

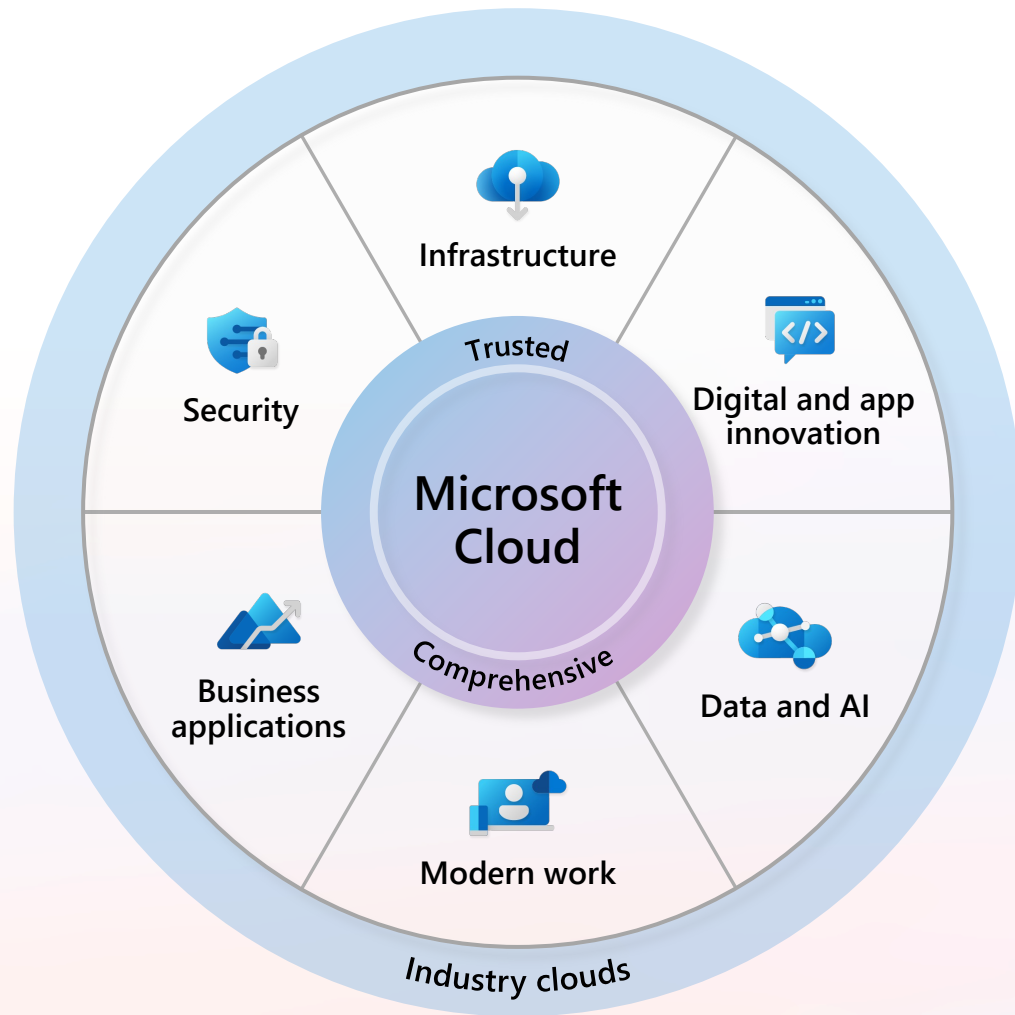
Copilot for Microsoft 365 Technical Overview



Table of contents

- What is Copilot for Microsoft 365?
- How Copilot for Microsoft 365 works
- Privacy
- Data residency and storage
- Compliance
- Security
- How to prepare for Copilot for Microsoft 365
- Technical questions
- Responsible AI
- Glossary

Select the underlined text to navigate to that section.



Microsoft Cloud

Radically accelerate your organization's productivity with built-in AI

Differentiate from the competition with intelligent apps using AI

Microsoft Cloud

Runs on trust

Your data is your data

Your data from any fine-tuning is not used to train the foundation AI models

Your data is protected by the most comprehensive enterprise compliance and security controls

What is Copilot for Microsoft 365?



