Adobe

SECURITY OVERVIEW

Adobe Experience Platform Agent Orchestrator and Agents

October 2025

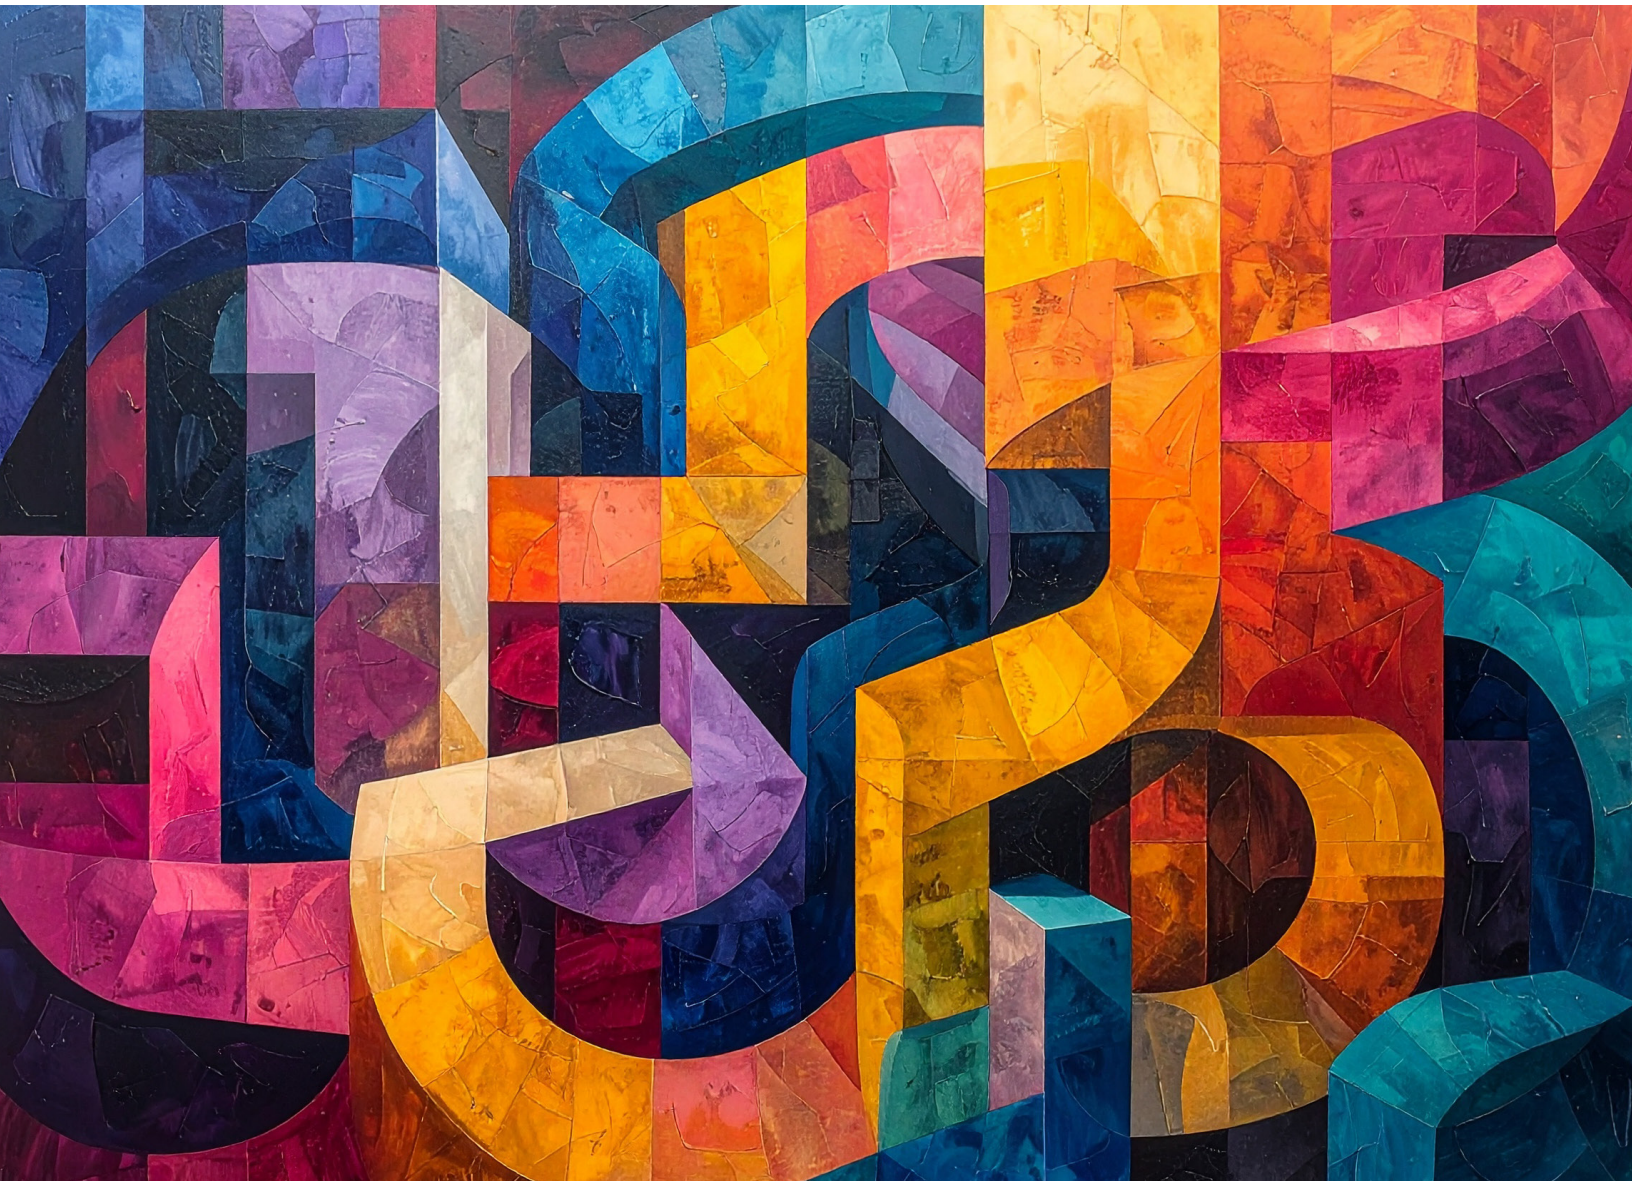

About Adobe Experience Platform Agent Orchestrator

Adobe Experience Platform Agent Orchestrator powers purpose-built Adobe Experience Platform Agents to work alongside marketing and customer experience teams, expanding their capacity to enable seamless customer experience orchestration and deliver personalization at scale. Agent Orchestrator and Experience Platform Agents perform complex decision-making and problem-solving tasks, understand context, retrieve precise information, plan multi-step actions, drive intelligent automation, refine responses, and automate key workflows.

Adobe Experience Platform Agents are a virtual grouping of functional agents that have the skills required to accomplish jobs within a certain customer experience domain. For a current list of Adobe Experience Platform Agents, please see Experience League.

Experience Platform Agent Orchestrator currently leverages either Microsoft Azure OpenAI or Meta Llama Large Language Model (LLM) service, depending on the customer's provisioned cloud hosting provider (see the *Data Processing and Storage Locations* section below).

Adobe Experience Platform Agent Orchestrator Components

Experience Platform Agent Orchestrator includes the following components:

- **Conversational interface** (AI Assistant) – Enables users to leverage generative AI and agentic AI capabilities through an intelligent natural language conversational experience.[1]

- **Reasoning engine** – Understands user intent, maintains context, breaks down complex requests into step-by-step plans, invokes the relevant agents, and brings together results from multiple agents into one clear, cohesive response.

- **Knowledge base** – Powers agents with secure access to customers' business intelligence, which is composed of structured and unstructured data sources, including Adobe product documentation, customer metadata about business objects, and analytics data.

- **Functional agents** – Specialize in one or more skills required to execute the steps involved in completing a job. Depending on the complexity of the user's prompt, one or more functional agents might be required to complete the job. Adobe groups functional agents into two categories:

  - **General purpose functional agents**: Designed to perform broad, reusable skills across multiple domains. Experience Platform Agent Orchestrator includes three (3) general purpose built-in functional agents:

    - **Product knowledge agent** – Identifies the appropriate set of documents (e.g., Adobe product documentation) to answer a given question, retrieves the relevant documents, and generates the appropriate answer based on the retrieved documents. The product knowledge agent also determines accurate source citations and verifies that responses are grounded in truth.

1 Note: Customer organizations must agree to additional legal terms to enable the use of AI Assistant. Please contact your Adobe account representative for more information.

