

Data Governance, Security, and Privacy with **Microsoft 365 Data**



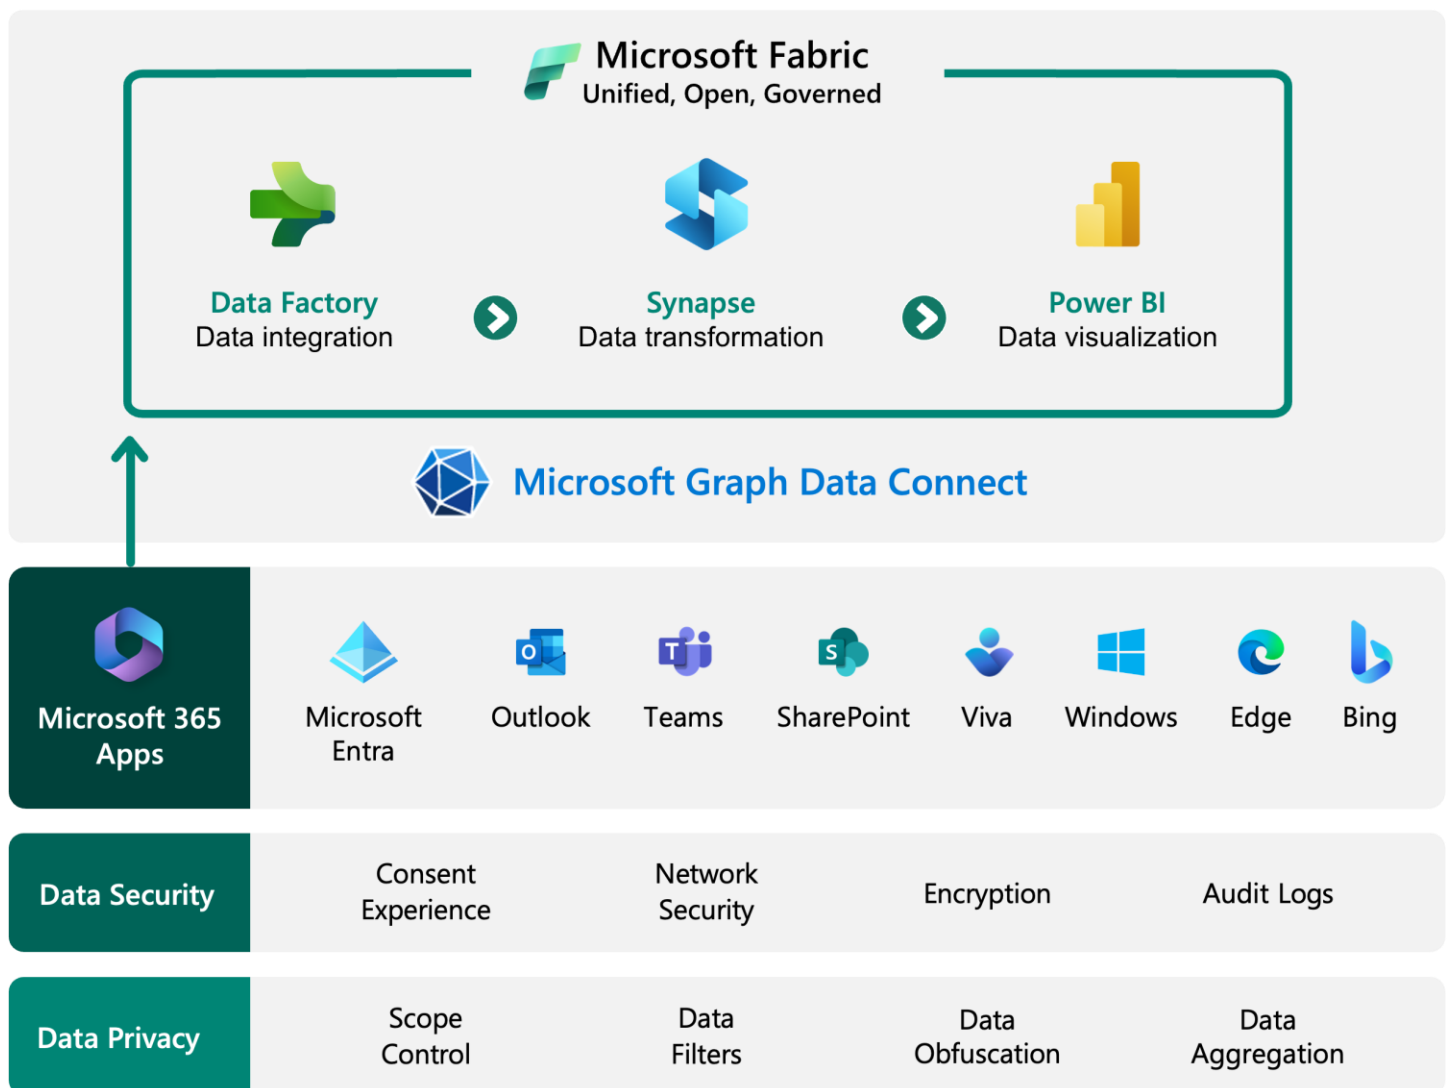
Table of Contents

Microsoft Graph Data Connect – Introduction	3	Conclusion	19
Compliance and Governance in Microsoft Graph Data Connect	5	Appendix 1 – Data-Protection Considerations	20
Microsoft Graph Data Connect Consent Experience	6	Roles and Responsibilities	20
Improved Consent Experience (Projected Availability CY23H2)	7	Determine if Data Protection Impact Assessment (DPIA) is Needed	23
Data Lineage in Microsoft Graph Data Connect	8	Support for handling data subject requests	24
Microsoft Graph Data Connect Available in 14 Regions	9		
Data Security Capabilities in Microsoft Graph Data Connect	10		
Network Security	10		
Encryption	11		
Geo Isolation for Data Access by Default	11		
Data Privacy and Access	12		
Data Privacy Principles	13		
Data Privacy Recommendations	13		
Data Privacy Capabilities in Microsoft Graph Data Connect	15		
Privacy Scrubbing with Allow and Deny list	15		
Scope and Filtering	16		
Data Obfuscation	17		

Microsoft Graph Data Connect – Introduction

Microsoft Graph Data Connect provides an intelligent way to access rich Microsoft 365 data (Microsoft 365 data) at scale. The available data includes how workers communicate and collaborate across all the applications and services in Microsoft 365.

Ideal for big data and machine learning, Microsoft Graph Data Connect lets you leverage Microsoft 365 data for analytics, intelligence, and AI/ML models by extending Microsoft 365 data into Microsoft Azure and Microsoft Fabric. By integrating in this way, you can take advantage of the vast suite of compute, storage, analytics, and BI services in Azure and Fabric while staying compliant with industry standards and helping to keep your data secure.