Copilot for Microsoft 365

Natural Language



Large Language
Models

+



Microsoft Graph
- Your Data -

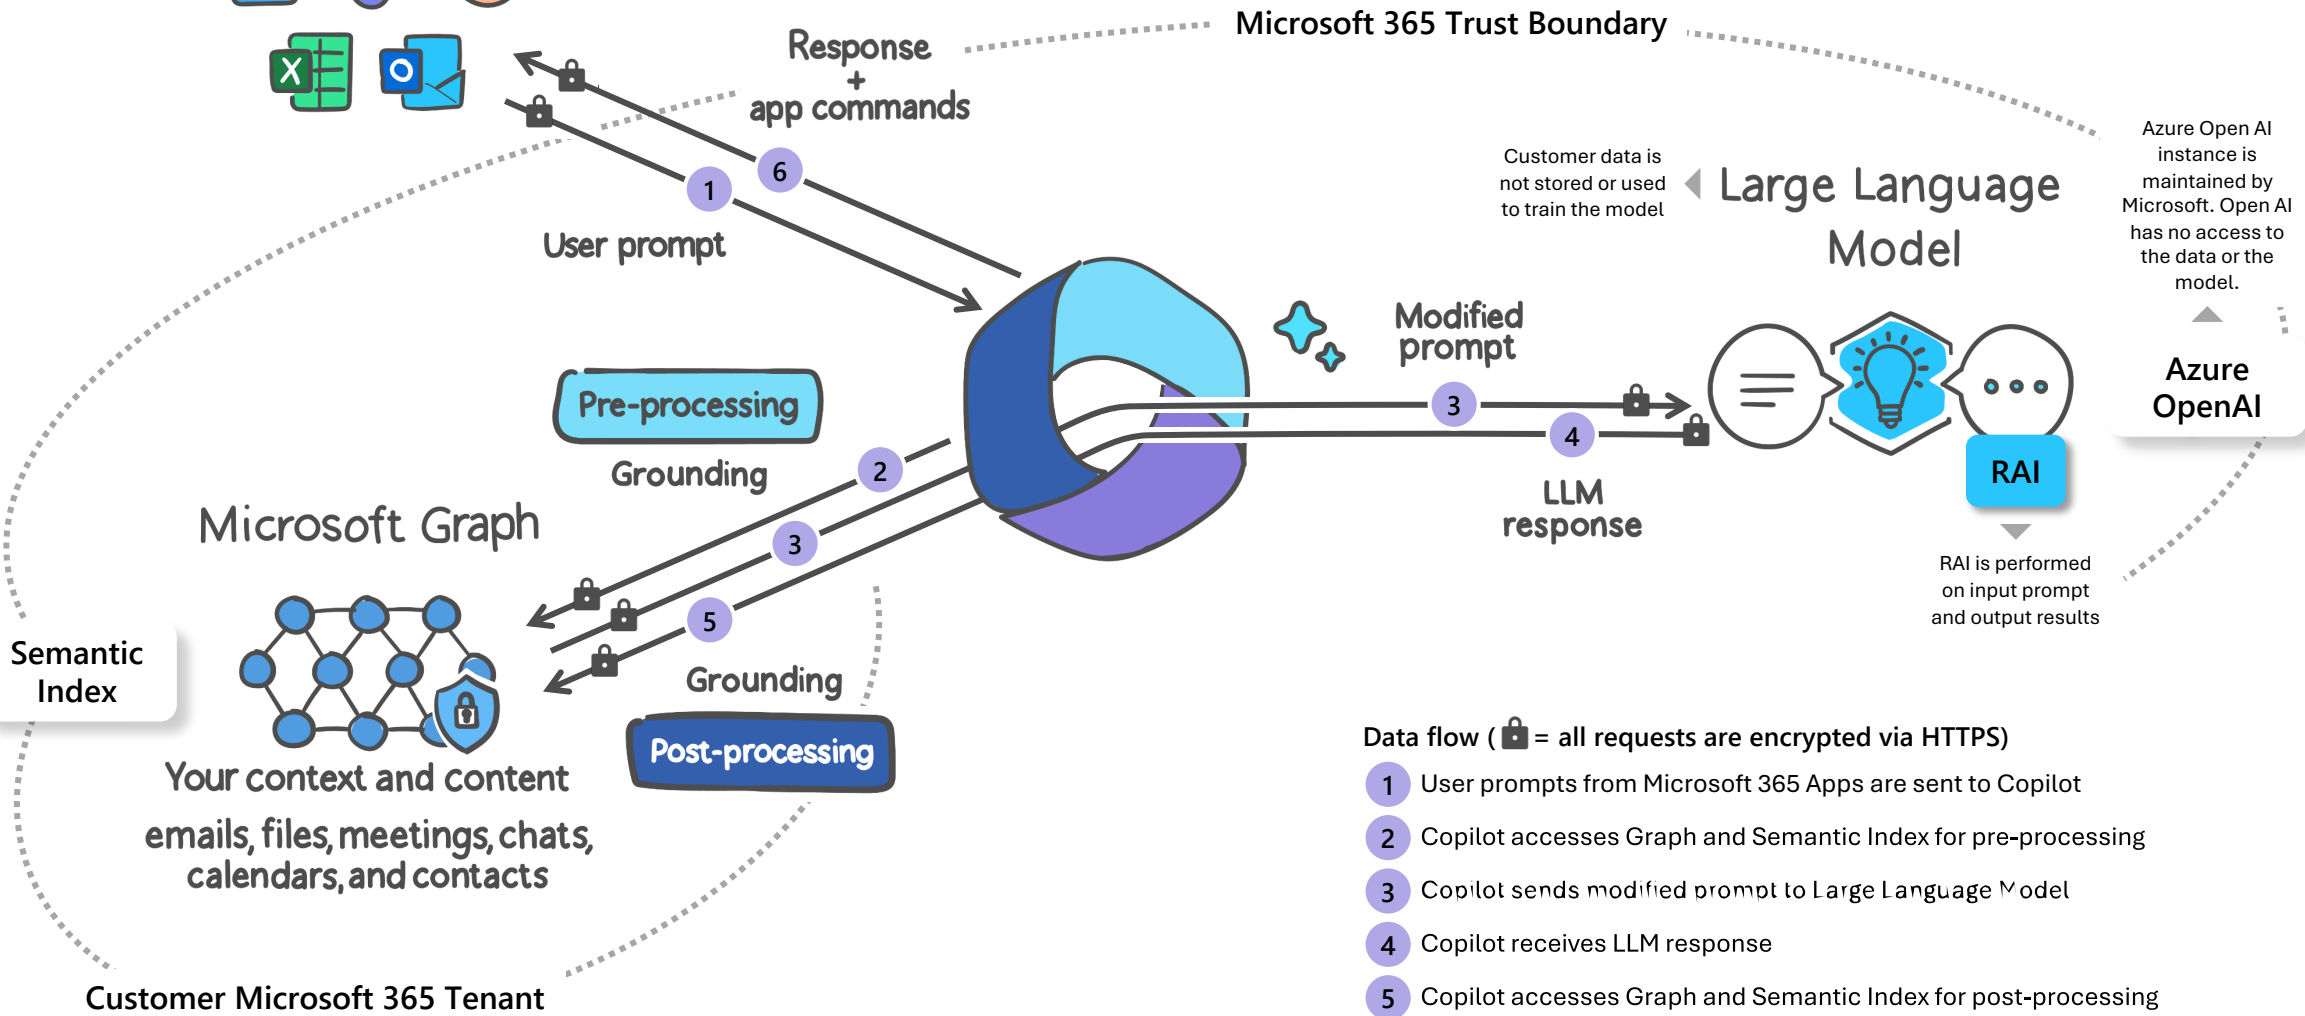
+



Microsoft 365
Apps

Copilot for Microsoft 365 basic architecture

Microsoft 365 Apps



- Data flow (🔒 = all requests are encrypted via HTTPS)
- 1 User prompts from Microsoft 365 Apps are sent to Copilot
 - 2 Copilot accesses Graph and Semantic Index for pre-processing
 - 3 Copilot sends modified prompt to Large Language Model
 - 4 Copilot receives LLM response
 - 5 Copilot accesses Graph and Semantic Index for post-processing
 - 6 Copilot sends the response, and app command back to Microsoft 365 Apps

Customer Microsoft 365 Tenant

**Microsoft adheres
to industry leading
privacy and security
commitments**

**Copilot for Microsoft 365 Early
Access Program falls under Preview
Terms**

**At GA, Copilot for Microsoft 365
will fall under [Product Terms](#)**

Copilot for Microsoft 365

Built on Microsoft's **comprehensive** approach



Security



Compliance



Privacy



Responsible AI

Copilot for Microsoft 365

Designed for the needs of the enterprise



Grounded in
your business data

Copilot for Microsoft 365 has real-time access to both your *content and context* in the Microsoft Graph.



Comprehensive
security, compliance, & privacy

Copilot inherits your security, compliance, and privacy policies set up in Microsoft 365.



Architected to protect
data

Your data doesn't leave the compliance boundary and isn't used to train the foundation model.



Integrated into the
apps you use every day

Word, Excel, PowerPoint, Outlook, Teams, and more.



Individual user & admin
always in control

Users decide what to use, modify, or discard.

