Microsoft Security

Security, IP & Compliance: Copilot for Microsoft 365

What you need to know for preparing your Copilot deployment

Martin Janisch
23 05 2024

Agenda

Data security and compliance

Privacy and data residency

How Copilot works (Copilot orchestration, In-app flows: Data egress and commanding apps)

Key takeaways and resources

Session recordings: Getting your enterprise ready for Copilot for Microsoft 365 and How Copilot for Microsoft 365 works

Preparing for AI

Preparing for the era of AI

Goals

Pain points

Current capabilities

Data strategy

Resource and adoption readiness

How to prepare for AI

"Security and risk management leaders must implement verifiable controls for AI data protection, privacy, application security and filtering of large language model content inputs and outputs."

- Gartner

Gartner, Quick Answer: How to Make Microsoft 365 Copilot Enterprise-Ready From a Security and Risk Perspective,, Avivah Litan, Matt Cain, Jeremy D'Hoinne, Nader Henein, Dennis Xu, 15 September 2023

Copilot stack

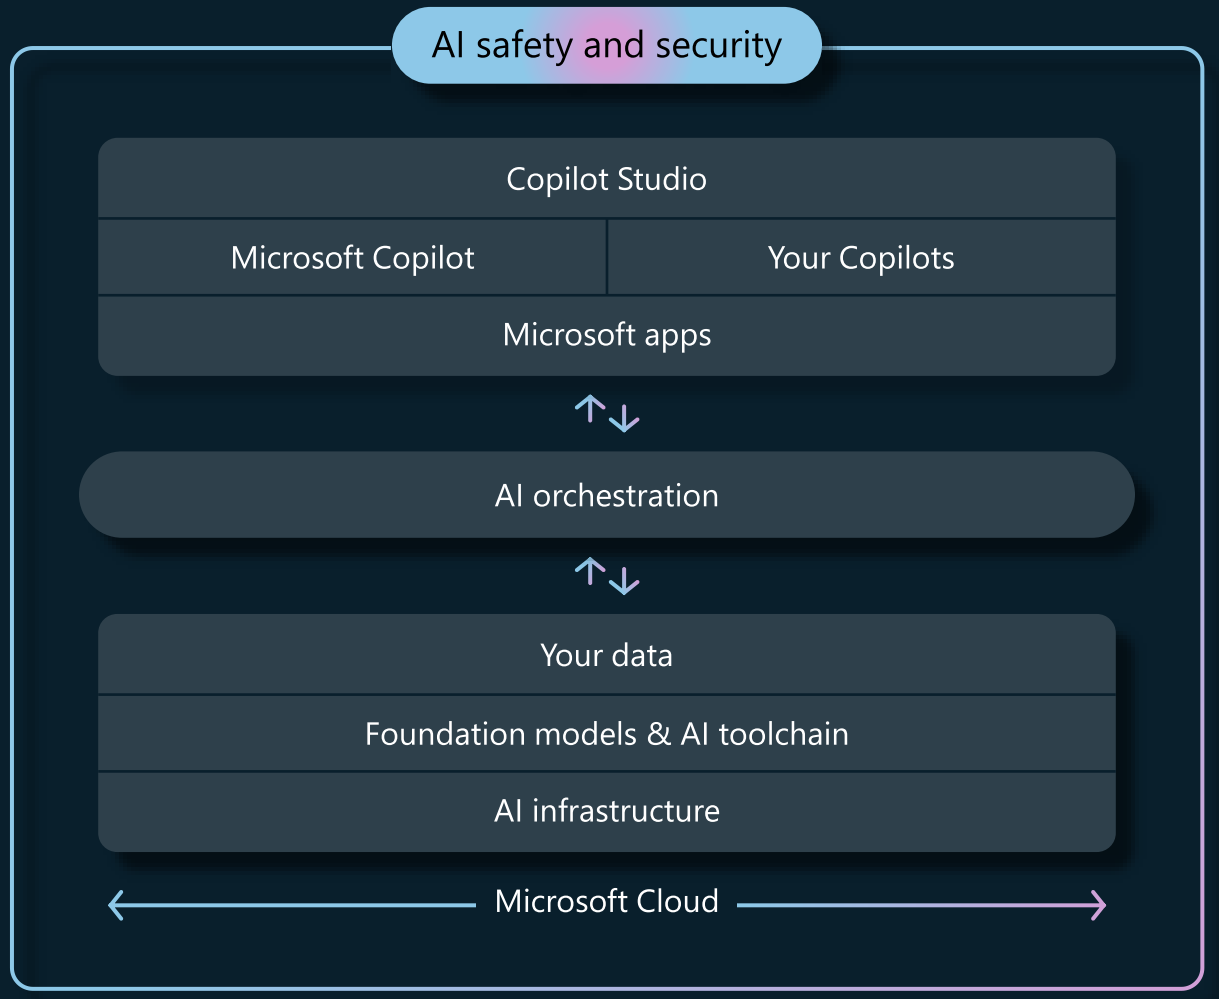

AI security shared responsibility model

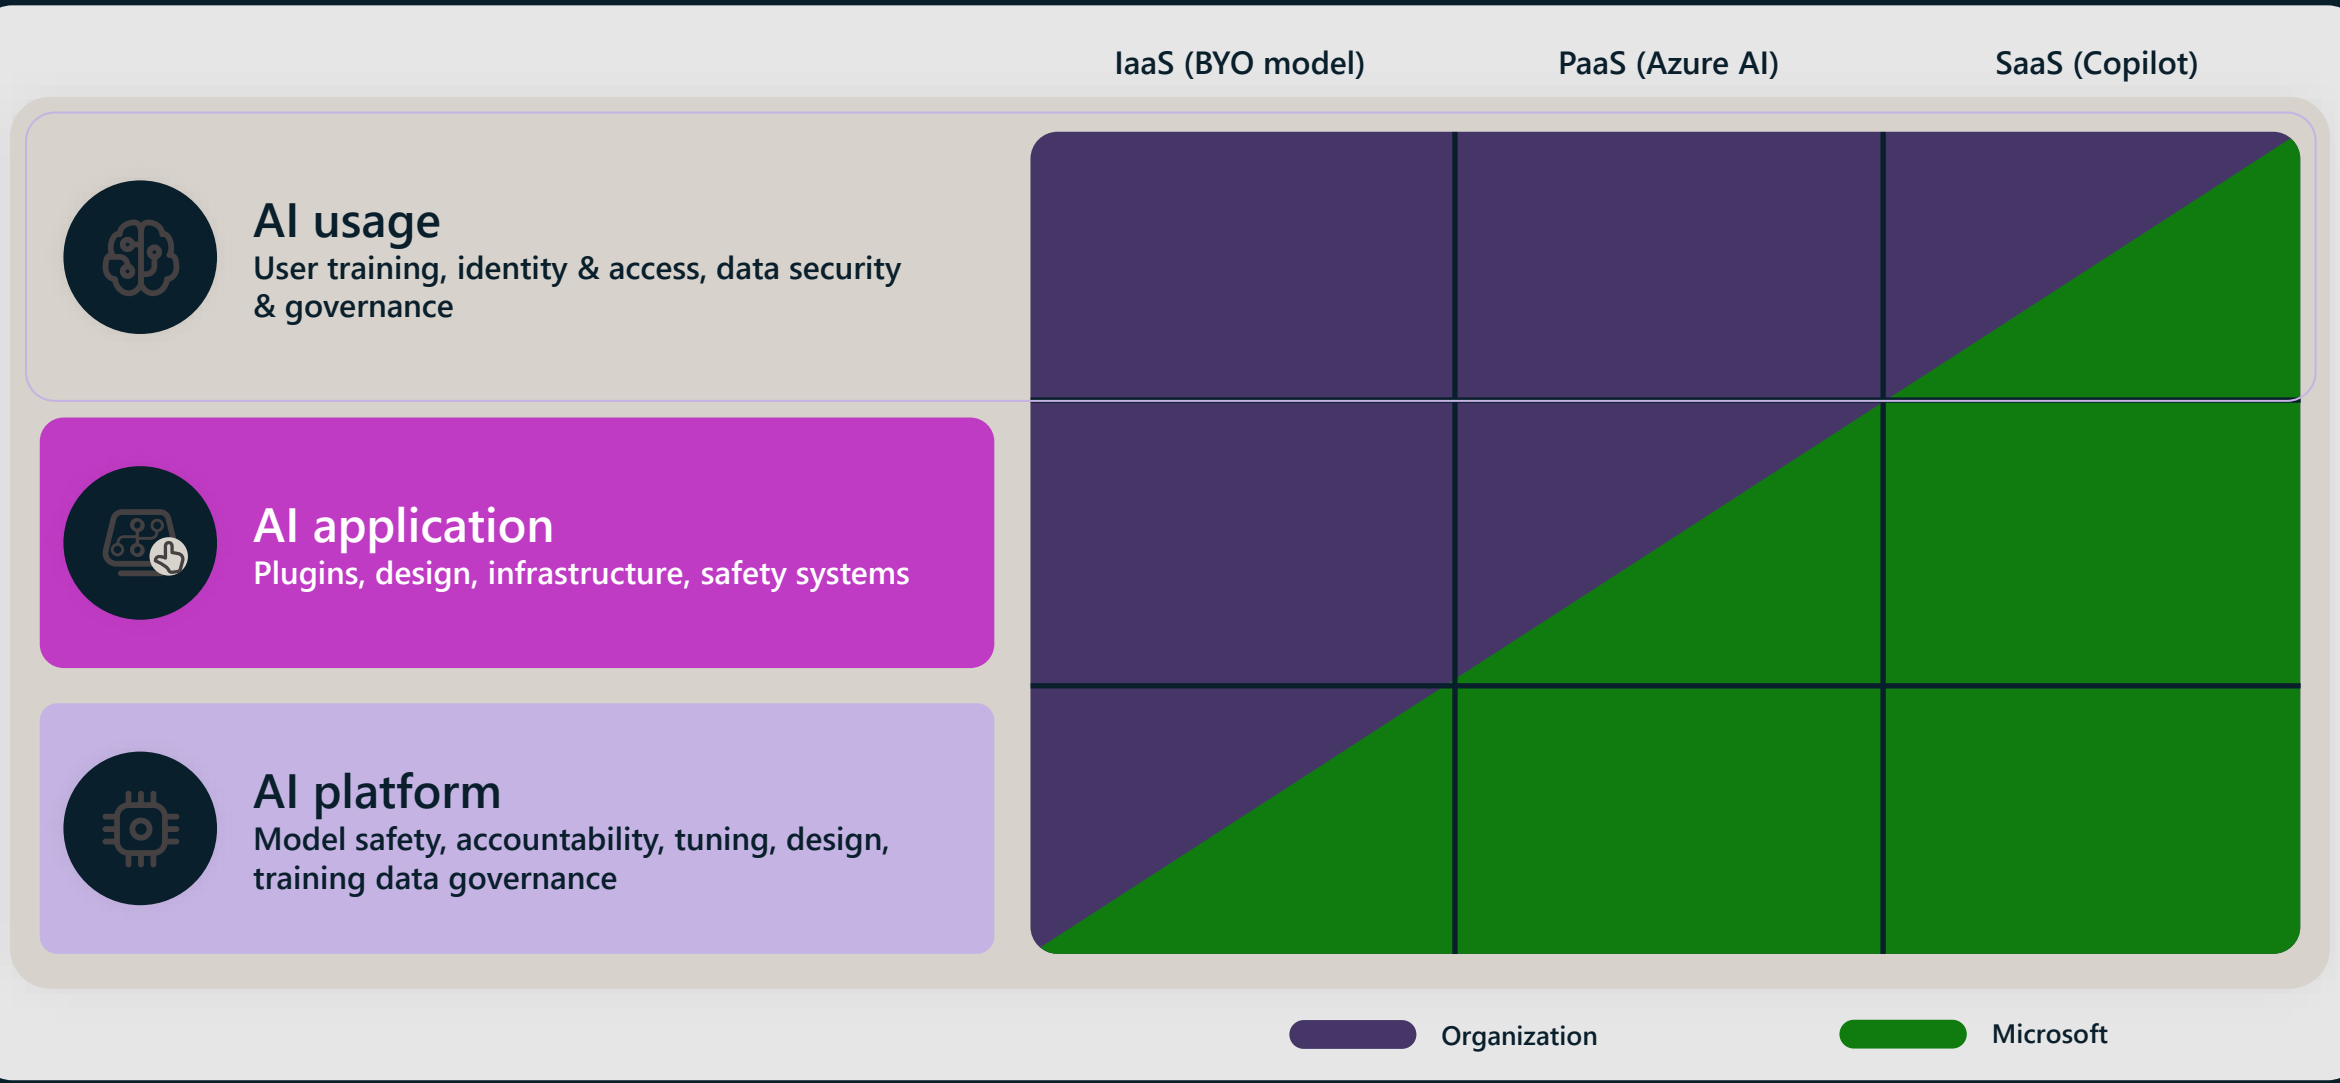

Microsoft's AI Principles

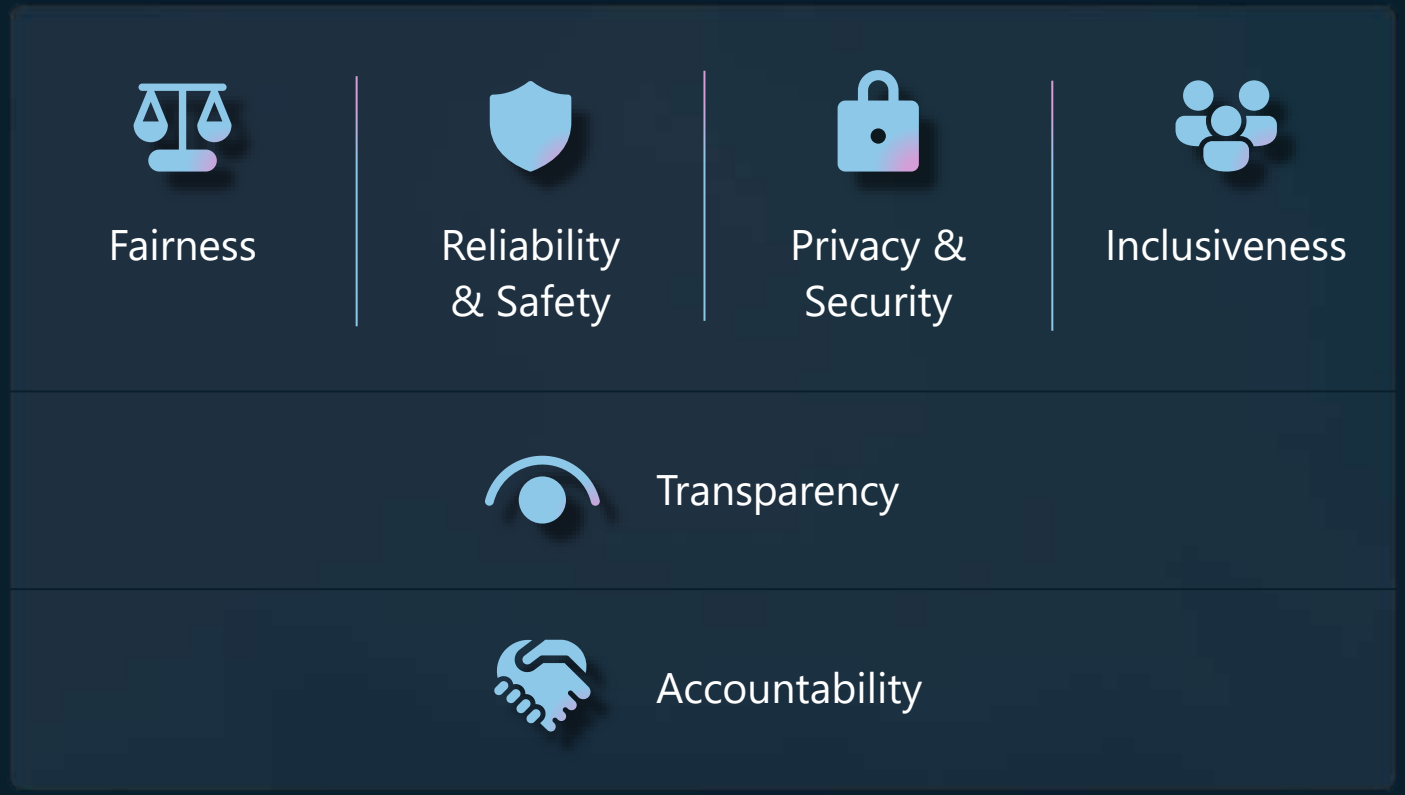

Microsoft Cloud — AI you can trust

Your data is **your** data.

