



SECURITY FACT SHEET

AI Assistant in Adobe Workfront

August 2025



About AI Assistant in Workfront

AI Assistant in Adobe Workfront is a generative AI tool designed to enhance productivity and help users expand product mastery, efficiently navigate enterprise data objects, and simplify tasks. Offering in-app information and suggestions in a natural-language conversation, AI Assistant in Workfront can help:

- Summarize work items or documents
- Find instructions or reference material for work processes
- Generate or check formulas for calculated fields

AI Assistant in Workfront Components

There are three (3) key components in AI Assistant in Workfront:

Workfront User Interface — Users interact with the AI Assistant by clicking on the icon in the upper right-hand corner of the Adobe Workfront UI, which reveals a right rail screen with a text box where the user can enter prompts.

Generative Experience Models (GEMs) — The primary “brains” behind AI Assistant in Workfront, GEMs include foundation agents that power AI Assistant use cases to provide engaging and natural responses.

- **AI Assistant Router GEM** — Orchestrates other models and directs user prompts to the correct model. The AI Assistant Router GEM is the core AI Assistant in Workfront conversational interface.
- **Product Knowledge GEM** — Identifies the appropriate Adobe Workfront product documentation to answer a given question, retrieves the relevant documentation, generates the appropriate answer based on the retrieved documentation, determines appropriate source citations, and verifies that responses are grounded.
- **Summarization GEM** — Translates a given question into a query against the underlying customer-specific operational object data and generates the appropriate summarization based on the data retrieved.
- **Formula GEM** — Uses Workfront-specific calculation definitions and fetches relevant customer-specific fields to use in generating a formula from natural language and verifies that formulas are grounded.
- **Workfront Actions GEM** — Leverages Workfront work object data, to find, create, update, and delete records in a Workfront portfolio, program, project, task, issue, document or planning record.

Data Services — API-based services invoked by GEMs to query the data stores that contain relevant data. Data in the data stores is organized, pre-joined, and indexed into a knowledge base, which then enables the GEMs to interact with it in an open-ended fashion. AI Assistant in Workfront includes two (2) data services:

- **Product Knowledge Data Service** — Queries a customer-agnostic product knowledge data store that contains general Adobe Workfront knowledge, such as product documentation.

- **Workfront API** — Certain GEMs will query a customer-specific Workfront application data store for work object data, partitioned by customer and environment. This data is the same customer data available via API for customers to access their own Workfront data and is used by the Summarization GEM, Formula GEM, and Workfront Actions GEM. The relevant GEMs query data associated with one or more of the object types can be found in the [AI Assistant overview](#) on Adobe Experience League.

Accessing AI Assistant in Workfront

AI Assistant in Workfront can be accessed by clicking the AI Assistant icon at the top of any Workfront page and typing a question or prompt into the panel at the right of the screen. The organization's Workfront administrator controls which users in the organization can access AI Assistant [using role-based access controls](#) in the Adobe Admin Console. Prerequisites to enable the AI Assistant in Workfront can be found in the [AI Assistant overview](#).

Please note: Each user granted permission to use AI Assistant in Workfront must accept the [Adobe Experience Cloud Generative AI User Guidelines](#) upon first use of the feature.