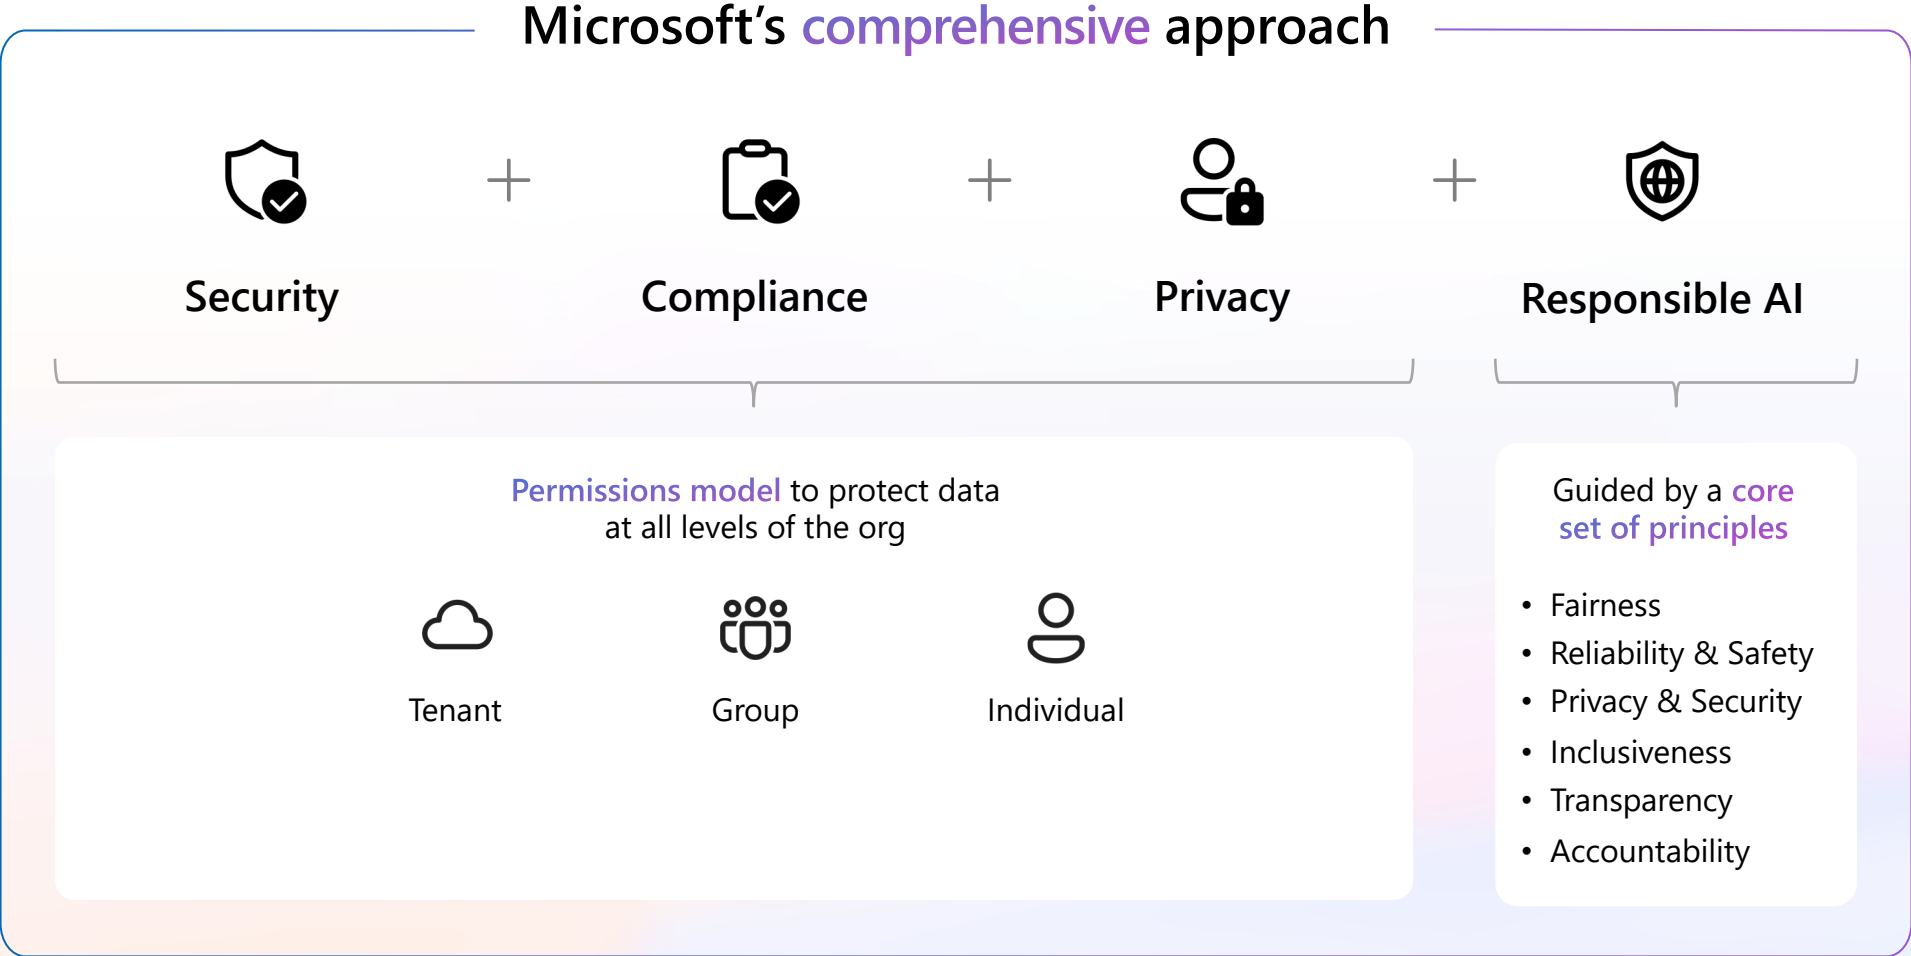


Designed to
integrate new skills

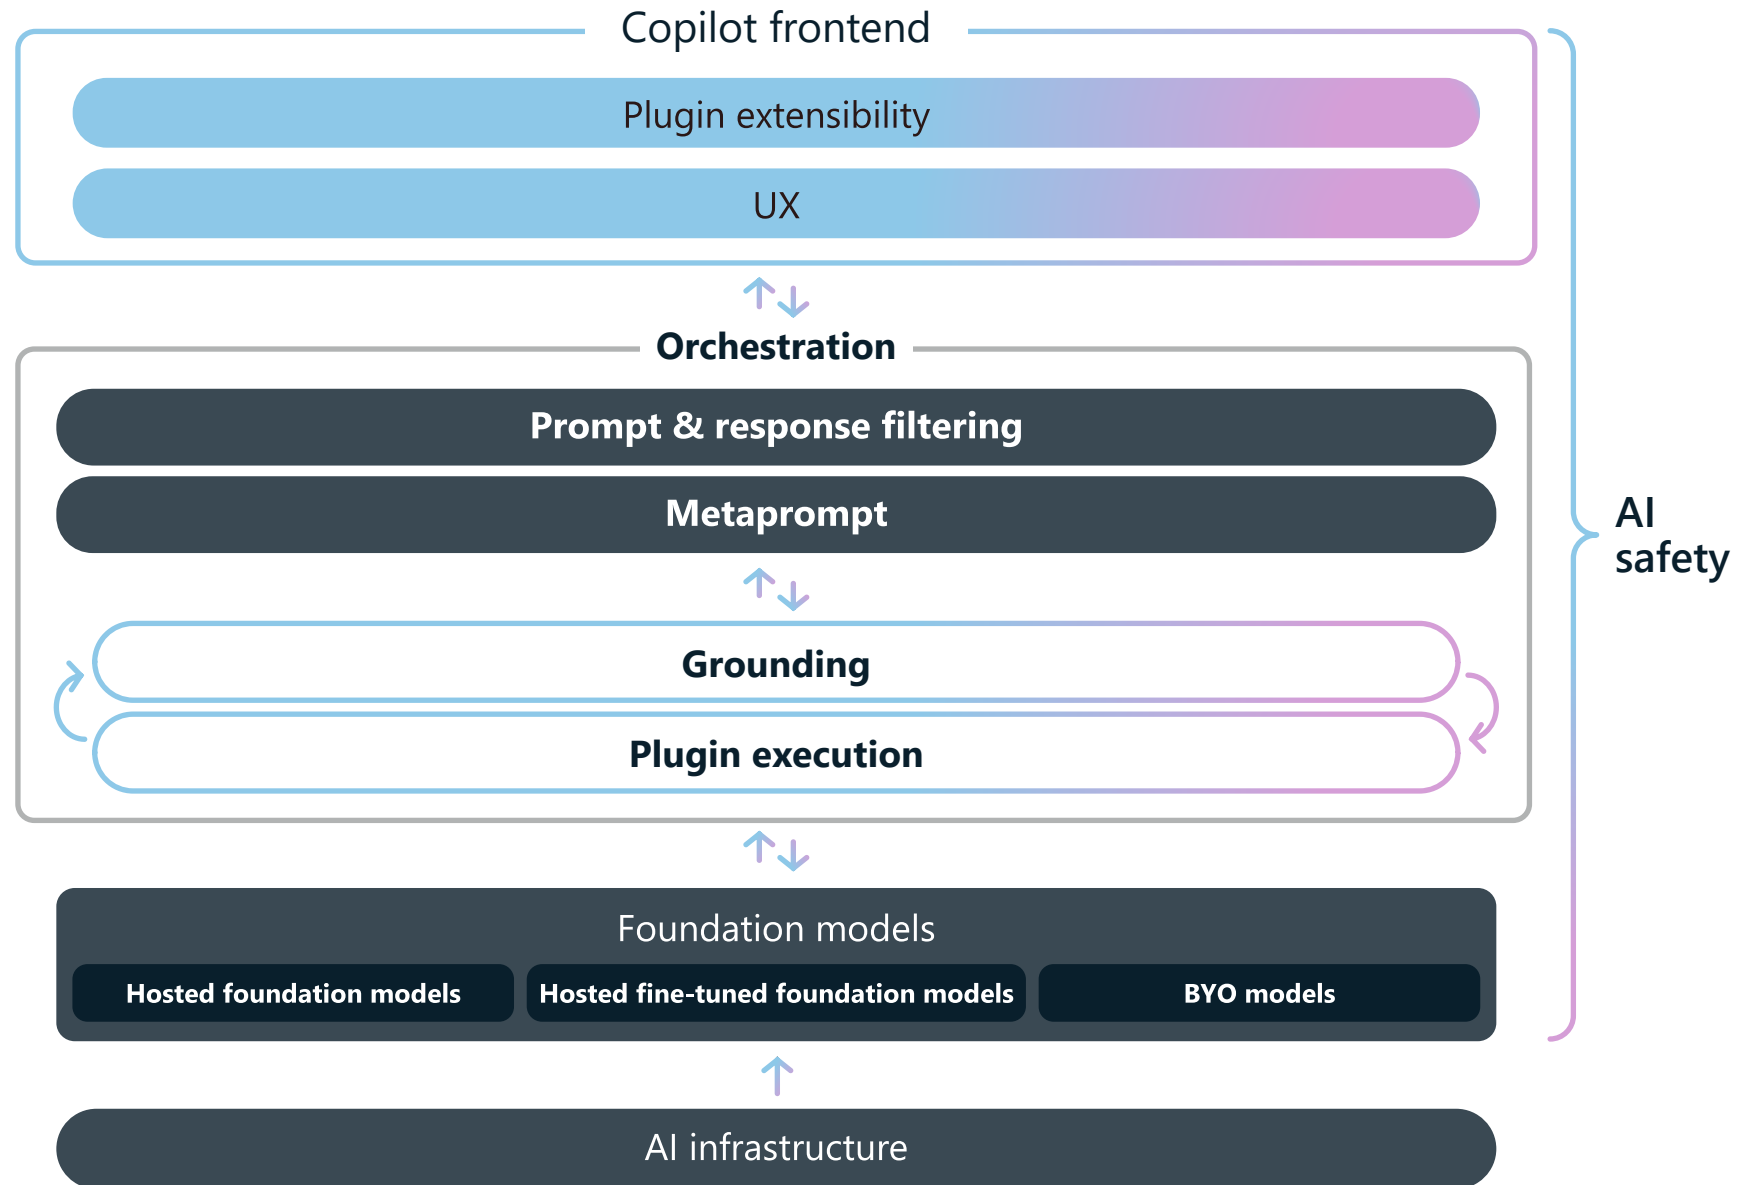
As Copilot learns about processes, it can perform