



SECURITY FACT SHEET

AI Assistant in Adobe Experience Platform

March 2025



About AI Assistant in Adobe Experience Platform

AI Assistant in Adobe Experience Platform is a generative AI tool integrated within native applications built on Adobe Experience Platform (AEP). Designed to enhance productivity and help users expand product mastery, efficiently navigate enterprise data objects, and simplify tasks while ensuring adherence to the customer organization's data security standards, AI Assistant in AEP can [answer questions](#) about:

- **Product Knowledge** – Learn about product functionality and troubleshoot unfamiliar concepts using Adobe Experience League documentation.
- **Operational Insights** – Discover usage, lineage, and hygiene insights about key business objects with which users interact, including datasets, segments, and destinations, based on operational data¹ in the organization's unique sandbox.

Adobe's agnostic approach to large language models (LLMs) enables us to choose the best-in-class technology for the task at hand. AI Assistant in AEP currently leverages Microsoft's Azure OpenAI Service to answer both product knowledge and operational insights questions.

AI Assistant in AEP Components

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

- **AEP User Interface** – Users interact with the AI Assistant by clicking on the icon in the upper right-hand corner of the Adobe Experience Platform 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 AEP, GEMs include foundation and custom models that power AI Assistant use cases. For more details on the specific models, please see:
 - **Dialog Management GEM** – Orchestrates other models to provide engaging and natural responses. Dialog Management is the core AI Assistant conversational interface.
 - **Product Knowledge GEM** – Identifies the appropriate Adobe 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.
 - **Operational Insights GEM** – Translates a given question into a query against the underlying customer-specific operational data, generates the appropriate answers based on the results returned by the query, and provides appropriate explanations for the query and the returned answer.

¹ Operational data is metadata about the application business objects that users create or configure via the AEP user interface or API within a product sandbox. It is descriptive data about the business object and not the underlying data in the business object itself. As an example, in the case of an audience, it is the name of the audience, the definition of the audience, and other associated metadata; it does not contain all the profiles within that audience. AI Assistant in AEP can currently query operational data about the following business objects: Attributes, Audiences, Dataflows, Datasets, Destinations, Schemas, Sources, and Journeys.

- **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 AEP includes two (2) data services:
 - **Product Knowledge Data Service:** Queries a customer-agnostic product knowledge data store that contains general Adobe knowledge, such as product documentation.
 - **Operational Insights Data Service:** Queries a customer-specific operational insights data store that contains centralized operational data, partitioned by the customer's AEP sandboxes and hydrated every 24 hours. This data is co-located with the customer's AEP instance.

Accessing AI Assistant in AEP

To enable a user to [access AI Assistant in AEP](#), the customer's Adobe Admin must grant specific permissions in both the [Adobe Admin Console](#) and the [Permissions UI](#).

- **For Real-time Customer Data Platform (Real-Time CDP) and Adobe Journey Optimizer users** – User must be granted "Enable AI Assistant" permission to ask product knowledge questions. To enable the user to ask operational insight questions, the Adobe Admin must additionally grant them "View operational insights" permission. These permissions are set by the Admin in the Permissions UI. [Learn more](#).
- **For Customer Journey Analytics users** – The Adobe Admin must grant permission for the user in the CJA Admin Console to access the AI Assistant in the context of Customer Journey Analytics, which enables the user to ask product knowledge questions. *Note: Operational Insights questions are not available for Customer Journey Analytics; therefore, no additional permissions apply.*

AI Assistant in AEP requests are authenticated using Adobe Identity Management Services (IMS) and authorizations are enforced by the [AEP Access Control Service](#). For more information about Adobe IMS, please see the [Adobe Identity Management Services Security Overview](#).

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 is encrypted at rest using AES 256-bit encryption.

AI Assistant in Adobe Experience Platform Security Architecture and Data Flow – Product Knowledge

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




Figure 1: AI Assistant in AEP data flow diagram for product knowledge questions

Data Flow Narrative – Product Knowledge

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

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

Step 2: AI Assistant authenticates the user with [Adobe Identity Management Services \(IMS\)](#) and checks that the user is entitled to use AI Assistant.

Step 3: User enters a product knowledge-type question in the prompt text box.

Step 4: AI Assistant UI sends the prompt text to the Dialog Management GEM, which classifies the prompt into the appropriate question type: product knowledge, operational insight, or out of scope. If the question is in scope for AI Assistant in AEP, the process moves to Step 5. If the question is out of scope, the user receives an error message.¹

Step 5: The Dialog Management GEM checks with the AEP Access Control Service to confirm that the user is entitled to ask product knowledge questions.

¹ Questions outside the scope of Adobe Experience Platform and its native applications, including questions about other Adobe products, such as Adobe Target and the Creative Cloud suite, cannot be answered by AI Assistant in AEP.

Step 6: If the user is entitled, the Dialog Management GEM applies a series of content filters to determine if the prompt adheres to [Adobe's Generative AI User Guidelines](#). If any part of the prompt violates these guidelines, the user receives an error message.

Step 7: The Dialog Management GEM then sends the prompt text to the Product Knowledge GEM, which uses semantic search to retrieve relevant snippets of documentation from the Product Knowledge Data Service to answer the question.

Step 8: The Dialog Management GEM combines the prompt text with the retrieved snippets of documentation from the Product Knowledge Data Service and sends them to the Azure OpenAI Service.

Step 9: Before sending the formulated answer back to the Dialog Management GEM, the Azure OpenAI content filtering service moderates generated responses that violate Azure Open AI user guidelines.