Data Encryption

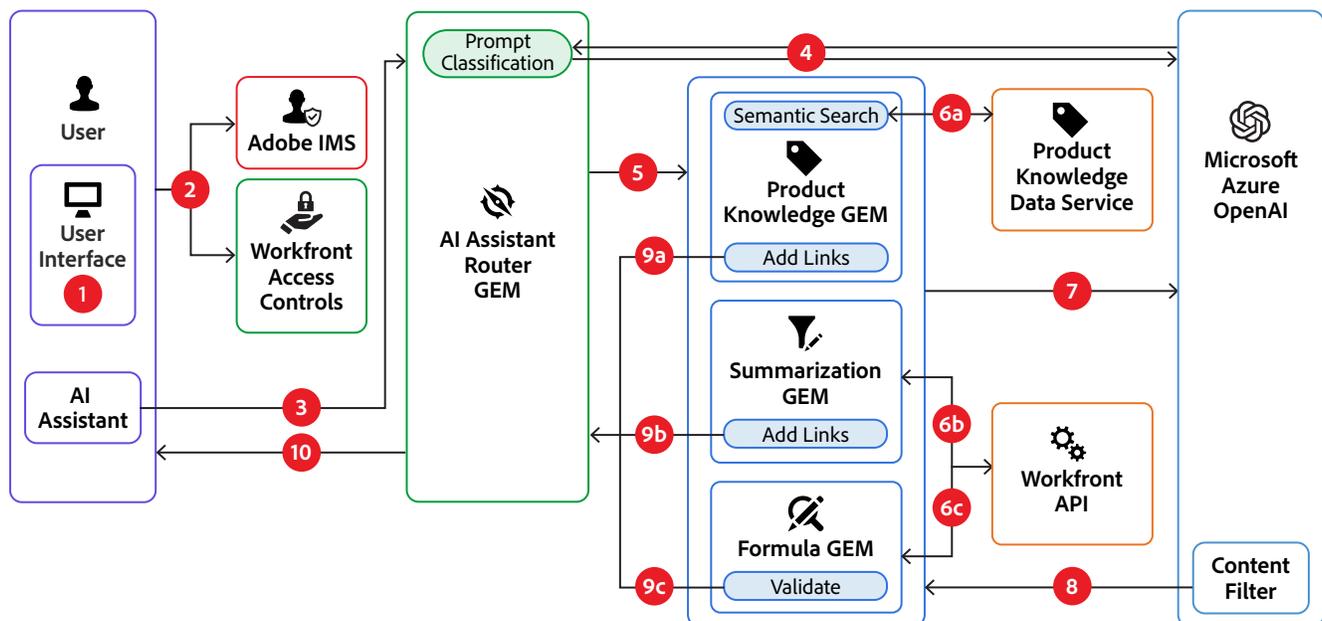
In Transit: All data is encrypted in transit over HTTPS using TLS 1.2 or greater.

At Rest: Any data stored by AI Assistant in Workfront is encrypted at rest using AES 256-bit encryption.

Security Architecture and Data Flow

General Workflow (Product Knowledge, Formula, Summarization)

The following example data flow illustrates how data flows in AI Assistant for product knowledge, formula and summarization questions:



Step 1: The user opens AI Assistant in the Adobe Workfront user interface.

Step 2: AI Assistant in Workfront authenticates the user with [Adobe Identity Management Services \(IMS\)](#), checks that the user is entitled to use AI Assistant and has been given access to use the AI Assistant by their Workfront Administrator.

Step 3: The user enters a product knowledge-type question in the prompt text box and the Workfront UI sends the prompt to the AI Assistant Router GEM.

Step 4: The AI Assistant Router GEM uses Azure OpenAI to help classify the user's prompt into a product knowledge question, a summarization request, or a formula help question, or determine if it is out of scope. If the question is out of scope, the user receives an error message, otherwise the process continues to step 5.

Step 5: The user's prompt is forwarded to the appropriate GEM based on the prompt classification.

Step 6a: If the AI Assistant Router GEM determines that the user's prompt is a product knowledge question, the Product Knowledge GEM uses semantic search to retrieve relevant snippets of documentation from the Product Knowledge Data Service to answer the question.

Step 6b: If the AI Assistant Router GEM determines that the user's prompt is a summarization question, the Summarization GEM makes a Workfront API call on the user's behalf to get the object that is intended to be summarized. The request is made with all the same access controls the user has in the Workfront application.

Step 6c: If the AI Assistant Router GEM determines that the user's prompt is a formula help question, the Formula GEM makes a Workfront API call on the user's behalf to retrieve the fields available on the user's current page. The request is made using the same access controls the user has in the Workfront application.

Step 7: The relevant GEM combines the prompt text with retrieved snippets of documentation or work objects from step 6 and sends them to Azure OpenAI.

Step 8: Before sending the formulated answer back to the relevant GEM, the Azure OpenAI content filtering service moderates generated responses that violate Azure Open AI user guidelines.

Step 9a: If the question was a product knowledge question, the Product Knowledge GEM cross-checks the answer provided by the Azure OpenAI Service, adds the appropriate links to the documentation snippet sources, and sends the complete answer and citations to the AI Router GEM.