Your data is **not** used to train the OpenAI foundation models without permission.

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

Copilot for Microsoft 365

Natural language

Large Language Models

+

Microsoft Graph - your data -

+

Microsoft 365 Apps

+

The web

Copilot for Microsoft 365

Built on Microsoft's comprehensive approach


Security

+


Compliance

+


Privacy

+


Responsible AI

Copilot top questions asked

Security/Compliance

How do we manage authentication and authorization of AI calls to the LLM to ensure data is safe?

How can I be aware when Copilot uses/returns sensitive information?

When and how will we be able to audit Copilot usage; see what content is being accessed?

Privacy/data residency

How and where are prompts stored and are they discoverable?

Where is my data processed?

How do we know our data is secure: how is our data encrypted?

Adoption/Readiness/Impact

What do we need to do to get ready?

What types of admin controls are available?

What type of success criteria / measures should I use to gauge the benefit and impact of giving users Copilot?

Data security and compliance

Start your data security and compliance journey today!

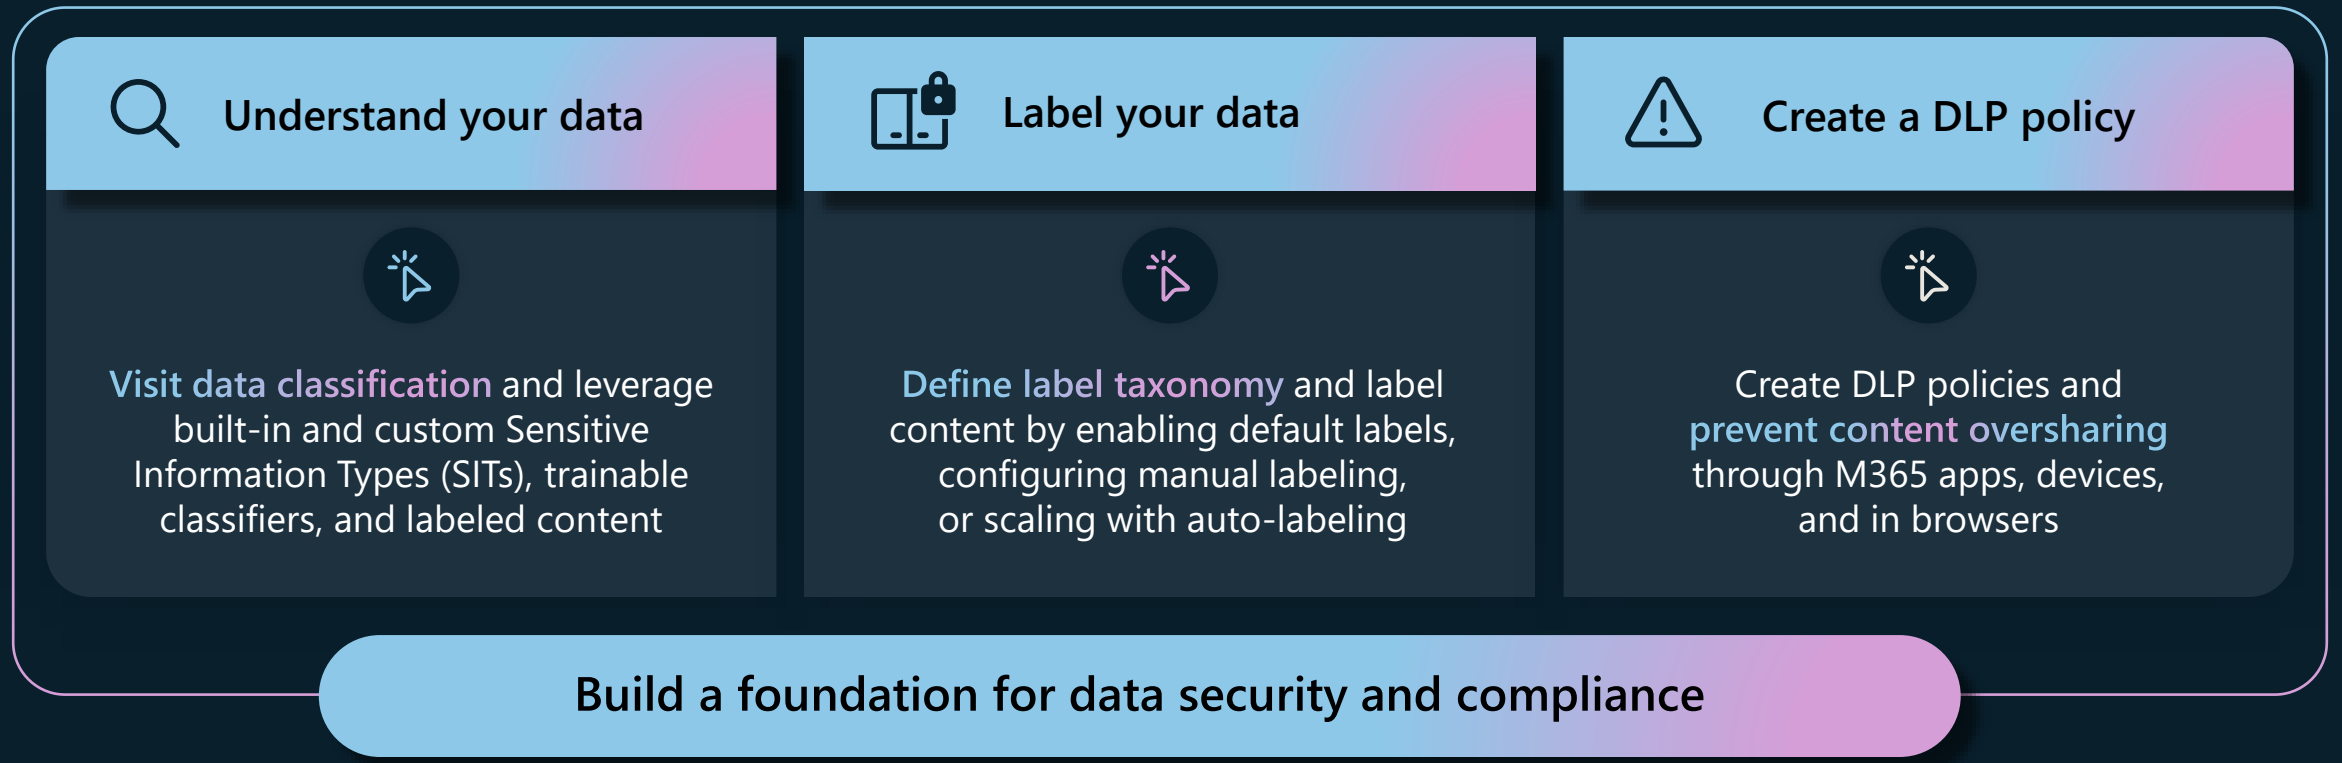

Demo

Sensitivity labels and Information Protection in Copilot

Admin config: labels, encryption, auto-labeling

Solutions
Catalog
Audit
Content search
Communication compliance
Data loss prevention
eDiscovery
Data lifecycle management
Information protection
Overview
Labels
Label policies
Auto-labeling
Information barriers
Insider risk management
Records management
Privacy risk management
Subject rights requests
Settings
More resources

When you turn this on, you'll be able to apply your sensitivity labels to files and schematized data assets in Microsoft Purview Data Map and Microsoft Defender for cloud. Learn more

Turn on

Sensitivity labels are used to classify email messages, documents, sites, and more. When a label is applied (automatically or by the user), the content or site is protected based on the settings you choose. For example, you can create labels that encrypt files, add content marking, and control user access to specific sites. Learn more about sensitivity labels

Create a label | Publish labels | Export | Refresh

6 items

Name	Priority	Scope	Created by	Last modified
Personal	0 - lowest	File, Email		Nov 7, 2023 4:12:14 PM
Public	1	File, Email		Nov 7, 2023 4:11:58 PM
General	2	File, Email		Nov 7, 2023 4:11:44 PM
Confidential	5	File, Email		Nov 7, 2023 4:11:10 PM
Anyone (unrestricted)	6	File, Email		Oct 17, 2023 9:55:21 AM
All Employees	7	File, Email, Site, UnifiedGroup		Nov 6, 2023 12:19:20 PM
Trusted People	8	File, Email		Oct 17, 2023 9:55:26 AM
Project Obsidian	9	File, Email, Meetings, Site, UnifiedGroup	MOD Administrator	Nov 6, 2023 12:21:49 PM
Highly Confidential	10	File, Email		Nov 7, 2023 4:10:54 PM

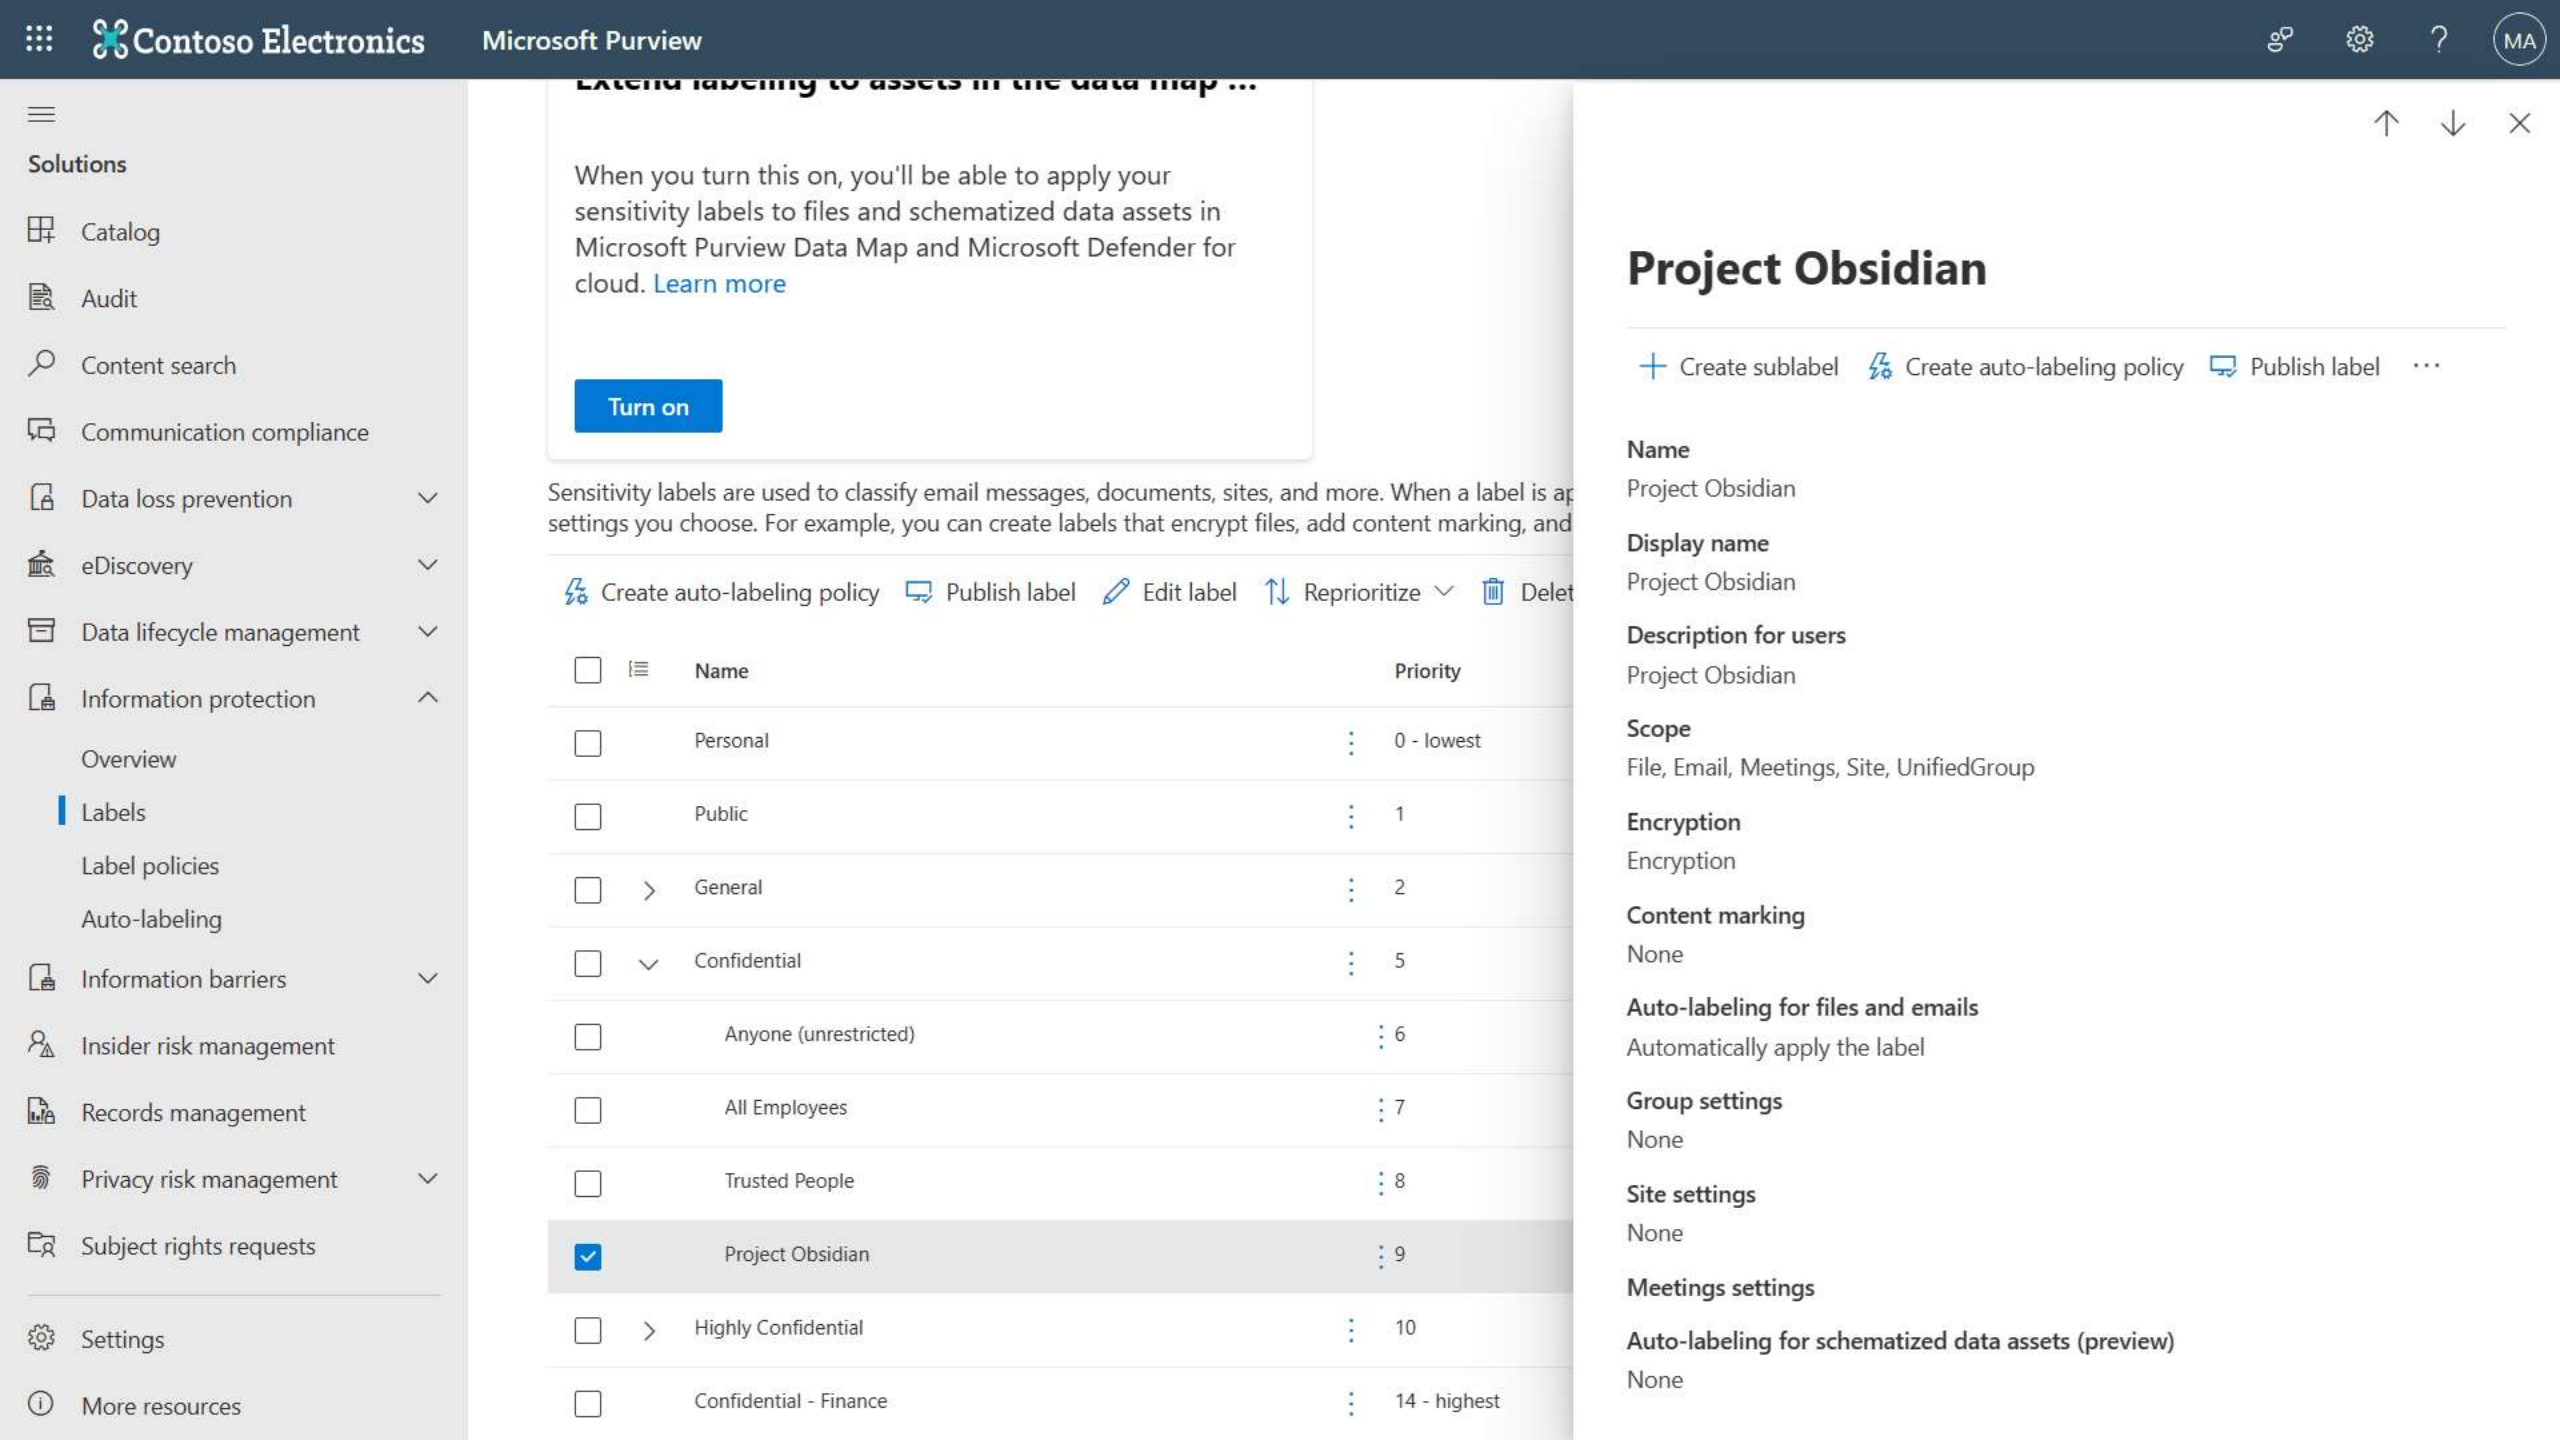
Contoso Electronics
Microsoft Purview
MA
Solutions
Catalog
Audit
Content search
Communication compliance
Data loss prevention
eDiscovery
Data lifecycle management
Information protection
Overview
Labels
Label policies
Auto-labeling
Information barriers
Insider risk management
Records management
Privacy risk management
Subject rights requests
Settings
More resources
When you turn this on, you'll be able to apply your sensitivity labels to files and schematized data assets in Microsoft Purview Data Map and Microsoft Defender for cloud. Learn more
Turn on
Sensitivity labels are used to classify email messages, documents, sites, and more. When a label is ap
settings you choose. For example, you can create labels that encrypt files, add content marking, and
Create auto-labeling policy
Publish label
Edit label
Reprioritize
Delet
Name
Priority
Personal 0 - lowest
Public 1
General 2
Confidential 5
Anyone (unrestricted) 6
All Employees 7
Trusted People 8
Project Obsidian 9
Highly Confidential 10
Confidential - Finance 14 - highest
Project Obsidian
Create sublabel
Create auto-labeling policy
Publish label
Name
Project Obsidian
Display name
Project Obsidian
Description for users
Project Obsidian
Scope
File, Email, Meetings, Site, UnifiedGroup
Encryption
Encryption
Content marking
None
Auto-labeling for files and emails
Automatically apply the label
Group settings
None
Site settings
None
Meetings settings
Auto-labeling for schematized data assets (preview)
None