- **Operational insights agent** – Translates a given question into a query against the underlying relational data stores (e.g., customer-specific operational data[2]), and generates the appropriate answer based on the results returned by the query. The operational insights agent also provides appropriate explanations for the query as well as the returned answer and respects role-based and attribute-level access control.

- **Field discovery agent** – Answers questions about the fields in the customer's schema by leveraging metadata, field-level embeddings, and vector similarity techniques to generate the appropriate answer based on this information. The field discovery agent also enables customers to perform natural language searches over XDM fields with fuzzy matching, leveraging both lexical and semantic similarity, to discover fields that best match their criteria.

– **Domain-specific functional agents**: Specialized and built for expert tasks within a specific domain. For domain-specific functional agents invoked by each Adobe Experience Platform Agent, please refer to the Adobe Experience Platform Agent-specific security fact sheet on the Adobe Trust Center.

User Authentication and Authorization

All AI Assistant access requests are authenticated using Adobe Identity Management Services (IMS) and authorizations are enforced by the Adobe Experience Platform access control service (or the CJA access control service).

To enable a user to access the AI Assistant conversational interface and use one or more Experience Platform Agents, the customer's Adobe Admin must grant relevant permissions in the Permissions UI:

- **For Real-Time Customer Data Platform (Real-Time CDP) and Adobe Journey Optimizer users** – The Adobe Admin must grant "Enable AI Assistant" permission to enable a user to access AI Assistant. To allow the user to ask operational insight questions, the Adobe Admin must additionally grant them "View operational insights" permission. Both permissions are set by the Admin in the Permissions UI.

- **For Customer Journey Analytics users** – The Adobe Admin must grant the user permission to access AI Assistant in the Adobe Admin Console, which enables the user to ask product knowledge and data insights questions. *Note: Operational insights questions are not available for Customer Journey Analytics; therefore, no additional permissions apply.*

To enable a user to access the AI Assistant conversational interface in Adobe Experience Manager (AEM), the customer's Adobe Admin must grant the user permission to use AI Assistant in the Adobe Admin Console. For more information, please see *Configure AI Assistant in AEM* on Experience League.

For more information about Adobe IMS, please see the Adobe Identity Management Services Security Overview.

2 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. The operational insights agent can currently query operational data about the following business objects: Attributes, Audiences, Dataflows, Datasets, Destinations, Schemas, Sources, and Journeys.

Data Encryption

- **In Transit** – All data is encrypted in transit over HTTPS using TLS 1.2 or greater.
- **At Rest** – Any data stored is encrypted at rest using AES 256-bit encryption.

Common Security Architecture and Data Flow

The following steps are common for all interactions with AI Assistant. The data flow for each Experience Platform Agent can be found in the agent-specific security fact sheets available on the Adobe Trust Center:

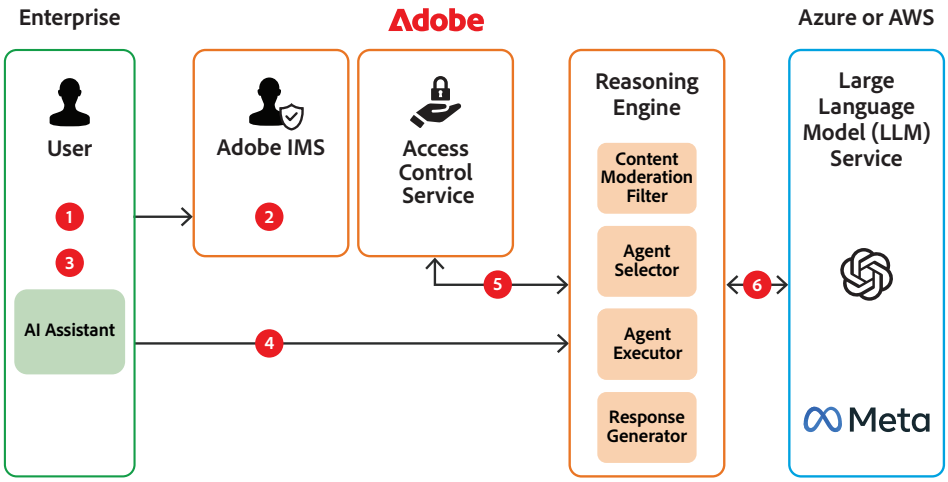

Figure 1: AI Assistant data flow diagram

Step 1: The user opens AI Assistant in the 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 with the access control service relevant to the product in use (e.g., Adobe Experience Platform, Adobe Experience Manager, or Customer Journey Analytics).