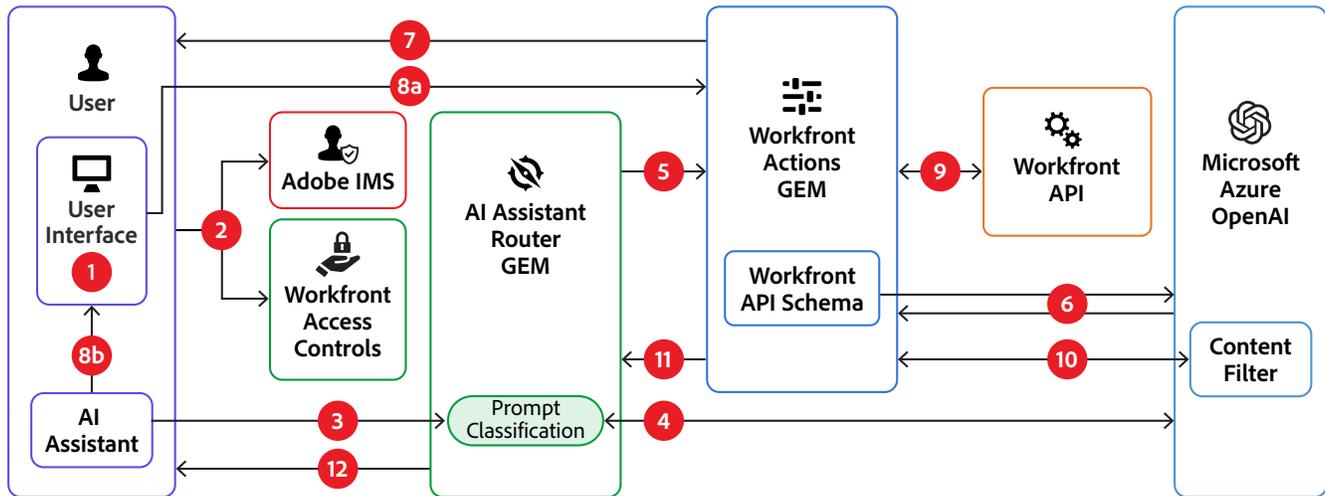
Step 9b: If the question was a summarization question, the Summarization GEM adds appropriate links to annotate the summarized Workfront objects and sends the response back to the AI Router GEM.

Step 9c: If the question was a formula question, the Formula GEM validates the formula being returned and returns the response to the AI Router GEM.

Step 10: The AI Router GEM returns the answer and the relevant citations to the user in the AI Assistant for Workfront UI.

Find, Create, Update or Delete Workfront objects Workflow

The following example data flow illustrates how data flows in AI Assistant for Workfront actions:



Step 1: User opens AI Assistant in the Adobe Workfront user interface.

Step 2: AI Assistant authenticates the user with [Adobe Identity Management Services \(IMS\)](#), checks that the user is entitled to use AI Assistant and has been given access to use the AI Assistant by their Workfront Administrator. The AI Assistant also checks the user's Workfront Access Controls and ensures that all actions taken on the user's behalf are done with the user's access in mind.

Step 3: User enters an action-based request in the prompt text box and the Workfront UI sends the prompt to the AI Assistant Router GEM.

Step 4: The AI Assistant Router GEM uses Azure OpenAI to help classify the user's prompt into an action request (find, update, create, or delete) or determine if it is out of scope. If the question is classified as out of scope, the user receives an error, otherwise the process moves to step 5.

Step 5: The user's prompt is forwarded to the appropriate GEM based on the prompt classification..

Step 6: The Workfront Actions GEM combines the prompt with the Workfront API schema and sends it to Azure OpenAI to generate an API-based execution plan for how to achieve the desired action in the user's request.

Step 7: If the request is a find action, the process moves to Step 9 because no user confirmation is needed to execute the plan. If the request is a create, update, or delete request, the Workfront Actions GEM sends the plan back to the user to confirm the action should be performed.

Step 8: The AI Assistant presents the user with what objects are going to be created, updated, or deleted, along with an option to approve or deny the plan.

Step 8a: If the user approves the plan, the workflow moves to step 9.

Step 8b: If the user denies the plan, the AI Assistant responds back to the user that the plan was cancelled, and the workflow does not move further.

Step 9: The Workfront Actions GEM executes the approved plan by sending a series of API requests to the Workfront API with the appropriate user access levels.