Access to Data at Scale

Microsoft Graph Data Connect provides access to large amounts of data, often from many users in your organization at once, accessing data in bulk and copying Microsoft 365 data to your Azure environment. Microsoft Graph Data Connect also lets you choose between accessing data from everyone in your organization or just specific groups of people.

Microsoft Graph Data Connect uses Azure Synapse or Azure Data Factory to copy Microsoft 365 data into your preferred storage at configurable intervals. It also provides a set of tools to streamline the delivery of this data to Microsoft Azure.

- [Azure Synapse](#) is a limitless analytics service that brings together enterprise data warehousing and big data analytics. It gives you the freedom to query data on your terms, using either serverless or dedicated resources at scale.
- [Azure Data Factory](#) is a managed cloud service that is built for complex hybrid extract-transform-load (ETL), extract-load-transform (ELT), and data integration projects. It provides the capability to create a pipeline calling Microsoft Graph Data Connect to access Microsoft 365 data.

Microsoft Graph Data Connect supports a variety of Microsoft 365 datasets, data regions, and storage locations in Microsoft Azure.

- Datasets include [Microsoft Entra, Outlook and Exchange Online, Microsoft Teams, Microsoft Groups, OneDrive and SharePoint Online](#) and more. The details on these data sets, including schemas and sample data are published for review in the linked locations.
- The current list of [supported Regions are available for reference](#).
- Storage locations (or sinks) are the output location used by Azure (Synapse or Data Factory); M365 data is delivered as JSON or Parquet format. To learn more, see the most up to date list of [supported sinks can be found here](#).

Below, you can see just some types of Microsoft 365 data that organizations get through Microsoft Graph Data Connect. There are datasets available from a wide variety of Microsoft 365 sources, including Microsoft Entra, Outlook, Teams, SharePoint, Viva Insight, and Microsoft Groups—and the [list](#) is growing.

To learn more: [Datasets, regions, and sinks supported by Microsoft Graph Data Connect – Microsoft Graph | Microsoft Learn](#)

Available 2023						
Microsoft Entra	Outlook	Teams	SharePoint	Viva Insights	Microsoft 365	Pre-built Solutions
User Profile	Messages, Sent	Meeting Chats	Sharing/Permissions	VI Person Report	Meeting Activity	Org. Network Analysis
Manager Info	Contact	1:1 and 1:n Chats	Groups	VE Messages	Email Activity	Information Oversharing
Direct Reports	Calendar View	Std. Channel Msgs.	Site		Office App Activity	Hybrid Work Effectiveness
Group Members	Mailbox Settings	Channel Details	File Usage		Teams Activity	Knowledge Mining
Group Details	Mail Folder	Call Records	File State			Security Controls
Group Owners	Events	Transcripts				
	Group Messages					
	Inbox Messages					
	To-Do Lists/Tasks					
	Conference Rooms					
Beyond						
Microsoft Entra	Outlook	Teams	SharePoint	Viva Insights	Microsoft 365	+
						Windows
						Edge
						Viva Topics
						Viva Learning
						Bing
						+
Expanding existing datasets				New datasets		

Compliance and Governance in Microsoft Graph Data Connect

Microsoft facilitates rich, connected communication between Microsoft Graph Data Connect and Azure that respects customer data. Microsoft Graph Data Connect supports all Azure-native service capabilities, such as encryption, geo-fencing, auditing, and policy enforcement.

To minimize data security and compliance management overhead with Microsoft Graph Data Connect, customers can specify a set of detailed policies on their Azure resources using [Azure Policy](#). Using Azure Policies, customers can continuously monitor their resources and services for policy adherence and set up controls that would limit or stop data flows if policy violations are observed.



Microsoft Graph Data Connect Consent Experience

Microsoft Graph Data Connect natively offers a granular control and consent model:

1. Microsoft 365 Global Administrators are required to explicitly opt-in to enable access to Microsoft 365 data through Microsoft Graph Data Connect.
2. Developers can request access to specific entities.
 - a. Developers choose only properties of the entity that they require for their use case.
 - b. The request specifies the scope of data access, the data policy enforcement and the reason for the request.
3. The Microsoft 365 Global Administrator manages and audits this data access.
4. Microsoft Graph Data Connect also allows for filtering the data in those entities by population scope, time duration and exclusion lists within the data set requested.

As a result, only required data is copied, and any unrelated content is excluded.

Data consent is achieved through integration with [Microsoft Privileged Access Management](#) (PAM). PAM is a solution that helps organizations restrict privileged access to protect organizations against cyberthreats by monitoring, detecting, and preventing unauthorized access to critical resources. PAM works through a combination of people, processes, and technology and gives you visibility into who is using privileged accounts and what they are doing while they are logged in. For Microsoft Graph Data Connect, this requires approval of data requests, and the approver and requester cannot be the same person. Each request always includes the following details about the [dataset](#) and the users about whom data is being extracted:

- **Requestor:** The user who requested the pipeline.
- **Duration:** If approved, how long the approval persists, which is always 4320 hours (6 months).
- **Reason:** Reason for the request, typically "An app installed for your organization requires approval for access to Office 365 Data."
- **Requested at:** The DateTime of the request.
- **Request id:** The ID of the request, used for approval purposes.
- **DataTable:** The data set being extracted (for example, Sent Items).
- **Columns:** The list of columns being extracted from the data table (for example, SentDateTime).
- **AllowedGroups:** The group or groups of users against whom the pipeline is extracting data. If the list of groups is empty, the pipeline is requesting access to data from all users in the tenant.
- **User Scope Query:** The predicate used to filter out users. This only applies if the request is for all users in the tenant. If this is empty, no filter is applied.
- **OutputUri:** The output path in which the extracted data is stored.
- **SourceTenantId:** The tenant ID from which data is being extracted.
- **InstallerIdentity:** The identity of the app installer.


Improved Consent Experience

The current PAM consent experience is based on the Microsoft admin opting into Microsoft Graph Data Connect and creating a group of approvers for pipeline and data extraction requests. This process involves several roles with global tenant or application administrator roles that are involved in the tenant's usage of Microsoft Graph Data Connect—allowing for accountability in the use of Microsoft Graph Data Connect.

Improvements to our consent experience allow application-level management of the consent. Microsoft Graph Data Connect aims to provide approvers the full context on how data is being accessed and why. The new streamlined consent process also allows app developers to request the consent before developing the data access pipelines, eliminating delays/failures at runtime.