Edit sensitivity label

- Label details
- Scope
- **Items**
- **Encryption**
- Auto-labeling for files and emails
- Groups & sites
- Schematized data assets (preview)
- Finish

Configure encryption settings

Turn on co-authoring Office desktop apps so multiple users can simultaneously edit documents that are labeled and encrypted by sensitivity labels. Learn more about this setting

Go to co-authoring setting

Assign permissions now or let users decide?

Assign permissions now

The encryption settings you choose will be automatically enforced when the label is applied to email and Office files.

User access to content expires

Never

Allow offline access

Always

Assign permissions to specific users and groups *

Assign permissions

1 item

Users and groups	Permissions	Edit	Delete
MODERNCOMMS382604.onmicrosoft.com	Viewer		

Back Next Cancel

Edit sensitivity label

- Label details
- Scope
- **Items**
- Encryption
- **Auto-labeling for files and emails**
- Groups & sites
- Schematized data assets (preview)
- Finish

Auto-labeling for files and emails

Since encryption is turned on, a large amount of content might be automatically encrypted when this label is applied. Turning on encryption impacts Office files (Word, PowerPoint, Excel) and PDF files that have this label applied. Because the files will be encrypted for security reasons, performance will be slow when the files are opened or saved, and some SharePoint and OneDrive features will be limited or unavailable. Learn more

Detect content that matches these conditions

Content contains

Group name *: Default

Group operator: Any of these

Sensitive info types

Sensitive terms: Low confidence — Instance count: 1 to Any

Add

Create group

Add condition

When content matches these conditions

Automatically apply the label

Back Next Cancel

Audit

Home
Compliance Manager
Data classification
Data connectors
Alerts
Policies
Roles & scopes
Trials

Solutions
Catalog
Audit
Content search
Communication compliance
Data loss prevention
eDiscovery
Data lifecycle management
Information protection
Information barriers
Insider risk management

Monday, Nov 6, 2023 12:00:00 AM to Wednesday, Nov 8, 2023 12:00:0

Export

Date	IP Address	User
Nov 7, 2023 12:41 PM	2001:4898:80e8:36:f986:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:40 PM	2001:4898:80e8:36:f986:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:36 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:25 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:24 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:20 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:19 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:19 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:16 PM	2001:4898:80e8:37:f985:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:11 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 10:20 AM	24.17.224.43	AlexW@MODERNCOMM

Users
AlexW@MODERNCOMMS382604.OnMicrosoft.com

Activity
Interacted with Copilot

Item

Details

CreationTime
2023-11-07T18:20:46

Id
8a2bfba6-c241-47fd-a6e5-6995b57590b0

Operation
CopilotInteraction

OrganizationId
b9ba404e-37f1-4363-bb0b-fc387ddfabe6

RecordType
261

UserKey
23f35b20-f05f-42f6-9ce8-d53c9edd3ce0

UserType
0

Version
1

Workload
Copilot

Close

Home
Compliance Manager
Data classification
Data connectors
Alerts
Policies
Roles & scopes
Trials

Solutions

Catalog
Audit
Content search
Communication compliance
Data loss prevention
eDiscovery
Data lifecycle management
Information protection
Information barriers
Insider risk management

Monday, Nov 6, 2023 12:00:00 AM to Wednesday, Nov 8, 2023 12:00:0

Export

Date	IP Address	User
Nov 7, 2023 12:41 PM	2001:4898:80e8:36:f986:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:40 PM	2001:4898:80e8:36:f986:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:36 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:25 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:24 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:20 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:19 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:19 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:16 PM	2001:4898:80e8:37:f985:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 12:11 PM	2001:4898:80e8:1:f9bb:fb39:1a3e:5fad	AlexW@MODERNCOMM
Nov 7, 2023 10:20 AM	24.17.224.43	AlexW@MODERNCOMM

1

Workload
Copilot

ClientIP
24.17.224.43

UserId
AlexW@MODERNCOMMS382604.OnMicrosoft.com

CopilotEventData


```
{
    "AccessedResources": [
        {
            "Id": "https://moderncomms382604.sharepoint.com/sites/(
            "Name": "kickoff.pptx",
            "SensitivityLabelId": "1f800ac5-34ff-40e6-aab6-2802e7f(
            "Type": "pptx"
        },
        {
            "Id": "https://moderncomms382604.sharepoint.com/sites/(
            "Name": "Design update.docx",
            "SensitivityLabelId": "1f800ac5-34ff-40e6-aab6-2802e7f(
            "Type": "docx"
        },
        {
            "Id": "https://moderncomms382604.sharepoint.com/sites/(
            "Name": "Next generation chip.docx",
            "SensitivityLabelId": "1f800ac5-34ff-40e6-aab6-2802e7f(
            "Type": "docx"
        }
    ],
    "AppHost": "bizchat",
    "Contexts": [],
    "MessageIds": [],
    "ThreadId": "19:qtOmIM5vzHCDQ1PGzya5KfTJfuhVOpYJcNbi1LDvqx81@th
}
```


Close

Data retention

Data lifecycle management

Overview | Retention policies | Labels | Label policies | Adaptive scopes | Policy lookup | Import

Your users create a lot of content every day, from emails to Teams and Yammer conversations. Use retention policies to keep the content y
policies.

If your role group permissions are restricted to a specific set of users or groups, you'll only be able to manage policies for those users or groups. Learn m

New retention policy | Edit | Delete | Disable policy | Export | Inactive mailbox | Refresh

Name	Created by
Copilot interactions	MOD Administrator
Employee Records	Megan Bowen
Personal Financial PII	Megan Bowen
Sensitivity	Megan Bowen
U.S. Financial Data Policy	Megan Bowen

MOD Administrator

Copilot interactions

Status
Enabled (Success)

Admin units (preview)
Full directory

Applies to content in these locations
Teams chats and Microsoft 365 Copilot interactions

Settings

Retention period
Keep content, and delete it if it's older than 5 years

Preservation lock
No

- Name
- Administrative Units
- **Type**
- **Locations**
- Retention settings
- Finish

Choose where to apply this policy

The policy will apply to content that's stored in the locations you choose.

You can set up data connectors to import content from non-Microsoft apps like Slack, WhatsApp and many more, for use with this solution. Set up now

Policies that apply to Teams chats or Teams channel messages can't include other locations.

Status	Location	Applicable Content	Included	Excluded
Off	Teams channel messages	Messages from channel conversations and channel meetings. Doesn't apply to Teams private channel messages. More details		
On	Teams chats and Microsoft 365 Copilot interactions	Messages from individual chats, group chats, meeting chats, bot chats, and Microsoft 365 Copilot interactions. More details	All users Edit	None Edit

Back Next Cancel

Secure and govern Copilot with Microsoft Security

Baseline

Copilot Microsoft 365
+ Business Standard / Office 365 E3

Multi-factor Authentication

Audit logging

Core

Copilot Microsoft 365
+ Microsoft 365 E3 / Business Premium
+ SharePoint Advanced Management

Conditional Access

Manual sensitivity labels

Data loss prevention policies

Advanced SharePoint sitewide access controls and reporting

Unified endpoint management

Best-in-class

Copilot Microsoft 365
+ Microsoft 365 E5
+ SharePoint Advanced Management

Conditional Access based on identity risk

Automatically apply sensitivity labels

Automatically remove inactive content

Prevent data leak on endpoint devices

Detect non-compliant usage

Copilot privacy and data location

Microsoft’s approach to privacy

You control your data

You know where your data is located

We secure your data at rest and in transit

We defend your data

You control your data

Microsoft does not use customer data to train foundation LLMs that Copilot uses.

