Microsoft

Microsoft Copilot for Microsoft 365 architecture & deployment

Copilot combines the power of large language models (LLMs) with your data in the Microsoft Graph — your calendar, emails, chats, documents, meetings, and more — and the Microsoft 365 apps to provide a powerful productivity tool.

Microsoft Copilot for Microsoft 365 logical architecture

Microsoft Copilot for Microsoft 365 or Copilot introduces several components to help users make use of content and data they already have access to. Note that only data a user has access to is returned in query responses (as illustrated).

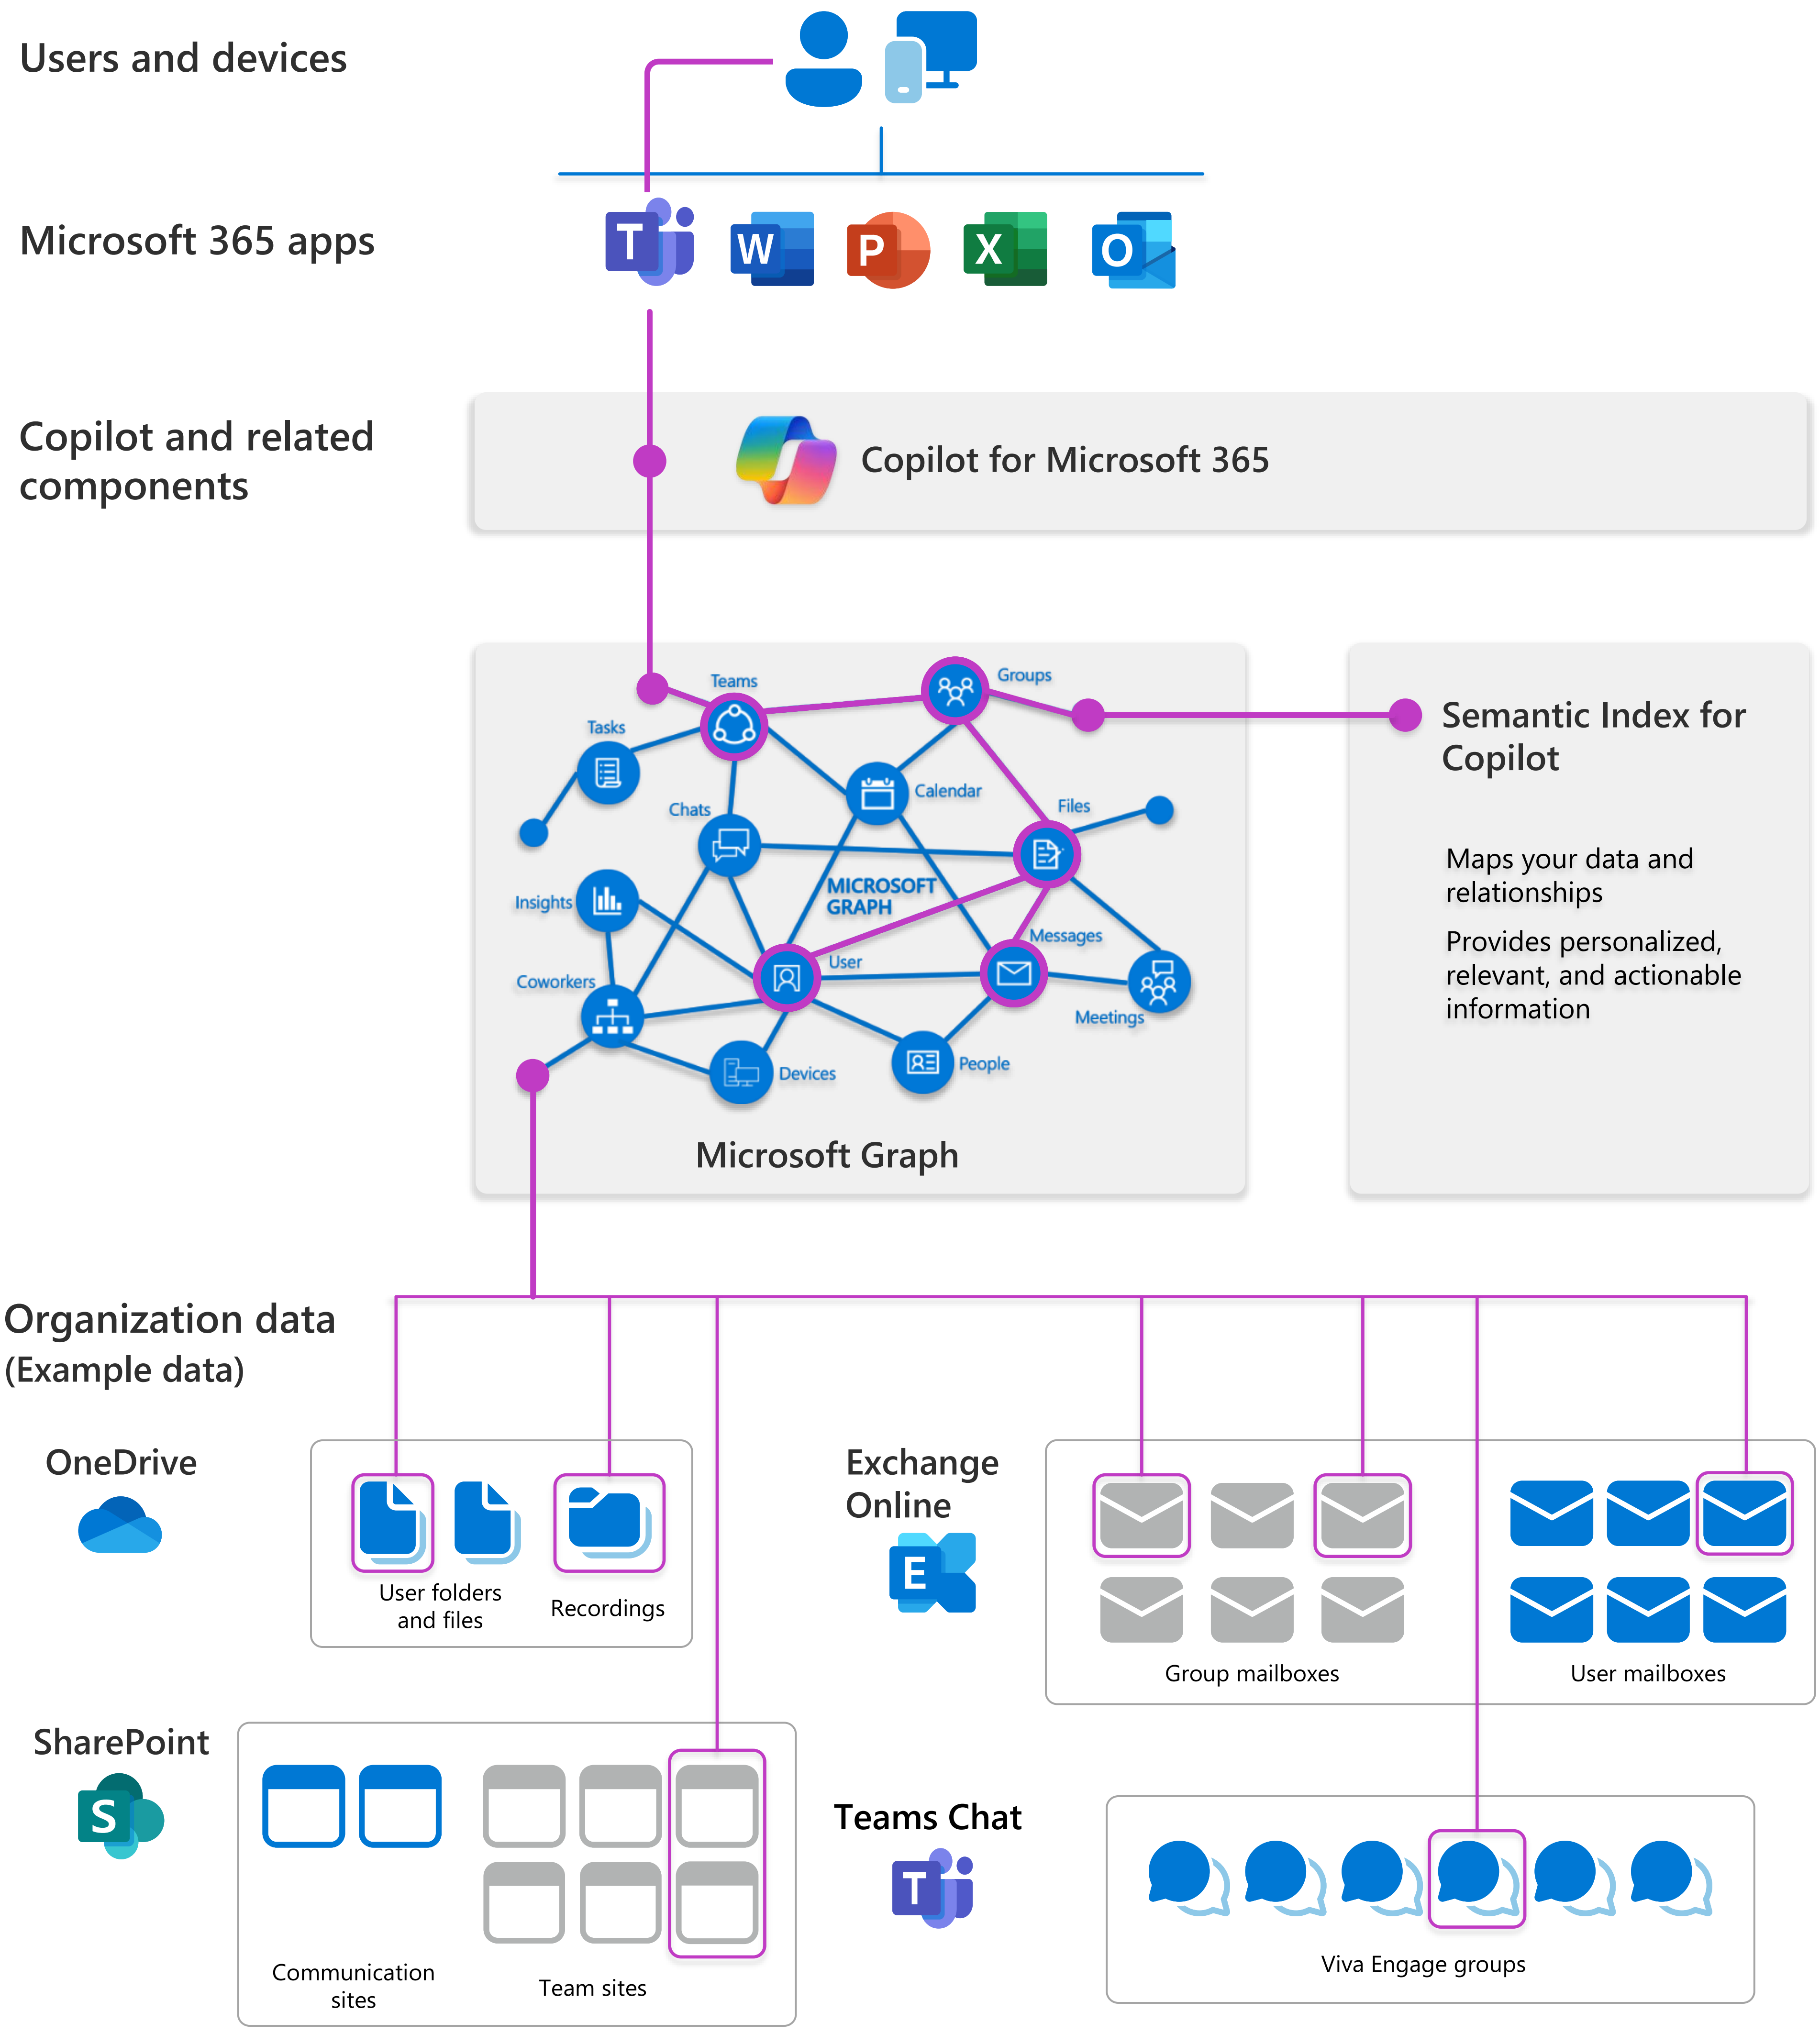

Users and devices

Microsoft 365 apps

Copilot and related components

Copilot for Microsoft 365

Microsoft Graph

Semantic Index for Copilot

Maps your data and relationships

Provides personalized, relevant, and actionable information

Organization data (Example data)

OneDrive

User folders and files

Recordings

Exchange Online

Group mailboxes

User mailboxes

SharePoint

Communication sites

Team sites

Teams Chat

Viva Engage groups

November 2023

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Microsoft Copilot for Microsoft 365 service and tenant logical architecture

