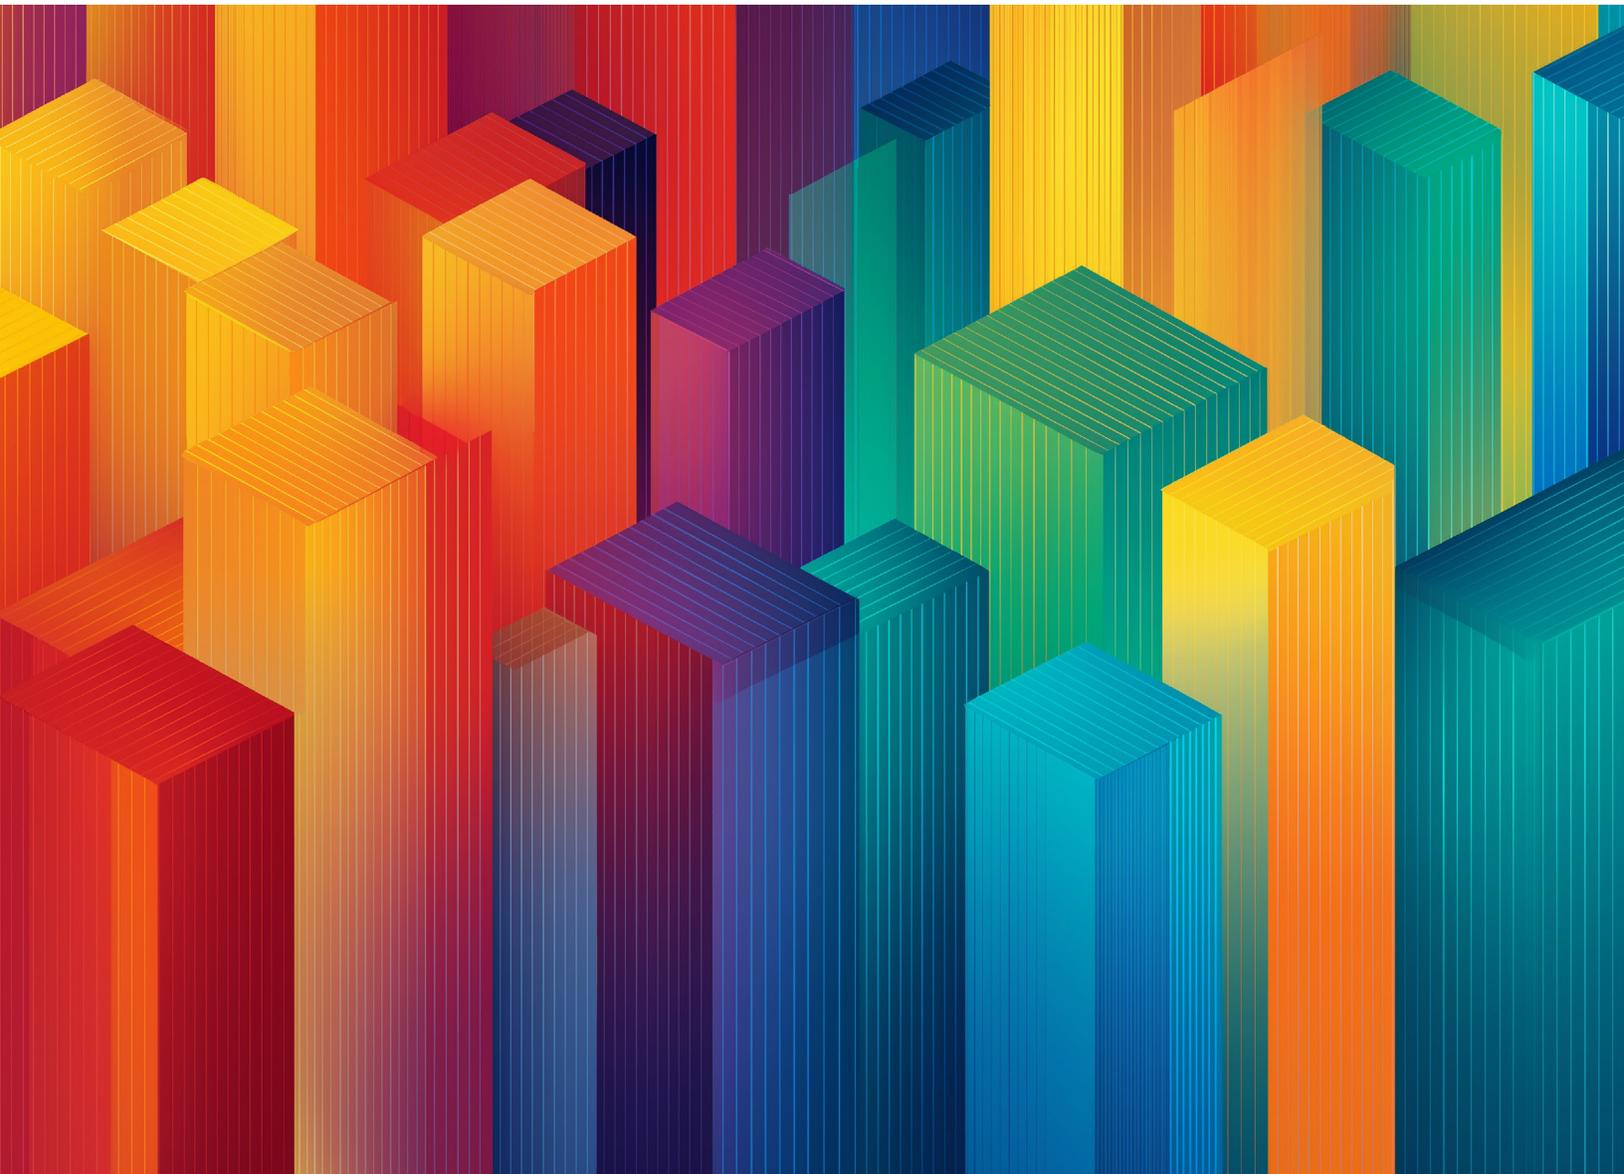




SECURITY OVERVIEW

Adobe Brand Concierge

March 2026



About Adobe Brand Concierge

Adobe Brand Concierge is an AI-powered conversational solution that improves customer experience by delivering personalized, context-rich, and brand-safe interactions. Grounded in first-party data and approved content to understand customer intent, provide accurate responses, and guide users through discovery, comparison, and decision-making, Brand Concierge is a multi-agent solution that enables organizations to deliver on-brand, intelligent interactions across owned digital channels to accelerate conversion.

Powered by [Adobe Experience Platform Agent Orchestrator](#), Brand Concierge supports both consumer and business buyer journeys with real-time intent understanding, dynamic responses, and low-code management for marketing teams to maintain brand integrity and agility.

Brand Concierge components

Brand Concierge consists of the following components:

- **Agent Composer** – Enables brands to configure, test, and modify Brand Concierge instances by uploading approved brand content; defining tuning parameters, such as brand voice, tone, response length, and conversation scope and behavior; and describing the look-and-feel of the end-user conversational experience.
- **Brand Concierge Knowledge Base** – Stores uploaded product catalogs and specified URLs and other brand information that powers the concierge's answers.
- **AEP Agent Orchestrator** – Invokes and coordinates specialized AI agents, dynamically routing conversations to the appropriate agent based on customer goals and intent. Agents used by Brand Concierge include:
 - **Product Advisory Agent** – Interprets customer intent; retrieves product data; delivers targeted recommendations, comparisons, and explanations; and offers behavioral insights to improve recommendations by leveraging a brand-specific product catalog to provide on-brand conversations.
 - **Site Advisory Agent** – Retrieves and summarizes uploaded brand-approved site content to answer customer questions quickly and accurately. Using this data to perform semantic and hybrid search, Site Advisory Agent supports natural language queries, ranks results by relevance and brand priority, and enables multi-turn conversational interactions.
- **Web Client** – Renders the brand's conversational experience in the browser or application using a JavaScript library and is responsible for capturing the customer's conversational input and displaying responses as free-form text or rich user interface components, such as product cards, carousels, and grids.
- **AEP Web and Mobile SDKs** – Handle the low-level communication between the browser and the Brand Concierge Runtime, providing support for gathering page contexts such as URL, browser/app details, and XDM identity map.
- **Brand Concierge Runtime** – Understands intent in real time to guide customers through discovery, comparison, and decision-making with accurate, context-rich interactions based on the configuration defined in Agent Composer. The runtime adapts dynamically while maintaining brand integrity, enabling marketing teams to manage the tone and content of AI-powered conversations without complicated coding.