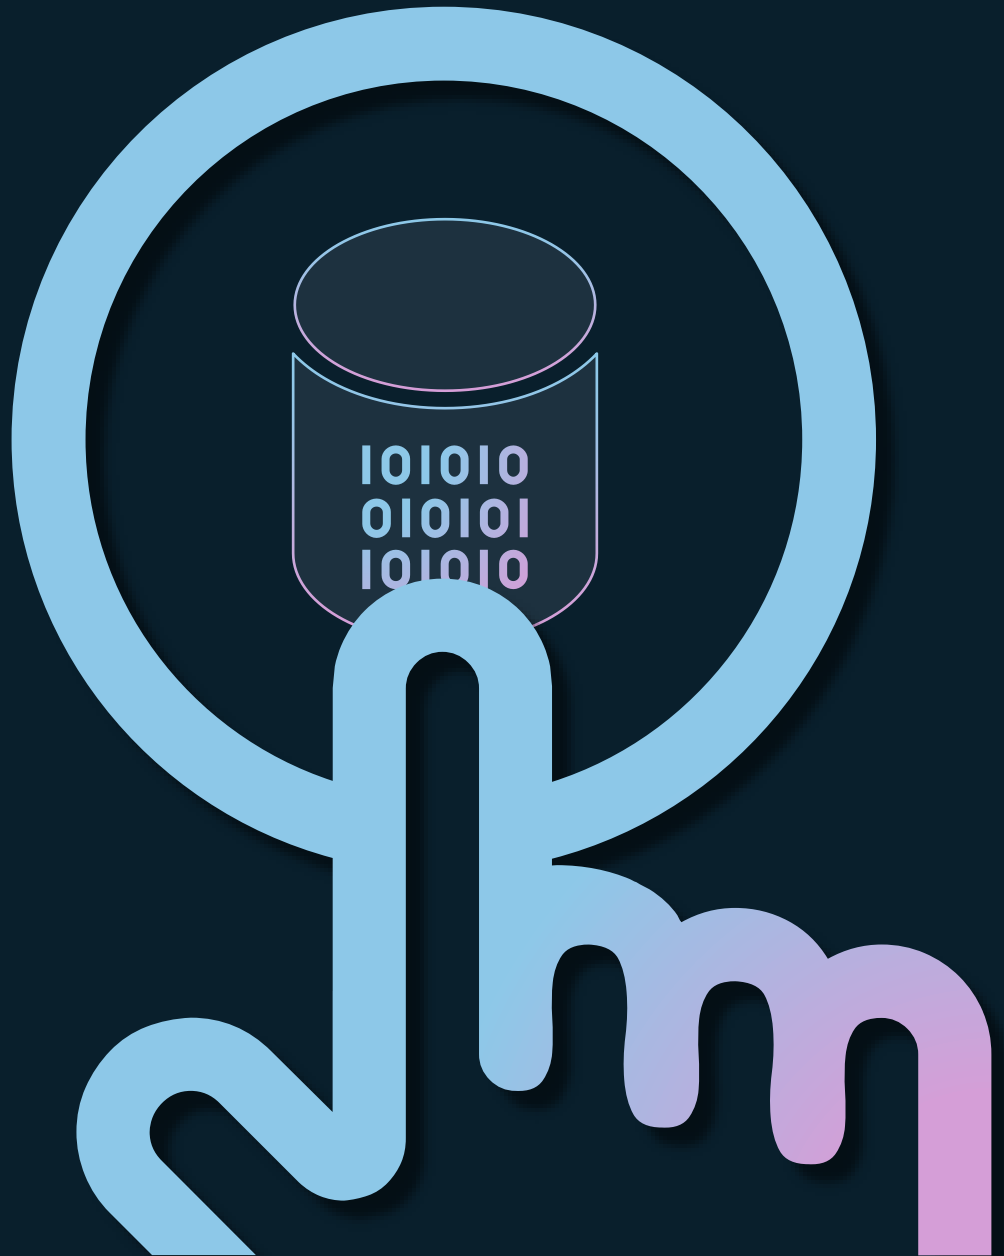

Microsoft encrypts customer data at rest and in transit

Encryption both at rest and in transit helps provided protections against unauthorized access, disclosure, sharing or usage of data.

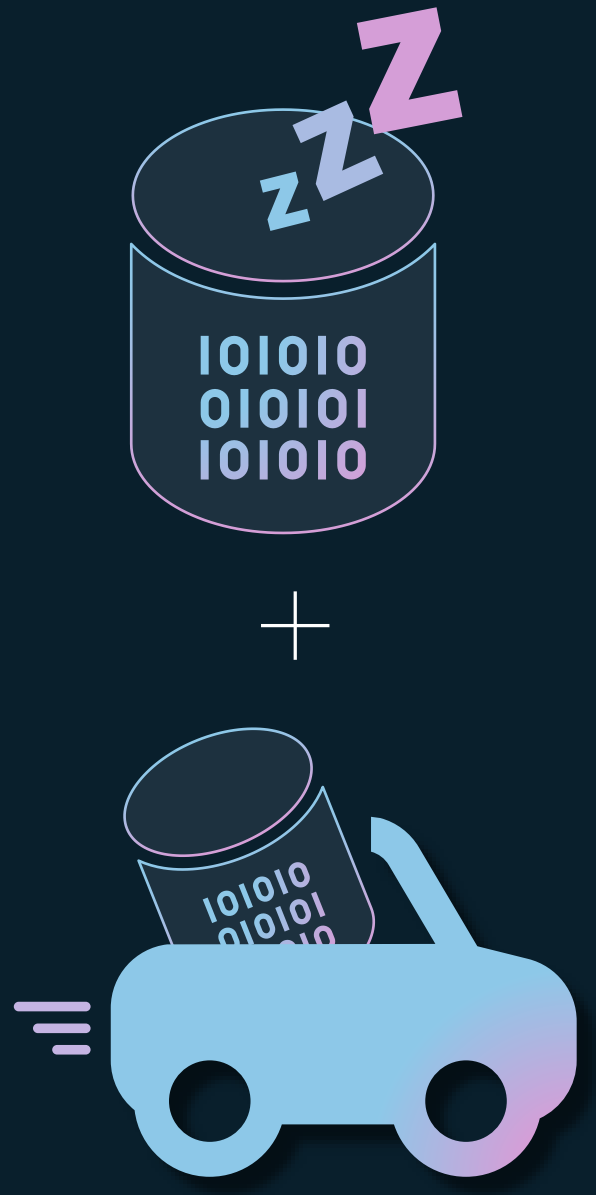

Where does Microsoft store customer data?

Customers specify a country or region while signing up for a new Microsoft 365 tenant. That selection determines the tenant's default geography. Microsoft makes decisions on where to store customer data by **combining that default geography with the available geographies for each service**.

The EU Data Boundary provides more expansive commitments

EUDB terms provide **contractual commitments for storing and processing EU + EFTA customer data** within the European Union.*

*NOTE: Some services are temporarily or permanently exempted from this boundary. Specific exemptions are listed on Microsoft Learn.

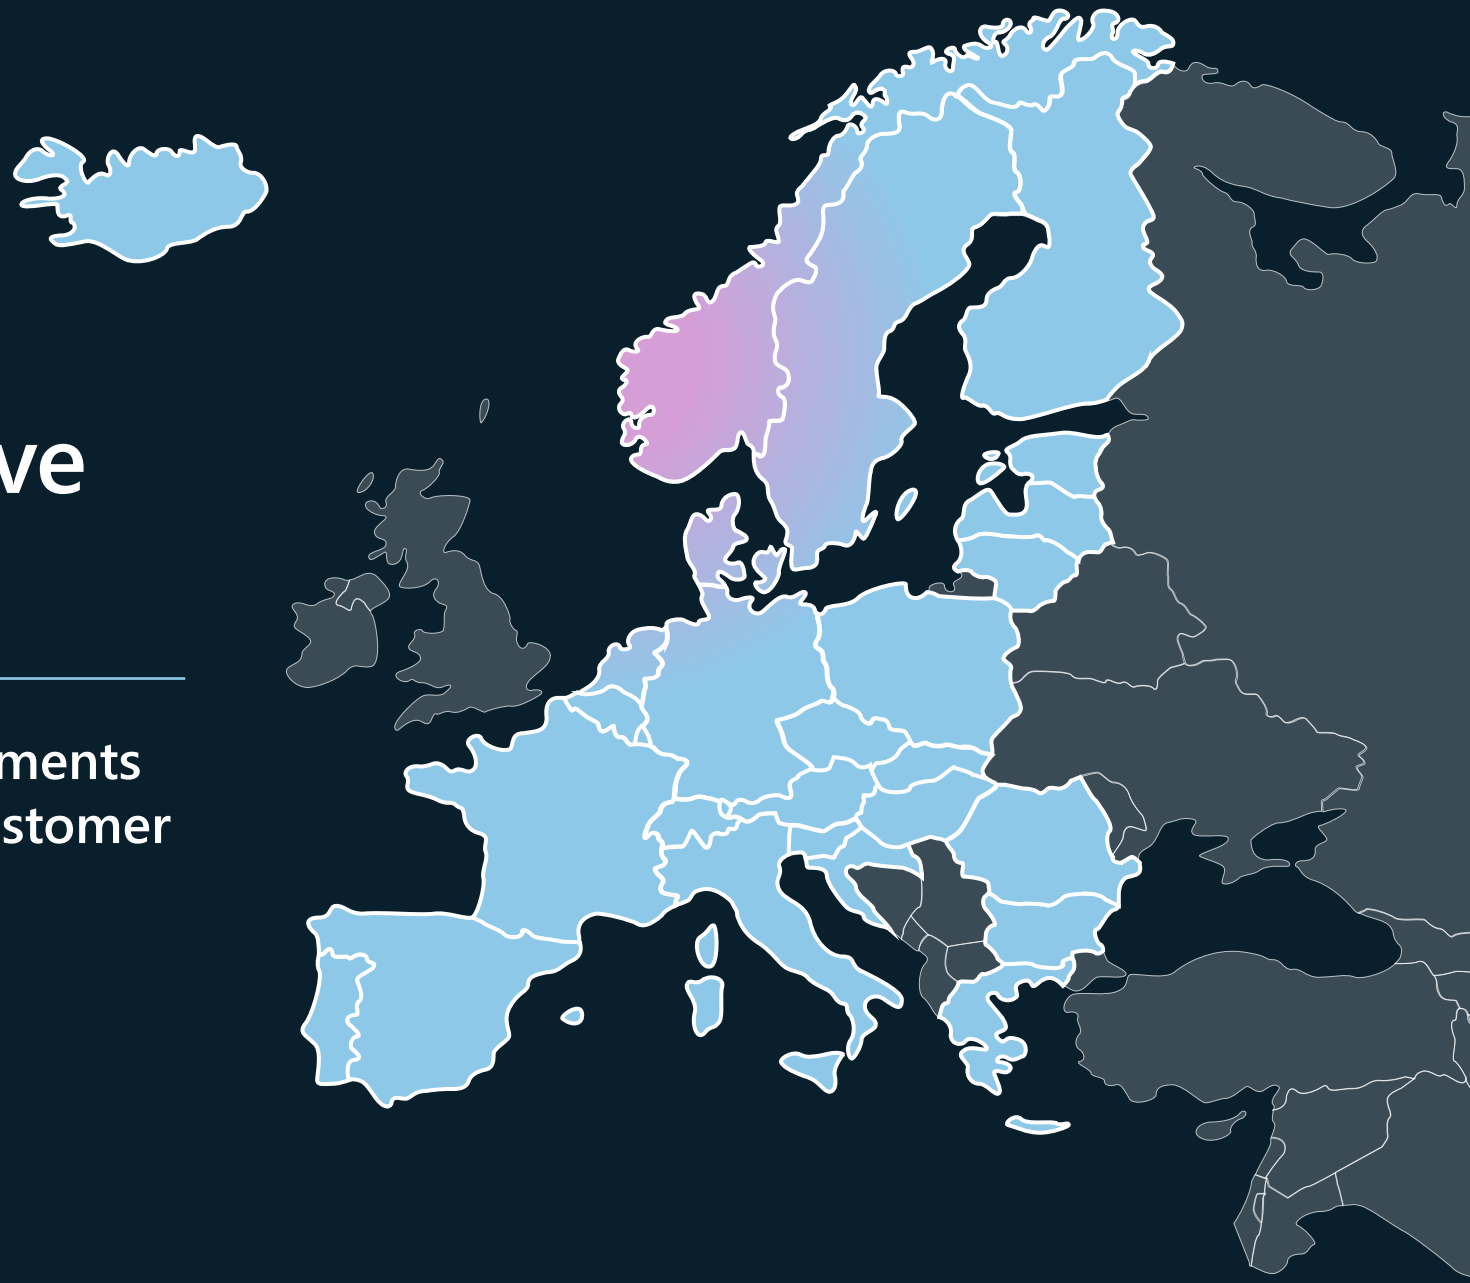

Global reach, local choices

Geography and availability

Data residency

Currently available in 19 regions (outside US)

Available to restricted and unrestricted industries

- Australia
- Brazil
- Canada
- France
- Germany
- India
- Israel
- Italy
- Japan
- Norway
- Poland
- Qatar
- Singapore*
- South Africa
- South Korea
- Sweden
- Switzerland
- United Arab Emirates
- United Kingdom

Adding 11 new regions by 2025

- Austria
- Chile
- Denmark
- Greece
- Indonesia
- Malaysia
- Mexico
- New Zealand
- Saudi Arabia
- Spain
- Taiwan