more sophisticated tasks and queries.

How Copilot for Microsoft 365 works

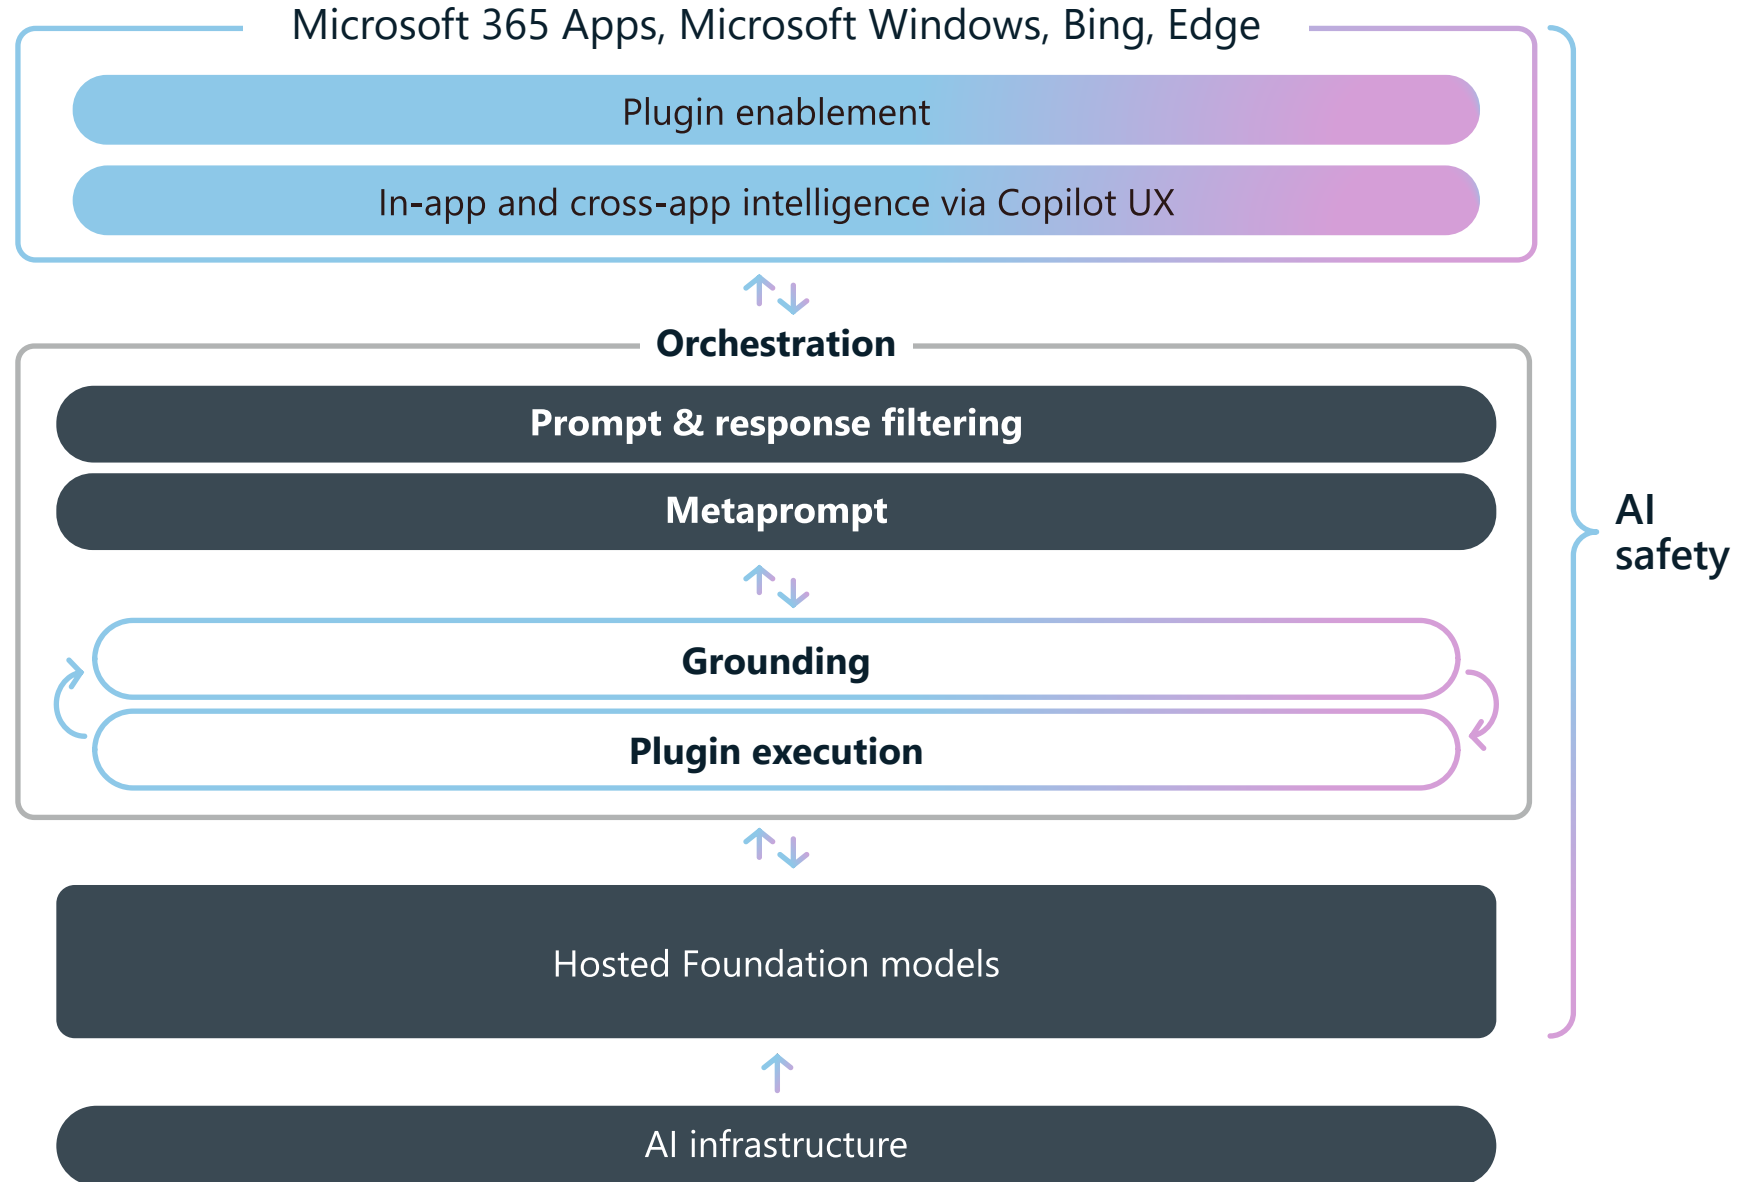
Copilot for Microsoft 365 foundational principles



