

# Anthropic Claude on Microsoft and Anthropic Platforms

Data residency and compliance briefing for Australian client engagements

---

**Prepared by FutureAbility**

Updated 11 May 2026

Classification: Client-presentable (subject to NDA)

## Executive Summary

This briefing assesses four deployment paths for Anthropic Claude that an Australian organisation might consider, and compares them against Australian data residency, privacy and regulatory requirements:

**Section 1** – Claude inside Microsoft 365 Copilot

**Section 2** – Claude inside Microsoft Copilot Studio and Power Platform

**Section 3** – Claude for Work (Anthropic's Claude.ai Teams and Enterprise SKUs)

**Section 4** – Anthropic Claude models in Azure AI Foundry

All four paths process data primarily in United States infrastructure. None offers an Australian-region inference guarantee as of May 2026. Microsoft now lists Anthropic as a subprocessor (January 2026) for M365 Copilot, Copilot Studio and Foundry, but Anthropic models are explicitly excluded from Microsoft's EU Data Boundary and in-country processing commitments.

**Bottom line:** for unregulated commercial workloads, all four paths are workable with appropriate disclosure. For APRA-regulated, IRAP-required or health (PHI) workloads, none of the four is clean today; engagements should be conditional on a formal third-party risk assessment, and in many cases the best answer will be Azure OpenAI in an Australian region instead.

# Section 1 – Microsoft 365 Copilot

## 1.1 What this is

Microsoft 365 Copilot is the AI assistant embedded across the M365 suite. Anthropic Claude is used as the default or alternative model for several Copilot experiences:

- **Copilot in Word, Excel, and PowerPoint** – Anthropic models power certain Copilot experiences. A separate admin toggle controls this independently.
- **Researcher** – Users can choose between Claude and OpenAI models for multi-document analysis.
- **Agent Mode in Excel** – Claude powers complex reasoning tasks.
- **Copilot Cowork** – A cloud-based agentic AI service that plans and executes multi-step tasks across Outlook, Teams, Excel, PowerPoint, and SharePoint. Critically, Cowork executes actions using the user's identity, not a dedicated AI agent identity.

## 1.2 Where does Claude actually run?

When Anthropic models are used in M365 Copilot, data processing occurs outside of the Microsoft EU Data Boundary. The same exclusion applies to in-country processing commitments, including Australia.

Microsoft made in-country Copilot processing available for Australia in late 2025, meaning standard Copilot interactions (using Azure OpenAI models) can stay within Australian borders. However, Anthropic models are explicitly excluded from those in-country commitments. Any request routed to Claude – including Cowork task execution – falls outside the Australian data residency guarantee.

## 1.3 Admin controls and default state

There are two independent admin settings that control Anthropic model use in M365 Copilot:

- **Global Anthropic subprocessor toggle** – Controls whether Anthropic models are available across M365 Copilot, Researcher, Copilot Studio, Power Platform, and Agent Mode in Excel.
- **Per-app Anthropic toggle** – A separate setting controlling Anthropic use specifically in Word, Excel, and PowerPoint. Changing one toggle does not affect the other.

**Default state for Australian tenants:** Global toggle is ON by default since 7 January 2026. Most organisations will not know the toggle is on.

## 1.4 What you lose if you disable Anthropic

- Researcher agent model choice (Claude vs. OpenAI)
- Copilot Studio agents powered by Claude
- Word, Excel, and PowerPoint agents (Frontier features)
- Agent Mode in Excel and related preview features

- Copilot Cowork (which relies on Claude for complex reasoning)

This is a material trade-off for organisations that want to use Frontier features but have data residency concerns.

## 1.5 Training use

Microsoft's Enterprise Data Protection documentation explicitly states that data is not used to train foundation LLMs, and this applies even when using Anthropic models. The commitment is from Microsoft's side – it means Microsoft does not provide data to Anthropic for training. Anthropic's Zero Data Retention (ZDR) is not available for M365 Copilot.

## 1.6 Retention

**Microsoft:** User prompts and Copilot responses are stored within the M365 service boundary, encrypted, and accessible via Content Search, Microsoft Purview, and eDiscovery.

**Anthropic:** Under standard terms, data is retained for abuse detection. If flagged for a policy violation, retention can extend up to 2 years. ZDR is not available for M365 Copilot.

## 1.7 Australian compliance verdict

- **Privacy Act 1988** – Workable with APP 8 cross-border disclosure in privacy notices and contractual safeguards via the Microsoft DPA.
- **APRA CPS 230** – Fourth-party risk assessment required. Anthropic is enabled by default – regulated entities may already be exposed.
- **IRAP** – Not assessed for the Anthropic layer. Government clients should default to Azure OpenAI in Australia East.
- **Health PHI** – Anthropic's HIPAA BAA does not extend to M365 Copilot. Not suitable for PHI.

