

Platform	Where data is typically processed/stored	Does the provider use your data for model training?	What level/account matters
Microsoft Copilot (M365 / Enterprise)	For Microsoft 365 Copilot, chats and related Copilot data are stored in the user's Microsoft 365 region. Australian tenants can have this data stored in Australia when the tenant and user location settings are configured correctly. Processing controls are improving, but storage location and processing location should still be checked separately.	Microsoft says prompts, responses and Microsoft Graph business data are not used to train foundation AI models.	Use Microsoft 365 Copilot in a business/enterprise tenant. Do not treat free/consumer Copilot as the same product.
ChatGPT Enterprise / Team	OpenAI offers data residency for eligible ChatGPT Enterprise and Edu customers, and some API customers. This covers supported customer content stored at rest. Inference/processing residency is available only for eligible Enterprise/Edu customers with data residency enabled in supported regions.	OpenAI says business data from ChatGPT Enterprise, Team, Edu and API is not used to train models by default.	ChatGPT Enterprise has the strongest controls. Team is better than consumer ChatGPT, but it does not provide the same governance and residency controls as Enterprise..
Claude (Anthropic Enterprise)	Claude data residency depends on how Claude is accessed. Anthropic only offers regional data residency and processing options for some enterprise/API deployments, including via AWS, Google Cloud Vertex AI and Microsoft Foundry. Australian residency should be confirmed in actual contract/deployment path.	Anthropic says business/API data is not used to train models unless the customer agrees. Consumer Claude is different and may use chats for training unless users opt out.	Use Claude Enterprise/API through an approved enterprise deployment. Do not use personal Claude accounts for client information.
Gemini (Google Workspace)	Gemini in Google Workspace follows Google Workspace enterprise controls. Google supports data-region controls for some Gemini features, but availability depends on Workspace edition and admin configuration. Processing may still depend on the specific Gemini product and region settings.	Google says Workspace content is not human reviewed or used to train generative AI models outside the customer's domain without permission.	Use Gemini for Google Workspace under a managed business account. Consumer Gemini should be treated differently.
NotebookLM	NotebookLM handling depends on account type. NotebookLM Enterprise runs inside Google Cloud controls and Google says data stays within the customer's Google Cloud project. Standard/personal NotebookLM does not provide the same enterprise controls.	In managed Workspace/Enterprise use, Google's Workspace/Cloud protections apply. Personal use should not be assumed to have the same protections.	Best used as NotebookLM Enterprise or within a managed Google Workspace business environment, not personal Gmail.
Perplexity	Perplexity is mainly a research/search tool. Its API has a zero-data-retention policy, but the general consumer product should not be assumed to provide Australian residency or broker-grade client data controls.	Perplexity says its Sonar API does not retain data or use customer data for training. Consumer and enterprise app settings may differ.	Use for general research only unless your organisation has approved the specific enterprise/API arrangement. Avoid identifiable client information.
Grok (xAI)	Grok is primarily US/global infrastructure. xAI has launched Business/Enterprise offerings, but Australian data residency is not clearly established publicly.	xAI says Grok Business/Enterprise data is not used for training. Consumer/X-related usage is different and has higher privacy risk.	Treat consumer Grok as unsuitable for client data. Enterprise may be considered only after legal/security review.