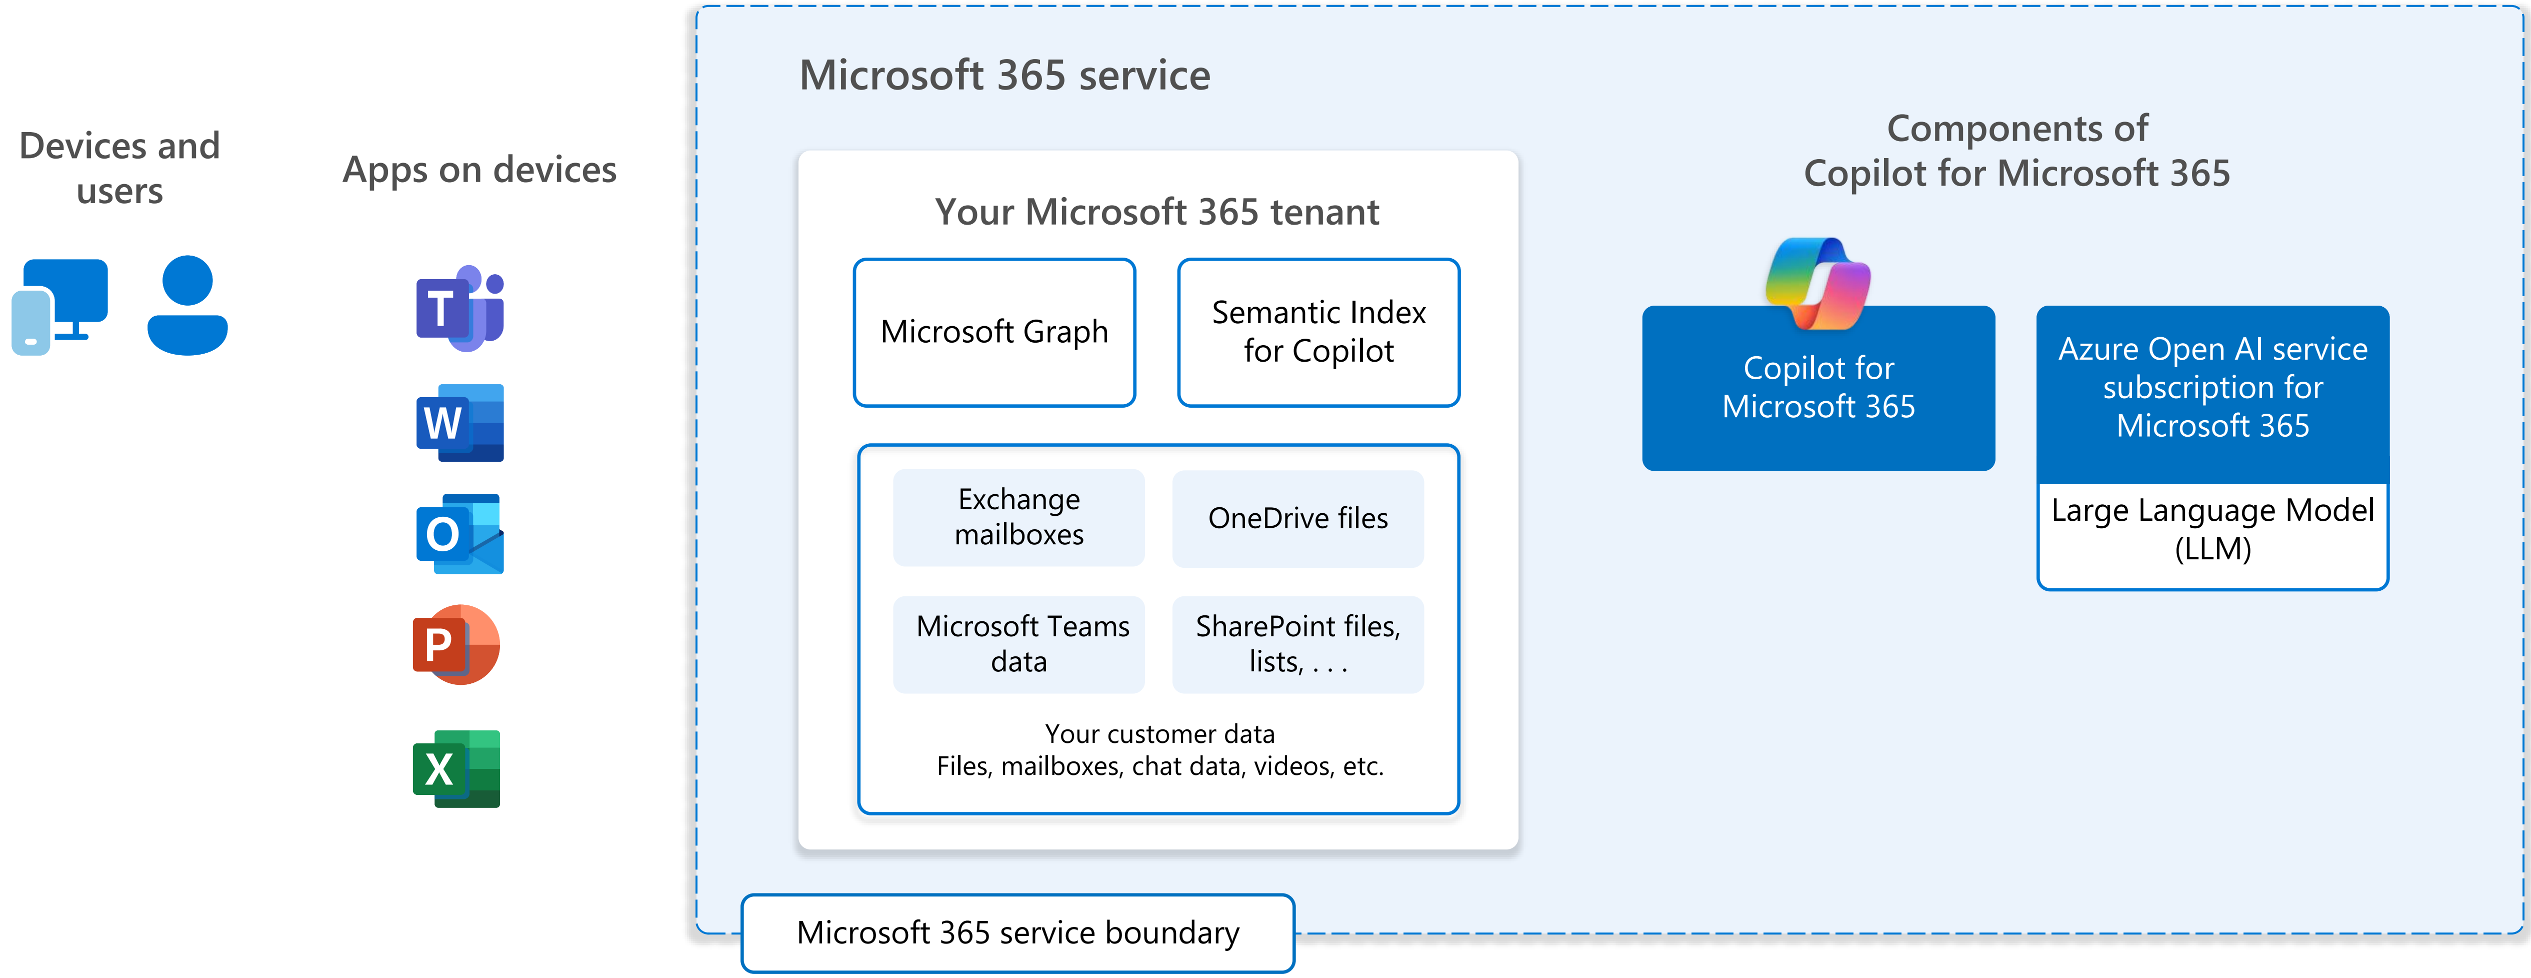
Devices and users
Apps on devices
Microsoft 365 service
Your Microsoft 365 tenant
Microsoft Graph
Semantic Index for Copilot
Exchange mailboxes
OneDrive files
Microsoft Teams data
SharePoint files, lists, . . .
Your customer data
Files, mailboxes, chat data, videos, etc.
Components of Copilot for Microsoft 365
Copilot for Microsoft 365
Azure Open AI service subscription for Microsoft 365
Large Language Model (LLM)
Microsoft 365 service boundary

