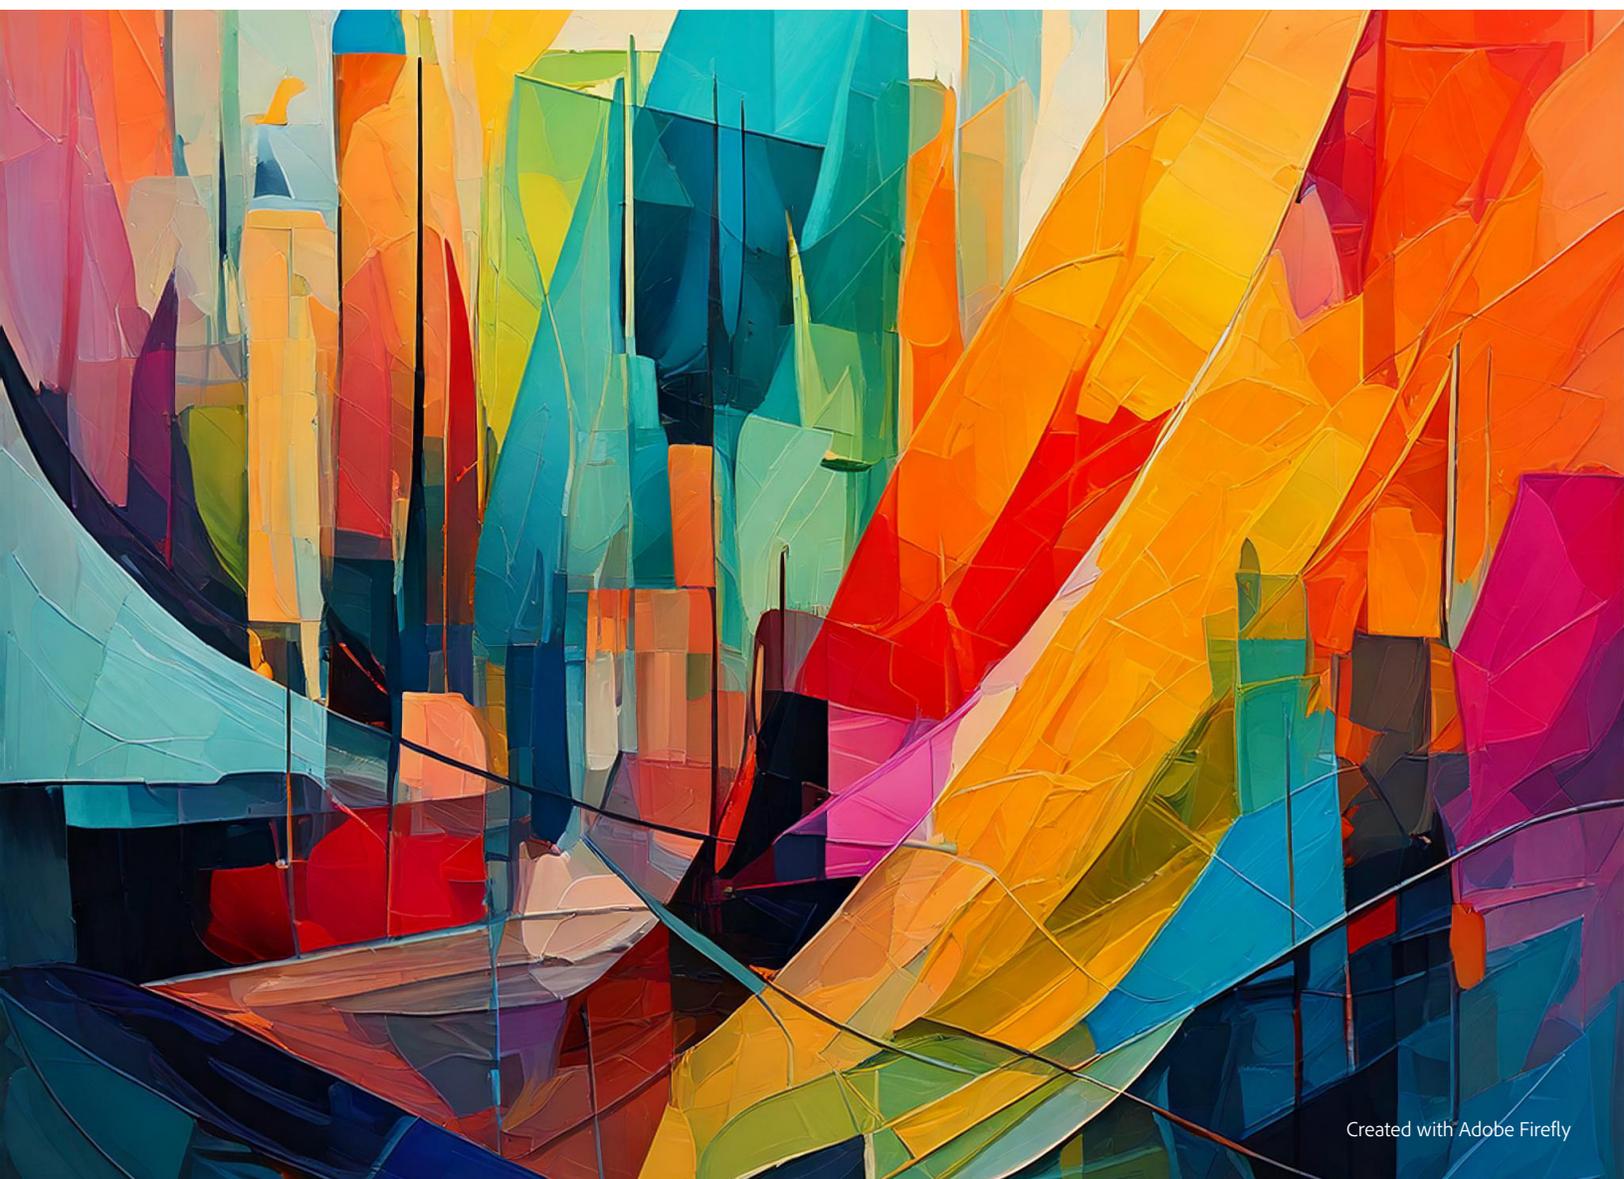




SECURITY FACT SHEET

AI Assistant Content Accelerator in Adobe Journey Optimizer

November 2024



About AI Assistant Content Accelerator in Journey Optimizer

AI Assistant Content Accelerator in Adobe Journey Optimizer enables customers to generate text and image content variations for use in email, SMS, push notifications, and web offers using generative AI (GenAI). With prompt-based text generation powered by the Azure OpenAI Service and image generation powered by Adobe Firefly, customers can optimize the impact of their marketing messages by experimenting with different main titles and images, varying the delivery content or subject, and defining multiple message treatments so they can measure performance with their target audience.

AI Assistant Content Accelerator Components

AI Assistant Content Accelerator includes four (4) components:

- **AI Assistant User interface** — Integrated into the Adobe Journey Optimizer content authoring screen, the AI Assistant Content Accelerator UI enables customers to upload and select existing brand assets, enter text prompts, and include other content features to generate and refine the content for different channels.
- **Web Application Service** — Provides REST APIs to accept requests for document processing and content generation. The Web Application Service authorizes and validates the request before passing it to one of the two dedicated “Workers” described below to complete the task.
- **Document Processing Worker** — Extracts the text chunks and image content from the brand documents uploaded by the customer and stores them as reference content for on-brand content generation.
- **Content Generation Worker** — Generates marketing content based on user-uploaded reference content (if used) and user-defined preferences. The Content Generation Worker finds the matching reference content and sends it to the Azure OpenAI Service to generate text content. If no matching image assets exist in the uploaded brand documents, the Content Generation Worker uses Adobe Firefly to generate image content.

Accessing AI Assistant Content Accelerator

Information about how to grant generative AI permission to users of AI Assistant Content Accelerator can be found on [Adobe Experience League](#).

AI Assistant Content Accelerator 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 Content Accelerator is encrypted at rest using AES 256-bit encryption.