Step 10: The Workfront Actions GEM sends the results of the action command to Azure OpenAI to create a natural language message describing what objects have been found or changed. The Azure OpenAI content filtering service moderates generated responses that violate Azure Open AI user guidelines.

Step 11: The Workfront Actions GEM combines the natural language message with the links of the Workfront objects that were found or changed (e.g., if 3 projects were found, the Workfront Actions GEM adds those customer-specific project links to the response), and returns the response to the AI Router GEM.

Step 12: The AI Router GEM returns the answer and the relevant citations to the user in the AI Assistant for Workfront UI.

AI Assistant in Workfront and Azure OpenAI

AI Assistant in Workfront currently leverages Azure OpenAI to answer customer questions. The following data may be passed to Azure OpenAI to facilitate answering product knowledge or operational insights questions:

- Adobe Experience League documentation
- Information related to the page the user is currently exploring
- User's conversation history (prompt and answer)

The following data may be passed to Azure OpenAI to facilitate answering summarization and Workfront Actions questions only.

- Example questions with ground truth queries (customer-agnostic)
- Attributes within application business objects (Projects, Tasks, Issues, Documents, Records), such as the name, description, customer's custom forms, relevant comments, etc. This information is retrieved with the access level of the user, so only data the user has access to can be retrieved and sent with the requests to Azure OpenAI.

Adobe has disabled logging in Azure OpenAI, helping ensure that Microsoft does not collect or review any data or documents sent for processing to Azure OpenAI by AI Assistant in Workfront. More information is available at [Azure OpenAI data privacy and security](#).

Adobe does not use any customer data or uploaded documents to train or fine-tune the Azure OpenAI Service.

Content Filtering

Adobe leverages Azure OpenAI's content filtering service to moderate the input (prompt) and output (answer) before it is returned to AI Assistant in Workfront. The service uses Microsoft's collection of proprietary models for content filtering that has both contextual and semantic understanding of text. Adobe has configured the content filter to filter "medium" and "high" severity for inputs and outputs to and from Azure OpenAI.

Testing

Adobe teams conduct testing to reduce the potential for biased and harmful outcomes in our generative AI products. For more information on the development and testing processes for our generative AI solutions, please see the [Generative AI Built for Business solution brief](#).

Data Retention

Chat History

Users can access their AI Assistant in Workfront chat history, including the prompt text and answer, for 30 days.

Adobe Workfront retains user chat history, system logs, and user feedback for up to 12 months to enable multi-turn query review tasks and for error analysis. Chat history is stored in the same region and cloud in which the customer's Workfront instance is provisioned.

If a customer would like to delete a user's chat history, the customer should contact their Adobe customer support representative to initiate the deletion process.

Data Usage

Adobe uses customer-aggregated data to enhance product features and configure LLM interactions to produce the best results for new and existing features. Customer data is not used to train or fine-tune any models for Workfront.

Document Usage

Adobe allows customers to upload documents to be analyzed by the AI Assistant. These documents are temporarily stored during the AI session. The documents are not inspected or accessible to any other customer or Azure OpenAI while they are stored. Upon completion of the AI session, the documents are deleted.

Data Processing and Storage Locations

Adobe Identity Management Services (IMS)

Regardless of the geographic location of the customer, all identity data is stored in multi-region, load-balanced, cloud infrastructure providers with data centers located in North America, Europe, and APAC. Identity data is replicated across all data centers for reliability reasons. All identity data is secured at-rest using AES-256-bit encryption in compliance with the Adobe Common Controls Framework (CCF) and meets our internal policies for encryption and storage of sensitive data.

AI Assistant in Workfront and Azure Open AI Service

All server-side components of AI Assistant in Workfront and corresponding data storage are co-located in the same region as the customer's Workfront service infrastructure, which is determined upon initial provisioning.

Data sent to the Azure OpenAI Service may be sent to a different Azure OpenAI region for processing, but no data will be persisted there.

Workfront Service Infrastructure and Data Storage	Azure OpenAI Data Center (Processing)
U.S. (Oregon, California, Iowa and Virginia)	U.S. East, U.S. West, Canada East
Europe (Ireland and Germany)	U.S. East, U.S. West, Canada East

Questions?

If you have any additional questions about the security posture and capabilities of Workfront, please contact your Adobe account manager. For all other questions about Adobe's security programs and processes and compliance certifications, please see the [Adobe Trust Center](#).