## Section 2 – Copilot Studio and Power Platform

### 2.1 Where does Claude actually run?

When a Copilot Studio agent invokes Claude, the request does not run entirely within Microsoft's infrastructure. Microsoft onboarded Anthropic as a subprocessor in January 2026. Anthropic models in Microsoft offerings are explicitly excluded from the EU Data Boundary and in-country processing commitments.

**Critical ambiguity:** Microsoft does not publicly state whether the Copilot Studio integration calls Anthropic's direct API, AWS Bedrock or a Microsoft-hosted instance.

### 2.2 Training use and retention

Microsoft commits that Copilot Studio prompts and responses are not used to train Microsoft foundation models. Anthropic's ZDR is not available for third-party integrations. Copilot Studio transcripts are stored in Dataverse for 30 days by default (configurable to 0). Anthropic retains inputs and outputs for up to 30 days for abuse detection, extending up to 2 years if flagged.

### 2.3 Australian compliance verdict

- **APRA CPS 230** – Requires formal board-approved fourth-party risk assessment.
- **IRAP** – Azure services in Australian regions are IRAP-assessed; Anthropic models are not.
- **Health PHI** – Anthropic's HIPAA BAA does not extend to Copilot Studio.
- **Privacy Act 1988** – APP 8 cross-border disclosure required. Must be disclosed in privacy notices.

## Section 3 – Claude for Work (Teams and Enterprise)

### 3.1 What this is

Claude for Work refers to Anthropic's commercial Claude.ai SKUs – Claude Team and Claude Enterprise – sold directly by Anthropic and accessed via Claude.ai (web, desktop, mobile). Organisations buy seats; users sign in with their work identity.

### 3.2 Data residency

Anthropic's Privacy Center states that traffic may be routed via servers in the US, Europe, Asia and Australia, but data is stored in the US by default. There is no published Australian data residency option as of May 2026.

### 3.3 Training use

Anthropic confirms commercial products (Team, Enterprise, API, Claude Gov) do not use customer conversations for model training by default. Thumbs up/down feedback can be retained for up to 5 years if submitted; admins can disable the feedback button.

### 3.4 Retention and security

Conversations persist until deleted by default. Enterprise plan offers custom data retention controls with a minimum 30-day retention period. Encryption at rest uses AES-256; encryption in transit uses TLS. Customer-managed keys are not currently standard.

### 3.5 Certifications

- SOC 2 Type II
- ISO/IEC 27001:2022
- ISO/IEC 42001:2023 (AI management systems)
- HIPAA-ready with BAA (API tier)

No published IRAP assessment for Claude.ai.

### 3.6 Australian compliance verdict

- **Privacy Act 1988** – Workable with a signed DPA and APP 8 cross-border disclosure.
- **APRA CPS 230** – No published assessment. Requires own fourth-party risk assessment.
- **IRAP** – Not certified. Not suitable for PROTECTED data.
- **Health PHI** – HIPAA-ready at API tier. Treat AU health info as overseas disclosure.

## Section 4 – Anthropic Claude in Azure AI Foundry

### 4.1 What this is

Anthropic Claude models are available in Azure AI Foundry as a Models-as-a-Service (MaaS) catalogue offering.

### 4.2 Hosting model – the critical finding

Anthropic's own documentation states clearly: "In this preview platform integration, Claude models run on Anthropic's infrastructure." You call an Azure endpoint, Microsoft authenticates and meters the call, but the actual model inference happens on Anthropic infrastructure outside the Azure trust boundary.

### 4.3 Data residency and retention

Australia East is not listed as a supported region. Microsoft's privacy notice states prompts and outputs may be processed anywhere globally. Anthropic retains inputs and outputs for up to 30 days for abuse detection. ZDR is not available (it remains an Azure first-party OpenAI feature only).

### 4.4 Australian compliance verdict

- **Privacy Act 1988** – Treat as overseas disclosure under APP 8.
- **APRA CPS 230** – Fourth-party assessment required.
- **IRAP** – Not assessed. Use Azure OpenAI in AU East for PROTECTED workloads.
- **Health PHI** – Not suitable by default.