Security Architecture and Data Flow – Document Processing

The following example illustrates the data flow for document processing in AI Assistant Content Accelerator:

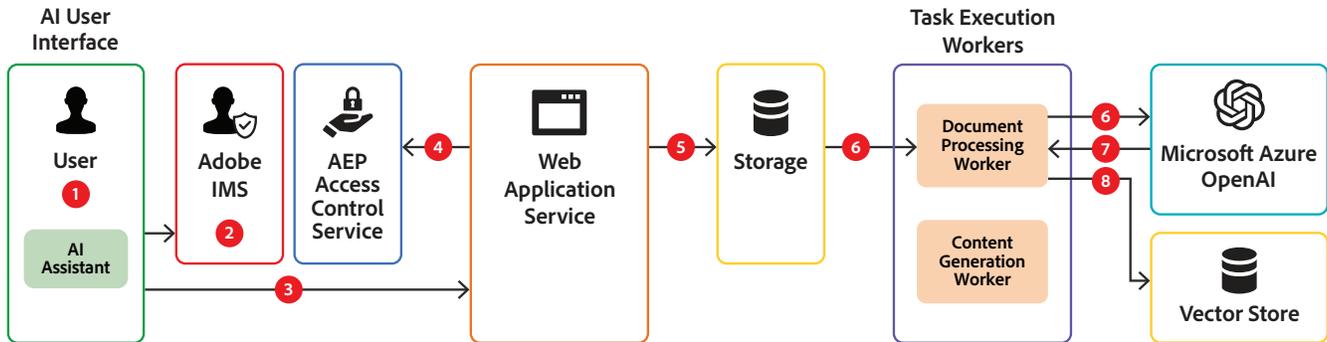


Figure 1: AI Assistant Content Accelerator data flow diagram for document processing

Data Flow Narrative – Document Processing

Step 1: User opens AI Assistant Content Accelerator in the Adobe Journey Optimizer authoring 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 Content Accelerator.

Step 3: User selects “upload content” in the authoring UI and optionally adds the brand assets and documents they want to use as reference content.

Step 4: The Web Application Service checks with the AEP Access Control Service to confirm that the user is entitled to ask upload content.

Step 5: If the user has the appropriate entitlement, the Web Application Service uploads the user-provided reference content in an encrypted database.

Step 6: The Document Processing Worker extracts relevant text and image content from the uploaded reference content in the database and sends this extracted content to the Azure OpenAI Service.

Step 7: The Azure OpenAI Service generates vector embeddings for the extracted content and returns the vector embeddings to the Document Processing Worker.

Step 8: The Document Processing Worker stores the extracted text and image vectors in encrypted storage to inform the generation of new marketing content.

Data Flow Narrative – Content Generation

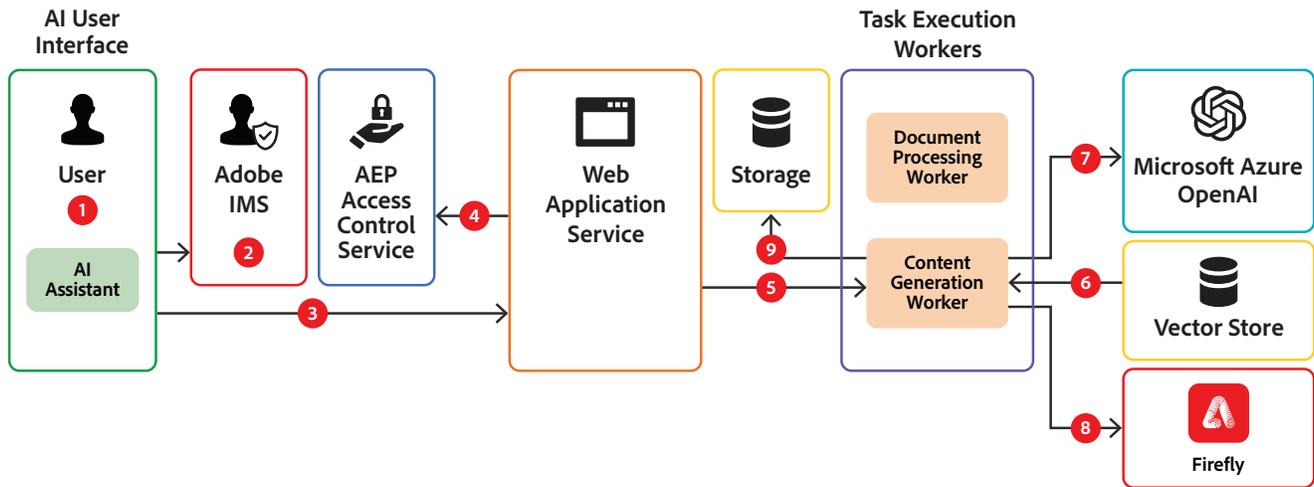


Figure 2: AI Assistant – Content Accelerator data flow diagram for content generation

Step 1: User opens AI Assistant Content Accelerator in the Adobe Journey Optimizer authoring 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 Content Accelerator.