Step 3: The user enters a question in the prompt text box.

Step 4: AI Assistant sends the prompt text to the reasoning engine, which determines if the prompt adheres to Adobe's Generative AI User Guidelines. If any part of the prompt violates these guidelines, the user will receive an error message.

Step 5: If the prompt does not violate the user guidelines, the reasoning engine checks with the access control service to confirm that the user is entitled to ask the question type in their prompt. If the user is not authorized, they will receive an error message. If they are authorized, the flow moves to Step 6.

Step 6: The reasoning engine sends the user's prompt to the LLM service, which determines the appropriate agent(s) to invoke based on the prompt. If the question is out of scope, the user receives an out-of-scope message. If it is in scope, the LLM informs the reasoning engine which agent/s it should invoke to answer the user's question.

Note: Each of the following security architecture diagrams and data flow narratives begin with each agent's specific Step 7, but steps 1-6 are applicable to all.

Security Architecture and Data Flow — Product Knowledge Agent

The following example data flow illustrates how data flows when a user asks a product knowledge question in AI Assistant:

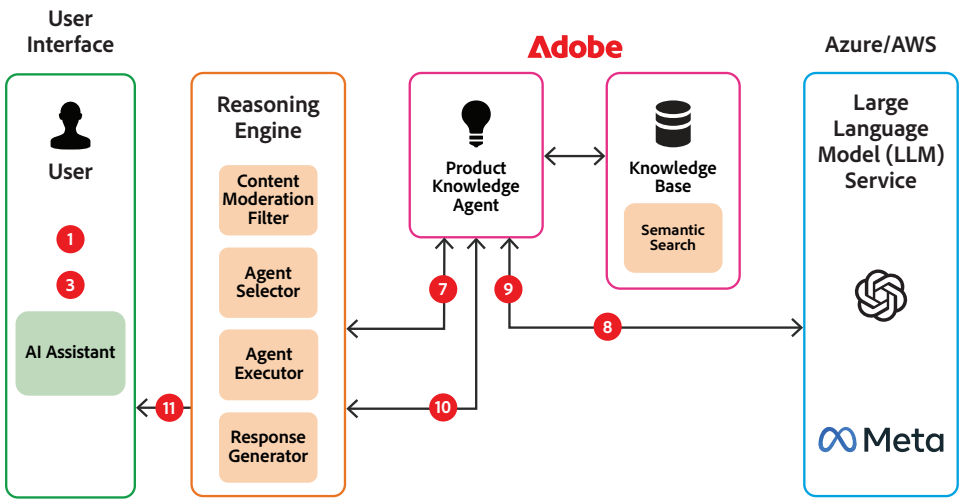

Figure 2: Product Knowledge Agent data flow diagram

Data Flow Narrative — Product Knowledge Agent

Step 7: The reasoning engine sends the prompt text to the product knowledge agent, which uses semantic search to retrieve relevant snippets of documentation from the knowledge base to answer the question.

Step 8: The product knowledge agent combines the prompt text with the retrieved snippets of documentation from the knowledge base and sends them to LLM service, which moderates the input before processing to provide guardrails against content that violates the LLM service's user guidelines. If the input violates the guidelines, the user receives an error message.

Step 9: Before sending the formulated answer back to the product knowledge agent, the content filter within the LLM service moderates the generated response to provide guardrails against content that violates the LLM service user guidelines. If the response violates the guidelines, the user receives an error message.

Step 10: The product knowledge agent cross-checks the answer provided by the LLM against the documentation snippets, adds the appropriate citations, and sends the complete answer and citations to the reasoning engine.

Step 11: The reasoning engine returns the answer and the relevant citations, along with suggested next prompts, to the user.