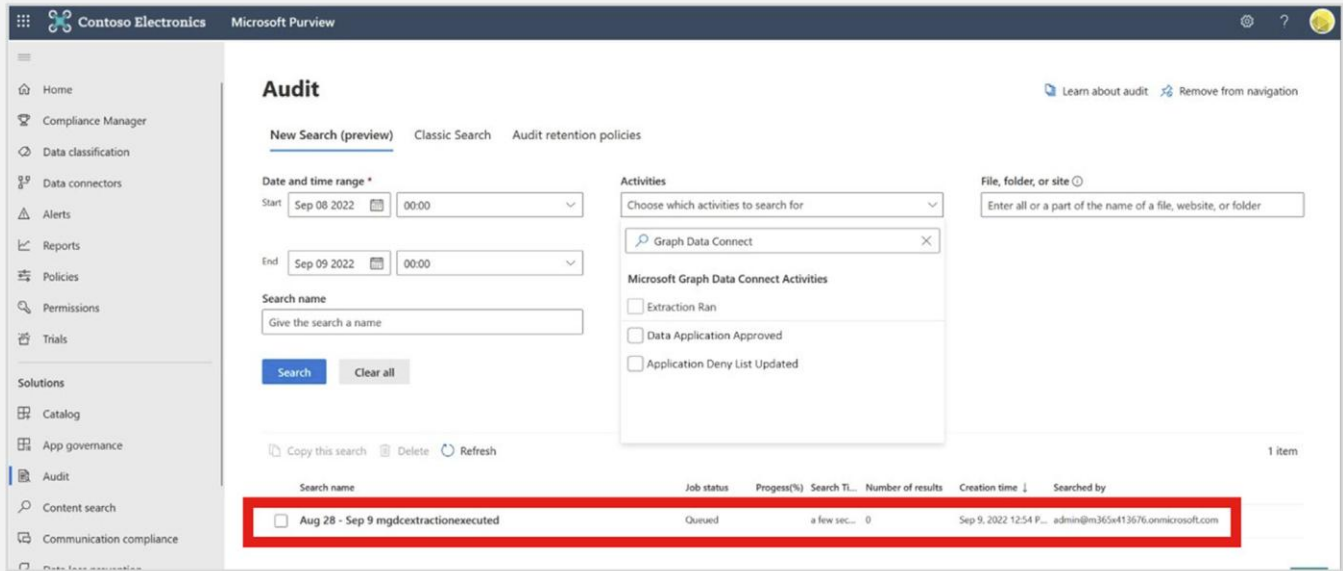
Key simplifications will include:

- Removing E5 license requirements for the approver
- Triggering consent before the pipeline run simplifies developer experience at runtime
- App-level consent provides one unified consent mechanism (instead of the existing per-dataset consent) allowing the admin to have a wholistic app view of data access. This helps the admin see a full picture of all the data the app intends to use and foresee any risks related to data access.



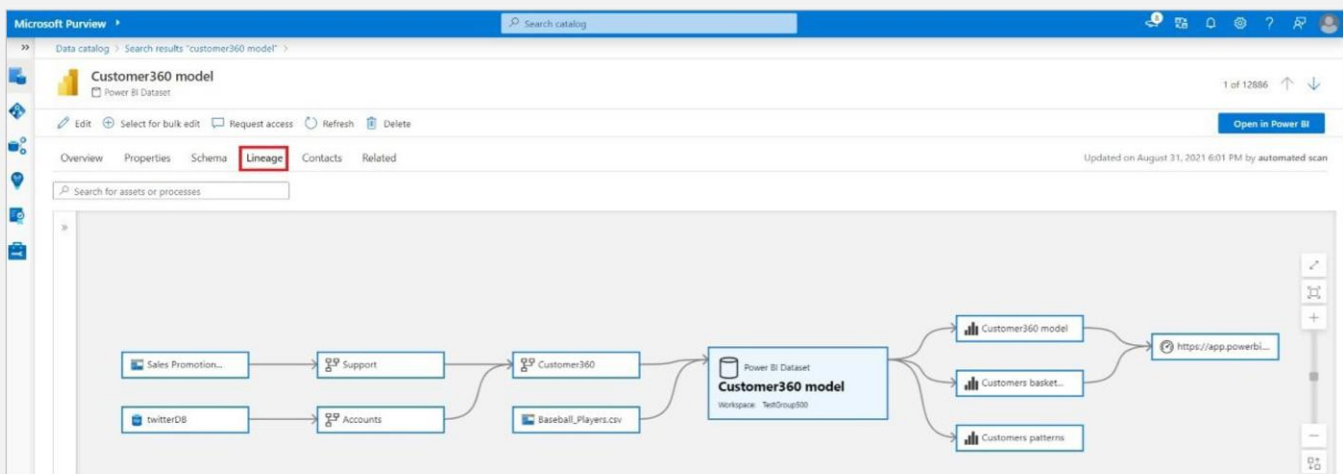
Data Lineage in Microsoft Graph Data Connect

Microsoft Graph Data Connect offers data lineage and tracking capabilities which tenant admins and authorized developers can access through [the compliance portal](#). Data tracking can help understand how data is being moved alongside recording who has access to it and which applications have authorized use of this data.



Customers can track data lineage movements for data quality analysis and troubleshooting purposes with [Microsoft Purview](#)

Microsoft Graph Data Connect offers governance capabilities such as audit log capabilities integrated with [Microsoft Purview \(compliance portal\)](#), which helps customers govern their data with granular controls. Tenant admins can authorize developers in their tenant to access audit logs in the compliance portal. Developers can track these logs in real time within the portal and export the data from the portal as well. These logs help developers understand consent runs (who authorized approvals and when) alongside details on pipeline runs such as datasets requested, status of runs, and requester. More functionality and information for audit logs can be found [here](#) and [here](#).



Admins can use Audit Logs above to record tenant creation and who triggers data movement.

Customers leverage [audit logs](#) to track how their data is being managed within Microsoft Graph Data Connect and enjoy easier data exporting functionality for their internal privacy and governance audit.

Microsoft Graph Data Connect Available in 14 Regions

Microsoft Graph Data Connect is now enabled in 14 regions (nine new regions since January 2023). The following table indicates which Microsoft 365 regions are supported and the corresponding Azure regions required for data movement (see below). Click here to learn more about [enabled regions](#).

Office Region	Azure region	
Asia-Pacific	• East Asia	• Southeast Asia
Australia	• Australia East	• Australia Southeast
Europe	• North Europe	• West Europe
North America	• Central US	• South Central US
	• East US	• West Central US
	• East US 2	• West US
	• North Central US	• West US 2
United Kingdom	• UK South	• UK West
Canada (CAN)	• Canada Central	• Canada East
Japan (JPN)	• Japan West	• Japan East
India (IND)	• South India	• Central India
Korea (KOR)	• Korea Central	• Korea South
Switzerland (CHE)	• Switzerland North	
Germany (DEU)	• Germany West Central	
Norway (NOR)	• Norway East	
France (FRA)	• France Central	
UAE (UAE)	• UAE North	

The alternate regions can be used when customers explicitly have a security requirement to lock down public IP access. To learn more, see our [security documentation](#).

Data Security Capabilities in Microsoft Graph Data Connect


Network Security

Our investments in network security give customers the ability to track and control how their data is being managed, accessed, and moved from Microsoft Graph Data Connect to their storage accounts. The requirements our Microsoft Graph Data Connect customers have will continue to shape the network security capabilities. Customer feedback is appreciated and directly influences the product roadmap. We would be happy to hear your additional requirements as well.