Step 10: The Product Knowledge GEM cross-checks the answer provided by the Azure OpenAI Service against the documentation snippets, adds the appropriate citations, and sends the complete answer and citations to the Dialog Management GEM.

Step 11: The Dialog Management GEM returns the answer and the relevant citations, along with suggested next prompts, to the user in the AI Assistant for AEP user interface.

AI Assistant in Adobe Experience Platform Security Architecture and Data Flow – Operational Insights

The following example data flow illustrates how data flows in AI Assistant in AEP for operational insights questions:




Figure 2: AI Assistant in AEP data flow diagram for operational insights

Data Flow Narrative – Operational Insights

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

Step 2: AI Assistant authenticates the user with [Adobe Identity Management Services \(IMS\)](#) and checks that the user is entitled to use AI Assistant in AEP.

Step 3: User enters an operational insights-type question in the prompt text box.

Step 4: AI Assistant sends the prompt text to the Dialog Management GEM, which classifies the prompt into the appropriate question type: product knowledge, operational insight, or out of scope. If the question is in scope for AI Assistant in AEP, the process moves to Step 5. If the question is out of scope, the user receives an error message.¹

Step 5: The Dialog Management GEM checks with the AEP Access Control Service to confirm that the user is entitled to ask operational insights questions.

Step 6: If the user is entitled, the Dialog Management GEM applies a series of content filters to determine if the prompt adheres to [Adobe's Generative AI User Guidelines](#). If any part of the prompt violates these guidelines, the user receives an error message.

Step 7: The Dialog Management GEM sends the prompt text to the Operational Insights GEM, which retrieves a customer-agnostic schema and sample queries relevant to the current prompt.²

Step 8: The Dialog Management GEM combines the prompt text with the customer-agnostic schema and sample queries and sends the data to the Azure OpenAI Service, which uses the information to formulate an answer.

Step 9: Before sending the formulated answer back to the Operational Insights GEM, the Azure OpenAI content filtering service moderates generated responses that violate Azure Open AI user guidelines.

Step 10: The Operational Insights GEM applies the relevant permissions on the business objects present in the query using role-based access control and object/attribute-level access control.

Step 11: The Operational Insights GEM runs the query in the context of the customer's Operational Insights Data Service and generates an intermediate response,³ which is typically a single- or multiple-row table.

Step 12: The Operational Insights GEM sends the query and the intermediate response to the Azure OpenAI Service, which generates the natural language description of the answer and provides the natural language explanation of the query. This step-by-step explanation helps the user to verify the query's accuracy.

Step 13: The Dialog Management GEM returns the answer to the user.

¹ Questions outside the scope of Adobe Experience Platform and its native applications, including questions about other Adobe products, such as Adobe Target and the Creative Cloud suite, cannot be answered by AI Assistant in AEP.

² The customer-agnostic schema represents the structure in the Operational Insights Data Service, including the tables it contains (e.g. audiences, schemas), details of each table (e.g., type of audience, audience definition, audience size), and the links between the tables.

³ Users can only query their own operational data because the operational insights datastore is partitioned by sandbox and the sandbox partition to query is determined based on the Adobe IMS and sandbox information contained in the API header.

AI Assistant in AEP and Azure OpenAI

AI Assistant in AEP 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:

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

The following data may be passed to Azure OpenAI to facilitate answering operational insights questions only:

- Schema of the tables being queried (customer-agnostic)
- Example questions with ground truth queries (customer-agnostic)
- Attributes within application business objects, such as the name, description, and counts (query results)

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

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

Content Filtering

As noted in the *Data Flow Narratives*, Adobe uses internally developed content filters to determine if the input (prompt) in AI Assistant in AEP adheres to [Adobe's Generative AI User Guidelines](#) before sending the prompt text to the Azure OpenAI Service.

Adobe leverages Azure OpenAI's content filtering service to moderate the output (answer) before it is returned to AI Assistant in AEP. 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 outputs from Azure OpenAI. Azure OpenAI's content filter does not filter any input to Azure OpenAI from AI Assistant in AEP.

In addition, AI Assistant in AEP uses Adobe's internally developed content filters to filter out any generated response that violates [Adobe's Generative AI User Guidelines](#) (e.g., hate speech and profanity) that was not moderated by Azure OpenAI's content filter.

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 AEP chat history, including the prompt text and answer, for 30 days. Chat history is stored in the same data center as the customer's Adobe data storage location (see the *Data Processing and Storage Locations* below).

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

Data Usage

Adobe uses customer-agnostic annotated data to fine-tune Adobe internal models (e.g., the linguistic models for documentation and the operational insights and classifier models for prompt classification or out-of-scope detection). The responses from these models are not shown directly to the users.

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 AEP and Azure Open AI Service

All server-side components of AI Assistant in AEP and corresponding data storage are co-located in the same region as the customer's AEP service infrastructure, which is determined upon initial provisioning. Data sent to the Azure OpenAI Service may be processed in a different data center but located in the same geographical region, per the table below.

AEP Service Infrastructure and Data Storage	Azure OpenAI Data Center (Processing)
U.S. East (Virginia)	Canada East (Quebec)
Canada Central (Toronto)	Canada East (Quebec)
Australia (Sydney)	Australia East (Girilambone)
Netherlands (Amsterdam)	UK South (London)

Questions?

If you have any additional questions about the security posture and capabilities of Adobe Experience Platform, native applications, or AI Assistant in Adobe Experience Platform, 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](#).