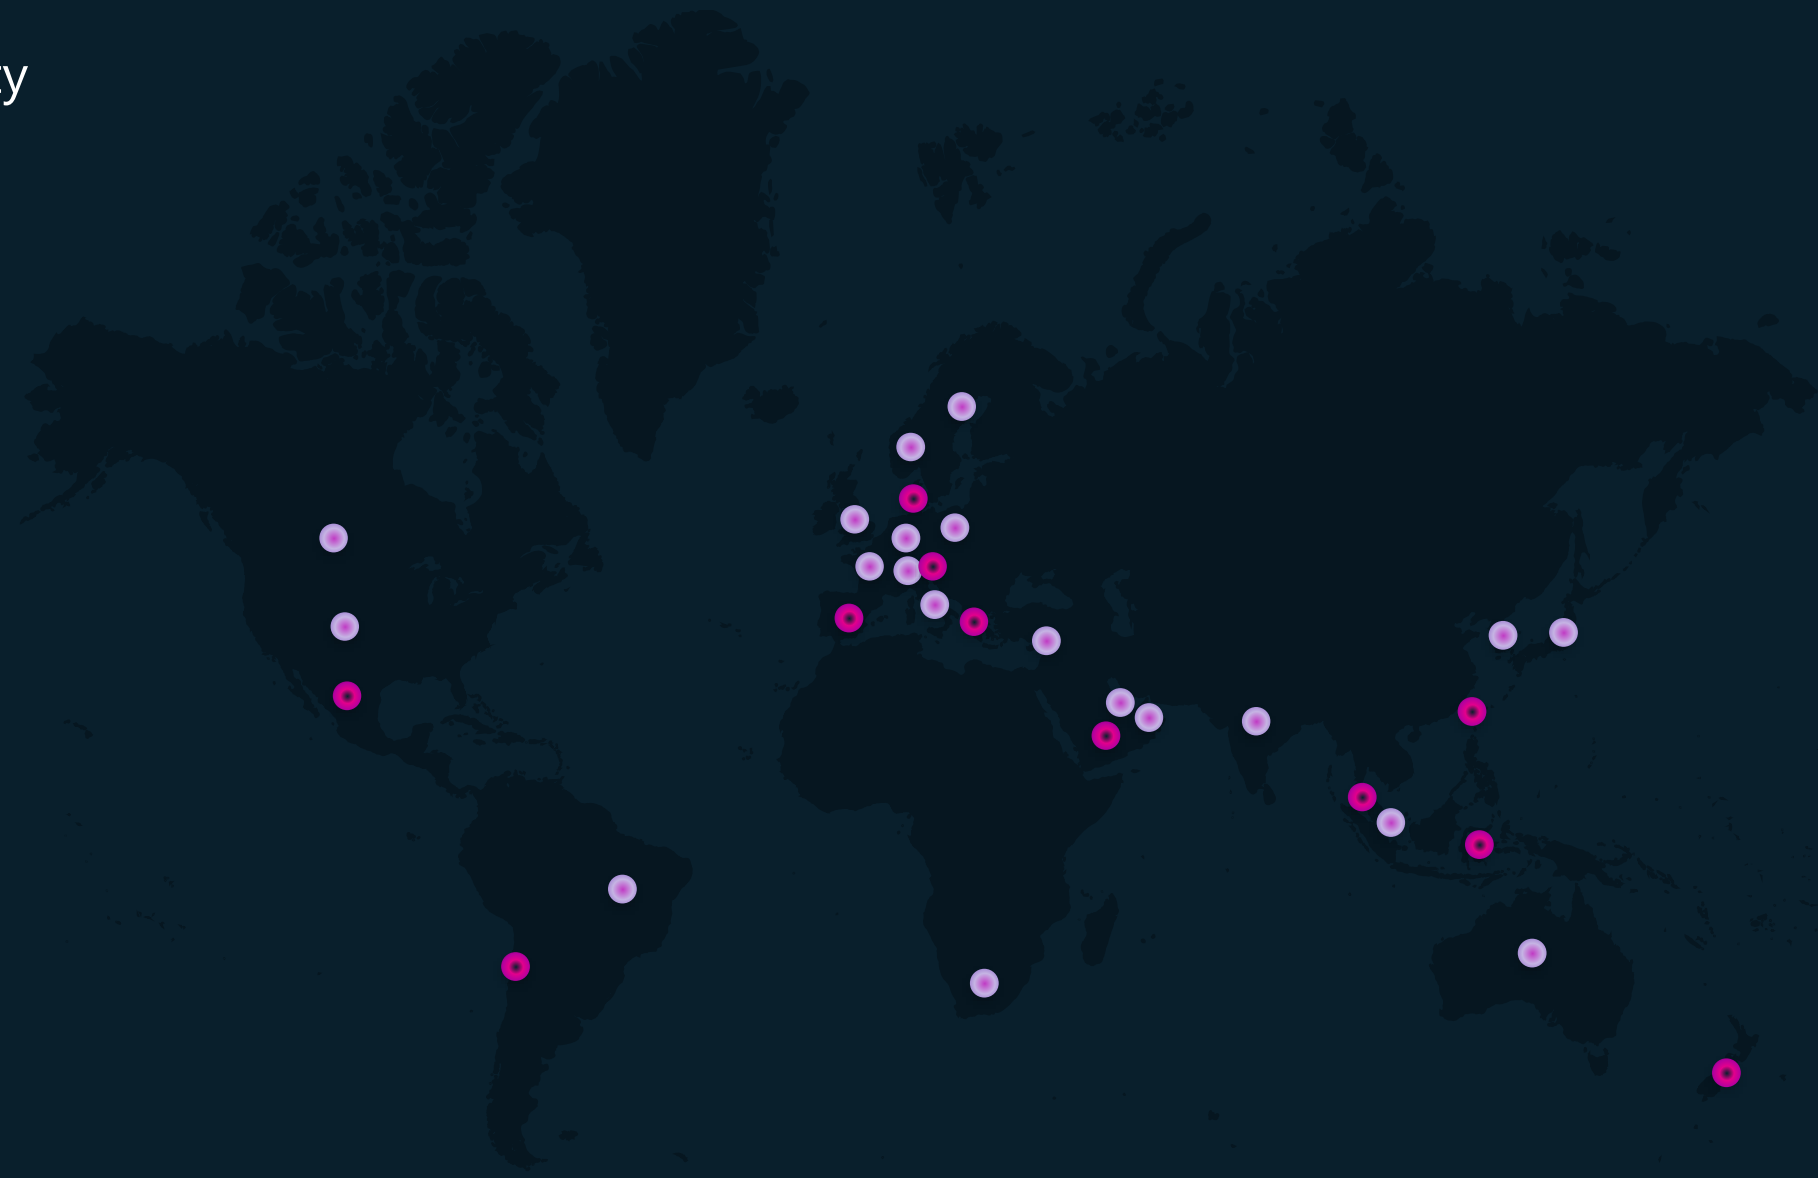

* Singapore local region geography is limited to specific customers

Where is Copilot for Microsoft 365 storing data

Data storage at rest

Copilot uses the same storage location as Microsoft 365 customer data content

Microsoft 365 tenant sign up country determines data location[1]

Single tenant data location is the same as default geo storage location

Multi-Geo add-on data location is determined by users' locations[2]

Customer Data content

Copilot for Microsoft 365 “content of interactions”

Stored customer data includes prompts, responses, and citations

Record of interactions is in the user’s Copilot interaction history

Users can delete their interaction history by going to My Account portal

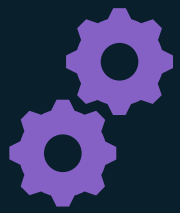

Other M365 Customer Data content

Covers in-scope Customer content in Exchange, OneDrive, SharePoint, Teams

Established commitments for data residency via Product Terms

Add-on SKUs: Advanced Data Residency (ADR) and Multi-Geo Capabilities

[1] Customer must define users’ *PreferredDataLocation* (PDL) in Microsoft Entra ID

2 Once Geo location is associated with signup Entra ID it cannot be altered at the tenant level.

New data residency treatment for Copilot for Microsoft 365

- Data residency commitments have been updated to include Copilot for Microsoft 365.
- Generally available as of 1 March 2024.
- Commitment is available in the following:
  - Product Terms
  - Advanced Data Residency (ADR) add-on
  - Multi-Geo Capabilities add-on

Product Terms

Location of Customer Data at Rest for Core Online Services

"Office 365 Services. If Customer provisions its tenant in Australia, Brazil, Canada, the European Union, France, Germany, India, Japan, Norway, Qatar, South Africa, South Korea, Sweden, Switzerland, the United Kingdom, the United Arab Emirates, or the United States, Microsoft will store the following Customer Data at rest only within that Geo: (1) Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), (2) SharePoint Online site content and the files stored within that site, (3) files uploaded to OneDrive for Business, and (4) Microsoft Teams chat messages (including private messages, channel messages, meeting messages and images used in chats), and for customers using Microsoft Stream (Classic) (on SharePoint), meeting recordings, ***and (5) any stored content of interactions with Microsoft Copilot for Microsoft 365 to the extent not included in the preceding commitments.*** *If Customer purchases an Advanced Data Residency subscription, then Microsoft will store certain Customer Data at rest in the applicable Geo in accordance with this section and the "Advanced Data Residency Commitments" section of the product documentation at https://aka.ms/adroverview.*

Published in Product Terms

ADR

"Content of Interactions" such as the user's prompt and Microsoft Copilot's response, including citations to any information used to ground Microsoft Copilot's response.

Published in Microsoft Learn docs in Advanced Data Residency in Microsoft 365 article.

Multi-Geo Capabilities

Multi-Geo capabilities in Microsoft Copilot for Microsoft 365 enable content of interactions with Microsoft Copilot for Microsoft 365 to be stored at rest in a specified Macro Region Geography or Local Region Geography location. Microsoft Copilot for Microsoft 365 uses the Preferred Data Location (PDL) for users and groups to determine where to store data. If the PDL isn't set or is invalid, data is stored in the Tenant's Primary Provisioned Geography location. The Geography where the content of interactions with Microsoft Copilot for Microsoft 365 are stored is determined by the PDL of the user interacting with Microsoft Copilot for Microsoft 365. This means that the storage of content of interactions for users in different regions will be based on their respective PDL configurations.

Published in Microsoft Learn docs in Microsoft 365 Multi-Geo article.

Data Location in Microsoft 365 admin center (MAC)

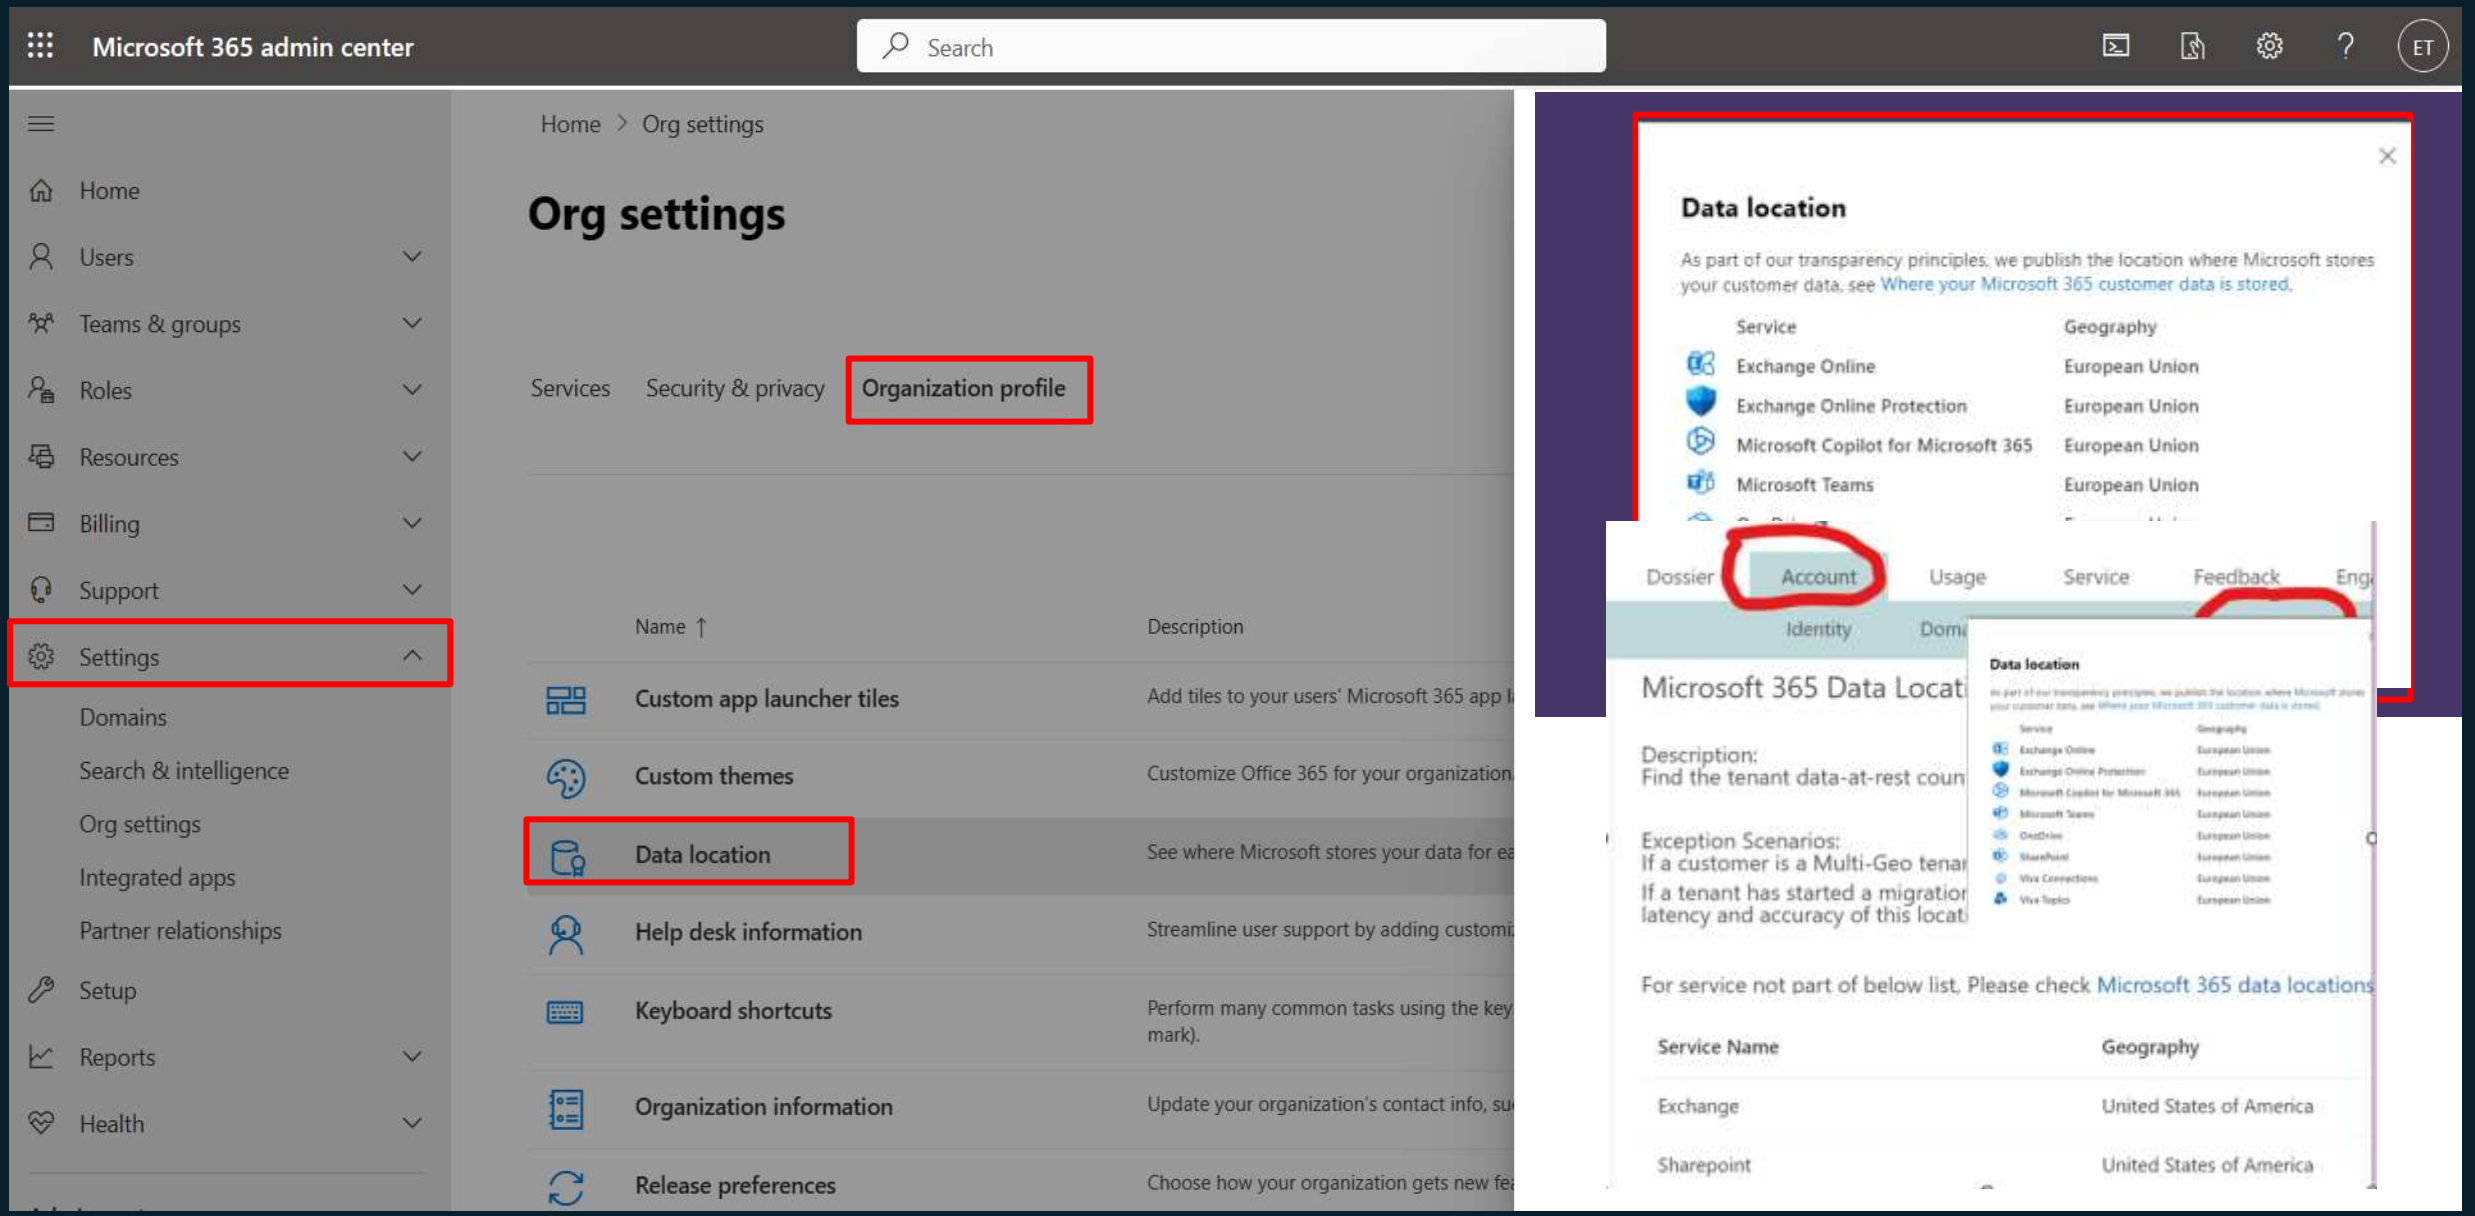

Note: Display of data location for Copilot for Microsoft 365 will be available in the coming months