1. **V-net IR:** Customers can leverage their [Azure V-net IR](#) capabilities with Microsoft Graph Data Connect. This is enabled in the Microsoft Graph Data Connect Mapping Data Flow capability (see [MDF in ADF](#)). The capability improves network level security during data egress in Microsoft Graph Data Connect. The projected availability is expected to be in CY23H1 however timelines are subject to change.

This includes private endpoint connectivity and the removal of using Azure Public IR. Customers no longer need to allow-list Public IPs when setting up data pipelines, ensuring for more security when data is in transit and reducing any overhead or risks involved with potentially exposing data publicly.

2. **Azure Key Vault:** Microsoft Graph Data Connect is reducing the efforts associated with [Service Principle \(SPNs\) with Managed Identity \(MSI\)](#). Customers no longer have to manage ownership and renewal of credentials when using Microsoft Graph Data Connect. They can instead leverage [Azure Key Vault](#) (AKV) credential management. With AKV, customers can store credentials needed for Microsoft Graph Data Connect to authenticate data movement. Using AKV allows developer peace of mind when it comes to managing their secrets and credentials as it's collectively secure in one place. Developers can use these credentials stores in AKV without any risk of accidentally exposing the content of the secrets to themselves or others.



Encryption

Microsoft Graph Data Connect offers out-of-the-box custom encrypted datasets alongside Azure encryption capabilities. Additionally, with [Azure Key Vault](#), customers can store their public and private keys in AKV and refresh them when needed.

1. **Encryption-in-transit:** Through custom encryption for datasets, Microsoft Graph Data Connect delivers datasets to multi-tenant applications to their storage accounts encrypted by default with a private key owned by the service. Multi-tenant application developers can then request the decryption of datasets for analysis. This feature allows customers the option to choose encrypted data delivery while giving them full control over their data and they can own the encryption key. Additionally, communications in data requests between Microsoft 365 and Azure resources while using Microsoft Graph Data Connect are securely authenticated and are using service standards from here alongside being audited [via SOC2](#).
2. **Encryption data-at-rest:** Customers can use [Azure's data-at-rest encryption feature](#) and [customer managed keys](#) when setting up their destination Azure storage account for Microsoft Graph Data Connect. Azure encrypts the data by default and allows for additional security measures.

Customers can now leverage encrypted datasets capabilities to secure their data, especially with the involvement of third-party independent software vendors (ISVs). This helps customers achieve the perfect balance of using Microsoft Graph Data Connect while safeguarding their data amongst their tenants and contracted ISVs.

Geo Isolation for Data Access by Default

Many multi-geo customers have their tenant data spread across different Microsoft 365 regions to align with their user locations. Thus, Microsoft Graph Data Connect enforces customers to adhere to different regional data security and privacy standards along with managing different Azure and Microsoft 365 regions.

Microsoft Graph Data Connect respects these geo boundaries by requiring data engineers to set up separate data pipelines per region with the corresponding Azure environments setup in each region. All data handling within Microsoft Graph Data Connect aligns with these geo boundaries to guarantee geo isolation of data.

More information on how geographic compatibility works when considering Azure and Microsoft 365 regions can be found [here](#).

Data Privacy and Access

A key component to ensuring a successful platform with Microsoft 365 data is being aware and planning for the privacy requirements our customers have and their internal policies. Microsoft Graph Data Connect functions provide controls to configure the data to meet privacy requirements. Below are key principles to consider when handling employee and Microsoft 365 data.



Data Privacy Principles and Recommendations

Clarify analysis goals and approach

There are many uses and specific data elements in Microsoft 365. Before you begin, it is recommended to determine the specific outcome that you desire. An analysis approach can be centered on a specific challenge that you want to address or a clear data-driven hypothesis. A clear analysis approach is useful in determining the employee/ mailbox scope, the specific data elements to use, other needed data sources, who will handle the data and who will see the results. It will also provide clarity about including other stakeholders and privacy security processes that your company has in place.

Determine the scope of included teams

You determine the scope of employee Microsoft 365 data that is required for your analysis. You can limit that scope to any specific groups that are relevant for your analyses. Microsoft Graph Data Connect allows you to create Microsoft Entra groups and you decide who to include. Microsoft Graph Data Connect also allows you to exclude specific individuals or teams that should not be part of the analysis through allow and deny lists. For a customer relationship analytics example, you might decide to include only the data from customer facing teams. For a SharePoint capacity case, you might not choose to include any other sensitive Microsoft 365 data.

Limit the Microsoft 365 data you need

Based on your analysis plan, you can determine which specific Microsoft 365 data elements are relevant. Microsoft Graph Data Connect allows you to use filters that let you define the specific mailboxes in scope, the time horizon of the data, and the selection specific data elements while discarding what you do not need to inform your analytics scenario. For most analytics cases, meta data, such as sender and receiver information, is sufficient. For other use cases, such as sentiment analysis, body of email or chat messages might be required. Based on your privacy considerations and desired outcome, you get to decide which specific elements from email, calendar, Teams chats, Teams meetings or SharePoint you would like to ingest.

Consider other data provided by your organization

You also control additional data sources that might be required for your analyses. Examples can include additional HR data, CRM data or LOB data. This data will come from other respective systems and can be ingested into your analytics environment and be merged with the Microsoft 365 data. Most approaches use the minimally required data to achieve the analysis goal. You can consider what type of data is necessary and if personal information should be included, or if there are other approaches that might avoid the use of personal data.

Determine who will use the data

You are in control of who will work with the data and analysis results. Like any other sensitive data in your organization, such as HR or financial data, you can apply the same considerations. The roles that usually participate in analyses include developers and data analysts. You can limit access to the data. You can also limit the exposure of the insights to the right constituencies via access controls in services like Azure or PowerBI. Like any other sensitive data, handling, storing and using Microsoft 365 should follow similar guidelines and requirements that you have established in your organization.

Consider using relevant privacy features and methods

Aligned with your desired analysis goal, you might consider additional data privacy protection mechanisms. Microsoft Graph Data Connect will soon allow to obfuscate data fields and elements (See below). Other techniques to consider are to aggregate the data to a certain group number and report it on that aggregation level. These are additional privacy methods you can consider in addition to the right user scope, the time period of the data and specific data fields that are needed.

Consult with subject matter experts

Your organization might have specific teams and processes that can help guide the handling of sensitive data, including from Microsoft 365. Subject matter experts in HR, privacy, IT, security and legal can help validate the analysis approach and data needs. You might consider the requirements of different countries, regulations like GDPR, or privacy, data security and data governance requirements of your company. Due diligence is particularly important in highly regulated jurisdictions like the European Union. Internal assessments, like a Data Protection Impact Assessment (DPIA) or a similar process, might be required in your company prior to handling sensitive data.