Your customer data stays within the Microsoft 365 service boundary. Your prompts, responses, and data in the Microsoft Graph is not used to train foundation LLMs that Copilot leverages. Your data is secured based on existing security, compliance, and privacy policies already deployed by your organization.

Your tenant sits inside the Microsoft 365 service boundary, where Microsoft's commitment to security, compliance, data location, and privacy are upheld.

Copilot is a shared service just like many other services in Microsoft 365. Communication between your tenant and Copilot components is encrypted.

For more information, see Data, Privacy, and Security for Copilot for Microsoft 365.

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Semantic Index for Microsoft Copilot for Microsoft 365

The Semantic Index for Copilot is a separate index or map of your user and company data — identifying relationships and connections. It works with Copilot and the Microsoft Graph to create a sophisticated map of all data and content in your organization to enable Copilot to deliver personalized, relevant, and actionable responses. The Semantic Index is part of the Microsoft 365 service and is created automatically.

Semantic indexing vs keyword indexing

Semantic Index builds upon keyword matching, personalization, and social matching capabilities within Microsoft 365 by creating vectorized indices to enable conceptual understanding, which helps determine your intent and helps you find what organizational content you need.

Unlike a standard keyword index, vectors are stored multi-dimensional spaces where semantically similar data points are clustered together in the vector space, enabling Microsoft 365 to handle a much broader set of search queries beyond "exact match."

Each dimension of a vector captures an aspect of semantic meaning of the data point being represented. This provides for fast and accurate search and retrieval of data based on vector distance or similarity. This means that instead of using traditional methods for querying based on exact matches or predefined criteria, the Semantic Index finds the most similar or relevant data based on the semantic or contextual meaning.

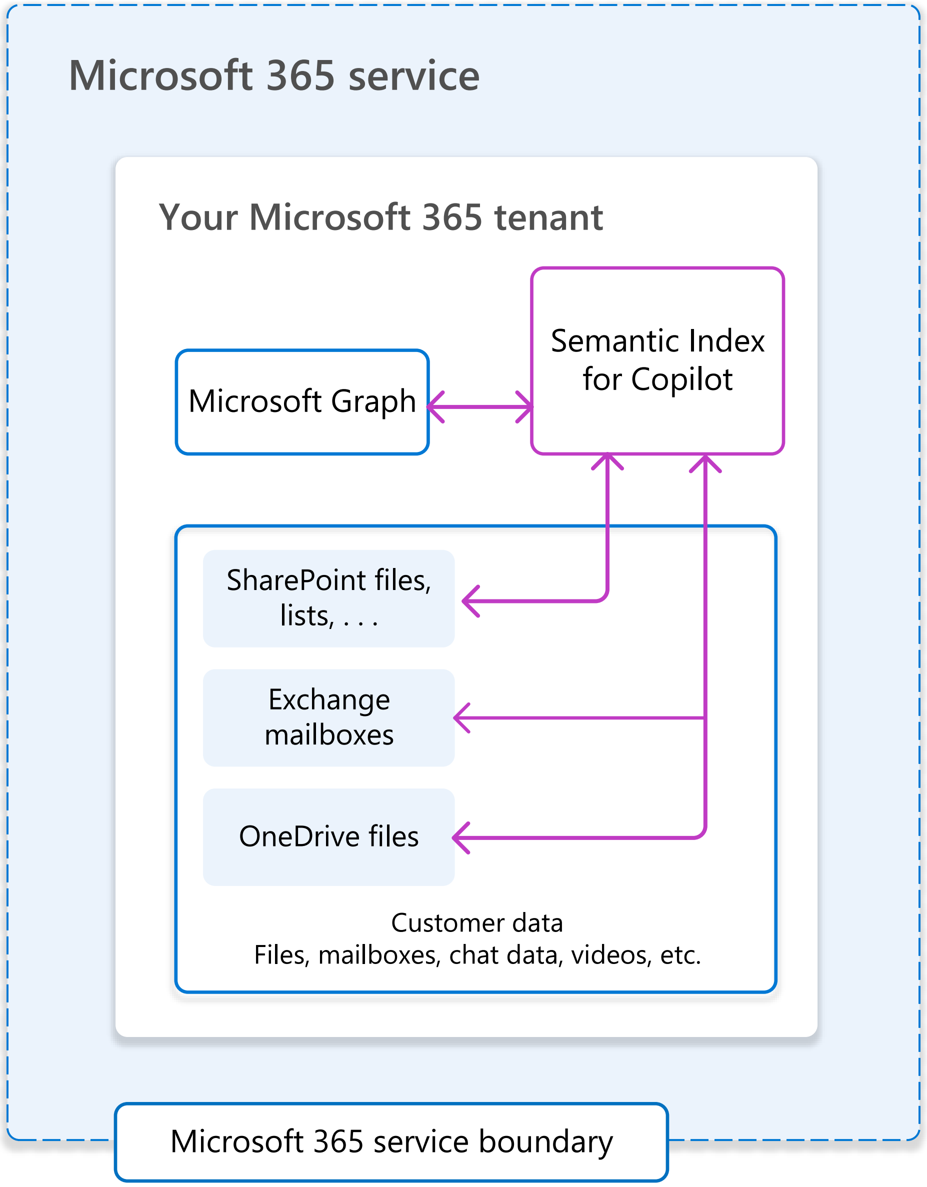
Microsoft 365 service
Your Microsoft 365 tenant
Microsoft Graph
Semantic Index for Copilot
SharePoint files, lists, . . .
Exchange mailboxes
OneDrive files
Customer data
Files, mailboxes, chat data, videos, etc.
Microsoft 365 service boundary