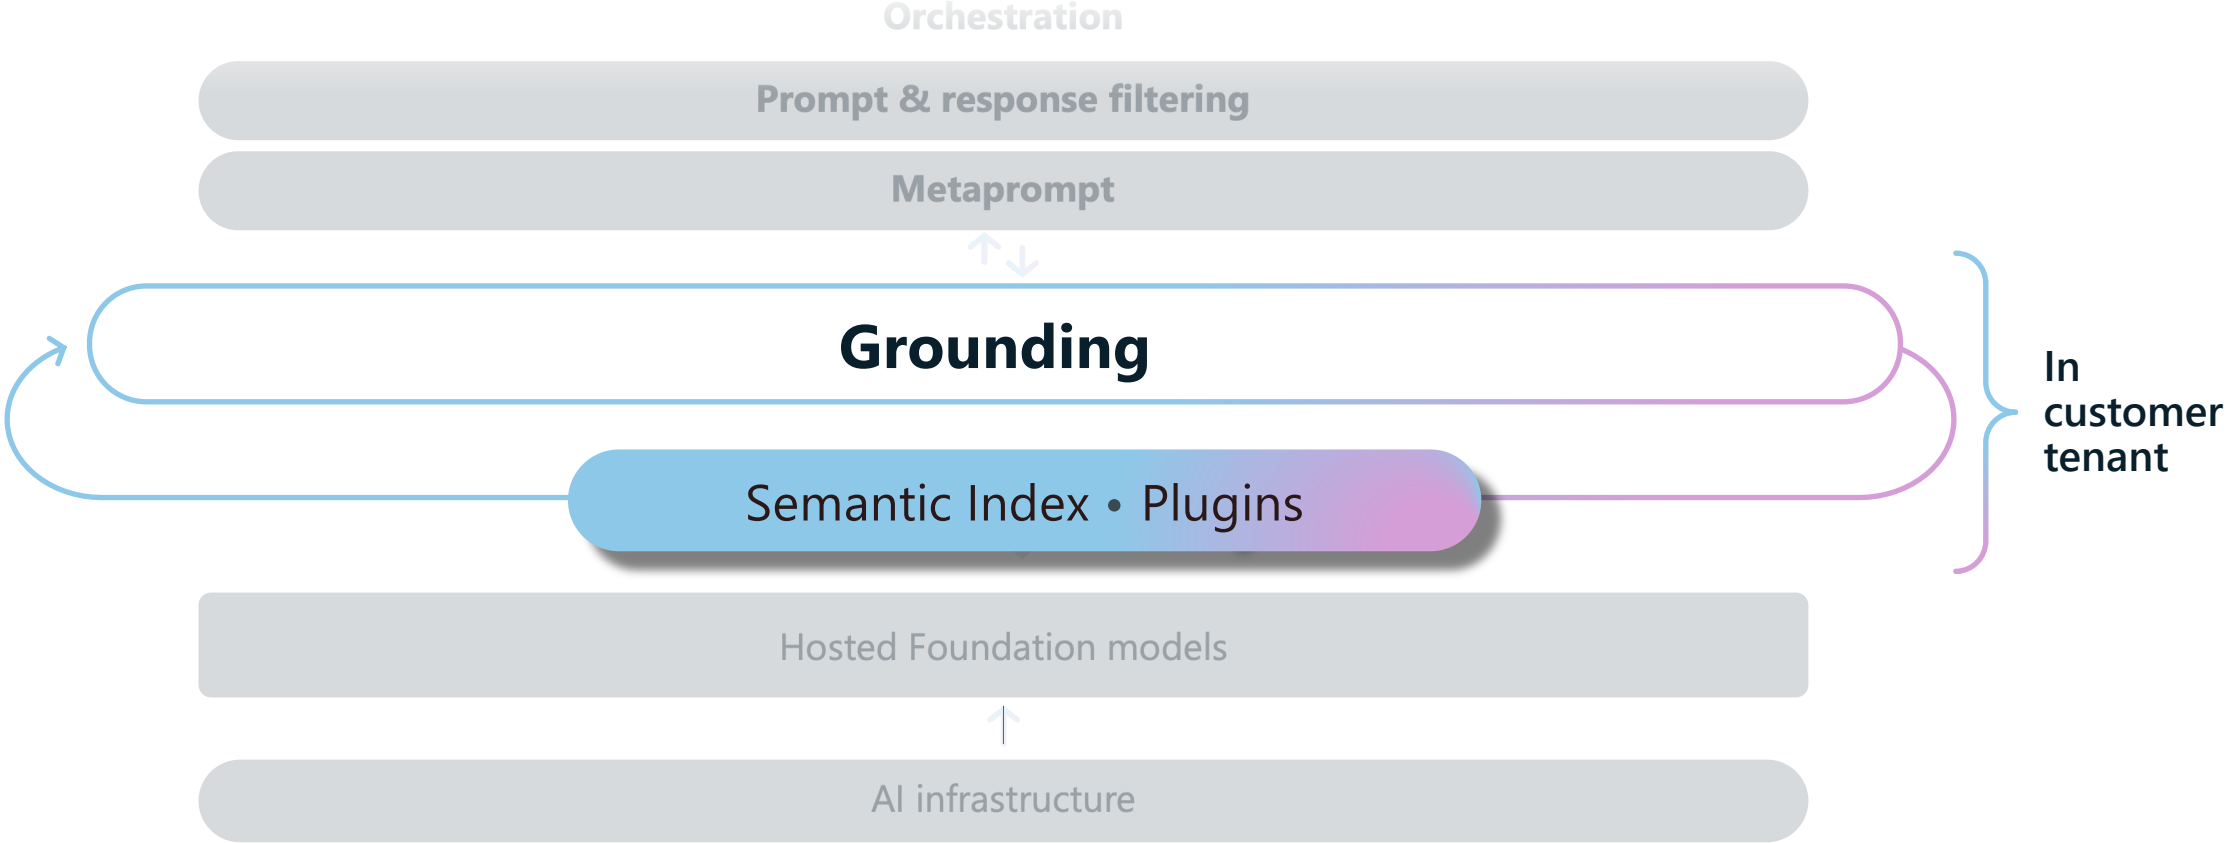
Copilot for Microsoft 365 stack



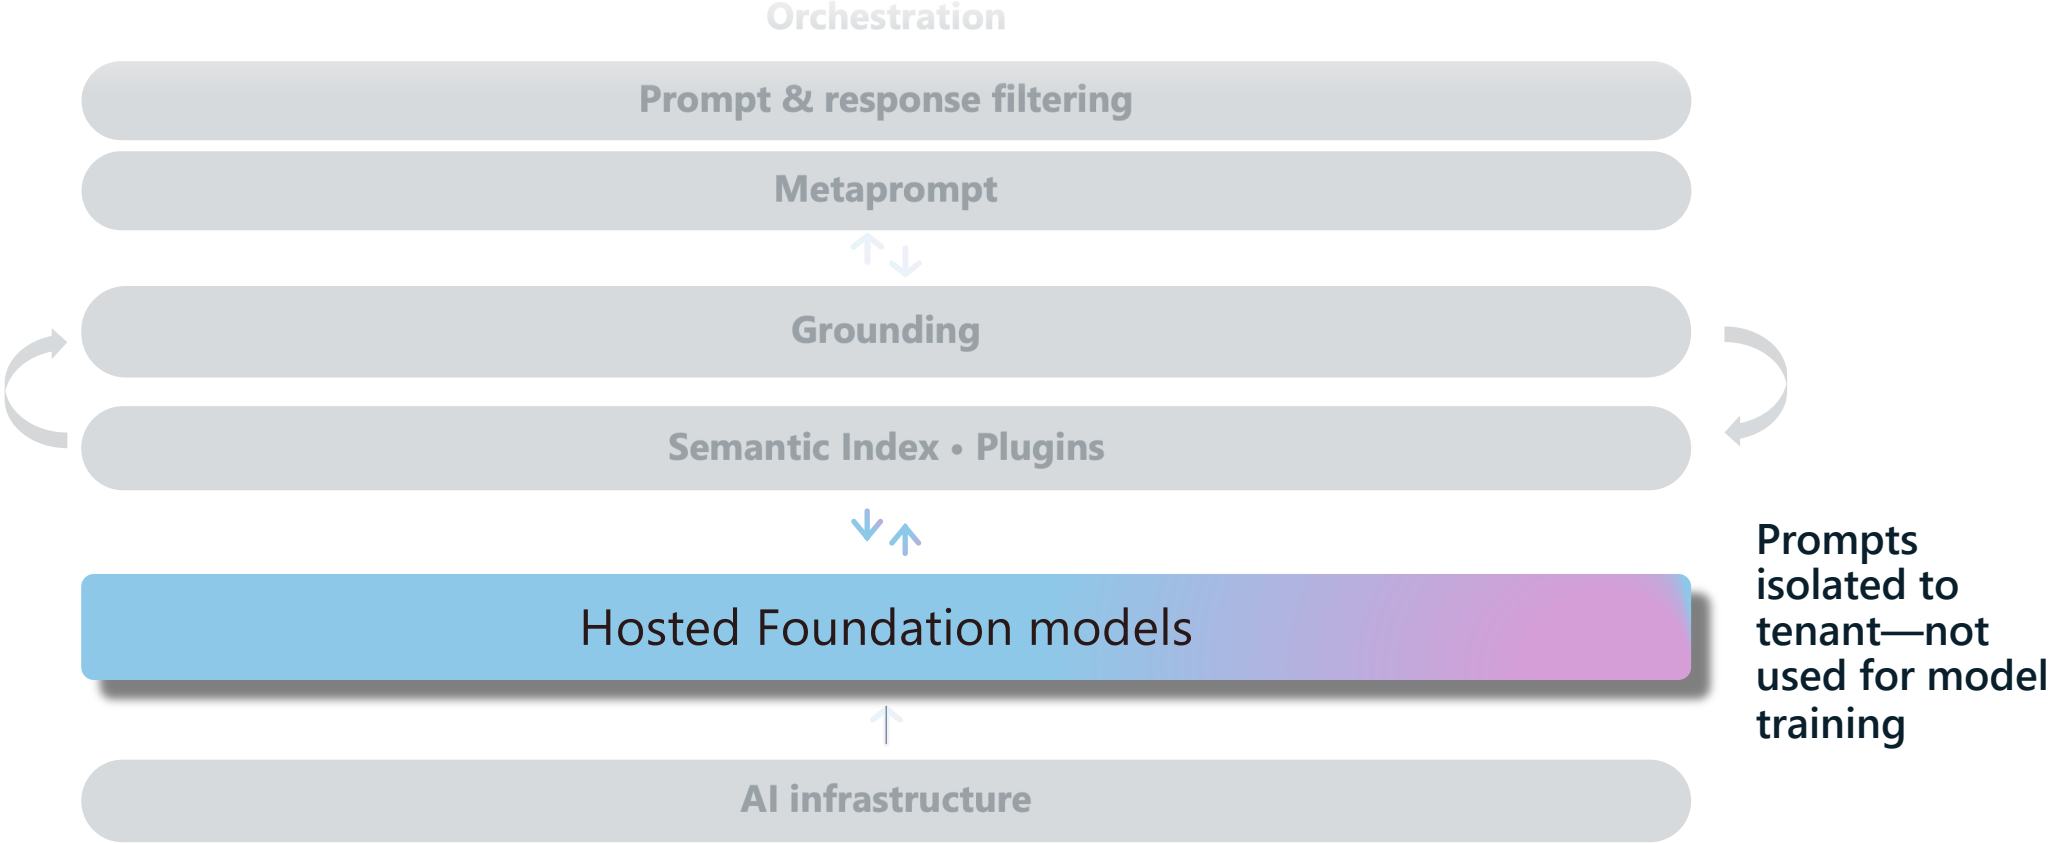
Copilot for Microsoft 365 stack



Retrieval Augmented Generation (RAG)



Natural language processing

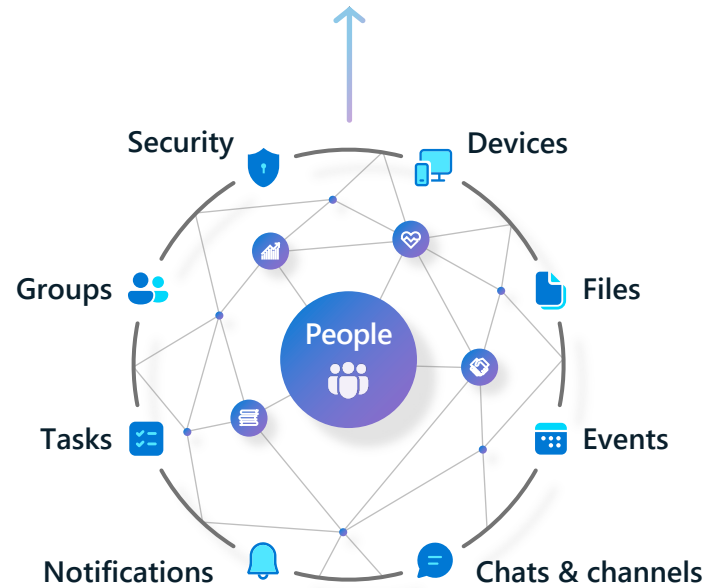


Copilot for Microsoft 365

Semantic Index

- Tenant-wide and user-scoped vector indices
- User and tenant signals & relationships
- Semantics and similarities around content and users
- Ranking and relevance
- Security and permissions
- Across documents, chats, emails, people, projects, meetings, & more