Microsoft offers several data residency options

Product Terms Data Residency

Advanced Data Residency (ADR) Add-on

Multi-Geo Capabilities Add-on

European Union Data Boundary (EUDB)

Copilot for Microsoft 365 commitments

Microsoft Product Terms

- Universal License Terms- Microsoft Generative AI services
- Microsoft Copilot copyright commitment
- EU Data Boundary Services
- Core Online Services*

Microsoft Products and Services Data Protection Addendum

Advanced Data Residency (ADR)* Add-on

Multi-Geo Capabilities* Add-on

*updates for 2024

How Copilot works

Microsoft Copilot for Microsoft 365 architecture

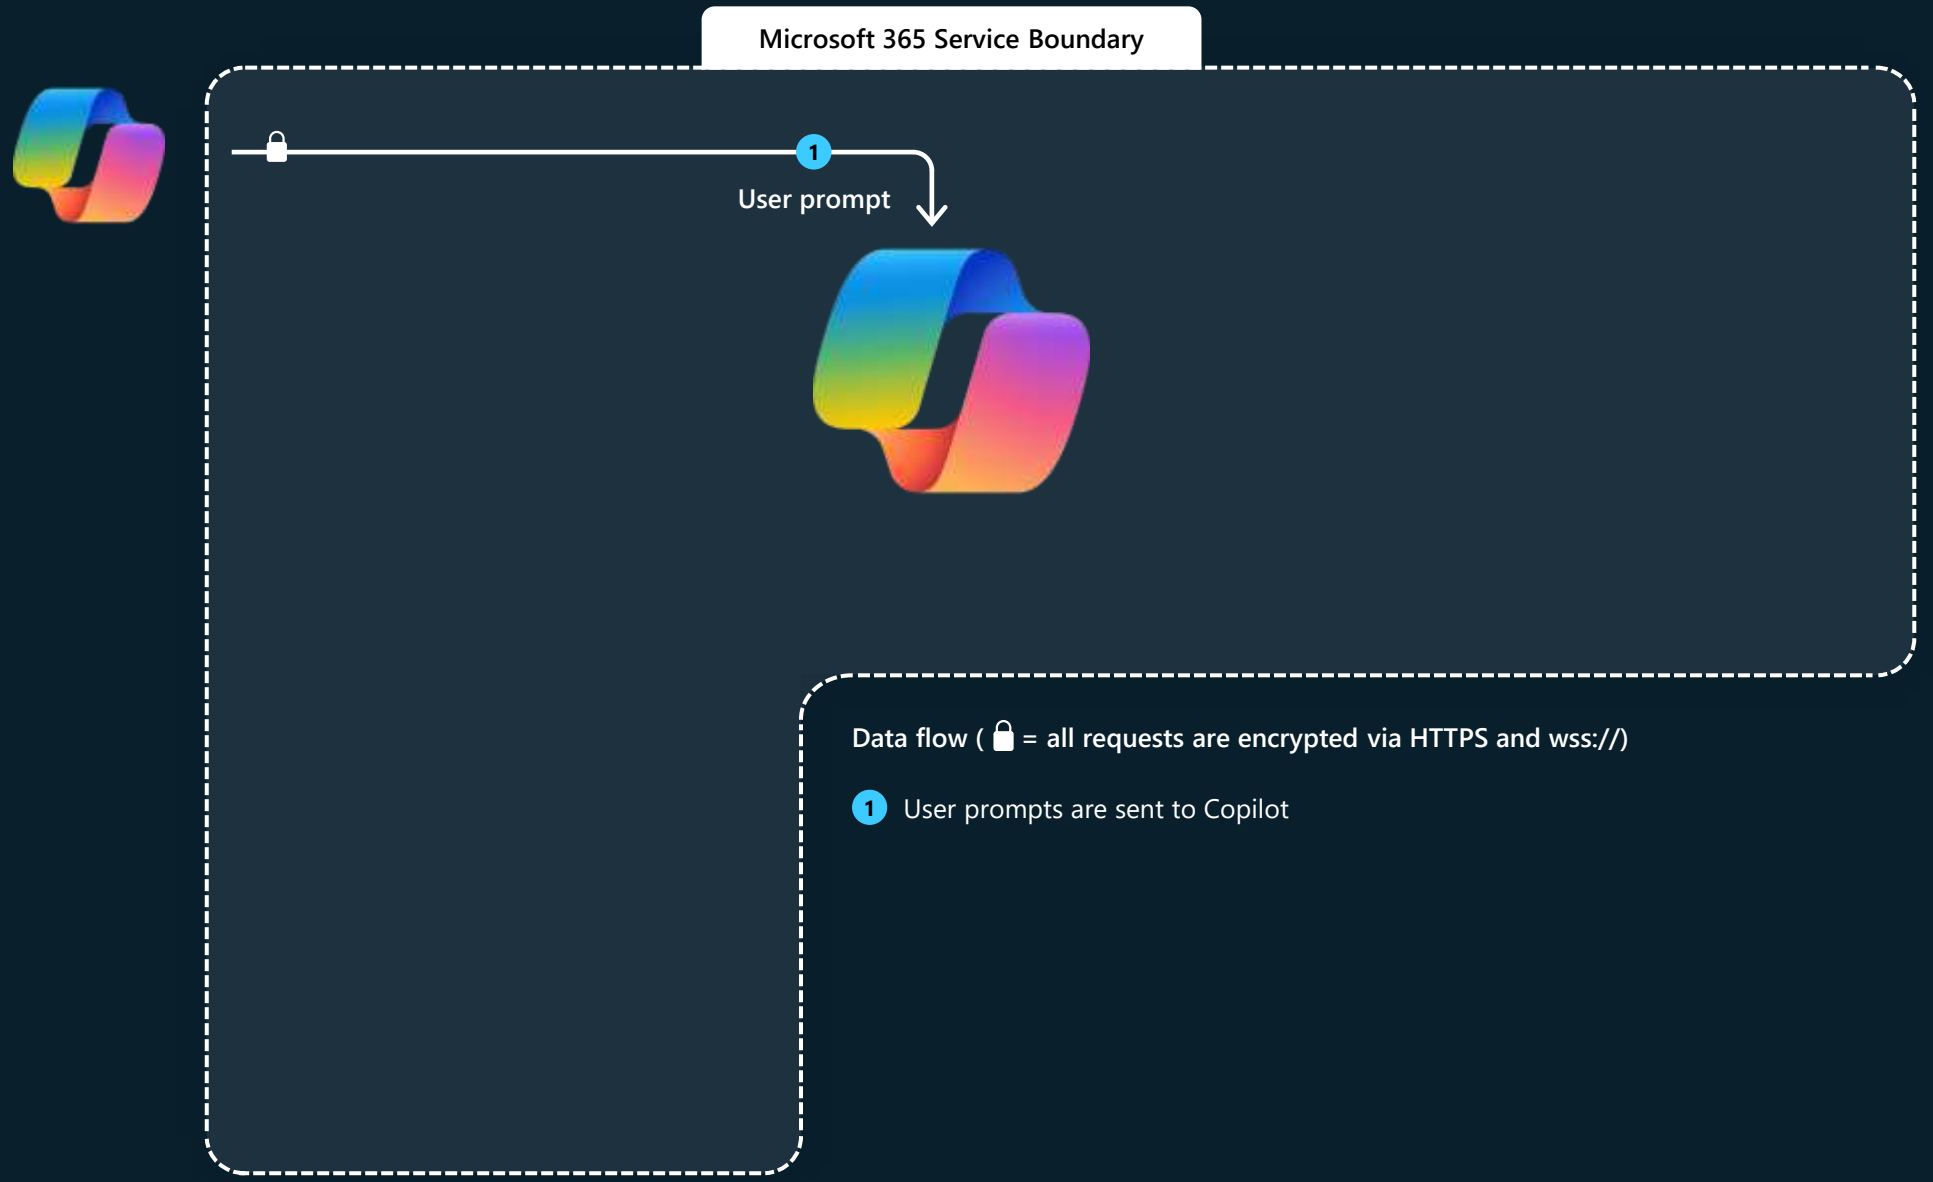

Microsoft Copilot for Microsoft 365 architecture

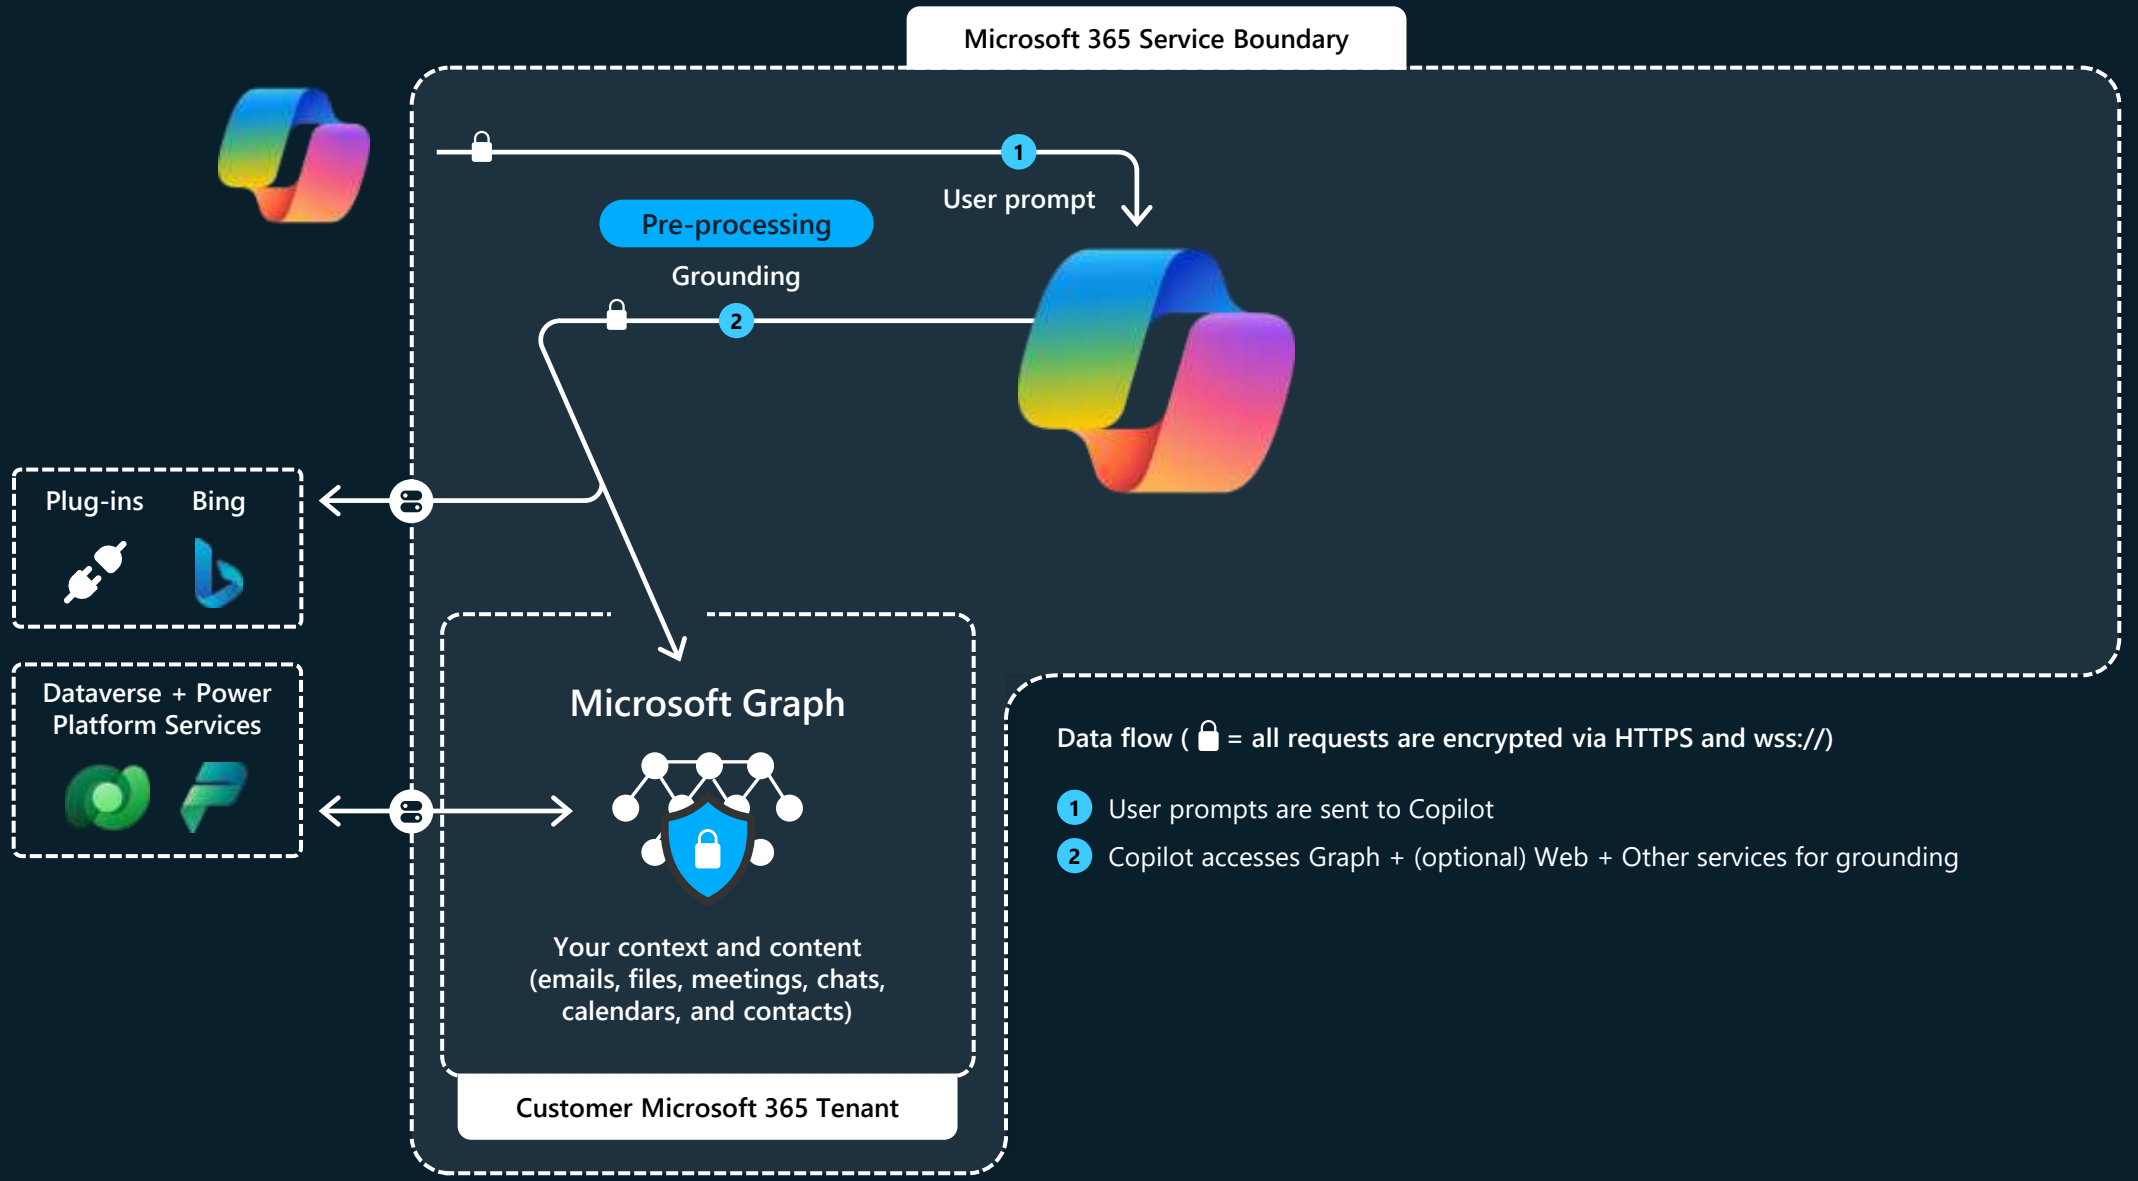

Microsoft Copilot for Microsoft 365 architecture

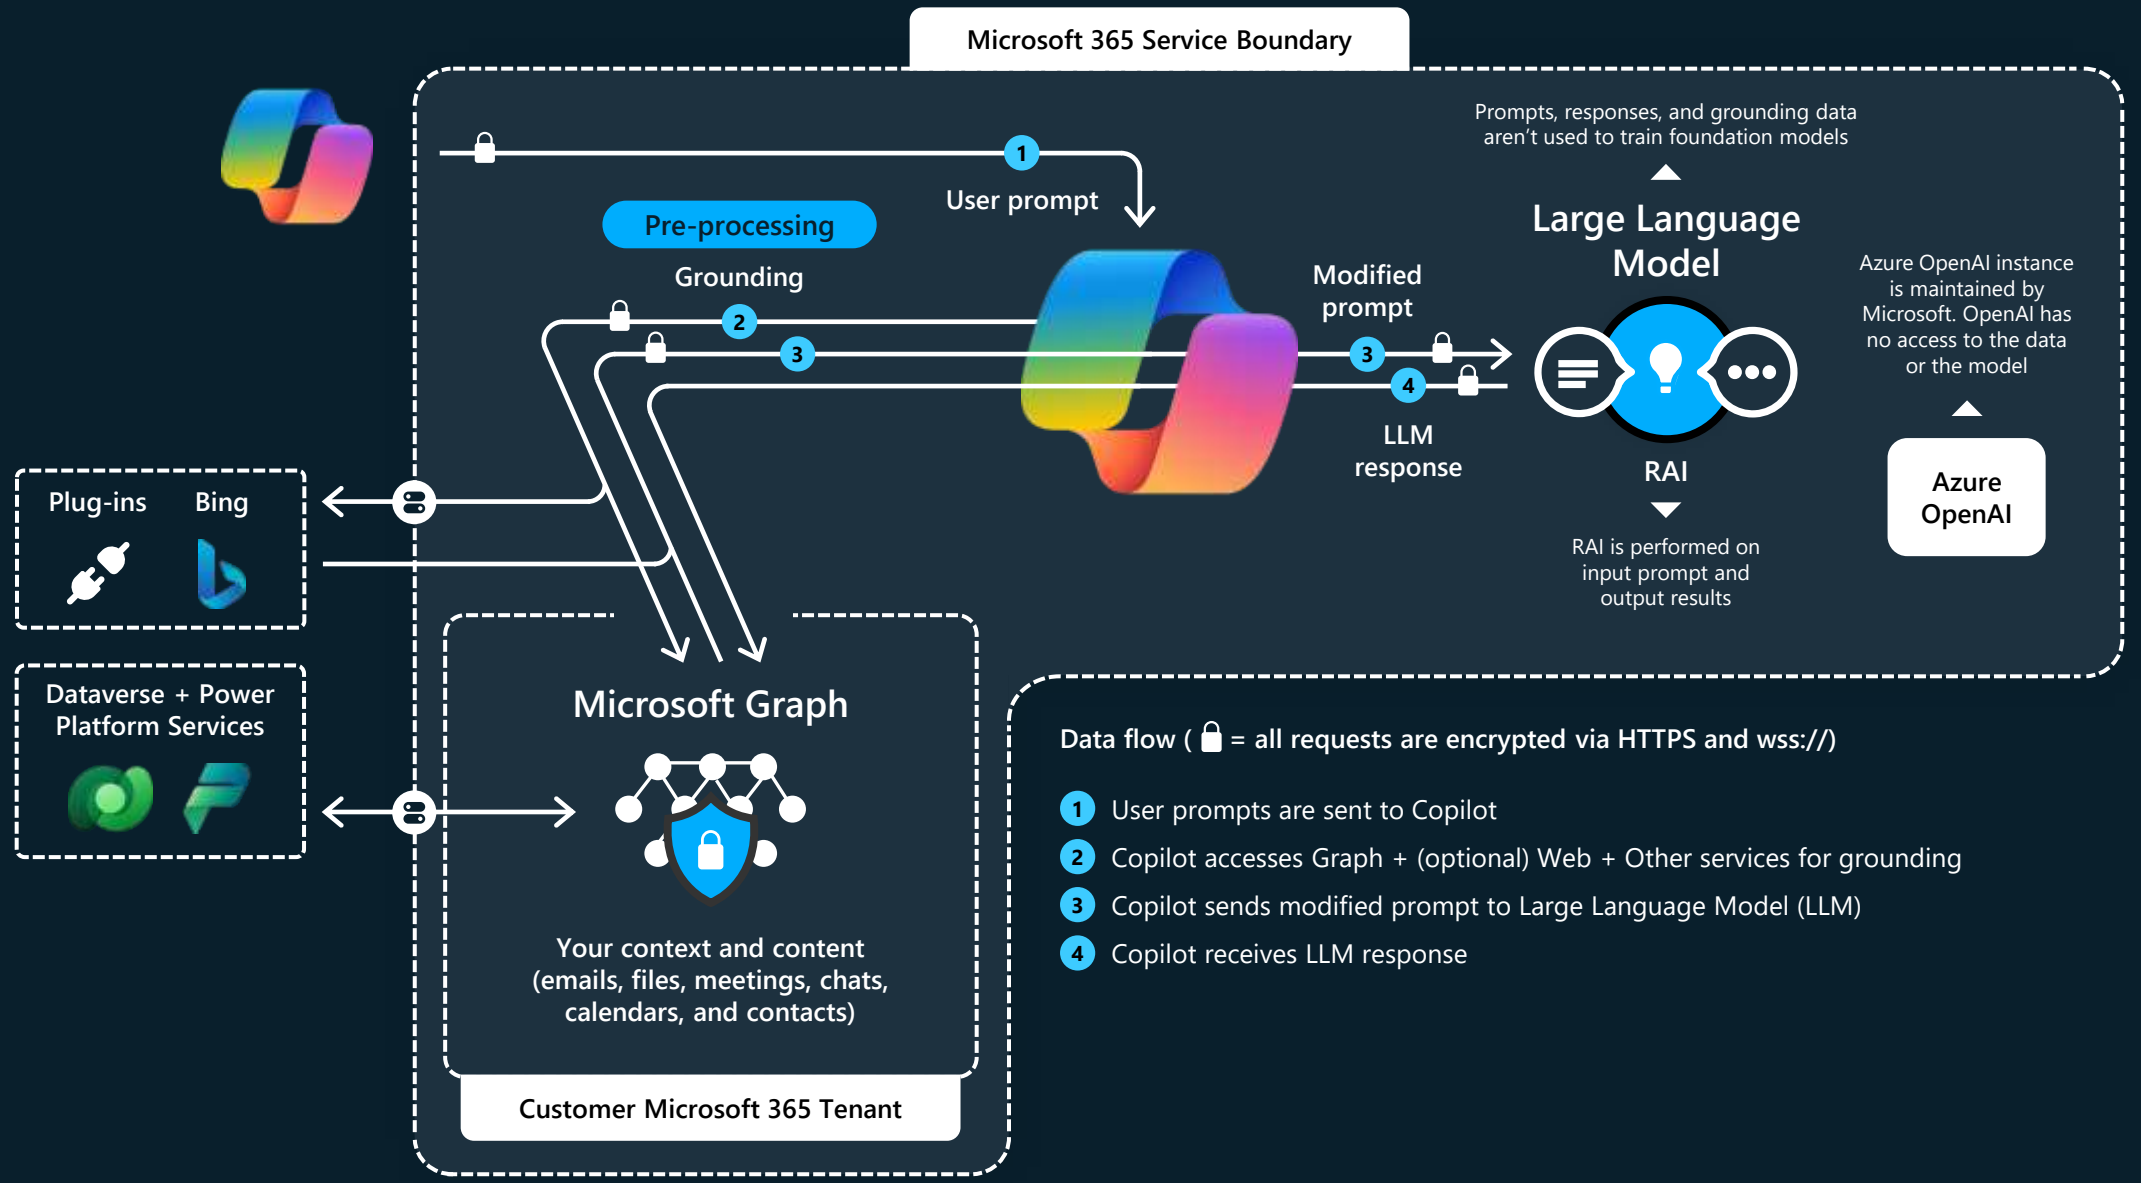

Data flow (= all requests are encrypted via HTTPS and wss://)

1. User prompts are sent to Copilot
2. Copilot accesses Graph + (optional) Web + Other services for grounding
3. Copilot sends modified prompt to Large Language Model (LLM)
4. Copilot receives LLM response

Microsoft Copilot for Microsoft 365 architecture

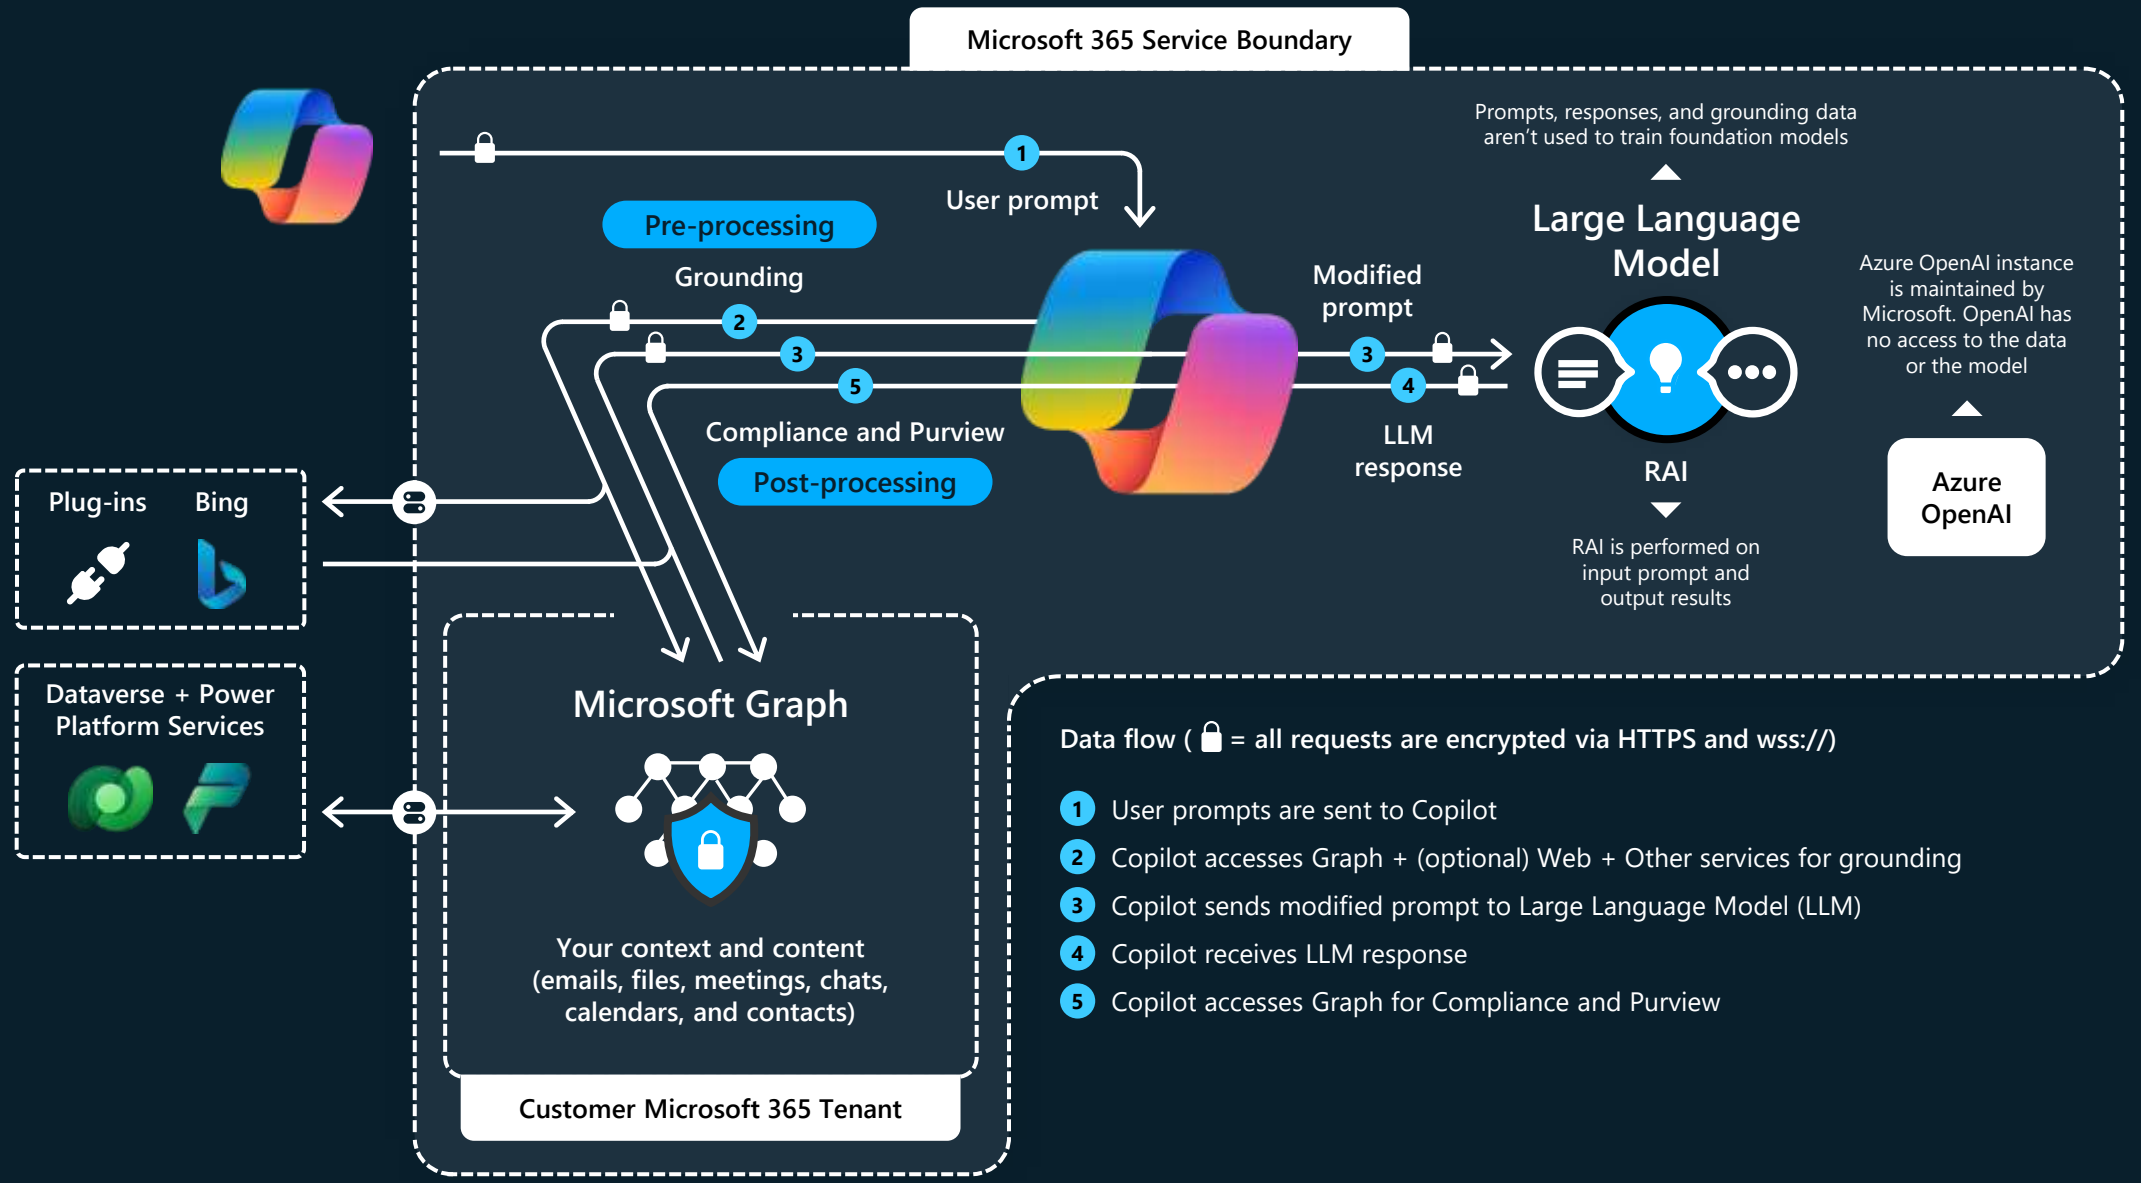

Data flow (= all requests are encrypted via HTTPS and wss://)

