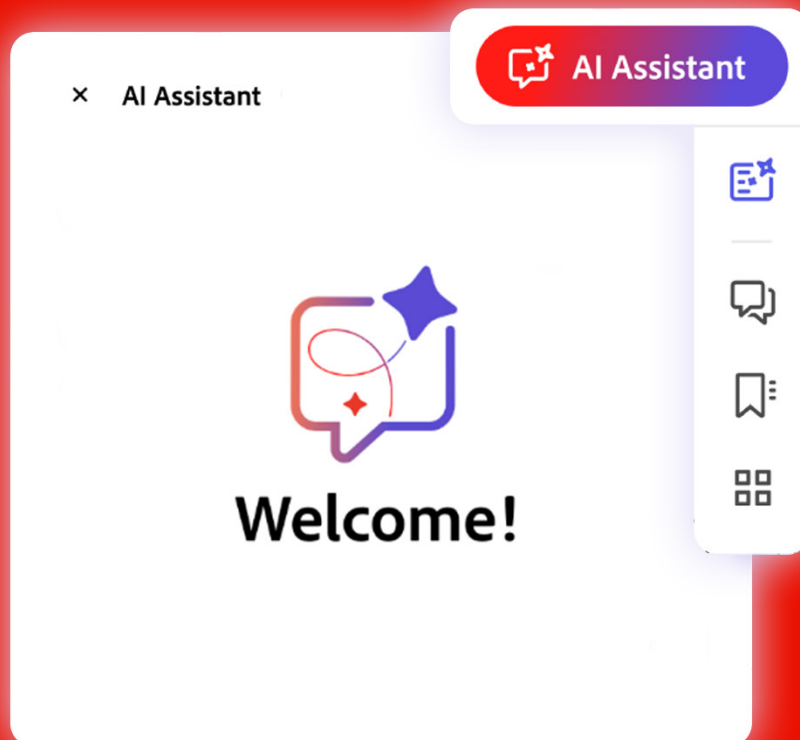




SECURITY FACT SHEET

Adobe Acrobat AI Assistant for enterprise

September 2024



About Adobe Acrobat AI Assistant

The generative AI capabilities in Adobe Acrobat desktop, web, and mobile enable users to ask questions and receive easy-to-understand answers within their PDF documents as well as automatically generate a summary of the document content. Acrobat AI Assistant for enterprise includes:

- **AI Assistant** – Quickly returns answers to questions based on the document content in a comprehensive and informative way. AI Assistant first provides a document overview along with suggested questions and links answers to text and data table sources. It also suggests follow-up questions based on the primary question.
- **Generative Summary** – Automatically get key points to easily find, discern, and navigate to essential information on demand. *Note: This feature is only available on Acrobat desktop and web.*

Acrobat AI Assistant Components

There are **two** primary components in Acrobat AI Assistant:

- **Adobe Acrobat desktop, web, and mobile apps** – On Acrobat desktop and web, users entitled to use Acrobat AI Assistant see the additional icons of AI Assistant and generative summary at the top of their right-hand panel when a document is open and can interact with the AI Assistant. When opening a document on Acrobat mobile, the user will see the AI Assistant icon floating button on the bottom right of their screen.
- **Acrobat Generative AI Service** – The Acrobat desktop, web, and mobile apps send the requested PDF to the Acrobat Generative AI Service, which is hosted in the cloud. The Acrobat Generative AI Service then extracts content from the user's PDF and generates output and attributions to answer questions or provide content summaries. It also caches content for seamless use of AI Assistant within a given period and blocks responses that include hate speech or profanity.

How it Works

The Acrobat Generative AI Service leverages Microsoft's Azure OpenAI Service and enhances it with Adobe proprietary technology to glean relevant information and insights. Adobe shares the PDF content with Azure OpenAI to generate an answer or summary. Adobe has disabled Azure OpenAI's abuse monitoring service so that Microsoft cannot access or use Adobe customer data to train LLMs. More information is available at [Azure OpenAI data privacy and security](#).

For seamless use of AI Assistant, the Acrobat Generative AI Service extracts all the content from the user-selected PDF and caches it with a corresponding session key for 12 hours. The session key uniquely identifies the PDF of a specific user. If the user makes subsequent requests of the Acrobat Generative AI Service within the 12-hour period, the extracted content will still be cached. If the user makes subsequent requests after the 12-hour period, the content is no longer in the cache and the Acrobat Generative AI Service begins the extraction and caching process again.