Semantic Index



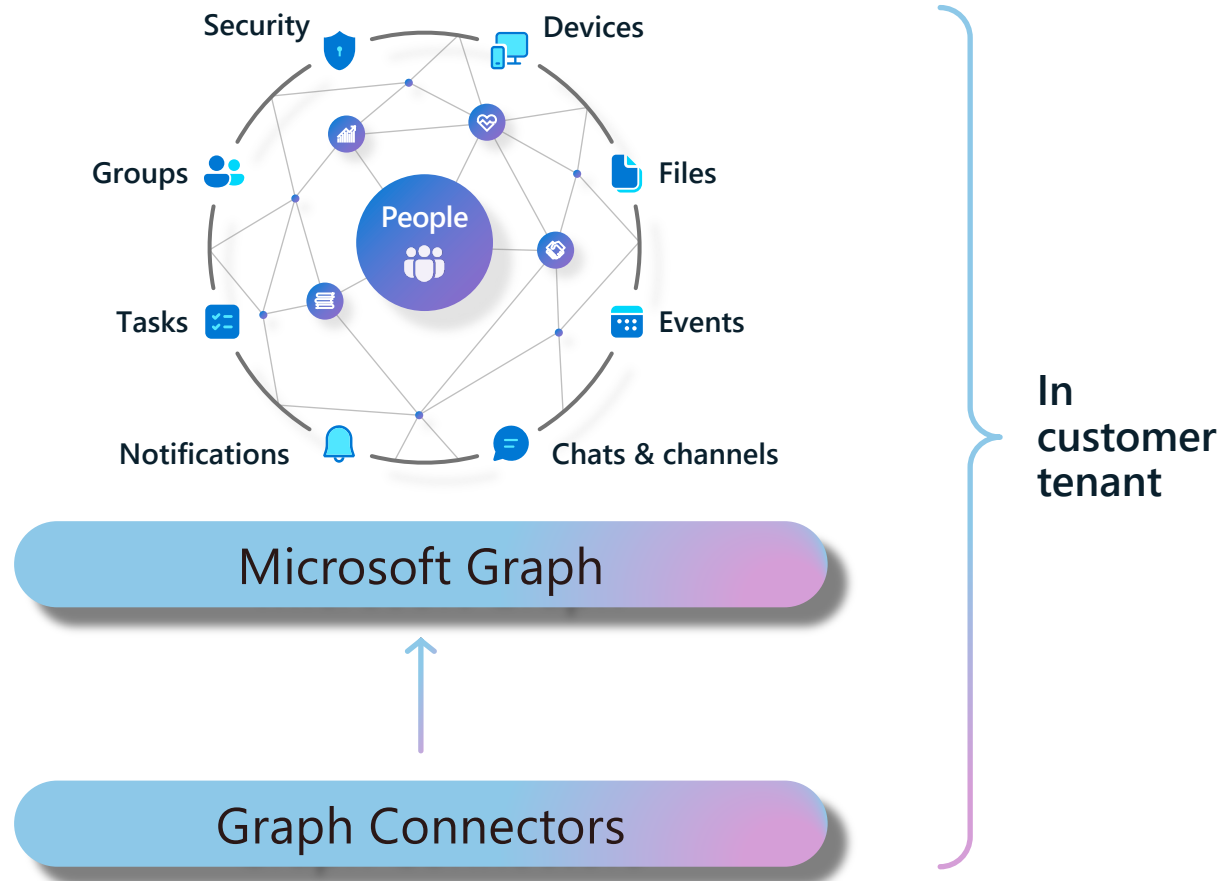
Microsoft Graph

In customer tenant

Copilot for Microsoft 365

Microsoft Graph

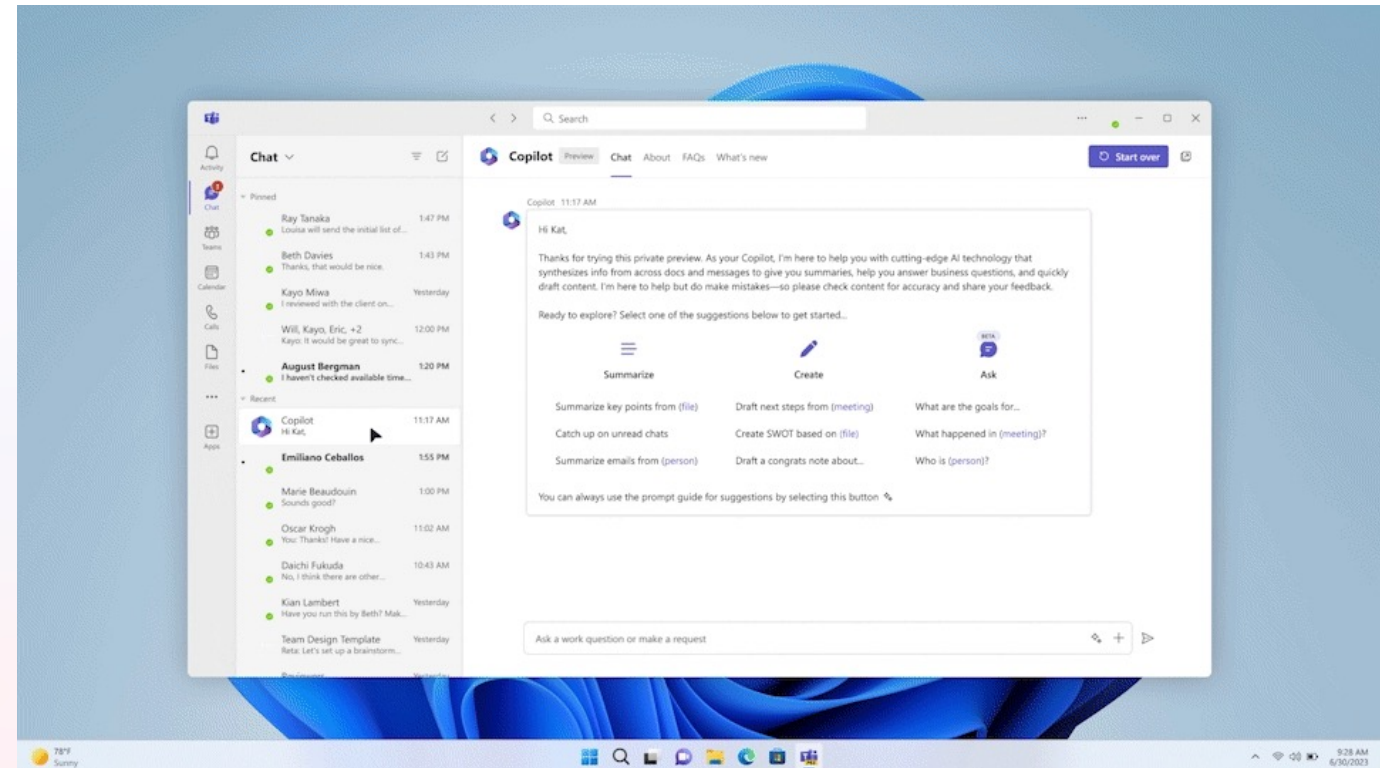
- Line-of-business and 3rd party data
- Security and permissions
- Search schema, filters and eDiscovery



Semantic Index for Copilot for Microsoft 365

To help every customer get AI-ready, we're rolling out Semantic Index for Copilot:

- The Semantic Index for Copilot creates a sophisticated map of your and your organizational data – identifying relationships and making important connections.
- It uses that conceptual understanding to determine your intent and help you find what you need
- The Semantic Index for Copilot enables Copilot for Microsoft 365 to deliver relevant, actionable responses to prompts – and do so in a secure, compliant, privacy-preserving way.



Semantic Index for Microsoft 365

keyword matched results & suggestions

Project Melbourne Shiroom

All Files Sites People Messages Images 2 more

Filters File type Last modified

Did you mean [project melbourne show room?](#)

- Melbourne - Upcoming Workshops**
Project Melbourne
Jodine Lingle modified on 1 March 2020
Project Melbourne is an ambitious and innovative endeavour aimed at transforming the technology...Project Melbourne focuses on lev...
- Blog Strategy - Project Melbourne**
Project Melbourne Blog
Youngji Kim modified on 12 October 2022
...address key challenges faced by hybrid work. The blog strategy envisions a holistic approach that integrates various aspects of...
- Melbourne phase 3 office build out**
Melbourne Office
Clodagh Fitzpatrick modified on 25 March 2017
Sustainability is a key pillar of the Melbourne office. We should aim to reduce carbon emissions, enhance energy efficiency, and promote renewable energy sources throughout...
- Facts About Melbourne**
Time and Places
Jodine Lingle modified on February 16, 2018
Founded in 1975, **Melbourne** is the worldwide leader in manufacturing, sales, services, devices and solutions that help...
- A New Life for Old Projects**
DCX Managed Service

Without Semantic Index for Copilot

Project Melbourne Shiroom

All Files Sites People Messages Images 2 more

Filters File type Last modified

Including results for **'Project Melbourne Shiproom'**
Search only for ['Project Melbourne Shiroom'](#)

Project Melbourne is a forward-thinking and comprehensive urban development initiative that envisions a sustainable, technologically advanced, and inclusive future for the organisation.

- PM Specification Document**
<https://microsoft.sharepoint.com/teams/SalesAndAccounts/Profiles/Melbourne/>
- PM Shiproom June 2023**
Project Melbourne
Jodine Lingle modified on 1 December 2021
Project Melbourne is an ambitious and innovative endeavour aimed at transforming the technology...Project Melbourne focuses on lev...
Karin Blair shared this with you 5 days ago
- Frequently Asked Questions**
 - What are the dependencies on Project Melbourne?
 - When is project Melbourne expected to finish?
 - What is the budget for project Melbourne?
- PM Quarterly Newsletter**

Was this answer useful?

With Semantic Index for Copilot

autocorrected spelling of organizational jargon

informative answers directly from documents

more relevant results, based on search intent

enhanced content sharing insights

AI generated questions and answers on topic

Semantic Index for Copilot for Microsoft 365 features

Organizational Terms creates a custom dictionary of generated terms that autocorrect misspelled jargon & acronyms unique to your organization.

Document Excerpts are extractive, informative answers to a query in the most relevant paragraph in the most relevant document.

Semantic Search understands the intent behind your prompts and commands, producing more relevant, transparent results

More Context provides collaboration, sharing and content insights to search results

Frequently Asked Questions provide AI generated questions and answers on a query to help you understand more about the query in a single search pane.

The screenshot shows the Microsoft 365 search interface for the query 'Project Melbourne Shiroom'. The interface includes a search bar at the top, a navigation pane on the left, and a main results area. The results are categorized into several sections:

- Search filters:** 'All', 'Files', 'Sites', 'People', 'Messages', 'Images', '2 more'. Below these are 'Filters', 'File type', and 'Last modified'.
- Search results:**
 - A box indicating 'Including results for *'Project Melbourne Shiproom'*' and 'Search only for *'Project Melbourne Shiroom'*'.
 - A document excerpt: 'Project Melbourne is a forward-thinking and comprehensive urban development initiative that envisions a sustainable, technologically advanced, and inclusive future for the organisation.'
 - A document result: 'PM Specification Document' with a URL: <https://microsoft.sharepoint.com/teams/SalesAndAccounts/Profiles/Melbourne/>
 - A presentation result: 'PM Shiproom June 2023' by 'Project Melbourne', modified by 'Jodine Lingle' on '1 December 2021'. The description reads: 'Project Melbourne is an ambitious and innovative endeavour aimed at transforming the technology...Project Melbourne focuses on lev...'
 - A sharing notification: 'Karin Blair shared this with you 5 days ago'.
 - A 'Frequently Asked Questions' section with three questions:
 - What are the dependencies on Project Melbourne?
 - How many products does Contoso produce?
 - Who is the CTO of Contoso Corporation?
 - A 'PM Quarterly Newsletter' result at the bottom.
- Related searches:** 'Project Melbourne', 'Melbourne office', 'Fabrikam Inc', 'Wide World Importers'. Below this is a feedback prompt: 'Were these work searches useful?' with thumbs up and down icons.

Arrows from the text on the left point to specific features in the screenshot: 'Organizational Terms' points to the search filters; 'Document Excerpts' points to the document excerpt; 'Semantic Search' points to the document results; 'More Context' points to the sharing notification; and 'Frequently Asked Questions' points to the FAQ section.



Microsoft Graph

Ground Copilot
in **your data**

Enforce **security**
& **permissions**
at runtime

Customer's **data**
- **not used** for
training

Power **Semantic**
Index for Copilot

Semantic Index for Copilot for Microsoft 365

Copilot for Microsoft 365 processing of customer data and prompts.

- Customer data sent to cloud model and results returned for user choice.
- Cloud model processes data without storing or retraining it.
- Customer data and prompts not stored by Microsoft.

Privacy

Privacy

Copilot inherits your Microsoft 365 policies and commitments.

Copilot for Microsoft 365 automatically inherits your organization's security, compliance, and privacy policies for Microsoft 365. Data is managed in-line with our current commitments.

Review the following for more information:

- [Privacy at Microsoft](#)
- [Microsoft Privacy Statement](#)
- [Microsoft Trust Center for data protection and privacy](#)

Privacy

Copilot for Microsoft 365 doesn't train with customer content.

We do not use customer data to train the large language models used to generate content.

Review [Microsoft's privacy policy and service documentation](#) for more information.

Privacy

Data in different geographies are stored and processed in compliance with GDPR and other relevant Privacy Laws.

The new Copilot for Microsoft 365 experiences do not change Microsoft's commitment and principles for data security and privacy. Copilot for Microsoft 365 is integrated into Microsoft 365 which adheres to all existing privacy and compliance obligations.

Review [Microsoft's privacy policy and service documentation](#) for more information.

Privacy

Microsoft commitments to its customers about the GDPR.

The General Data Protection Regulation (GDPR) introduces new rules for organizations that offer goods and services to people in the European Union (EU), or that collect and analyze data for EU residents no matter where you or your enterprise are located.

The GDPR requires controllers (such as organizations using Microsoft's enterprise online services) only use processors (such as Microsoft) that provide sufficient guarantees to meet key requirements of the GDPR. **Microsoft has taken the proactive step of providing these commitments to all Volume Licensing customers as part of their agreements.**

Learn more about Microsoft's commitment to GDPR from our [full summary](#) and [short summary](#).

Data residency and storage

Data residency and storage

Copilot for Microsoft 365 data remains within the Microsoft cloud.

When you input prompts using our AI tools, the information contained within your prompts, the data they retrieve, as well as the generated responses **remain within the Microsoft cloud.**

Data residency and storage

Customer content data is written in home region datacenters.

- **LLM calls will be routed to the closest data centers in the regions.**
- The LLM call will process customer data and return the response to the home region.
- **NO customer content data is written in region outside the home region for the user.**
- For EU users we have additional safeguards to comply with EUDB.