Step 3: User selects “generate content” in the authoring UI and types in a prompt that describes what they want to generate and uploads a content template. Optionally, the user may also upload brand assets and other parameters to guide the content generation.

Step 4: The Web Application Service checks with the AEP Access Control Service to confirm that the user is entitled to generate content.

Step 5: If the user has the appropriate entitlement, the Web Application Service sends the template and the user prompt to the Content Generation Worker.

Step 6: The Content Generation Worker retrieves the previously uploaded relevant text brand assets from the vector store based on the user-provided input and generation parameters.

Step 7: The Content Generation Worker sends the user input/prompt and the retrieved text from the vector store to Azure OpenAI to generate the new text content and provide suggested descriptions for relevant images.

Step 8: Using the suggested image descriptions, the Content Generation Worker searches the uploaded brand assets for images that match the description. If it does not find any images that match the description, the Content Generation Worker uses Adobe Firefly to generate images for the corresponding image placeholders.

Step 9: The Content Generation Worker stores the newly generated brand assets in encrypted storage. The user can then access the content through the Journey Optimizer authoring UI.

AI Assistant Content Accelerator and Azure OpenAI

AI Assistant Content Accelerator leverages Azure OpenAI to generate text-based content. The following data may be passed to Azure OpenAI to facilitate content generation:

- Prompt text and other generation parameters provided by the user
- Text and images from user-uploaded brand documents

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. 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.

Please note: Each user granted permission to use AI Assistant Content Accelerator is subject to the [Adobe Experience Cloud Generative AI User Guidelines](#).

Content Filtering

AI Assistant Content Accelerator leverages Azure OpenAI's content filtering service to moderate the input before it is processed by Azure OpenAI and the output after generation by Azure OpenAI. Azure OpenAI uses Microsoft's collection of proprietary models for content filtering that includes both contextual and semantic understanding of text. Adobe has configured the Azure OpenAI's content filtering service to filter for "medium" and "high" severity output from Azure OpenAI.

AI Assistant Content Accelerator additionally filters any input and output that violates Adobe's Generative AI User Guidelines (e.g., hate speech and profanity) that was not moderated by Azure OpenAI's content filter.

When sending content to Azure OpenAI, AI Assistant Content Accelerator also sends system prompt instructions to help provide guardrails that enforce intended behavior.

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

Brand Assets

Uploaded reference content and brand assets are stored in an encrypted database located in the same region in which the customer's Adobe Journey Optimizer service infrastructure is provisioned. By default, Adobe stores this data for 30 days after which time it is automatically deleted. If the user wishes to delete their content prior to 30 days, they may delete it at any time.

Data Usage

While Adobe does not use customer data to fine-tune or train any model, a human evaluation team observes and uses the input and output data to evaluate and improve the quality of generated content.

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 Content Accelerator and Azure OpenAI

Customer data sent to Adobe Journey Optimizer and AI Assistant Content Accelerator is processed and stored in regions as indicated in the table below. Customer data sent to Azure OpenAI is processed but not stored in Azure data centers listed here.

Journey Optimizer Service Infrastructure and Data Storage	AI Assistant Content Accelerator Service Infrastructure and Data Storage	Azure OpenAI Data Center (Processing)
U.S. East (Virginia)	U.S. East (Virginia)	U.S. East 2 (Virginia), US West California)
Canada Central (Toronto)	Canada Central (Toronto)	U.S. East 2 (Virginia), US West (California)
Australia (Sydney)	Australia (Sydney)	U.S. East 2 (Virginia), US West (California)
Netherlands (Amsterdam)	Netherlands (Amsterdam)	Sweden Central (Gävle and Sandviken)
Central India	Central India	Sweden Central (Gävle and Sandviken)
UK South (London)	UK South (London)	Sweden Central (Gävle and Sandviken)

Questions?

If you have any additional questions about the security posture and capabilities of AI Assistant Content Accelerator, 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](#).

© 2024 Adobe. All rights reserved.

Adobe, the Adobe logo, Acrobat, the Adobe PDF logo, Adobe Premiere, After Effects, Audition, Behance, Creative Cloud, the Creative Cloud logo, Dreamweaver, Illustrator, InCopy, InDesign, Lightroom, Photoshop and Prelude are either registered trademarks or trademarks of Adobe in the United States and/ or other countries.

The Adobe logo, consisting of the word "Adobe" in a bold, red, sans-serif font.