Note: Other than caching content, the Acrobat Generative AI Service only stores user content when a user reports harm or bias or provides feedback on a response received from the Acrobat Generative AI Service. For more information, please see the *Content Storage and Processing* section below.

Data Encryption

- **In Transit** – All traffic is encrypted in transit over HTTPS using TLS 1.2 or greater.
- **At Rest** – Any data stored by the Acrobat Generative AI Service is encrypted using AES-256. For more information, please see the *Content Storage and Processing* section below.

AI Assistant Security Architecture and Data Flow

The following steps describe the flow and security of data when the user clicks the AI Assistant icon for the first time. The steps in this narrative map to the numbers in Figure 1, below:




Figure 1: Data Flow and Security in Acrobat AI Assistant

Data Flow Narrative

Step 1: After launching Acrobat and signing in, the user opens a PDF file and clicks on the AI Assistant button.

Step 2: The Acrobat app sends an auth token to [Adobe Identity Management Services](#), which authenticates the user and their entitlements. If authenticated, the PDF is then passed to the Acrobat Generative AI Service, which creates a unique session key.

Step 3: The Acrobat Generative AI Service extracts the PDF content and prepares it for responding to user prompts.

Step 4: The Acrobat Generative AI Service caches content for seamless use of GenAI features for 12 hours and sends the content to Azure OpenAI.

Step 5: Azure OpenAI generates a document overview and suggested questions.

Step 6: Before sending the generated content back to the Acrobat Generative AI Service, the Azure OpenAI content filtering service moderates generated responses that violate Adobe's [Generative AI User Guidelines](#).

Step 7: Utilizing a block list, the Acrobat Generative AI Service further blocks any generated responses that violate Adobe's [Generative AI User Guidelines](#) and were not moderated by the Azure OpenAI content filter.

Step 8: The Acrobat Generative AI Service returns the document overview and suggested questions to the user.

Step 9: If the user asks a follow-up question in the prompt window, the Acrobat Generative AI Service prepares relevant content and sends the information to Azure Open AI, which generates the answer and suggested questions, runs the results through the content filter, and returns the results to the Acrobat Generative AI Service. The Acrobat Generative AI Service then attributes the answer in the source document and checks the block list again for any generated content that contains hate or profanity.

Step 10: The Acrobat Generative AI Service returns the answer to the prompt with attributions and suggested questions to the user.

For users of Acrobat desktop and mobile, chat history (both request and response) is stored on the user's device. For users of Acrobat web, chat history is stored in Adobe cloud storage.

Generative Summary Security Architecture and Data Flow

The following steps describe the flow of data when a user clicks the generative summary icon for the first time. The steps in this narrative map to the numbers in Figure 2, below:

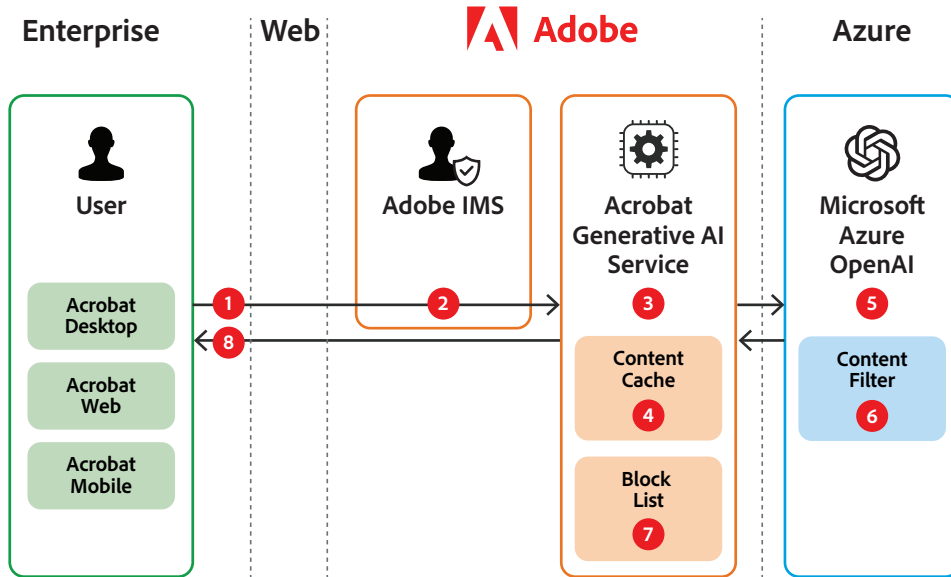


Figure 2: Data Flow and Security in Acrobat AI Assistant Generative Summary

Data Flow Narrative

Step 1: After launching Acrobat and signing in, the user opens a PDF file and clicks on the generative summary icon in the right-hand panel of the app.