Accessing Brand Concierge

To enable a marketing professional or administrator to provision, configure, and update Brand Concierge instances, the brand's Adobe Admin must assign the user to the product profile for Adobe Experience Platform and grant the "manage-brand-concierge" role in the Adobe Admin Console.

Data encryption

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

Security architecture and data flow – configuration

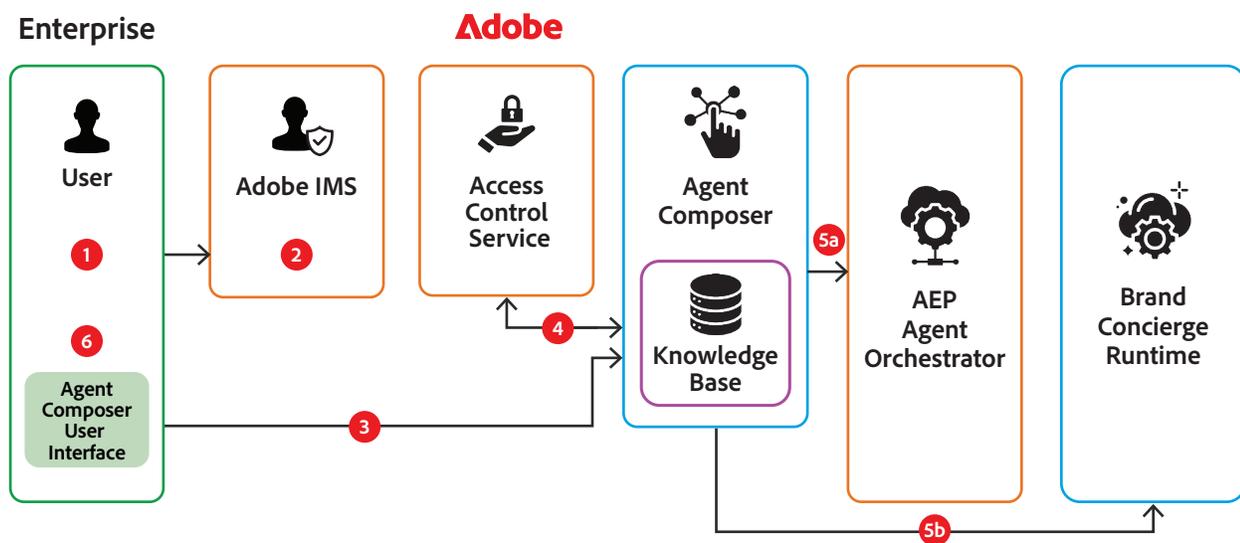


Figure 1: Brand Concierge configuration data flow diagram

Initial configuration

Step 1: The marketer opens the Agent Composer user interface.

Step 2: Agent Composer authenticates the marketer with Adobe Identity Management Services (IMS) and checks that the user is entitled to access Brand Concierge with the access control service.

Step 3: The marketer specifies the URLs for ingestion, uploads the product catalog, adds brand guidelines, selects tuning parameters, and defines the look-and-feel of the customer's conversational experience in Agent Composer. For more information, see the "Data Sources" section below.

Step 4: Agent Composer validates that the marketer is authorized to create the Brand Concierge instance with the access control service, including verifying the required entitlements and role assignments, and creates a Brand Concierge instance with the information provided in Step 3 by populating the Brand Concierge Knowledge Base.