Compliance

Compliance

Microsoft adheres to industry leading privacy, security, and compliance commitments.

At GA, Copilot will fall under Product Terms.

Learn more about our commitment to compliance through our [compliance website](#) and [Service Trust Portal](#).

Understand [compliance offering definitions](#).

Security

Security

Built on Microsoft's comprehensive approach to security, compliance, and privacy.

Copilot is integrated into Microsoft 365 apps and **automatically inherits all your company's valuable security, compliance, and privacy policies and processes.** Two-factor authentication, compliance boundaries, privacy protections, and more make Copilot the AI solution you can trust.

Security

Copilot for Microsoft 365 models ensure data is not exposed to other organizations.

Copilot protects data in multiple ways:

- **Logical isolation** of customer content within each tenant with AAD authorization and role-based access control.
- Microsoft uses **physical security, screening, and a multi-layered encryption** strategy to protect customer content.
- Microsoft 365 encrypts customer content at rest and in transit.

Learn more about [Microsoft 365 isolation controls](#).

Security

Data is protected with tenant controls and permissions models.

The permissions model within your Microsoft 365 tenant will ensure that data will not unintentionally leak between users and groups. **Copilot presents only data that each individual can access** using the same underlying controls for data access used in other Microsoft 365 services.

Review [Microsoft's privacy policy and service documentation](#) for more information.

Security

Architected to protect tenant, group, and individual data.

LLMs are not further trained on, or learn from, your tenant data or your prompts. Copilot presents only data you can access using the same technology that we've been using for years to secure customer data.

Security

Data protected with Purview or AIP labeling remains protected.

Data protected using Microsoft Purview Information Protection (MPIP or AIP) labeling will continue to be protected according to those policies.

How to prepare for Copilot for Microsoft 365

Technical requirements and recommendations

- ✓ Microsoft 365 E3 or E5 are licensing prerequisites.
- ✓ AAD - based account is required¹.
- ✓ Several features require users to have a OneDrive account.
- ✓ Copilot in Word, Excel, and PowerPoint: first available in the web versions at <https://microsoft365.com> and later expand to desktop
- ✓ For Outlook: Users need to be using the new Outlook for Windows (which is publicly called “the new Outlook for Windows” and it is currently in Preview. Users can load the new Outlook by toggling or Outlook Mobile. [Getting started with the new Outlook for Windows – Microsoft Support](#).
- ✓ For Teams: Users need to be using Teams desktop client or the web app. Both current and the new versions of Teams supported.
- ✓ For Loop: Tenants must have Loop enabled. [Learn more how to enable Loop](#). *(only required if customer want to use Copilot in Loop)*
- ✓ The richness of the cross-app experience will depend on the data sources indexed by Microsoft 365. Tenants with the richest data in Microsoft 365 (Exchange, OneDrive, SharePoint, Teams) will get the best results.

Some Copilot experiences require a websocket connection from the device where the Microsoft 365 app is running to a Microsoft service.

How to prepare for Copilot for Microsoft 365

- ① Understand file permissions:
 - Overview of [sharing and permissions available in SharePoint in Microsoft 365](#)
 - [Sensitivity labels](#) from Microsoft Purview
- ② Implement data loss prevention policies:
 - Learn about [data loss prevention](#)
 - Watch: [automatically classifying and protecting documents and data](#) and [data loss prevention in Microsoft Purview](#)
- ③ Use reporting to stay updated:
 - [Usage reports](#) in the Microsoft 365 admin center
 - [Data access governance reports](#) in SharePoint admin center

Review this [blog post](#) on getting tenant ready for Copilot for Microsoft 365 or watch this Microsoft Mechanics [video](#).

Permissions policies

Preparing your organization for data governance in Copilot for Microsoft 365.

Data permissions will be consistent, and users will only be able to search the information they already have access to.

Review our blog about content management as you prepare your organization to use Copilot: [New era in content management and security in SharePoint, OneDrive, and Teams](#).

Technical questions

How accurate is Copilot for Microsoft 365?

Responses that generative AI produces aren't guaranteed to be 100% factual.

While we continue to improve responses, **people should still use their judgment when reviewing the output before sending them to others.**

Our Copilot capabilities provide useful drafts and summaries to help you achieve more while giving you a chance to review the generated AI rather than fully automating these tasks.

We continue to improve algorithms to proactively address issues, such as misinformation and disinformation, content blocking, data safety, and preventing the promotion of harmful or discriminatory content in line with our [responsible AI principles](#).

Where does Copilot for Microsoft 365 get its responses?

Copilot for Microsoft 365 users can view where data is being sourced from in order to validate the accuracy of the data.

Our design principles include giving users clarity into where responses are sourced, where possible.

- Outlook provides a summary of an email thread it will include footnotes throughout, linking to the places in the thread that support Copilot's summary.
- Cross-app intelligence features (chatting with Copilot in Teams or Bing), responses include links to the documents, emails, etc. that support Copilot's response.

What languages does Copilot for Microsoft 365 support?

Copilot for Microsoft 365 will initially support these languages during the Early Access Program.

English, Spanish, French, German, Italian, Japanese, Portuguese, and Simplified Chinese.

Who owns content generated by Copilot for Microsoft 365?

Microsoft does not claim ownership of the output of the service.

Microsoft does not own Customer's Output Content. Customer understands and acknowledges that generative AI systems may produce similar responses to similar prompts or queries from multiple customers.

What's the difference between ChatGPT and Copilot for Microsoft 365?

ChatGPT is a general-purpose LLM that works like it a chatbot. It uses machine learning algorithms to answer user questions or prompts in a way that mimics human conversations.

ChatGPT cannot search the internet, and its knowledge base does not extend beyond September 2021. This means it can “hallucinate” and provide wrong information. Users also cannot custom-tune ChatGPT.

Popular ChatGPT use cases include text compositions, lists, coding, and content summaries.

Copilot for Microsoft 365 is an AI-powered digital assistant that combines ChatGPT with data from the Microsoft Graph and Microsoft 365 apps. It uses data from user calendars, emails, chats, documents, meetings, and more to deliver contextual, relevant solutions directly in Microsoft 365 apps.

This approach only improves the context for a user's specific scenario; Microsoft does not leverage user data to train the Copilot model.

Since Copilot works directly in Microsoft 365 apps, it can do things like create new content in PowerPoint, suggest email replies in Outlook, and recap a Teams meeting.

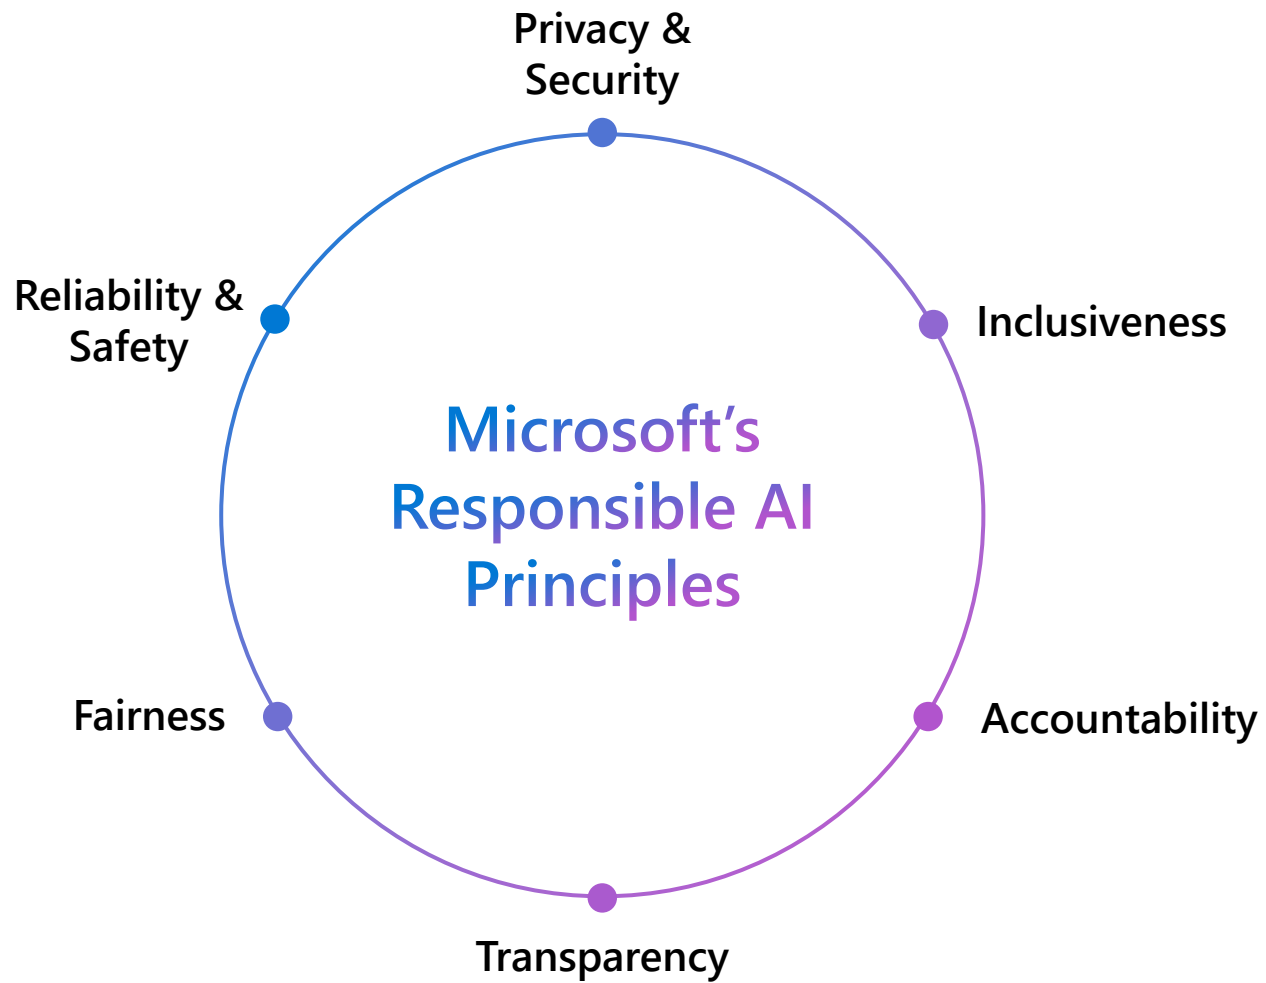
Read more about the [difference between ChatGPT and Copilot for Microsoft 365](#).