## Comparative Analysis – The Four Paths

This matrix compares the four deployment paths against the dimensions an Australian client will care about.

| Dimension                    | M365 Copilot                               | Copilot Studio                              | Claude for Work                                 | Azure AI Foundry                         |
|------------------------------|--------------------------------------------|---------------------------------------------|-------------------------------------------------|------------------------------------------|
| <b>Inference location</b>    | Anthropic infra (US default)               | Anthropic infra (US default)                | Anthropic infra; storage US default             | Anthropic infra (confirmed)              |
| <b>AU region available</b>   | No                                         | No                                          | No (traffic may route via AU; storage US)       | No (AU East not listed)                  |
| <b>Training use</b>          | No (EDP commitment). Abuse monitoring only | No (Microsoft). Anthropic side not explicit | No by default. Feedback if submitted            | No (contractually committed)             |
| <b>Zero Data Retention</b>   | Not available                              | Not available                               | Available for Claude Code Enterprise on request | Not available (OpenAI-only)              |
| <b>Default retention</b>     | Purview-governed; Anthropic 30 days        | Dataverse 30 days; Anthropic 30 days        | Indefinite; Enterprise custom retention         | Anthropic 30 days, then auto-delete      |
| <b>Anthropic toggle (AU)</b> | ON by default. Two separate toggles        | Linked to global toggle                     | N/A (purchased directly)                        | N/A (provisioned in Foundry)             |
| <b>Privacy Act (APP 8)</b>   | Cross-border disclosure required           | Cross-border disclosure required            | Cross-border disclosure required                | Cross-border disclosure required         |
| <b>APRA CPS 230</b>          | Fourth-party assessment + board approval   | Fourth-party assessment + board approval    | Fourth-party assessment + board approval        | Fourth-party assessment + board approval |
| <b>IRAP (PROTECTED)</b>      | Not assessed for Anthropic                 | Not assessed for Anthropic                  | Not certified                                   | Not assessed; use Azure OpenAI           |
| <b>Health PHI</b>            | BAA does not flow through                  | BAA does not flow through                   | API tier has BAA                                | Not suitable by default                  |

## Combined Risk Summary

| Risk                                                     | Severity | Mitigation                                                                                 | Owner           |
|----------------------------------------------------------|----------|--------------------------------------------------------------------------------------------|-----------------|
| Data leaves Australia (all four paths)                   | High     | Disclose in privacy notice; APP 8 contractual safeguards; consider Azure OpenAI in AU East | Client / FAb    |
| No published AU inference region for Anthropic           | High     | Wait for AU support, or fall back to Azure OpenAI / native Foundry models                  | Microsoft       |
| Anthropic enabled by default in M365 Copilot (AU)        | High     | Audit both global and per-app toggles immediately. Scope to Entra security groups          | Client admin    |
| APRA CPS 230 fourth-party risk                           | High     | Formal board-approved assessment of Anthropic and AWS/Azure as material service providers  | Client          |
| IRAP not assessed for Anthropic layer                    | High     | Do not use Anthropic for PROTECTED data; use Azure OpenAI in AU East                       | ASD / Microsoft |
| HIPAA BAA does not flow through M365/Studio/Foundry      | High     | Use Anthropic direct API with BAA for PHI                                                  | Client          |
| Anthropic training use under Copilot Studio not explicit | Medium   | Get written confirmation from Microsoft                                                    | Microsoft       |
| Copilot Cowork identity governance                       | Medium   | Map audit trail implications. No dedicated AI agent identity                               | Client          |
| Indefinite Claude.ai retention by default                | Medium   | Enterprise plan + custom retention policy                                                  | Client          |
| ZDR not available for any Microsoft path                 | Medium   | Use Anthropic direct API + ZDR for sensitive workloads                                     | Client          |



# Decision Guide for Prospect Conversations

Use the workload's regulatory profile to choose the path:

**Unregulated commercial workload, AU residency a preference but not a hard rule** – Any of the four paths is workable. Choose by developer/user experience. Note that M365 Copilot's Anthropic toggle is on by default – raise this proactively.

**AU residency is a hard contractual rule** – None of the four is suitable today. Recommend Azure OpenAI in Australia East with a roadmap watch on Anthropic AU regions. Disable the Anthropic toggle in M365 Copilot.

**APRA-regulated entity (banking/insurance)** – Any path is conditional on a CPS 230 fourth-party risk assessment with board sign-off. Lead with Azure OpenAI in AU East if residency is the gating concern.

**Australian Government – PROTECTED data** – Do not use any of the four Anthropic paths for PROTECTED workloads. Use Azure OpenAI in AU East.

**Health PHI** – Only Anthropic's direct API with a signed BAA gives a defensible posture.

## Questions to Put to Microsoft and Anthropic

1. Can Anthropic models be constrained to a specific region (US-only, AU-only) when called from M365 Copilot, Copilot Studio or Foundry?
2. Can Microsoft cite the contractual clause that prevents Anthropic from using customer data for training under M365 Copilot, Copilot Studio and Foundry?
3. Is AWS Bedrock anywhere in the M365 Copilot or Copilot Studio inference chain?
4. Does Microsoft offer a third-party risk addendum for CPS 230-regulated clients?
5. When will Anthropic models in Microsoft services complete IRAP assessment for PROTECTED?
6. When will an Australian-region inference option be available for Anthropic?
7. Is there any roadmap to extend Anthropic ZDR to the M365/Studio/Foundry channels?
8. Does data fetched via MCP follow Claude.ai's US storage default?
9. Can Microsoft expose a per-environment Anthropic retention setting?
10. What audit trail and governance controls exist to distinguish user actions from Cowork actions?
11. What data does Anthropic retain for abuse detection, and can this be reduced for enterprise customers?