Step 5: When the marketer saves the configuration, Agent Composer sends the required details to AEP Agent Orchestrator, which uses the information to inform agent-level reasoning and orchestration. At the same time, Agent Composer sends the configuration information to the Brand Concierge Runtime, which applies the experience and presentation configuration to inform conversations and render responses in the web or mobile client.

Step 6: After the marketer configures the Brand Concierge instance, they can preview the experience using the Preview functionality in the Agent Composer user interface.¹

Changing a configuration

Step 1: The marketer opens the Agent Composer user interface.

Step 2: Agent Composer authenticates the marketer with Adobe Identity Management Services (IMS) and checks that the user is entitled to access Brand Concierge with the access control service.

Step 3: The marketer makes the required changes to the configuration in Agent Composer.

Step 4: Agent Composer validates that the marketer is authorized to update the Brand Concierge instance with the access control service, including verifying the required entitlements and role assignments, and updates the Brand Concierge instance with the changes entered in Step 3 by updating the Brand Concierge Knowledge Base.

Step 5: When the marketer saves their changes, Agent Composer sends the new configuration details to AEP Agent Orchestrator, which updates the agent-level reasoning and orchestration. At the same time, Agent Composer sends the new details to the Brand Concierge Runtime, which updates the experience and presentation configuration to inform conversations and render responses in the web and mobile client.

Step 6: The marketer can preview the changes in Agent Composer and validate that the changes match their desired outcome.

¹ The brand organization is solely responsible for all configuration and testing of their Brand Concierge instances.

Security architecture and data flow – customer interaction

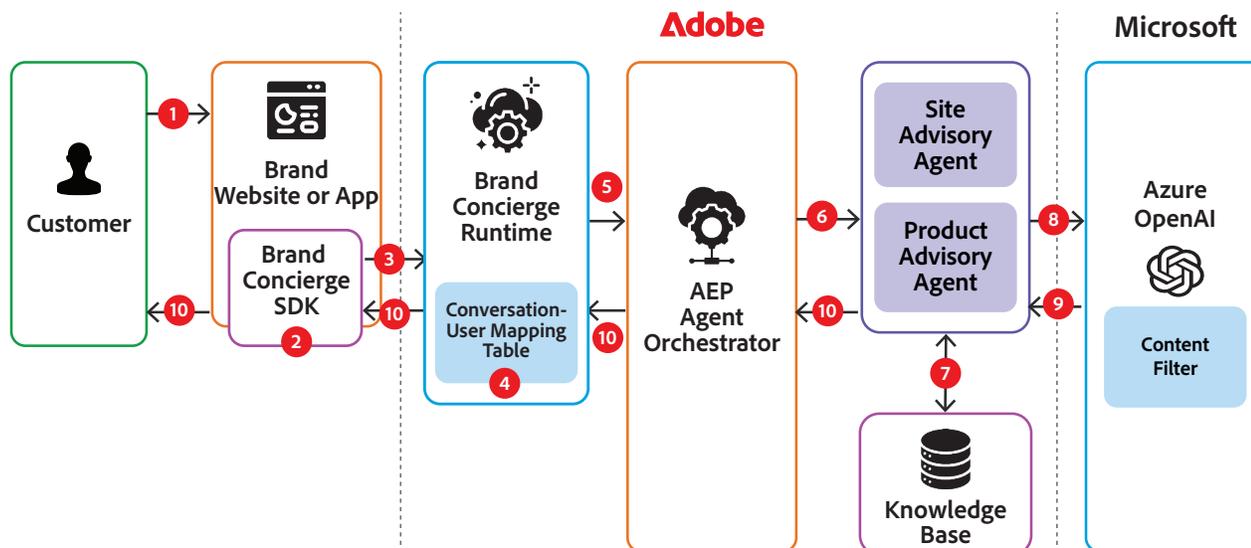


Figure 2: Brand Concierge customer interaction data flow diagram

Step 1: The customer interacts with the brand's website or app and initiates a conversational session by engaging with the conversational bar embedded in the experience.