Is Copilot connected to the internet?

Copilot's large language models are not trained on your tenant data or on your prompts. And importantly, with Copilot your data—including your prompts—stay within the compliance boundary. At an individual user level, Copilot works like enterprise search today: it can only access the data to which the user already has access, and all that data—including search terms—stay within the tenant.

You may see previously available public data in your Copilot results. Copilot for Microsoft 365 gets its data from the **LLMs, like ChatGPT or GPT 3.5 that have been trained on the publicly-available data sources** before being brought to the Microsoft Cloud.

Responsible AI



Building blocks to enact principles



Tools and processes



Training and practices



Rules



Governance

Tools for a Responsible AI approach



Fairness

Assess fairness and mitigate fairness issues to build models for everyone.



Explainability

Understand model predictions by generating feature importance values for your mode.



Counterfactuals

Observe feature perturbations and find the closest datapoints with different model predictions.



Causal analysis

Estimate the effect of a feature on real-world outcomes.



Error analysis

Identify dataset cohorts with high error rates and visualize error distribution in your model.



Responsible AI scorecard

Get a PDF summary of your Responsible AI insights to share with your technical and non-technical stakeholders to aid in compliance reviews.

Microsoft's Responsible AI Standard

Guidance for designing, developing, deploying, and using AI, which Microsoft's own engineering teams follow too



Identify

Identify and prioritize potential harms that could result from your AI system through iterative red-teaming, stress-testing, and analysis.



Measure

Measure the frequency and severity of those harms by establishing clear metrics, creating measurement test sets, and completing iterative, systematic testing (both manual and automated).



Mitigate

Mitigate harms by implementing tools and strategies such as [prompt engineering](#) and using our [content filters](#). Repeat measurement to test effectiveness after implementing mitigations.



Operate

Define and execute a deployment and operational readiness plan.

Review our [Responsible AI practices](#) to learn more

Next Steps

- ✓ Join our [Copilot Readiness Community](#)
- ✓ Take the [Copilot Readiness Assessment](#)
- ✓ [Contact CodeBlue](#) to discuss your Copilot strategy

Glossary

Glossary of terms for Copilot for Microsoft 365 (1 of 2)

- **Microsoft 365 apps:** Solutions like Word, Excel, PowerPoint, Outlook, Teams, and Loop that operate with Copilot to support users in the context of their work. For example, Copilot in Word is designed to assist users specifically in the process of creating, comprehending, and editing documents. In a similar way, Copilot in the other apps help users in the context of their work within those apps.
- **Chat:** Copilot has cross-app intelligence, which allows users with broad needs a simpler way to work with multiple apps. Users access cross-app intelligence by chat in the same way they would interact using open prompts with ChatGPT or Microsoft Copilot. Those prompts access the core training data in the LLM as well as users' business data and apps to surface the information and insights they need from their organization's data. Prompts work with Copilot across a range of experiences, including Teams (chat), Bing, Edge, and the Microsoft 365 app.
- **The Microsoft Graph:** A foundational part of Microsoft 365, the Graph includes information about the relationships between users, activities, and your organization's data, working together with the Semantic Index for Copilot, as well as orchestrating information retrieval steps using search. The Microsoft Graph API brings additional context from customer signals into the prompt, such as information from emails, chats, documents, meetings, and more.

Glossary of terms for Copilot for Microsoft 365 (2 of 2)

- **The Semantic Index for Copilot:** A sophisticated map of your user and company data. It uses multiple large language models that sit on top of the Microsoft Graph, which interpret user queries and produce sophisticated, meaningful, and multilingual responses that help you to be more productive. It allows Microsoft 365 E3 and E5 customers to search through billions of vectors (mathematical representations of features or attributes) and return the most related results in tens of milliseconds. Combined with enhancements across the Microsoft Graph, the Semantic Index for Copilot connects you with the most relevant and actionable information in your organization and is built on Microsoft's comprehensive approach to security, compliance, privacy, and respects all organizational boundaries within your tenant.
- **The Copilot System:** The common underlying AI stack that connects Microsoft 365 apps, chat, the Microsoft Graph, and the Semantic Index. It includes baseline LLM, AI platform, skills repository and runtime that powers end user experiences Microsoft Copilot, Copilot in Microsoft 365 apps, and cross-app intelligence.

Glossary of terms for Responsible AI (1 of 2)

- **Fairness:** Fairness is a core ethical principle that all humans aim to understand and apply. This principle is even more important when AI systems are being developed. Key checks and balances need to make sure that the system's decisions don't discriminate or run a gender, race, sexual orientation, or religion bias toward a group or individual.
 - Microsoft provides an [AI fairness checklist](#) that offers guidance and solutions for AI systems. These solutions are loosely categorized into five stages: envision, prototype, build, launch, and evolve. Each stage lists recommended due diligence activities that help to minimize the impact of unfairness in the system.
 - Fairlearn integrates with Azure Machine Learning and supports data scientists and developers to assess and improve the fairness of their AI systems. The toolbox provides various unfairness mitigation algorithms and an interactive dashboard that visualizes the fairness of the model. Use the toolkit and closely assess the fairness of the model while it's being built; this should be an integral part of the data science process.
 - Learn how to [mitigate fairness in machine learning models](#).
- **Reliability & Safety:** AI systems need to be reliable and safe in order to be trusted. It's important for a system to perform as it was originally designed and for it to respond safely to new situations. Its inherent resilience should resist intended or unintended manipulation. Rigorous testing and validation should be established for operating conditions to ensure that the system responds safely to edge cases, and A/B testing and champion/challenger methods should be integrated into the evaluation process. An AI system's performance can degrade over time, so a robust monitoring and model tracking process needs to be established to reactively and proactively measure the model's performance and retrain it, as necessary, to modernize it..

Glossary of terms for Responsible AI (2 of 2)

- **Privacy & Security:** A data holder is obligated to protect the data in an AI system, and privacy and security are an integral part of this system. Personal needs to be secured, and it should be accessed in a way that doesn't compromise an individual's privacy. [Azure differential privacy](#) protects and preserves privacy by randomizing data and adding noise to conceal personal information from data scientists.
- **Inclusiveness:** Inclusiveness mandates that AI should consider all human races and experiences, and [inclusive design](#) practices can help developers to understand and address potential barriers that could unintentionally exclude people. Where possible, speech-to-text, text-to-speech, and visual recognition technology should be used to empower people with hearing, visual, and other impairments.
- **Transparency:** Achieving transparency helps the team to understand the data and algorithms used to train the model, what transformation logic was applied to the data, the final model generated, and its associated assets. This information offers insights about how the model was created, which allows it to be reproduced in a transparent way. Snapshots within [Azure Machine Learning workspaces](#) support transparency by recording or retraining all training-related assets and metrics involved in the experiment.
- **Accountability:** Accountability is an essential pillar of responsible AI. The people who design and deploy the AI system need to be accountable for its actions and decisions, especially as we progress toward more autonomous systems. Organizations should consider establishing an internal review body that provides oversight, insights, and guidance about developing and deploying AI systems. While this guidance might vary depending on the company and region, it should reflect an organization's AI journey.

Glossary of terms for **Azure OpenAI** (1 of 3)

- **OpenAI:** A research organization focused on developing artificial intelligence in a safe and beneficial manner.
- **Generative AI:** Generative AI is a type of artificial intelligence that involves the creation of new content or information, such as images, videos, or text, by an algorithm. Unlike other AI technologies, such as predictive or prescriptive analytics, which use historical data to make predictions or recommendations, generative AI is focused on the creation of new content that does not necessarily rely on past data.
- **GPT:** Generative Pre-trained Transformer. This is a deep learning algorithm that can generate human-like language and has been used for tasks such as language translation and text completion.
- **GPT-3:** Generative Pre-trained Transformer 3 is a language model developed by OpenAI, which uses deep learning techniques to generate natural language text. It is capable of generating coherent, context-sensitive responses to a wide range of prompts.
- **ChatGPT:** A variant of the GPT-3 model developed specifically for use in conversational AI applications. ChatGPT has been optimized for generating human-like responses in a conversational context and can be fine-tuned on specific domains or use cases to improve its performance.
- **Hyper-personalization:** The practice of using data and algorithms to tailor products, services, and content.
- **Fine-tuning a model:** this is the process of taking a pre-trained machine learning model and adjusting its parameters to better suit a specific task or domain. This can involve training the model on a new dataset, adjusting its architecture, or tweaking its hyperparameters.

Glossary of terms for **Azure OpenAI** (2 of 3)

- **Transformer:** A type of neural network architecture that allows for parallel processing of inputs and outputs.
- **Prompt engineering:** The practice of designing natural language prompts that can effectively guide a language model to generate desired responses. This involves careful consideration of the language used, the context of the prompt, and the potential responses that the model might generate.
- **Neural Network:** A computational system that simulates the behavior of the human brain, used in machine learning and artificial intelligence.
- **Natural Language Processing (NLP):** The ability of computers to understand, interpret, and generate human language.
- **Reinforcement Learning:** A type of machine learning where an agent learns to take actions in an environment to maximize a reward signal.
- **Machine Learning:** A type of artificial intelligence where machines learn from data, rather than being explicitly programmed.
- **Deep Learning:** A subfield of machine learning that uses neural networks with many layers to learn complex patterns in data.
- **Transfer Learning:** The ability of a model to apply knowledge learned in one domain to a different domain.
- **Language Model:** A statistical model that predicts the probability of the next word in a sentence or sequence of words.

Glossary of terms for Azure OpenAI (3 of 3)

- **Supervised Learning:** A type of machine learning where a model is trained on labeled data, where the desired output is known.
- **Unsupervised Learning:** A type of machine learning where a model is trained on unlabeled data, where the desired output is unknown.
- **Hyperparameters:** Parameters that are set before training begins, such as learning rate and batch size, which can greatly affect the model's performance.
- **Training Data:** The data used to train the model, typically labeled with the desired output.