Step 2: The Acrobat app sends an auth token to [Adobe Identity Management Services](#), which authenticates the user and their entitlements. If authenticated, the PDF is then passed to the Acrobat Generative AI Service, which creates a unique session key.

Step 3: The Acrobat Generative AI Service extracts the PDF content and prepares it for summarization.

Step 4: The Acrobat Generative AI Service caches content for seamless use of Acrobat's GenAI features for 12 hours and sends the content to Microsoft Azure OpenAI.

Step 5: Azure OpenAI summarizes the content.

Step 6: Before sending the generated content back to the Acrobat Generative AI Service, the Azure OpenAI Service content filter moderates generated responses that violate Adobe's [Generative AI User Guidelines](#).

Step 7: Utilizing a block list, the Acrobat Generative AI Service further blocks any generated responses that violate Adobe's [Generative AI User Guidelines](#) and were not moderated by the Azure OpenAI content filter.

Step 8: The Acrobat Generative AI Service returns the generated summary to the user.

For users of Acrobat desktop and mobile, chat history (both request and response) is stored on the user's device. For users of Acrobat web, chat history is stored in Adobe cloud storage.

Content Storage and Processing

When a user interacts with the Acrobat Generative AI Service, Adobe may do the following:

- **When using Acrobat AI Assistant** – Chat history remains in the user's control for review or query continuation. For users of Acrobat desktop and mobile, chat history is persistently stored on the user's device. For users of Acrobat web, chat history is stored in Adobe cloud storage.

Except for chat history as explained above, Adobe extracts, encrypts, and stores the PDF content, user questions, and generated responses in a cache for 12 hours to enable seamless use of Acrobat AI Assistant. The cache is stored in the same servers as the Acrobat Generative AI Service.

- **When a user flags content as being harmful, illegal, or offensive** – Adobe stores the flagged content only for 30 days unless we have other legal reasons to retain it in accordance with the [Adobe Privacy Policy](#). You can always contact us via the form in the [Adobe Privacy Policy](#) if you want Adobe to delete this data sooner.
- **When a user provides feedback on a response** – Enterprises that have opted into enabling the in-app feedback mechanism will see a thumbs-up or thumbs-down within Acrobat AI Assistant.

Enterprises also have the option to provide explicit consent to Adobe to access their document and its associated data for product improvement. If an enterprise has chosen this option, Adobe stores the related session data (including the PDF file, prompt, and response) for 30 days for product improvement purposes. Feedback data is stored in an encrypted and secured Amazon S3 bucket.

Adobe strictly limits who can access this information to a small number of trained Adobe employees directly involved in the development of the Adobe Generative AI Service. These employees use algorithmic and manual processes to review the session data and address the reported issue within the Acrobat Generative AI Service. Session data is automatically deleted after 30 days but can be deleted sooner upon request using the form in the [Adobe Privacy Policy](#).

Except for the chat history retention mentioned above and any content reported by the user, Adobe automatically deletes the uploaded documents, submitted prompts, and generated responses from Adobe cloud storage after 12 hours.

Content Filtering

Adobe leverages Azure OpenAI's content filtering service to moderate hate, sexual, violent, and self-harm content. 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 the model but not to filter any input.

To further ensure we provide ethically responsible responses, the Acrobat Generative AI Service also leverages a defined block list as an additional mechanism to filter hate and profanity. Responses are only provided to the user if they pass both checks.

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 Processing and Storage Locations

Adobe Identity Management Services

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.

Acrobat Generative AI Service

Adobe hosts the Acrobat Generative AI Service in Amazon Web Services (AWS) data centers in US-East-1; EU-West-1; and AP-Northeast-1.

Azure OpenAI Service

Microsoft hosts the Azure OpenAI Service in Microsoft Azure data centers. For location information, please refer to Azure geographies page.

Questions?

If you have any additional questions about the security posture and capabilities of Acrobat AI Assistant, please reach out to your Adobe account manager or visit the [Adobe Trust Center](#).