What is currently indexed

- Semantic Index indexes text-based files in SharePoint that are shared with two or more people.
- At the user level, Semantic Index indexes all email. It also indexes all text-based files in a user's OneDrive that have been shared, interacted with (even just by the user), or commented on.
- Current supported file types include:
  - Word documents (doc/docx)
  - PowerPoint (pptx)
  - PDF
  - Web pages (html/aspx)
  - OneNote (one)
- Semantic Index leverages the Microsoft Graph to better correlate relationships and understand permissions.

Excluding SharePoint Online sites

You can exclude a SharePoint Online site from being indexed by both Microsoft Search and the Semantic Index (such as payroll data or financial information). Site administrators select this option within the Site Settings page.

Additional resources

Video — Semantic Index for Copilot: Explained by Microsoft

Training article — Examine how Copilot uses the Semantic Index (MS-012 Prepare your organization for Copilot for Microsoft 365)

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Microsoft Copilot for Microsoft 365 query flow

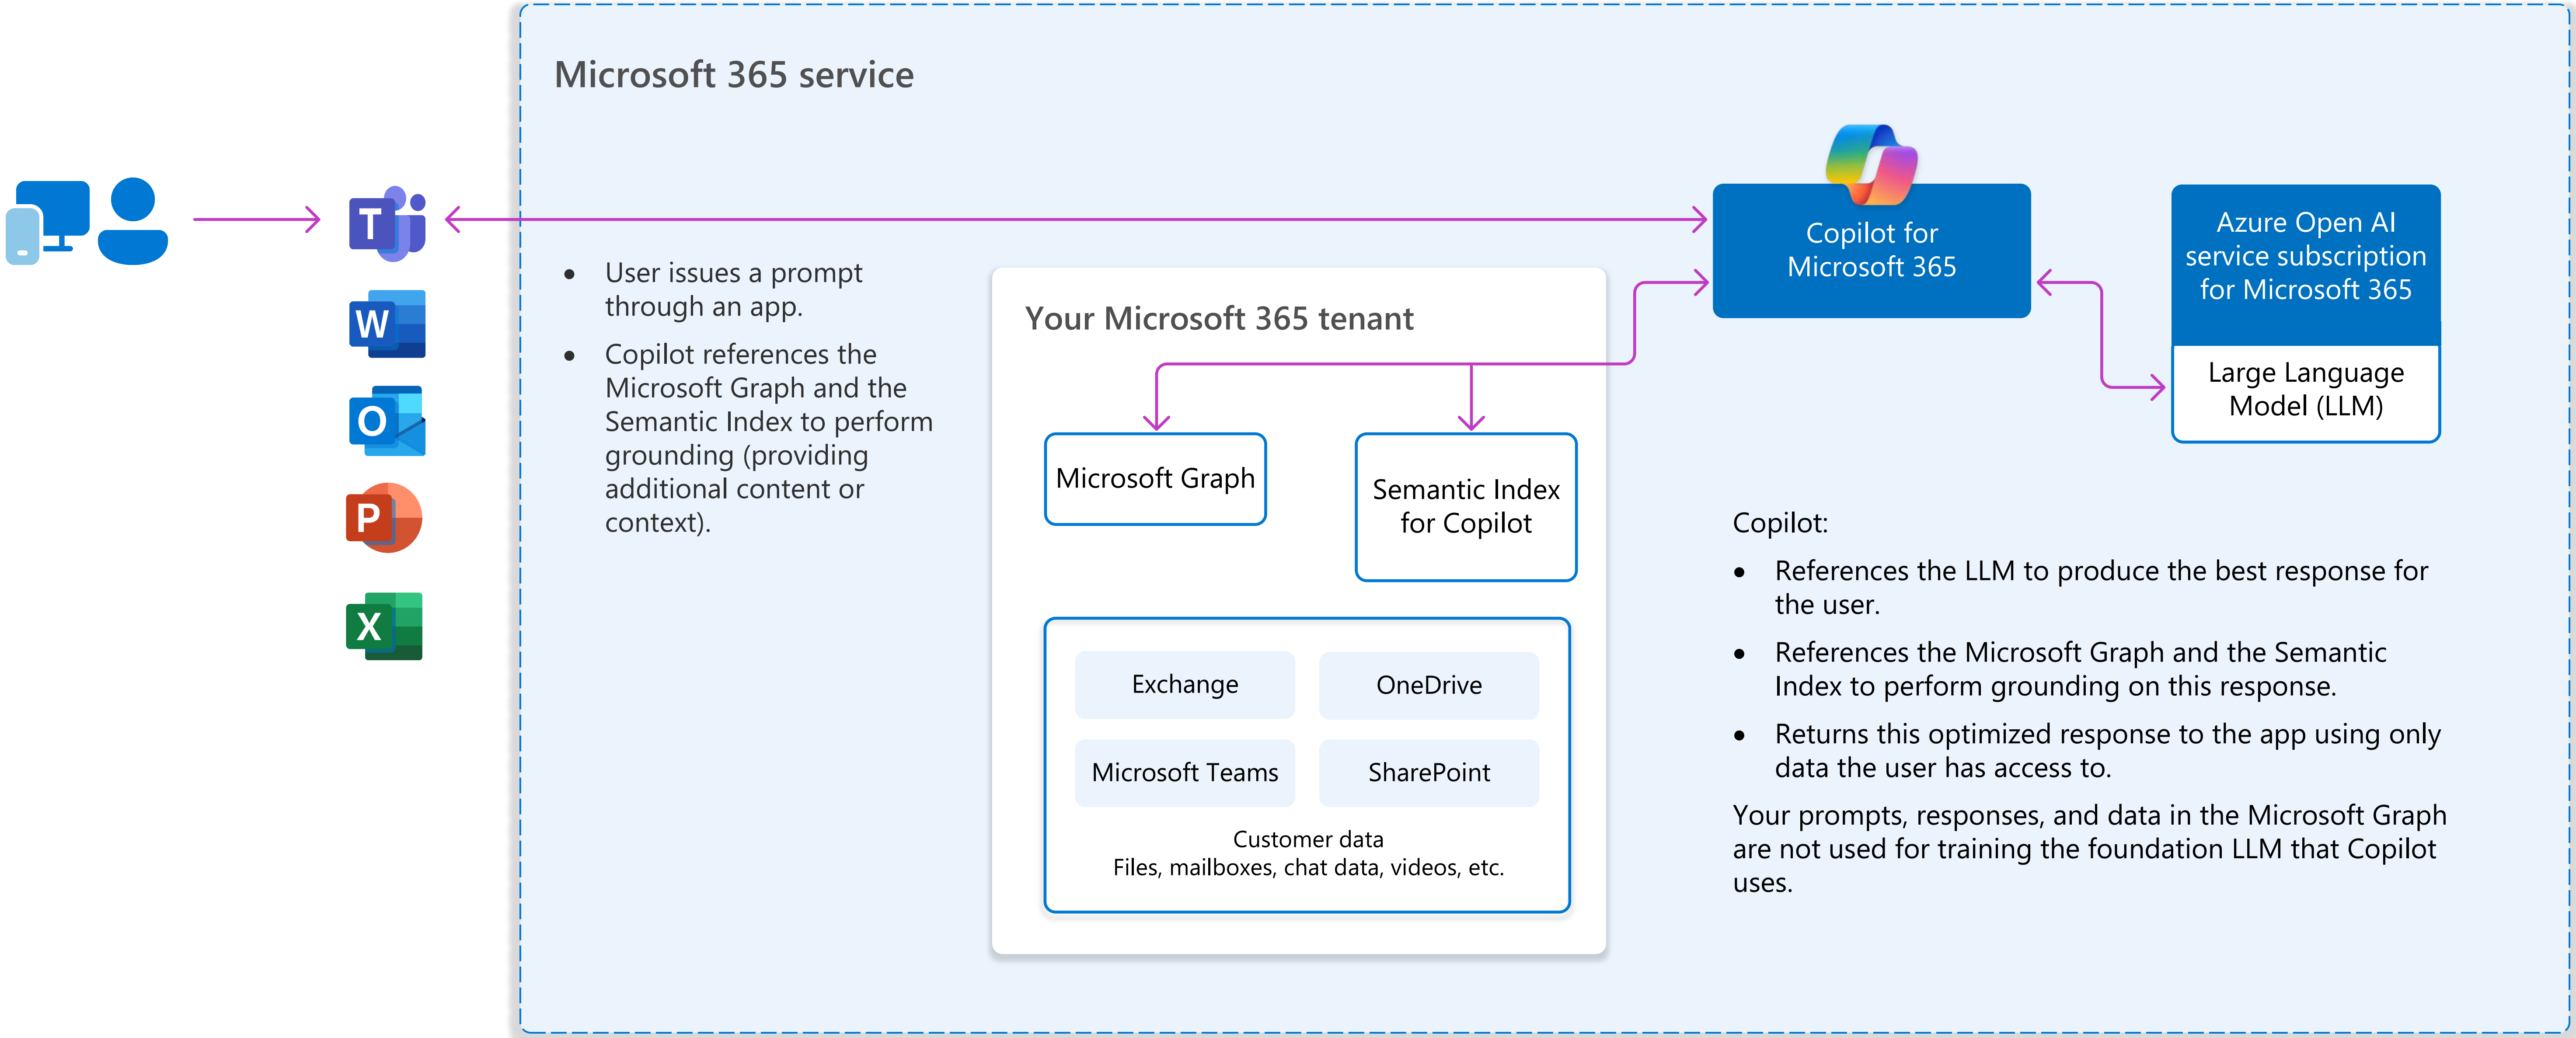

Microsoft 365 service

- User issues a prompt through an app.
- Copilot references the Microsoft Graph and the Semantic Index to perform grounding (providing additional content or context).

Copilot for Microsoft 365
Azure Open AI service subscription for Microsoft 365
Large Language Model (LLM)
Your Microsoft 365 tenant
Microsoft Graph
Semantic Index for Copilot
Exchange
OneDrive
Microsoft Teams
SharePoint
Customer data
Files, mailboxes, chat data, videos, etc.

Copilot:

- References the LLM to produce the best response for the user.
- References the Microsoft Graph and the Semantic Index to perform grounding on this response.
- Returns this optimized response to the app using only data the user has access to.

Your prompts, responses, and data in the Microsoft Graph are not used for training the foundation LLM that Copilot uses.

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Prepare your environment for Microsoft Copilot for Microsoft 365

Security and information protection recommendations

Microsoft recommends building a foundation of secure productivity to get AI-ready, including Microsoft Copilot for Microsoft 365 or Copilot.