Data Privacy Capabilities in Microsoft Graph Data Connect

Privacy Scrubbing with Allow and Deny list.

Microsoft Graph Data Connect offers privacy scrubbing to allow customers to exclude specific constituencies and data elements. Customers can utilize Allow or Deny lists to protect sensitive data such as executive emails, calendars, communications and ensure privacy amongst a specific group in their tenant. Customers can use this capability to extract the data they need while avoiding extracting data containing content they wish to protect.

1. Customers can use allow lists to set up the scope of certain employees and mailbox related to those employees and can check datasets are enabled and compatible for [privacy scrubbing](#).
2. Customers can leverage the deny list to exclude certain individuals or departments that might have sensitive information. For example, customers can use this feature filter out sensitive communications from their C-suite, Legal, or HR teams by adding those departments to the deny list. Microsoft Graph Data Connect will not extract data for individuals on the deny-list during extraction or pipeline runs.

The table below outlines how allow and deny lists work. Deny lists are currently not available for Teams chats.

Options to Exclude	Email – Messages Data Set	Email – SentbData Set	Teams – Chat Messages
Allow List (Microsoft Entra Group): Excluding Executive mailboxes	<ul style="list-style-type: none"> • Messages (send and received) from excluded mailboxes will not be extracted • Messages data set of included mailboxes will be included. This can include: <ul style="list-style-type: none"> – Emails received from excluded mailboxes – Emails sent to excluded mailboxes 	<ul style="list-style-type: none"> • Messages sent from excluded mailboxes will not be extracted <ul style="list-style-type: none"> – “Sent” data set only contains sent emails. Therefore, emails from excluded mailboxes will not show in the recipient’s data set • Messages sent to excluded mailboxes will be included 	<ul style="list-style-type: none"> • Chat data set (send and receive) from excluded Teams users will not be extracted • Chat data set of included Teams users will be included. This can include: <ul style="list-style-type: none"> – Chats received from excluded Teams users – Chats sent to excluded Teams users
Deny List – Including executive mailboxes on deny list by Microsoft Graph Data Connect pipeline approver	<ul style="list-style-type: none"> • Emails sent from executives are excluded • Emails with executives in to, cc and bcc are excluded 	<ul style="list-style-type: none"> • Messages sent from excluded mailboxes will not be extracted <ul style="list-style-type: none"> – “Sent” data set only contains sent emails. Therefore, emails from excluded mailboxes will not show in the recipient’s data set • Messages sent to excluded mailboxes will be included 	<ul style="list-style-type: none"> • Not available for Teams chat

Scope and Filtering

Microsoft Graph Data Connect has scoping and filtering capabilities that narrow down the amount of data needed to focus on for insight and analyses. For example, this allows customers to focus only on meta data required such as sender or receiver while leaving out sensitive data such as email bodies if they are not needed. Filtering capabilities help the customer reduce their overall risk of data exposure, protect the data of users who are filtered out, and reduce the monetary cost of data.

Note: SharePoint (ODSP) datasets are not available for filtering yet. Data filters exist for Microsoft Entra, Outlook (Emails and Meetings) and Teams datasets.

1. **Mailbox and Employee Filter:** Is limited to the right employees in scope through creating a Microsoft Entra group for mailboxes in scope.
2. **Data Elements:** The filter focuses on the specific elements in Microsoft 365 data sets. Only selected data elements will be exported through Microsoft Graph Data Connect. Customers can scope down on metadata elements like content for email, meetings, teams.
3. **Time period:** Customers can select for a range of time when requesting data. Customers can leverage this for insights which may rely on historical data and data exports by referring to the very specific time range of extraction. Time period filters are not available for SharePoint data sets as SPO data sets retain 21 days of history after which the data will be deleted due to data regulations.

The screenshot displays the configuration for a 'Copy data' activity in Microsoft Graph Data Connect. The 'Source' tab is selected, showing the source dataset 'Office365Table1'. The 'Scope' is set to 'All users in the Microsoft 365 tenant'. The 'Scope filter' is empty. The 'Date filter' is set to 'createdDateTime' with a start time of '1/1/22 00:00:00' and an end time of '4/1/22 00:00:00'. The 'Output columns' section shows a list of columns with checkboxes for selection. The 'receivedDateTime' column is selected. A red box highlights the 'Scope' and 'Date filter' sections, with a red arrow pointing to the 'Scope filter' field labeled 'Scoping M365 data'. Another red box highlights the 'Output columns' table, with a red arrow pointing to the 'receivedDateTime' row labeled 'Filtering columns'.

Column name	Type
<input checked="" type="checkbox"/> receivedDateTime	ANY Default
<input type="checkbox"/> sentDateTime	ANY Default
<input type="checkbox"/> hasAttachments	ANY Default
<input type="checkbox"/> internetMessageId	ANY Default
<input type="checkbox"/> subject	ANY Default
<input type="checkbox"/> importance	ANY Default
<input type="checkbox"/> parentFolderId	ANY Default
<input type="checkbox"/> sender	ANY Default

Scope Selection and Filtering for Microsoft 365 Data in Microsoft Graph Data Connect