1. User prompts are sent to Copilot
2. Copilot accesses Graph + (optional) Web + Other services for grounding
3. Copilot sends modified prompt to Large Language Model (LLM)
4. Copilot receives LLM response
5. Copilot accesses Graph for Compliance and Purview

Microsoft Copilot for Microsoft 365 architecture

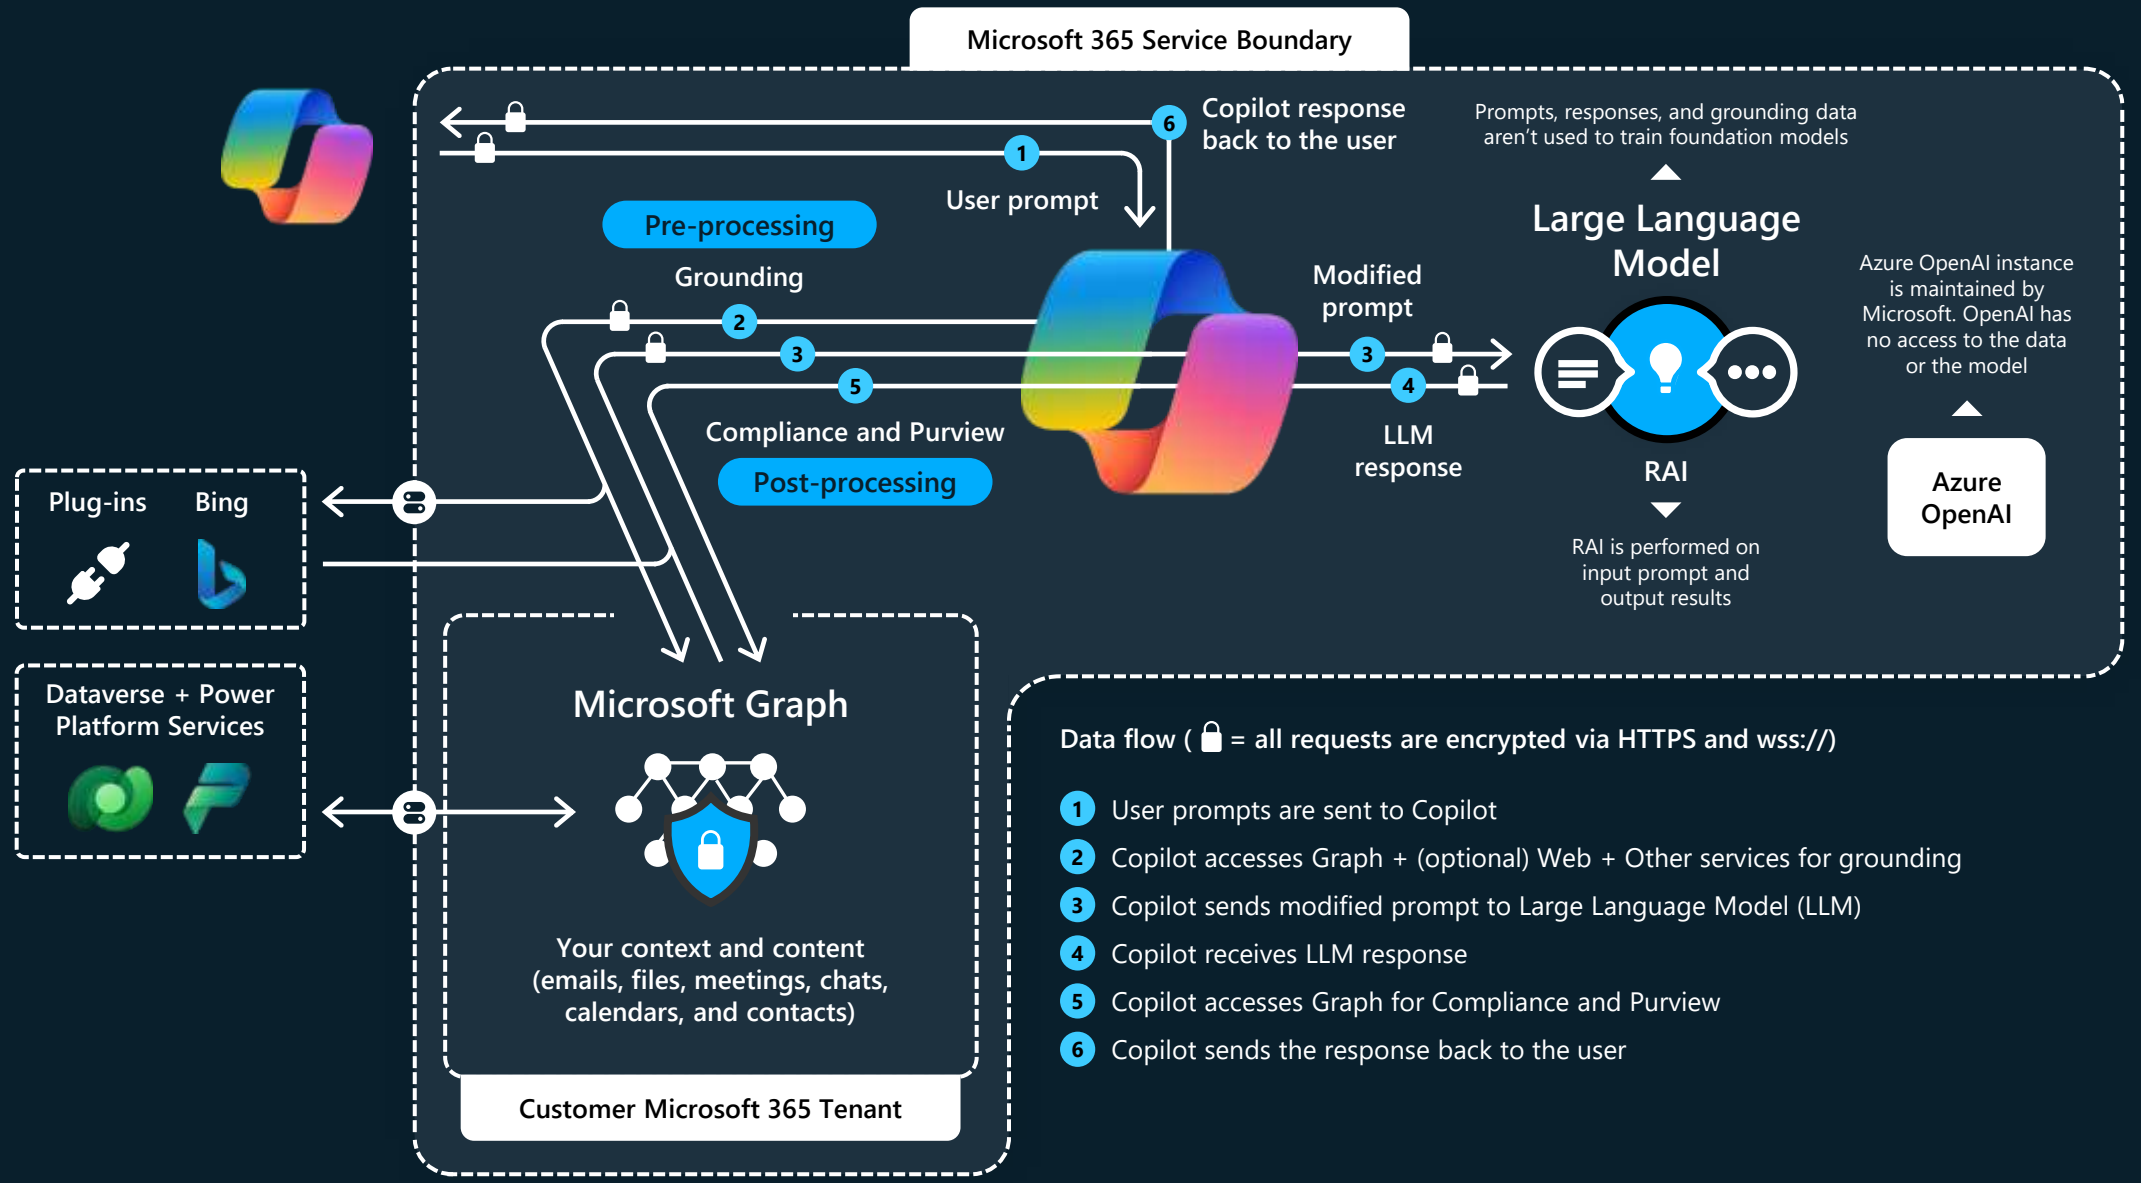

Data flow (= all requests are encrypted via HTTPS and wss://)

1. User prompts are sent to Copilot
2. Copilot accesses Graph + (optional) Web + Other services for grounding
3. Copilot sends modified prompt to Large Language Model (LLM)
4. Copilot receives LLM response
5. Copilot accesses Graph for Compliance and Purview
6. Copilot sends the response back to the user

Microsoft Graph

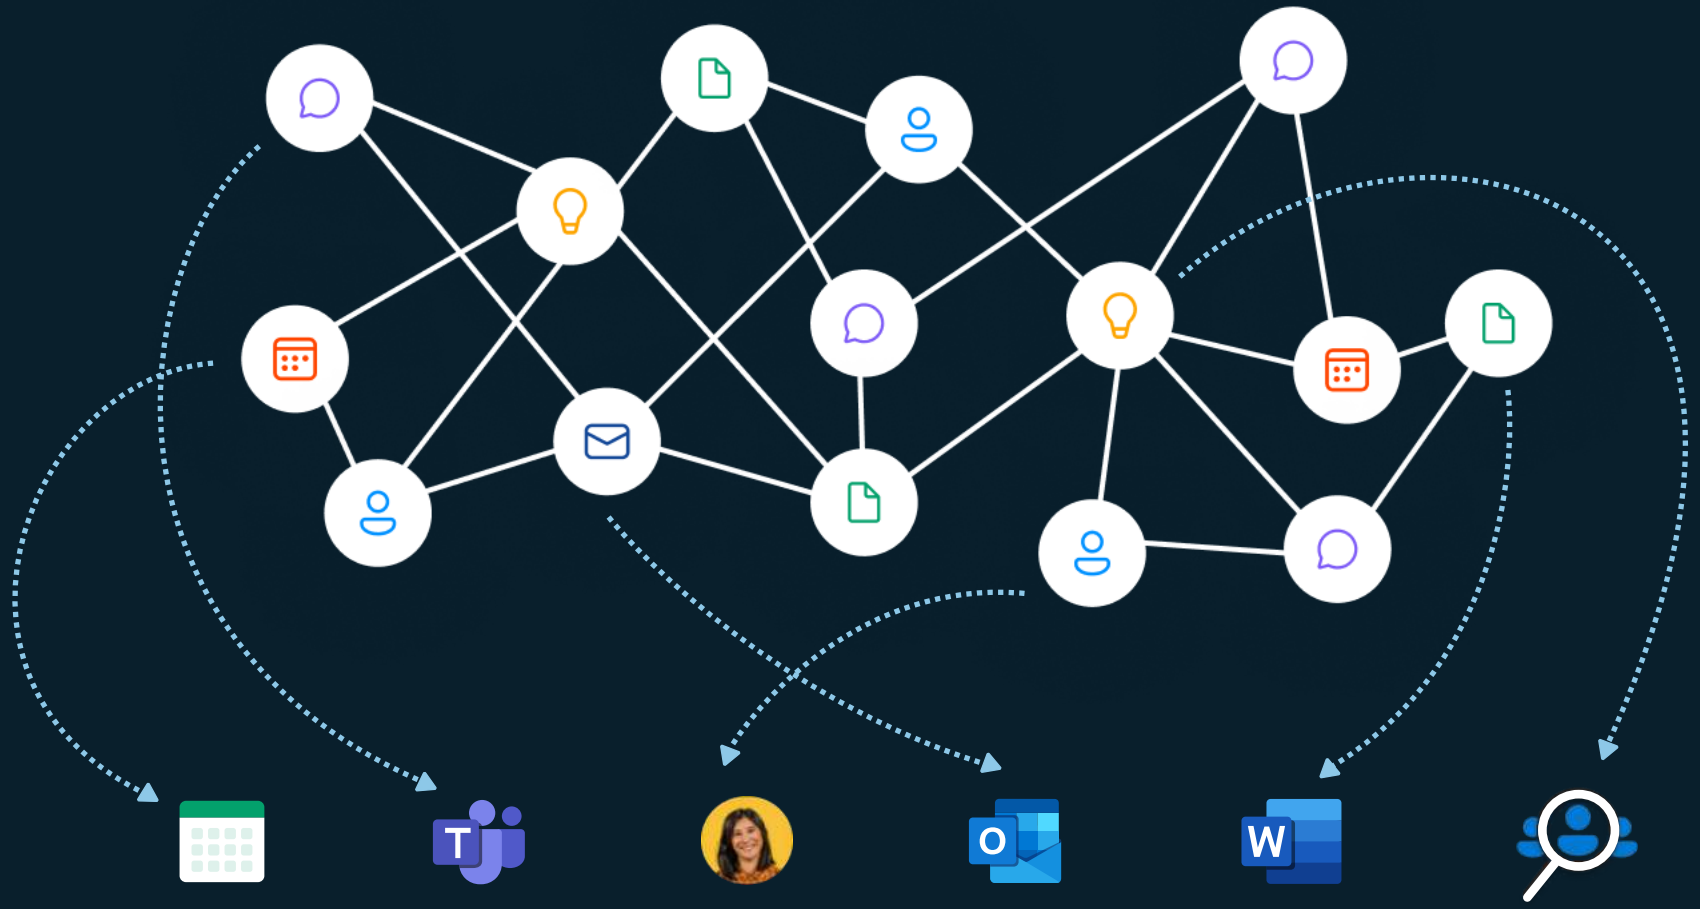

Knowledge & Insights via Microsoft Search

Embeddings for all Microsoft 365 entities

Key takeaways

Microsoft protects your data and enables you to control it

Inheriting Microsoft 365 policies and controls

- Data access & permissions
- User-tenant focus
- Customer data protection
- Data processing and residency

Protecting data processed through LLMs

- Security/ Compliance
- Data usage
- Committed to responsible AI

Resources and calls to action

Watch Ignite Copilot sessions

Getting your Enterprise ready for Copilot for Microsoft 365

Extend Copilot for Microsoft 365

Building Plugins for Copilot for Microsoft 365 (lab)

How Copilot works

Next steps

Develop your comprehensive data strategy

Watch explainer graphics and demos

Learn more about Copilot and technical details

Learn more about security, compliance, privacy and data residency

Learn how to develop successful adoption strategy

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 Bing chat. 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 Bing chat, 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.