of new employees, increases efficiency between silos and departments, and minimises administrative burdens. This capability lets public sector staff focus on their mission priorities, reducing burnout and allowing them to do more with less.

- **Deep data:** Large language models (LLMs) can tackle the intersection between vast troves of data, which may have been previously analysed separately and manually. Simple prompts to the AI can yield both typical and unexpected connections between topics and domains that can help to spur the analytic process.

Insightful and succinct summaries of vast amounts of media coverage or public feedback can be generated in seconds. Generative AI helps to objectively challenge conventional wisdom – raising new angles, questions, or counterarguments that may have been implicitly screened by the bias of the author. This approach ultimately yields stronger and more comprehensive output.



- **Creative aid:** No more writer's block! Generative AI can provide helpful initial drafts of abstracts, outlines, speeches, simple correspondence, memos, frequently asked questions, white papers and citizen guides. While official communications should always require a human in the loop to verify accuracy, apply human "voice," and ensure that the information is complete and not misleading, generative AI as a creative writing aid can accelerate the process dramatically and help light the creative spark while reducing time-to-completion for common writing tasks.
- **Enhanced security:** Generative AI can support cybersecurity teams and protect your organisation from threats. Generative AI models can be trained to review applications and code for weaknesses using a dynamic model that evolves to keep pace with threat. This can also be used to review and deploy new code more quickly by automating vulnerability detection which will help security professional scale workloads by freeing them up from lower-value tasks.

By improving citizen services, increasing efficiency, better managing and analysing data and serving as a creative aid, generative AI can help to create a more effective, inclusive and responsive government.

Generative AI can also help create a more efficient, productive and rewarding work environment for public sector employees. Governments should carefully consider the implications of using AI in their operations and take appropriate measures to ensure the technology is used ethically and responsibly. Now is the time for public sector organisations to begin leveraging and adopting generative AI capabilities and they can and should do so from a position of engagement and experimentation.



# Appendix 3:

## Frequently asked questions (FAQs)

### How is my organisation's data protected when I use Microsoft's generative AI services?

Microsoft runs on trust. We are committed to security, privacy and compliance across everything we do and our approach to generative AI is no different.

Privacy is built into our approach to Responsible AI and we will continue to uphold our core values of privacy, security, fairness, accountability, transparency, reliability, inclusiveness and safety in our AI products and solutions.

In [Part 2 of this paper](#), we outline seven commitments that demonstrate our continued commitment to protecting our customers' data when they use our Generative AI services:

- We will keep your organisation's data private.
- You are in control of your organisation's data.
- Your access control and enterprise policies are maintained.
- Your organisation's data is not shared without your permission.
- Your organisation's data privacy and security are protected by design.
- Your organisation's data is not used to train foundation models without your permission.
- Our products and solutions continue to comply with global data protection regulations.

### What is generative AI and what are the different types of AI models Microsoft uses?

Generative AI is a type of artificial intelligence that can create new things, like pictures, text, or speech, that are similar to examples it has seen before. It does this by learning from a set of examples, figuring out the patterns and rules that make them similar and then using those patterns and rules to make new examples that are similar to the ones it learned from. It's different from other types of AI because it can create new

things, instead of just recognising or classifying things it has seen before.

Microsoft's Azure OpenAI Service and Copilot for Microsoft 365 allow customers to leverage OpenAI's models, including GPT-3, GPT-4, and Codex in the Microsoft environment. These models are commonly referred to as "foundation models," which are generally understood to be large-scale AI models that are trained on vast quantities of primarily unlabelled data at scale (usually by self-supervised learning) and can be adapted with minimal fine-tuning for a range of different downstream tasks.

### What are the differences between procuring cloud and generative AI services from a Privacy Act perspective?

The obligations under the Privacy Act that apply to procuring and using cloud computing services are the same as those that apply to procuring and using generative AI services. Best practice under the Privacy Act suggests a risk-based approach should be taken towards the implementation and use of any new technologies.

The level of risk involved will depend on the nature, scope, content and purpose for which personal information will be used. When procuring and using cloud services and/or generative AI services, an organisation will need to consider what technical and organisational measures are in place to protect and safeguard the use of personal information and ensure that it has appropriate contractual commitments and operational processes so it can comply with its obligations under the Privacy Act.

[Find out more about how Microsoft can assist customers in undertaking this assessment when they are looking to use Copilot for Microsoft 365 and/or Azure OpenAI Service in Part 2 of this paper.](#)



## What are the key obligations of the Privacy Act that apply to procuring and using generative AI systems?

The obligations under the Privacy Act and the APPs will apply whenever a generative AI system uses or otherwise processes personal information.