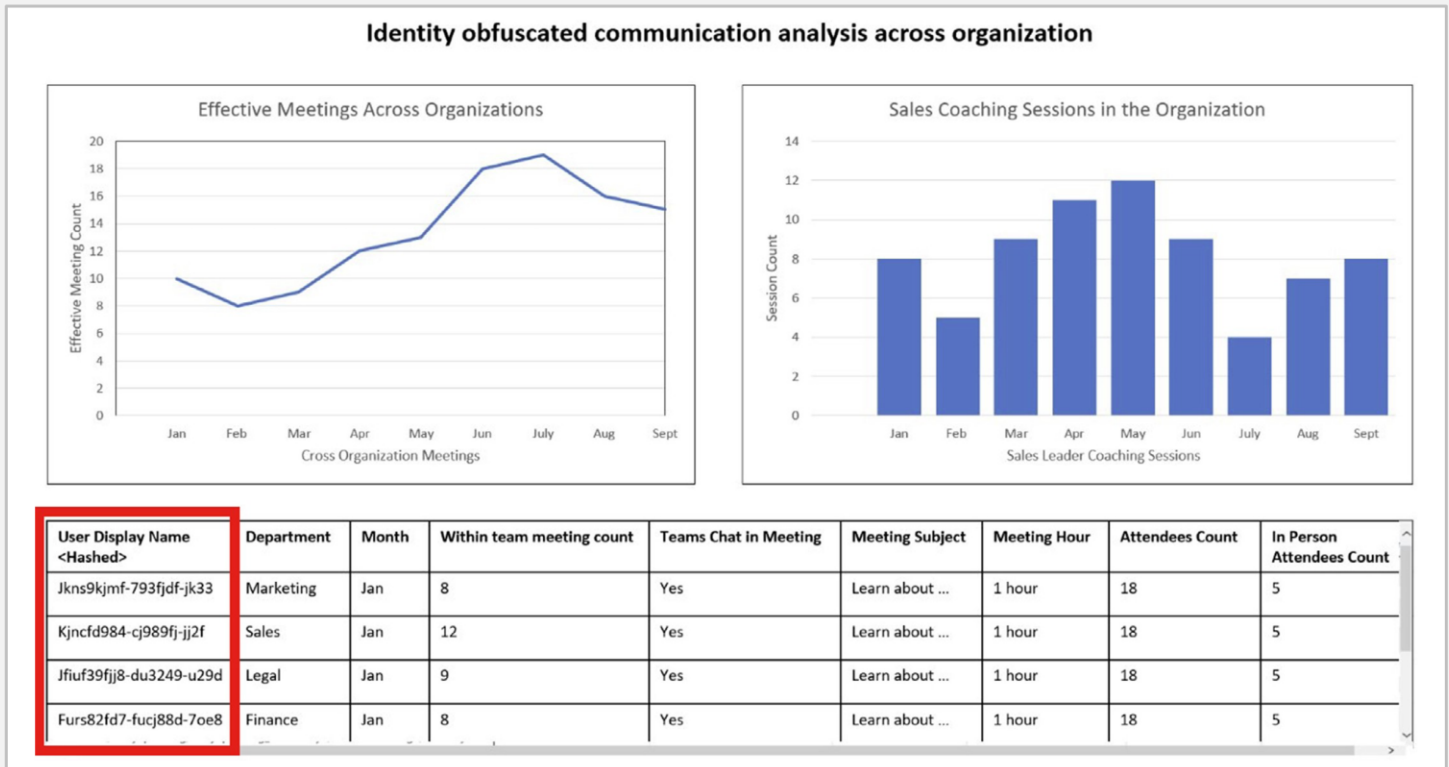
In addition to the filtering capabilities in Microsoft Graph Data Connect, customers can leverage additional filtering once the Microsoft 365 data has been egressed into Azure Synapse:

- 1. Domain filtering:** Filter collaboration data by domain to exclude external communication. [Synapse "endsWith" Expression function](#) allows customers to implement this approach as needed. Note that doing so prevents the ability to deliver on certain use cases (e.g. Customer Relationship analytics, Supply Chain optimization, etc.). This functionality is available on Microsoft Graph Data Connect data using Synapse after extraction.
- 2. Term based filtering:** Using [Synapse Expression functions](#) like "contains" provides customers with term-based filtering and removal of data. All message and events are filtered out based on terms present in the content body like "private", "ACP – Attorney Client Privileged", etc. This functionality is available on Microsoft Graph Data Connect data using Synapse after extraction.

Data Obfuscation

Obfuscation is a measure to help protect any user-identifiable data, including personal identifiable information (PII) across datasets. The projected availability is expected to be in CY23H2 however timelines are subject to change.

Data obfuscation in Microsoft Graph Data Connect allows customers to mask data and to maintain data privacy. In this case, Microsoft Graph Data Connect scrubs out any PII related to first and last name that occur within fields such as "email address", "first name", "last name", etc. in Microsoft Graph Data Connect datasets.



Customers can still gain valuable insights while protecting PII through obfuscation.

Microsoft Graph Data Connect provides only one-way de-identification obfuscation. Once developers and approvers enable obfuscation for datasets for their application, the data will be hashed consistently across the application. However, there is no reversing the hash, not only for the data once it reaches the customer storage account, but also for the pipeline run once the request is approved.

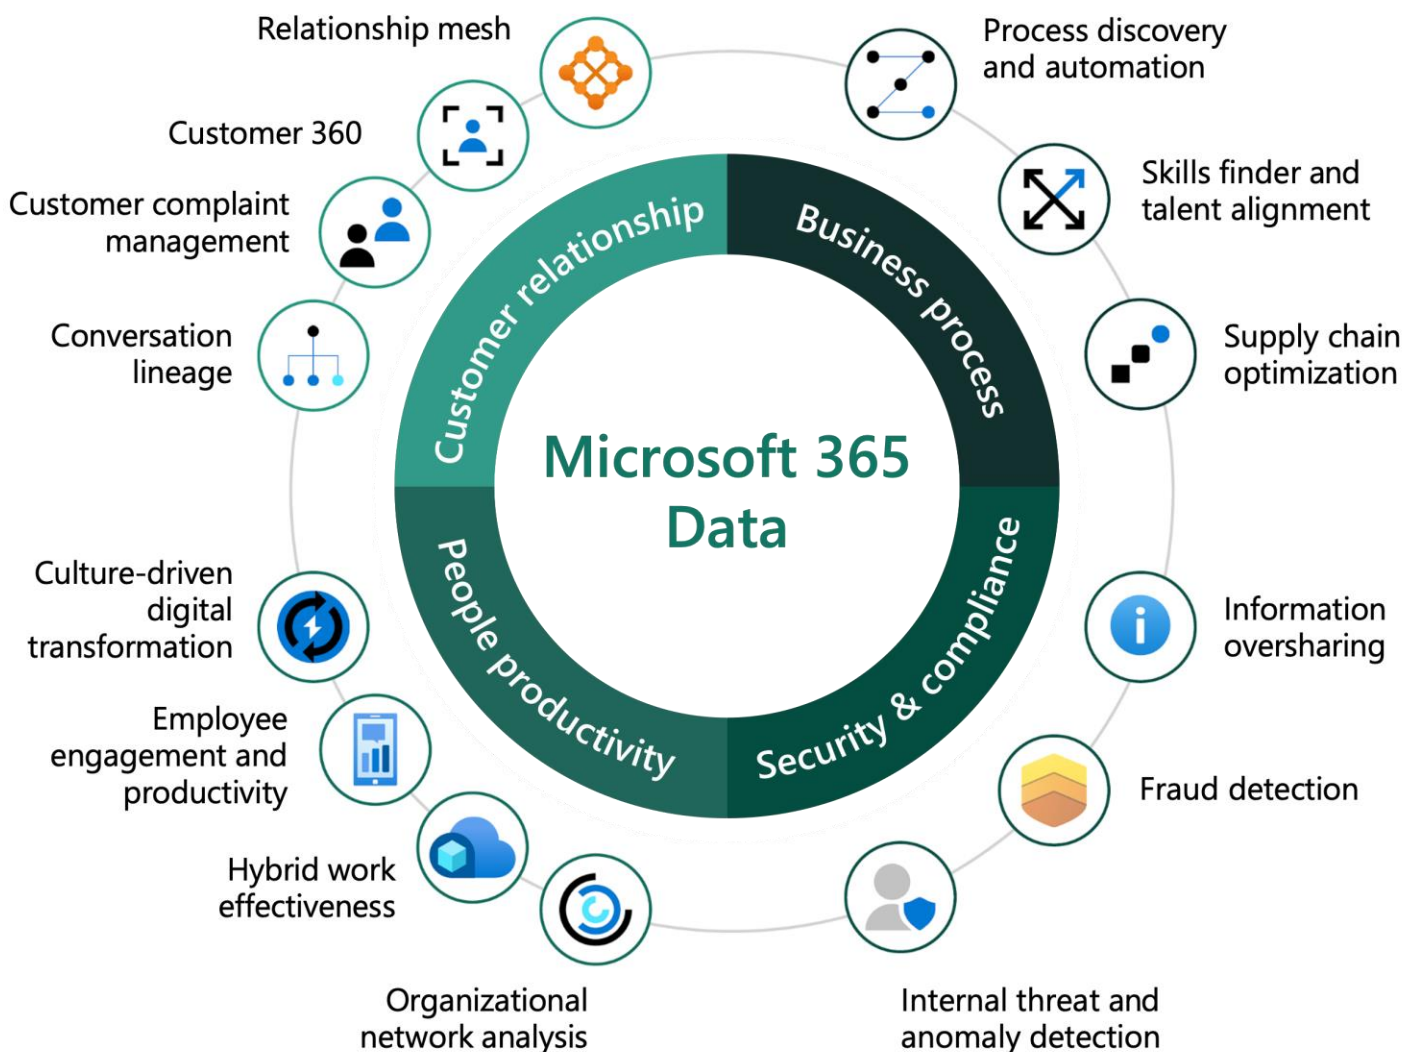
1. Microsoft Graph Data Connect provides customer choice with an “all or nothing” approach to obfuscation as the first iteration of its release. In this release, the app developer and approvers can choose to enable obfuscation for their datasets in the application during application registration. This ensures all datasets in the application are hashed consistently and can be joined simultaneously without any risk of privacy loss. Microsoft Graph Data Connect will carry out post processing to remove basic PII such as fields related to first and last name. The obfuscated data will then be delivered to the customer’s destination storage account for scenario use. Microsoft Graph Data Connect aims to provide out-of-the-box capabilities with obfuscation to take the workload off customers when it comes to meeting their internal privacy requirements—minimizing risk and development cost for the customer to build these capabilities by themselves.
2. Over time, we plan on extending this capability to column obfuscation where developers can choose per dataset, which columns to obfuscate and which remain as is—allowing for a better balance between joining data together and limiting privacy loss. This gives a more flexible approach to the developer and tenant on how they choose to enact their privacy principles.
3. ISVs can also leverage obfuscation on behalf of their customers similar to the custom encryption work enabled on Microsoft Graph Data Connect—allowing for more security and trust between customers, ISVs and Microsoft Graph Data Connect.