Security Architecture and Data Flow — Operational Insights Agent

The following example data flow illustrates how data flows when a user asks an operational insights question in AI Assistant:

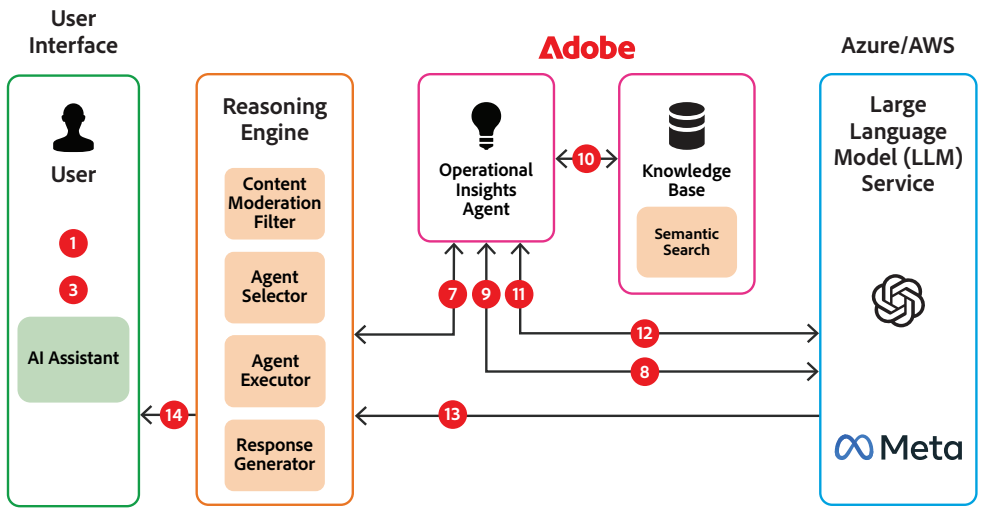

Figure 3: Operational Insights Agent data flow diagram

Data Flow Narrative — Operational Insights Agent

Step 7: The Reasoning Engine sends the prompt to the operational insights agent, which retrieves customer-agnostic schema definitions and example queries relevant to the prompt. The definitions and queries are stored within the operational insights agent.[3]

Step 8: The operational insights agent combines the prompt text with the customer-agnostic schema and sample queries and sends this information to the LLM service, which moderates the input before processing to provide guardrails against content that violates the LLM service's user guidelines. If the input violates the guidelines, the user will receive an error message.

Step 9: Before sending the formulated answer back to the operational insights agent, the content filter within the LLM service moderates the generated response to provide guardrails against content that violates the LLM service's user guidelines. If the output violates the guidelines, the user receives an error message.

Step 10: The operational insights agent applies the relevant permissions on the business objects present in the query using role-based and object/attribute-level access control and runs the query in the context of the customer's knowledge base using semantic search and generates an intermediate response,[4] which is typically a single- or multiple-row table.

3 The customer-agnostic schema represents the structure in the knowledge base, 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.

4 Users can only query their own operational data because the knowledge base 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.

Step 11: The operational insights agent sends the query and the intermediate response to the LLM service, which generates the natural language description of the answer and provides the natural language explanation of the query, helping the user verify the accuracy of the query.

Step 12: The LLM service returns the answer to the operational insights agent.

Step 13: The operational insights agent returns the answer to the reasoning engine.

Step 14: The reasoning engine returns the answer to the user in the AI Assistant conversational interface.

Security Architecture and Data Flow — Field Discovery Agent

The following example data flow illustrates how data flows when a user asks a field discovery-type question in AI Assistant:

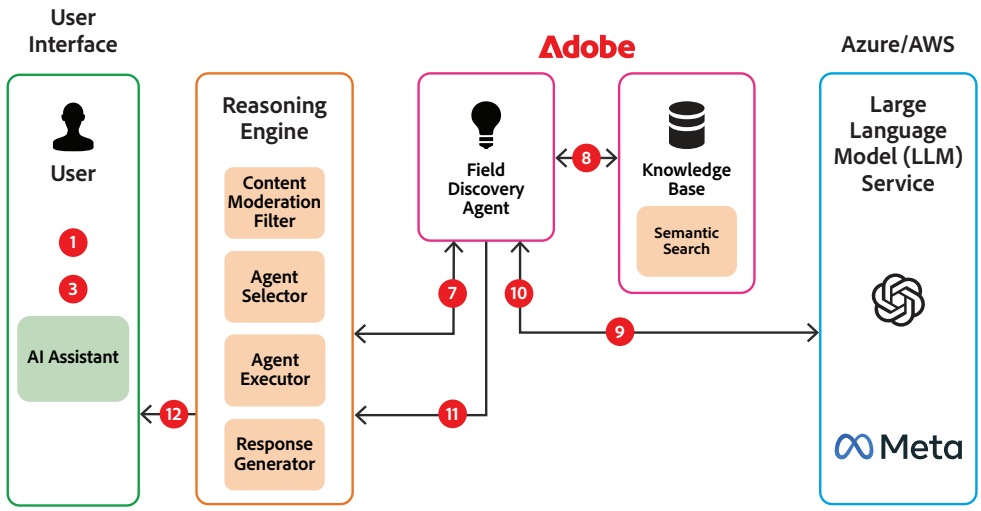

Figure 4: Field Discovery Agent data flow diagram

Data Flow Narrative — Field Discovery Agent

Step 7: The reasoning engine sends the user's prompt text to the field discovery agent. The prompt should contain Experience Data Model (XDM) entities, including name/alias, semantic description, or value.