Area to protect	Getting started with E3	Next steps with E5
Identity and access	**Configure common conditional access policies** With Microsoft Entra ID P1, configure the following policies to use multi-factor authentication (MFA): - Require MFA for administrators - Require MFA for all users - Block legacy authentication See Common Conditional Access policies. Be sure Microsoft 365 Services and your other SaaS apps are included in the scope of these policies. If your environment includes hybrid identities, also enforce on-premises Microsoft Entra Password Protection for Active Directory Domain Services.	**Configure recommended policies for Zero Trust** With Microsoft Entra ID P1, configure the following policies to use multi-factor authentication (MFA): - Require MFA when sign-in risk is medium or high - Block legacy authentication - Require high risk users to change their password See Common security policies for Microsoft 365 organizations. Also configure Privileged Identity Management.
Microsoft 365 Apps	**Implement Intune App Protection policies (APP)** With APP, Intune creates a wall between your organization data and personal data. Policies ensure corporate data in the apps you specify cannot be copied and pasted to other apps on the device, even if the device is not managed. See Implement App Protection policies.	
Devices	**Manage devices** After devices are enrolled, set up compliance policies and then require healthy and compliant devices. Finally, deploy device profiles to manage settings and features on devices. Enroll devices into management Set up compliance policies Require healthy and compliant devices Deploy device profiles	**Monitor device risk and compliance to security baselines** Integrate Intune with Defender for Endpoint to monitor device risk as a condition for access. For Windows devices, monitor compliance of these devices to security baselines. See Monitor device risk and compliance to security baselines.
Threat protection	**Configure Exchange Online Protection and endpoint protection** Exchange Online Protection (EOP) helps protect your email and collaboration tools from phishing, impersonation, and other threats. You can rapidly apply these protections by configuring preset security policies. Microsoft Defender for Endpoint P1 includes Attack surface reduction and Next generation protection for antimalware and antivirus protection. See Overview of Microsoft Defender for Endpoint Plan 1.	**Pilot and deploy Microsoft 365 Defender** For more comprehensive threat protection, pilot and deploy Microsoft 365 Defender, including: - Defender for Identity - Defender for Office 365 - Defender for Endpoint - Defender for Cloud Apps See Evaluate and pilot Microsoft 365 Defender.
Organization data	**Develop your classification schema and get started with sensitivity labels and other policies** Sensitivity labels form the cornerstone of protecting your data. Before you create the labels to denote the sensitivity of items and the protection actions to be applied, understand your organization's existing classification taxonomy and how it will map to labels that users will see and apply in apps. Create data loss prevention policies Create retention policies Use context explorer (to review results)	**Extend policies to more data and begin using automation with data protection policies** Sensitivity labeling expands to protecting more content and more labeling methods. For example, labeling SharePoint sites and Teams by using container labels, and automatically labeling items in Microsoft 365 and beyond. For more information, see a list of common labeling scenarios and how they align to business goals.

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Set up secure file sharing and collaboration with Microsoft Teams

Configure Microsoft Teams with appropriate protection

Microsoft provides guidance for protecting your Teams at three different levels – baseline, sensitive, and highly sensitive. Guidance includes sensitivity labels and site sharing settings.

Introducing Copilot is a good time to review your environment and ensure that appropriate protection is configured.

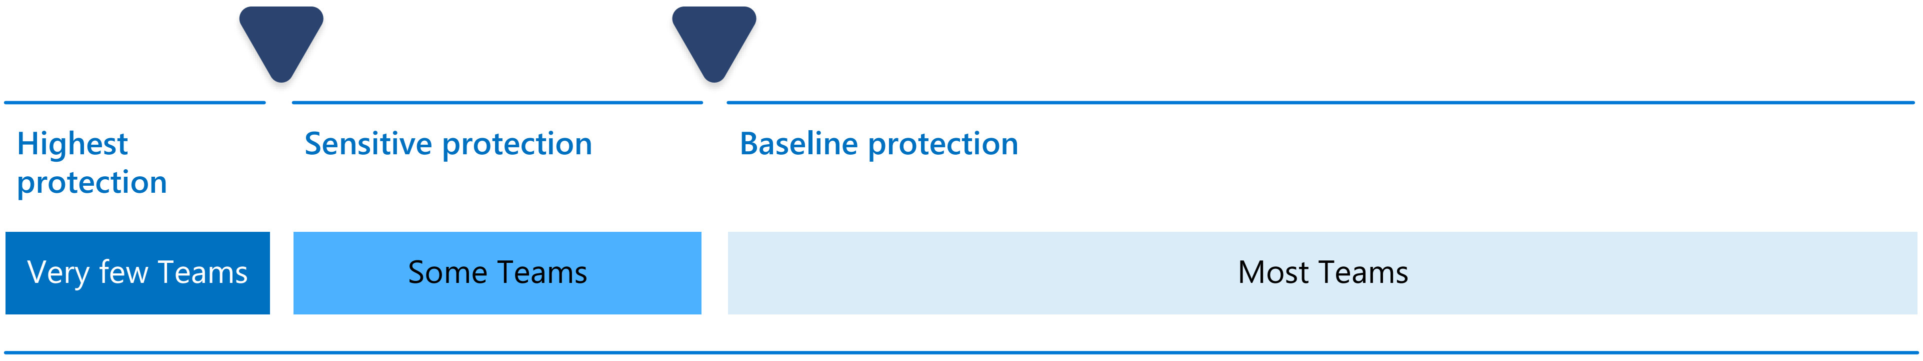

Highest protection

Very few Teams

Sensitive protection

Some Teams

Baseline protection

Most Teams

1. First, identify Teams or projects that warrant highly sensitive protection. Configure protections for this level. Many organizations don't have data that requires this level of protection.
2. Next, identify Teams or projects that warrant sensitive protection and apply this protection.
3. Finally, ensure all Teams and projects are configured for baseline protection, at a minimum.

Resources

- Set up secure file sharing and collaboration with Microsoft Teams
- Compare levels of protection
- Configure Teams with three tiers of protection

Configure external sharing with appropriate security

Introducing Copilot is a good time to review your policies for sharing files with people outside your organization and for allowing external contributors. Note that guest accounts are not licensed to use Copilot.

Sharing with people outside your organization

You may need to share information of any sensitivity with people outside your organization. Use these resources:

- Apply best practices for sharing files and folders with unauthenticated users
- Limit accidental exposure to files when sharing with people outside your organization
- Create a secure guest sharing environment

Collaborating with people outside your organization

Use these resources for setting up your environment for collaborating with people outside your organization:

- Collaborate on documents — share individual files or folders
- Collaborate on a site — collaborate with guests in a SharePoint site
- Collaborate as a team — collaborate with guests in a team
- Collaborate with external participants in a channel — collaborate with people outside the organization in a shared channel

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.

Pilot and deploy Microsoft Copilot for Microsoft 365

Onboard users to Microsoft Copilot for Microsoft 365

Microsoft recommends deploying Copilot as you phase in protections for access, devices, and data. Whether you're starting with Microsoft 365 E3 or you're taking next steps with Microsoft 365 E5, use the following sequence:

1. Apply identity and access protection.
2. Apply device protection.
3. Assign Copilot licenses to users with these protections.

Ongoing — Continue your deployment of information protection capabilities.