Conclusion

Today's organizations need a new set of security tools to take control of their data and get organizational insights to effectively adapt to the complexity of the modern work environment. Microsoft Graph Data Connect protects sensitive [Microsoft 365 data](#) through our security, privacy, and governance capabilities that come ready to use, out of the box.

Microsoft ensures that customer trust is at the center of all our product and services. To learn about Microsoft's complete approach to end-to-end security, please see the following related to [Microsoft 365 data](#) and [Microsoft cloud](#). We are constantly expanding our security, privacy, and governance portfolio based on customer feedback. For any further questions or feedback, please reach out to us at dataconnect@microsoft.com.



Key data insight scenarios

Appendix 1

Data-Protection Considerations

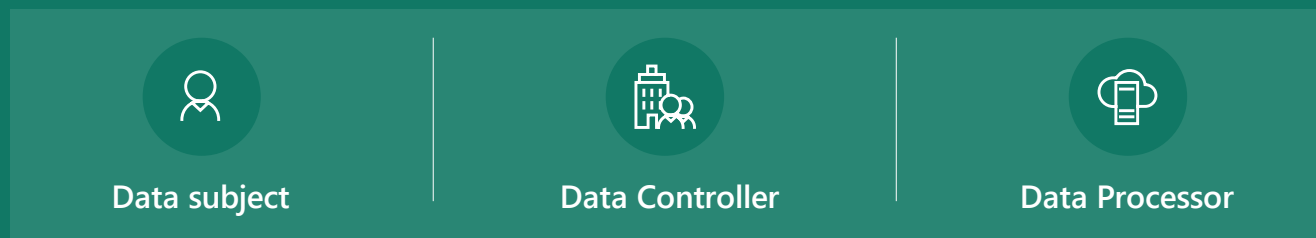
Using data generated from everyday work in Microsoft 365 provides the single largest data set about how work actually gets done in an organization by tapping into the actual collaboration data that fuels productivity.

The following includes a basic overview of the roles, responsibilities, types of data, and data-privacy recommendations. The general suggestions offered here are a starting point for planning your data-protection strategy and deployment. These are not intended as a substitute for addressing your organization's unique needs by engaging with legal, privacy, human-resources, and other subject matter experts within your organization.

Roles and Responsibilities

The concepts of [data controller](#), [data processor](#), and [data subject](#) originate in European privacy law. Regardless of where your organization is located or whether any personal data of European Union citizens is involved, these concepts provide a useful framework for thinking about data protection.

The following graphic shows the central position of the data controller (your organization) between the data subject (left) and the data processor, Microsoft (right):



Before accessing Microsoft 365 data, first consider the respective roles and responsibilities of your organization and of Microsoft to protect personal data and honor the rights of data subjects.

Data controller

The data controller is a party that determines the purposes and means of processing a data subject's personal data.




As a data controller, your organization should:

- Determine the scope of data to analyze and the purpose and objectives of the analysis.
- Use controls to direct what data will be analyzed, how data will appear in results, and who will have access to both raw data and the results of analysis.
- Work with your organization's legal, privacy, and human resources teams for the following tasks:
 - Determining whether you should obtain consent from employees in your company.
 - Determining what information you provide to employees about how your organization will process their personal data in Microsoft 365 data.
 - Accounting for local or country-specific considerations.

Note:

Some countries require employers to consult with employee representatives or seek approval from a works council before deploying certain information technology in the workplace, while others restrict when and how employers can process certain employee data. For example, if your company has employees in Germany or Netherlands, then you should consider if works council engagement or approval is required. Moreover, you choose if Microsoft 365 data includes employee communications, which could be considered "communications data" (including "traffic data") in Finland. Thus, if your company has employees in Finland, then you should understand how Finnish laws apply to the processing of employee personal data and communications or traffic data to determine if use is permissible.



Data Processor

The data processor is a party that processes personal data on behalf of the data controller. When your organization uses Microsoft 365, Microsoft is the data processor.