# Citations

## Microsoft Sources

- Anthropic as a subprocessor for Microsoft Online Services  
<https://learn.microsoft.com/en-us/copilot/microsoft-365/connect-to-ai-subprocessor>
- Move data across regions for Copilots, AI agents and generative AI features  
<https://learn.microsoft.com/en-us/power-platform/admin/geographical-availability-copilot>
- Copilot in M365 apps with Anthropic models  
<https://learn.microsoft.com/en-us/microsoft-365/copilot/copilot-anthropic-apps>
- Control how transcripts are retained and accessed – Copilot Studio  
<https://learn.microsoft.com/en-us/microsoft-copilot-studio/admin-transcript-controls>
- Enterprise data protection in Microsoft 365 Copilot  
<https://learn.microsoft.com/en-us/copilot/microsoft-365/enterprise-data-protection>
- Data, Privacy, and Security for Microsoft 365 Copilot  
<https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>
- Australia IRAP – Azure Compliance  
<https://learn.microsoft.com/en-us/azure/compliance/offerings/offering-australia-irap>
- Data privacy for Anthropic Claude in Azure Foundry  
<https://learn.microsoft.com/en-us/azure/foundry/responsible-ai/claude-models/data-privacy>
- Foundry models from partners and community  
<https://learn.microsoft.com/en-us/azure/foundry/foundry-models/concepts/models-from-partners>
- Are Anthropic models deployed under Microsoft? (Q&A)  
<https://learn.microsoft.com/en-gb/answers/questions/5843526/are-the-anthropic-models-deployed-under-microsoft>
- Microsoft offers in-country data processing for M365 Copilot  
<https://www.microsoft.com/en-us/microsoft-365/blog/2025/11/04/microsoft-offers-in-country-data-processing-to-15-countries-to-strengthen-sovereign-controls-for-microsoft-365-copilot/>
- Claude is now inside your M365 tenant, mind the data residency gap (IDM Magazine)  
<https://idm.net.au/article/0015519-claude-now-inside-your-m365-tenant-mind-data-residency-gap>
- Anthropic Models On by Default: Admin Action Plan (2toLead)  
<https://www.2tolead.com/insights/anthropic-models-on-default-copilot-admin-action-plan-and-risks>

## Anthropic Sources

- Anthropic Trust Center  
<https://trust.anthropic.com/>
- Subprocessors list  
<https://trust.anthropic.com/subprocessors>
- Where are your servers located?  
<https://privacy.claude.com/en/articles/7996890-where-are-your-servers-located-do-you-host-your-models-on-eu-servers>
- Is my data used for model training?  
<https://privacy.claude.com/en/articles/7996868-is-my-data-used-for-model-training>
- How long do you store my data?  
<https://privacy.claude.com/en/articles/10023548-how-long-do-you-store-my-data>
- Custom data retention controls for Enterprise  
<https://support.claude.com/en/articles/10440198-custom-data-retention-controls-for-enterprise-plans>
- What is the Enterprise plan?  
<https://support.claude.com/en/articles/9797531-what-is-the-enterprise-plan>
- Access audit logs  
<https://support.claude.com/en/articles/9970975-access-audit-logs>
- What certifications has Anthropic obtained?  
<https://support.claude.com/en/articles/10015870-what-certifications-has-anthropic-obtained>

- How do I view and sign your DPA?

<https://privacy.claude.com/en/articles/7996862-how-do-i-view-and-sign-your-data-processing-addendum-dpa>

- Claude Code – data usage

<https://code.claude.com/docs/en/data-usage>

- Claude in Microsoft Foundry

<https://platform.claude.com/docs/en/build-with-claude/claude-in-microsoft-foundry>

- Claude on Amazon Bedrock

<https://platform.claude.com/docs/en/build-with-claude/claude-in-amazon-bedrock>

- Data residency – Claude API

<https://platform.claude.com/docs/en/build-with-claude/data-residency>

- API and data retention

<https://platform.claude.com/docs/en/build-with-claude/api-and-data-retention>

- Privacy Act 1988 (Cth), section 16A (Australian Privacy Principle 1.2)

---

Document prepared: 15 April 2026 | Last updated: 11 May 2026

Scope: Australian client engagement risk assessment – M365 Copilot, Copilot Studio, Claude for Work, Azure AI Foundry

Classification: Client-presentable (subject to NDA or engagement agreement) | Expiry: 30 June 2026