Begin by building a plan and then testing the plan. Then roll out new configurations and capabilities incrementally. This provides the opportunity to improve on the plan while lessons are learned.

The following diagram illustrates the recommendation to start a project with a small group to evaluate the changes. This small group can be members of your IT team or a partner team. Then, pilot the changes with a larger group. Full deployment is accomplished by gradually increasing the scope of the deployment until your whole organization is covered.

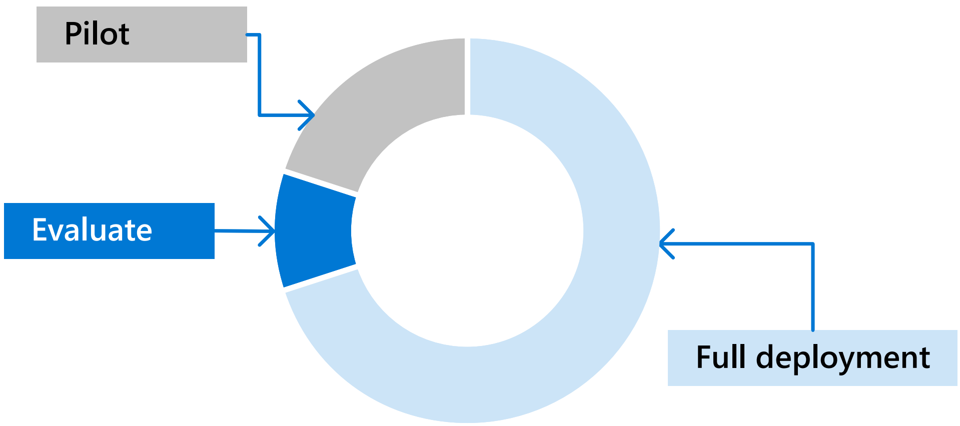

Applying protections and deploying Copilot in parallel

	Evaluate	Pilot	Full deployment
Identity and access	Identify 50 users for testing	Identify the next 50-100 users in the production environment	Apply protections to the rest of the users in larger increments
Devices	Test device protections with the same 50 users	Apply device protections to the same users	Enroll the rest of the endpoints in larger increments
Copilot	Assign Copilot licenses to users AFTER their account and devices are protected		

Technical adoption of information protection

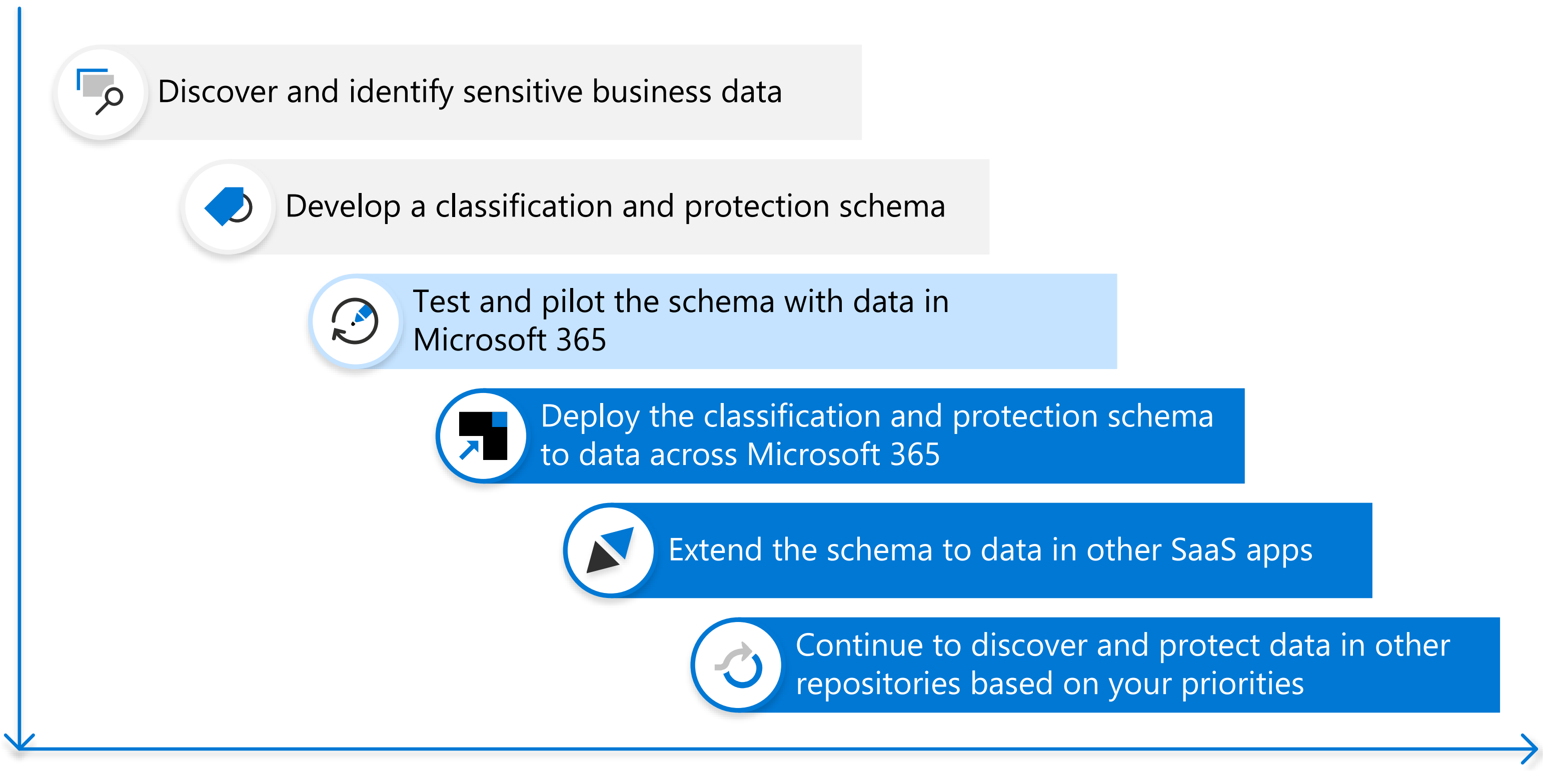

Discover and identify sensitive business data

Develop a classification and protection schema

Test and pilot the schema with data in Microsoft 365

Deploy the classification and protection schema to data across Microsoft 365

Extend the schema to data in other SaaS apps

Continue to discover and protect data in other repositories based on your priorities

© 2023 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at CloudAdopt@microsoft.com.