Step 2: The brand's chatbot, powered by Brand Concierge, sends the customer's question or request to the Brand Concierge SDK.

Step 3: The Brand Concierge SDK forwards the request to the Brand Concierge Runtime.

Step 4: The Brand Concierge Runtime maps the customer's identity to a conversation using a mapping table, enabling customers to remain in-conversation across multiple devices.

Step 5: The Brand Concierge Runtime then forwards the customer's question to AEP Agent Orchestrator, which uses Azure OpenAI to determine the best agent/s to answer the customer's query. For this data flow, see the [AEP Agent Orchestrator and Agents security overview](#).

Step 6: If the question is related to site navigation or general information about the brand, the reasoning engine routes the query to the Site Advisory Agent. If the question is about the brand's products or services, the reasoning engine routes the query to the Product Advisory Agent.

Step 7: In each case, the relevant agent retrieves the relevant contextual information—such as the customer's question—from the Brand Concierge Knowledge Base.

Step 8: The relevant agent sends the contextual information retrieved from the Brand Concierge Knowledge Base along with tuning parameters defined during configuration to Azure OpenAI for response generation.

Step 9: Azure OpenAI uses the data and the prompt sent by the agent to generate a response and returns the response to the relevant agent.

Step 10: The relevant agent returns the response to Agent Orchestrator, which forwards it to the customer via the Brand Concierge Runtime and the Web and Mobile SDK.

Data sources

Brand Concierge can ingest data from the following sources:

- Structured data, such as a product catalog, uploaded in the form of JSON or CSV files by the customer during product configuration. This data is ingested into AEP datasets, which Brand Concierge Knowledge Base uses to hydrate its database.
- Unstructured data, such as a list of the brand's URLs, provided by the customer during product configuration. Agent Composer crawls these URLs, scrapes and cleans content from the URLs, and stores it in the Brand Concierge Knowledge Base.
- Third-party applications, such as product search and catalog browsers, can be integrated via a REST API provided in Agent Composer.

Brand Concierge and Azure OpenAI

Brand Concierge currently leverages Azure OpenAI to process answers to customer questions in the conversational experience. The following data may be passed to Azure OpenAI:

- Structured and unstructured data stored in the knowledge base
- End-user queries, profile attributes, and segment memberships
- Prompts and tuning parameters modified in Agent Composer
- Data retrieved from integrated third-party systems (to guide a conversation)
- Past conversations for the purposes of mining signals and intents

Adobe has disabled logging in Azure OpenAI, helping ensure that Microsoft does not collect or review any data sent for processing to Azure OpenAI. 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

Adobe uses internally developed content filters to determine if the customer input (prompt) adheres to [Adobe's Generative AI User Guidelines](#) before sending the prompt text to the LLM.

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 Brand Concierge. 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.

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

Brand Concierge stores conversational events generated in each turn of the customer's conversation as standard experience events in AEP dataset(s). These events include metadata and interaction signals (such as conversation and turn identifiers, agent context, and engagement outcomes). Prompt and/or response text may also be stored only if the brand provides Adobe with explicit permission to store this information.

Brand Concierge stores conversation history (to keep the user in-context with LLM interactions) in AEP Agent Orchestrator. Brands may access this history for 30 days, after which time it will be automatically deleted.

Conversation history is stored in the AEP Data Lake for analytics purposes only if the user has provided consent (e.g., cookie consent) for their data to be used for analytics. Data stored in the data lake is integrated with the AEP privacy service.

Data usage

Adobe does not use customer data to train Brand Concierge or any of its components.

Adobe collects and stores Brand Concierge user prompts and AI-generated responses without direct user identifiers. This data is siloed and stored on a per-customer basis and is used solely to monitor quality and improve Brand Concierge product performance.

Data Processing and Storage Locations

Brand Concierge

All Brand Concierge components are hosted in an Adobe-managed data center in the US-East region.

Azure OpenAI

Brand Concierge currently routes all Azure OpenAI requests to the Azure data center in the US-East region.

Questions?

If you have any additional questions about the security posture and capabilities of Brand Concierge, 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.