Key obligations organisations should consider when procuring and/or implementing generative AI systems include:

- Consider whether you need to update your privacy notices to reflect any new processing activity or to clarify activities (APP 1 and APP 5).
- Consider what information will be used (or generated) by the generative AI system and, where that information is personal information, you have steps in place to ensure that any personal information you collect, use or disclose is done in accordance with the Privacy Act (APP 3, 6 and 9).
- Ensure you have processes in place to enable you to comply with requests from individuals to exercise their rights under the Privacy Act (APP 12 and APP 13).
- Consider whether you have to, or should, conduct a privacy impact assessment (e.g. Australian Government Agencies Privacy Code).

[Learn more about how Microsoft assists customers in meeting these obligations in Part 2 of this paper.](#)

## How does Microsoft comply with applicable law?

Microsoft's AI products and solutions are designed and built for compliance with applicable data protection and privacy laws today, including the Privacy Act.

Microsoft's approach to protecting privacy in AI is underpinned by a commitment to compliance with existing and emerging regulatory and legal obligations globally. We will continue to support meaningful privacy and AI regulation, and believe the best way to make rapid progress on needed guardrails for AI is to make the best use of existing legal protections, approaches and regulatory tools that could be applied to protecting privacy and safety in these systems today.

## Does Microsoft share Customer Data with OpenAI/ChatGPT?

No. Your organisation's Customer Data, including prompts (inputs) and completions (outputs), your embeddings and any training data you might provide to the Microsoft Online Services are not available to OpenAI.

Azure OpenAI Service is fully controlled by Microsoft; Microsoft hosts the OpenAI models in Microsoft's Azure environment and Azure OpenAI Service does not interact with any services operated by OpenAI (e.g., ChatGPT or the OpenAI API). OpenAI is not a sub-processor to Microsoft.

[Learn more about the underlying OpenAI models that power Azure OpenAI Service.](#)

## Can I share confidential information with Microsoft's generative AI services?

Yes. When using Azure OpenAI or Copilot for Microsoft 365, customers may confidently share their confidential information. The foundation models that are accessed via Azure OpenAI Service and Copilot for Microsoft 365 do not use Customer Data for training without permission. These foundation models are stateless and do not store any data, including prompts that a customer inputs and completions that the model outputs. Customers can also trust that their confidential information will not be transmitted to other customers.

## How does Microsoft protect security in this new era of AI?

Security is built-in throughout the development life cycle of all of our enterprise services (including those that include generative AI technology), from inception to deployment.

Azure OpenAI Service and Copilot for Microsoft 365 are hosted in Azure infrastructure and protected by some of the most comprehensive enterprise compliance and security controls in the industry. These services were built to take advantage of the security and compliance features that are already well-established in Microsoft's hyperscale cloud. This includes prioritisation of reliability, redundancy, availability and scalability, all of which are designed into our cloud services by default.

Because generative AI systems are also software systems, all elements of our Security Development Lifecycle (SDL) apply: from threat modelling to static analysis, secure build and operations, use of strong cryptography, identity standards and more.

We've also added new steps to our Security Development Lifecycle to prepare for AI threat vectors, including updating the Threat Modelling SDL requirement to account for AI and machine-learning-specific threats. We put our AI products through AI red teaming to look for vulnerabilities and ensure we have proper mitigation strategies in place.

Learn more about Security for Copilot for Microsoft 365 in [Part 3 of this paper](#), and about Security for Azure OpenAI Service in [Part 4 of this paper](#).

## Are data transfers to countries outside of Australia allowed under the Privacy Act?

Yes – under APP 8, APP Entities can generally disclose personal information to overseas recipients provided that the APP Entity ensures that the overseas recipient will handle an individual's personal information in accordance with the APPs.

If you provision your Microsoft 365 tenant in Australia:

- Any stored content of interactions with Copilot for Microsoft 365 and Azure OpenAI Service is stored at rest in Australia.
- As outlined in [Microsoft's Data Protection Addendum](#) (which forms part of your agreement with Microsoft), Microsoft will comply with all laws and regulations applicable to its providing the Products and Services, including the Privacy Act.

[Find out more about how Microsoft approaches data transfers to third countries in Part 2 of this paper.](#)

## Where will my data be stored and processed?

Your data residency choices will be respected when you use Microsoft's generative AI products and services that offer local storage and/or processing capabilities.

## How can customers set up their procurement and use of generative AI services to be compliant with the Privacy Act?

The Privacy Act requires entities that collect, hold, use and disclose personal information to consider data protection issues at every stage of their information-handling activities, from the initial design (including during the procurement phase) to final implementation.