Step 8: The field discovery agent retrieves the entities relevant to the query from the Knowledge Base, using keyword and semantic search. Role-based, field-level access controls set by the customer's Adobe Admin in the Adobe Admin Console ensure users only see content they are permitted to access.

Step 9: The field discovery agent sends the retrieved entities along with the user's prompt to the LLM service.

Step 10: The LLM service ranks the entities according to relevance to the user's prompt question and returns them to the field discovery agent along with entity details, such as attribute data type and sample data, and an explanation of the ranking relevance so users can understand the results.

Step 11: The field discovery agent returns the results to the reasoning engine.

Step 12: The reasoning engine returns the results to the user in the AI Assistant conversational interface.

Large Language Model Services

Experience Platform Agent Orchestrator leverages either Microsoft Azure OpenAI or Meta Llama, depending on the provisioned cloud hosting provider (see the *Data Processing and Storage Locations* section below).

The following data may be passed to the LLM service to facilitate answering product knowledge questions:

- Adobe Experience League documentation
- Information related to the page in Adobe Experience League documentation that the user is viewing, such as Page Name
- User's conversation history (prompt and answer)

The following data may be passed to the LLM 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)

The following data may be passed to the LLM to facilitate answering field discovery questions only:

- User prompt text
- Field names (customer-specific)
- Field values, type, descriptions, aliases, if existent
- Augmented descriptions and augmented aliases, generated via semantic enrichment

Adobe has disabled logging in connection with third-party providers that host the LLM services, ensuring that data is not collected or reviewed by such third parties.

Content Filtering

As noted in the Data Flow Narratives, Adobe uses internally developed content filters to determine if the customer's input (prompt) adheres to Adobe's Generative AI User Guidelines before sending the prompt text to the LLM.

- **For Azure OpenAI:** Adobe leverages Azure OpenAI's content filtering service to moderate both input (prompts) before processing by Azure OpenAI and output (responses) before returning the response to Experience Platform Agent Orchestrator. 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. Adobe has disabled logging for Azure OpenAI content moderation.

- **For Meta Llama:** Adobe leverages the hate, abuse, and profanity (HAP) detector (also known as a HAP filter) in IBM WatsonX to moderate both input (prompts) before processing by Llama and output (responses) before returning the response to Experience Platform Agent Orchestrator.

In addition, Experience Platform Agent Orchestrator 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 the Azure OpenAI or IBM WatsonX content filters.

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 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* section 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, including linguistic models and various classification models. The responses from these models are not shown directly to the users.

Adobe does not use any customer data to train or fine-tune LLMs.

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.

Adobe Experience Platform Agent Orchestrator and LLM Services

All server-side components of Experience Platform Agent Orchestrator and corresponding data storage are co-located in the same region as the Experience Platform service infrastructure, which is determined upon initial provisioning.

For customers whose Platform instances are hosted on Azure, generative AI requests are processed by Azure OpenAI. For customers whose Experience Platform instances are hosted on Amazon Web Services (AWS), generative AI requests are processed by Llama models hosted within Adobe's AWS environment, unless the customer agreement specifies otherwise.

Data sent to the LLM service is processed in the same geographical region as the data is stored but may be processed in a physically different data center.

Cloud Hosting Provider	Adobe Experience Platform Service Infrastructure and Data Storage	LLM Service	LLM Service Data Center (Processing only)
Azure	North America	Azure OpenAI	North America
Azure	Australia	Azure OpenAI	Australia
Azure	The Netherlands	Azure OpenAI	Great Britain
Azure	Great Britain	Azure OpenAI	Great Britain
Azure	India	Azure OpenAI	Australia
Amazon AWS	North America	Meta Llama	North America

Questions?

If you have any additional questions about the security posture and capabilities of Adobe Experience Platform, native applications, Adobe Experience Platform Agent Orchestrator and Experience Platform Agents, 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.