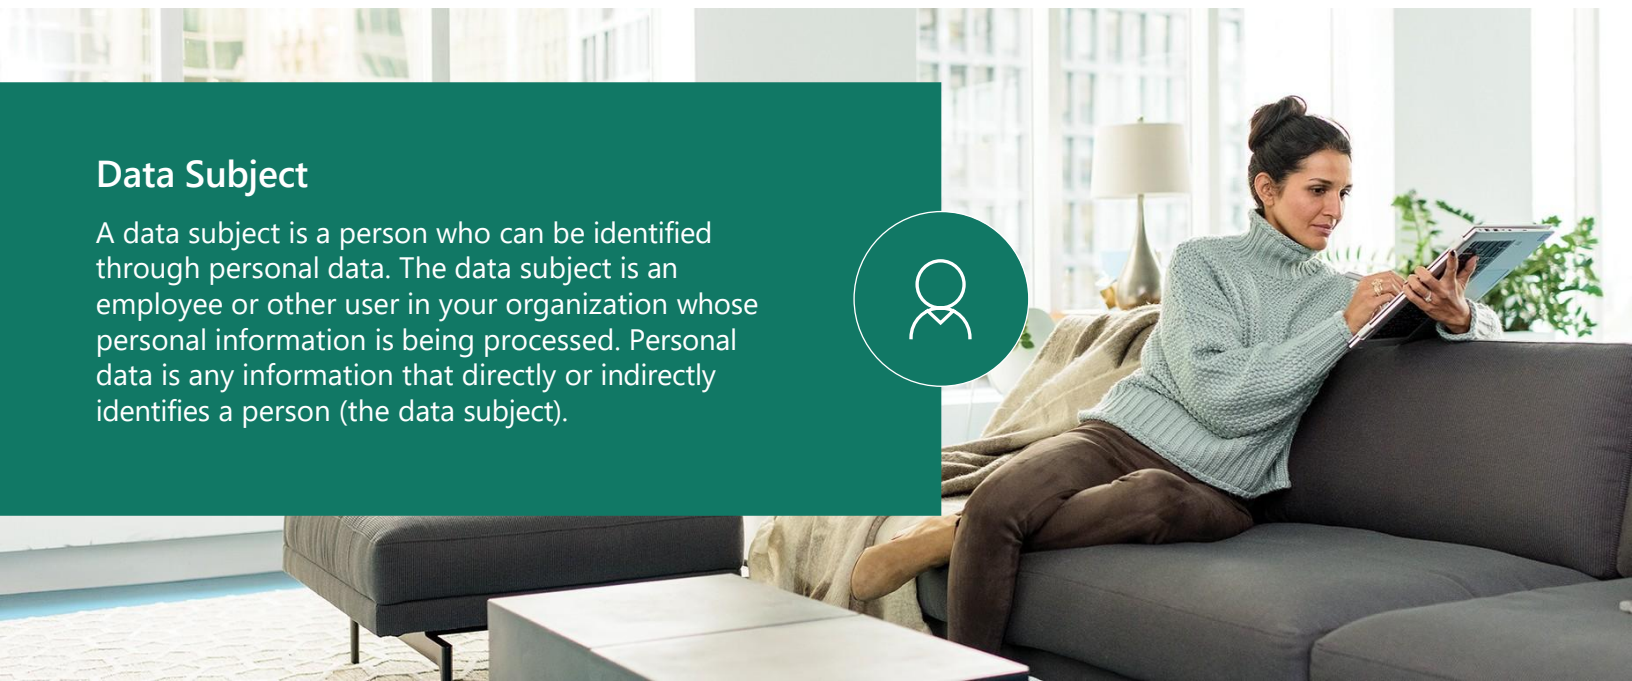
As a data processor, Microsoft will:

- Process personal data in accordance with your organization's instructions as directed through your settings configuration.
- Process all data provided to Microsoft (including personal data) according to the same general privacy and security terms in the Online Services Terms (OST) as Microsoft 365.
- As part of Microsoft's commitments under the OST, remain certified under the EU-U.S. and Swiss-U.S. Privacy Shield Frameworks and the commitments that these frameworks entail to legitimize transfers of personal data from the EU and Switzerland to the U.S.
- Contractually commit to abide by applicable provisions of the European Union General Data Protection Regulation (GDPR), effective starting May 25, 2018.
- Provide features that help organizations meet their data-controller obligations and honor data-subject rights under the GDPR, including the right of exclusion from processing, access, and erasure, and including the right of transparency regarding methods of processing.
- Implement technical and organizational security measures to protect the confidentiality of your organization's (and employees') data.

In addition, Microsoft does not use data for advertising, nor does it volunteer to provide data to law enforcement.

Data Subject

A data subject is a person who can be identified through personal data. The data subject is an employee or other user in your organization whose personal information is being processed. Personal data is any information that directly or indirectly identifies a person (the data subject).



Determine if Data Protection Impact Assessment (DPIA) is Needed

The degree of privacy risk to employees and other users in your organization is largely within your control. The risk depends primarily on the Microsoft 365 data leveraged and how you will use that data.


After you have developed an analysis plan but before you begin processing data, determine whether you need to complete a DPIA. If your proposed use involves processing personal data in a manner that could lead to high risks to the rights of employees and other users in your organization, completing a DPIA might be warranted. If you are unsure whether a DPIA is required, consult your organization's privacy subject matter experts, such as legal or HR personnel.

Higher risk data includes sensitive demographic data, such as racial or ethnic origin, sex or gender, and trade union membership. Higher risk uses include using the service for profiling or making automated decisions or predictions about employees.

If you determine that a DPIA is necessary, you will need to document several aspects of how you will use the service to process personal data, including how the data will be collected; how it will be processed; the necessity of the processing in relation to the purpose; what risks the processing presents to employees; the data flows; feedback that you received from employees regarding the proposed data processing; and any other information that your organization's data-protection officer (or equivalent) deems necessary for the DPIA. As a data controller, your organization is entirely responsible for determining your organization's purposes for using Microsoft 365 data. As a data processor, Microsoft informs you how the product functions and processes data pursuant to your configuration of the services.

Note:

Microsoft [offers more details on DPIAs, relevant Microsoft 365 information](#) – including a [DPIA Service Elements Matrix](#) and [a customizable DPIA document](#) to help you get started.



Support for handling data subject requests

Under GDPR, CCPA and other emerging US State legislation, data subjects may have rights to request exclusion from processing, access, correction, or deletion of their personal data. It is your organization's role as data controller to evaluate whether a particular data subject request is valid and, if appropriate, to take action to fulfill the request. As a data processor, Microsoft provides mechanisms for your organization as the data controller to honor data subject rights through controls.

- **Exclusion from processing** – Data subjects have the right to have their personal information excluded from processing.
- **Access** – Data subjects have the right to demand what personal information is being processed, and Microsoft 365 gives you the ability to export the raw data, which may contain personal data.
- **Correction** – Data subjects have the right to rectify their personal data. This data is not corrected through Microsoft 365 data controls.
- **Deletion** – Microsoft supports the GDPR [Right to erasure](#). Additionally, if necessary, customers themselves can also delete reports that identify the data subject.

Note:

More information on data subject [requests are available](#), including specifics for Microsoft 365.