The risks associated with the use of generative AI will vary depending on the specific use case and related nature, sensitivity and volume of personal information that will be used in connection with that use case.

One way you can demonstrate compliance with the Privacy Act and the APPs is to complete a privacy impact assessment (PIA) relating to specific use cases for generative AI solutions. A PIA helps organisations identify and reduce the data protection risks. A PIA is legally required for Australian Federal Government agencies where the processing activity is of high privacy risk (i.e. it involves new or changed ways of handling personal information that are likely to have a significant impact on the privacy of individuals). Even if it is not legally required, a PIA is good practice and can help you work through the specific data protection risks associated with how you wish to implement generative AI for a specific use case.

[Find out more about PIAs in Part 2 of this paper.](#)

## Can customers comply with the Privacy Act when using a public cloud to use generative AI services?

Microsoft's public cloud services have been developed to ensure they can be used by customers in compliance with the Privacy Act (and many customers already make use of these services). The information set out in this paper and contained in the [Product Terms](#) and [Data Protection Addendum](#) can be used by you to undertake an appropriate risk-based assessment of any proposed use of Copilot for Microsoft 365 and Azure OpenAI Service so as to demonstrate compliance with the relevant requirements of the Privacy Act.

## How can organisations comply with their transparency obligations under the Privacy Act when deploying AI technologies?

APP 1 and APP 5 of the Privacy Act require organisations to provide individuals with certain key information about how their personal information will be used. This information is often provided in the form of privacy notices. If you deploy a new technology

(such as Copilot for Microsoft 365 or Azure OpenAI Service) and intend to use such technology in a way that is not reflected in your existing privacy notices, then you will need to update these notices to reflect these new processing activities.

The information set out in this paper is intended to assist you to understand how Copilot for Microsoft 365 and Azure OpenAI Service use data and to determine what information needs to be communicated to individuals.





# Appendix 3:

## Additional resources

Microsoft is committed to providing our customers with clear information about how we use and share data and choices they have in managing their data. This Appendix sets out additional resources which you can reference to supplement and expand on the information set out in this paper.

### Responsible AI

- [Empowering responsible AI practices](#)
- [Governing AI: A Blueprint for the Future](#)
- [Microsoft's principles and approach to Responsible AI](#)
- [Microsoft Responsible AI Standard](#)

### Microsoft's Customer Commitments

- [AI Assurance Program and AI Customer Commitments](#)
- [Customer Copyright Commitment](#)
- [Protecting the data of our commercial and public sector customers in the AI era](#)
- [FAQ: Protecting the Data of our Commercial and Public Sector Customers in the AI Era](#)

### Understanding Generative AI

- [The underlying LLMs that power Microsoft's generative AI solutions](#)
- [The art and science of prompting \(the ingredients of a prompt\)](#)
- [Prompting do's and don'ts](#)

### Data Protection Addendum and Product Terms

- [Data Protection Addendum](#)
- [Microsoft Product Terms](#)

### Copilot for Microsoft 365

- [Copilot for Microsoft 365](#)
- [Copilot Lab](#)
- [Copilot for Microsoft 365 Documentation](#)
- [Data, Privacy and Security for Copilot for Microsoft 365](#)
- [FAQs for Copilot data security and privacy](#)
- [Microsoft 365 isolation controls](#)
- [Encryption in the Microsoft Cloud](#)

### Azure OpenAI Service

- [Azure OpenAI Service - Documentation, quickstarts and API reference guides](#)
- [Configure usage rights for Azure Information Protection](#)
- [Data, privacy and security for Azure OpenAI Service](#)
- [Prompt Engineering](#)
- [Azure OpenAI On Your Data](#)
- [Azure OpenAI fine tuning](#)
- [Content filtering](#)
- [Abuse monitoring](#)
- [Enterprise security for Azure Machine Learning](#)



Disclaimer:

© Microsoft Corporation 2024. All rights reserved. Microsoft, the Microsoft logo, Azure, Bing, BitLocker, Dynamics 365, Excel, Exchange Designs, Microsoft 365, Microsoft Copilot Designs, Microsoft Entra, Microsoft Purview, OneDrive, Outlook, Power Platform and SharePoint are all trademarks or registered trademarks of Microsoft in the United States and/or other countries.

Microsoft makes no warranties, express or implied, in this document. This document is for informational purposes only and provided "as-is." The document may not contain the most up to date information or guidance. Information and views expressed in this document including references to any of our terms, URL and other references may change without notice. You bear the risk of using it. This document is not legal or regulatory advice and does not constitute any warranty or contractual commitment on the part of Microsoft. You should seek independent legal advice on your legal and regulatory obligations. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.