



SECURITY OVERVIEW

Adobe Acrobat Studio for Enterprise

August 2026



About Adobe Acrobat Studio for Enterprise

Adobe Acrobat Studio for Enterprise is an AI-powered workspace built on the industry-standard Acrobat platform that teams already rely on for business-critical PDFs. Because it is built on Acrobat, it inherits enterprise-grade security, administrative controls, and governance, enabling fast adoption at scale without retraining or added risk.

Bringing together PDF management, document intelligence, collaboration, and content creation in one unified environment, Acrobat Studio understands the structure, layout, sources, and intent of complex documents to deliver fast, traceable insights teams can trust. Thanks to a centralized AI-powered workspace, Acrobat Studio reduces tool switching, accelerates decisions, and drives measurable impact.

Acrobat Studio for Enterprise components

Acrobat Studio for Enterprise includes the following components:

- **Adobe Acrobat Pro** – Allows users to create, edit, convert, protect, and collaborate on PDF documents, request and collect e-signatures, and redact sensitive information.
- **AI Assistant in Acrobat** – Lets users ask questions about their files, generate summaries, and extract insights through a conversational interface. Using AI Assistant, users can also perform context-aware PDF actions through natural language, such as organizing pages, adding headers, and preparing documents for printing, as well as generate presentations and other content from information in their documents.
- **PDF Spaces** – Provides AI-powered shared workspaces in which users can organize files, links, notes, and insights in a single location. Using a pre-built or custom AI assistant, users can ask questions and surface insights from the content in a PDF Space. Users can also generate presentations or podcasts from the workspace content and collaborate with others.
- **Adobe Express** – Enables teams to create polished, professional, on-brand multi-format content without design bottlenecks. Adobe Express powers several workflows in Acrobat Studio, including generating presentations, stylizing or translating a PDF, and editing existing or generating new images in a PDF.

All components in Acrobat Studio are powered by the **Acrobat Generative AI Service**, which extracts content from the user's uploaded document/s, maps text and layout structures, generates output and attributions to answer user prompts and provide content summaries, and orchestrates interactions with Microsoft Azure OpenAI. It also caches content for seamless use of AI Assistant within a given period. See the Content Storage and Processing section for details.

Administrative controls in Acrobat Studio for Enterprise

To enable a user to access Acrobat Studio, the customer's Adobe Admin must add the user to the Acrobat Studio product profile in the organization's [Adobe Admin Console](#). For more information on administering Acrobat Studio and role-based access, see this [HelpX](#).

Additionally, administrators may configure access to specific services within Acrobat Studio, including PDF Services, AI Assistant in Acrobat, AI Assistant chat history on cloud, and Adobe Express. Additional settings may be required in the Admin Console depending on the service/s. Detailed information is in this [HelpX](#).

Administrators may also grant access to third-party models in Firefly to specific users and/or roles within the organization. More information is available in this [HelpX](#).

For information about administrative controls to manage cloud storage, see this [HelpX](#).

Acrobat Studio users are authenticated using Adobe Identity Management Services (IMS). For more information, see the [Adobe Identity Management Services security overview](#).

Data encryption

- **In transit** – All traffic is encrypted in transit over HTTPS using TLS 1.2 or greater.
- **At rest** – Data stored by the Acrobat Generative AI Service is encrypted using AES-256.¹

For information about encryption in transit and at rest in Adobe Acrobat, see the [Acrobat and Acrobat Services security overview](#).

For information about encryption in transit and at rest for Adobe Express, see the [Adobe Express for enterprise security overview](#).

¹ Enterprise customers may request that Adobe generate a dedicated encryption key for their organization. More information on this option is available on [HelpX](#).

Acrobat Studio for Enterprise security architecture and data flows

The following sections describe the primary security architecture and data flows for each solution included in Acrobat Studio, other than Adobe Acrobat itself. For information about the security architecture of Acrobat and Acrobat Services, see the [Acrobat and Acrobat Services security overview](#).

AI Assistant in Acrobat

The following steps describe the flow and security of data when the user clicks the AI Assistant icon in Acrobat Studio for the first time. The steps in the narrative map to the numbers in the diagram below:




Figure 1: AI Assistant data flow diagram

Step 1: After signing into Acrobat, the user clicks the AI Assistant icon and selects a PDF (or other supported format) document.

Note: If the user selects more than one document, they are prompted to create a new PDF Space. For the data flow related to PDF Space creation, see the PDF Spaces section below.

Step 2: Acrobat sends an authentication token to Adobe Identity Management Services, which authenticates the user and their entitlements.

Step 3: If authenticated, Acrobat passes the PDF document or other supported document format to the Acrobat Generative AI Service, which creates a unique session key.

Step 4: The Acrobat Generative AI Service extracts the content from the PDF document or other supported document format and caches the content for 12 hours.

Step 5: The Acrobat Generative AI Service sends the content to Azure OpenAI, which generates suggested prompts and returns them to the Acrobat Generative AI Service.

Step 6: The Acrobat Generative AI Service returns the suggested prompts to the user in the AI Assistant conversational interface.

Step 7: The user selects one of the suggested prompts or writes their own in the AI Assistant prompt box.² Based on the prompt, the Acrobat Generative AI Service prepares the relevant content and sends the information to Azure Open AI, which generates the answer and additional suggested prompts and returns the results to the Acrobat Generative AI Service. The Acrobat Generative AI Service then generates attributions for the answer from the source document.

Step 8: The Acrobat Generative AI Service returns the answer to the prompt with attributions and additional suggested prompts to the user in AI Assistant.

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. For information on how to configure Acrobat Studio, please see [HelpX](#).

Use the Generative Summary feature in AI Assistant

The following steps describe the flow and security of data when a user clicks the *Generative Summary* icon. The steps in the narrative map to the numbers in the diagram below:




Figure 2: Generative Summary data flow diagram

Step 1: After signing into Acrobat, the user opens a PDF (or other supported format) document, clicks on the Generative Summary icon, and asks for a document summary.

Step 2: Acrobat sends an authentication token to Adobe Identity Management Services, which authenticates the user and their entitlements.

Step 3: If authenticated, the PDF document or other supported document format is passed to the Acrobat Generative AI Service, which creates a unique session key.

² If the user enters a manipulation-related prompt (for example, "edit this PDF" or "organize these pages"), see *Edit/manipulate PDFs with chat for the corresponding data flow*.

Step 4: The Acrobat Generative AI Service extracts the content from the PDF document or other supported document format, prepares it for summarization, and caches the content for 12 hours.

Step 5: The Acrobat Generative AI Service sends the content to Azure OpenAI, which summarizes the content and returns it to the Acrobat Generative AI Service.

Step 6: The Acrobat Generative AI Service returns the generated summary to the user in AI Assistant.

Edit/manipulate PDFs with chat in AI Assistant

Users can use natural language prompts to complete a task in a PDF without navigating menus or seek help on a specific feature or tool using the AI Assistant in Acrobat. For information on the supported tools and actions users can leverage to edit or otherwise manipulate PDFs in chat, including multi-tool and content-aware actions, see [HelpX](#).

The steps in the narrative map to the numbers in the diagram below:




Figure 3: Edit a PDF in chat data flow diagram

Note: Steps 1 through 4 are the same as in the AI Assistant in Acrobat data flow above.

Step 5: The user submits a prompt in AI Assistant, describing the action they wish to take on a PDF or seeking guidance on how to complete that task.

Step 6: AI Assistant sends the prompt to the Acrobat Generative AI Service, which determines the user's intent.

Step 7: The Acrobat Generative AI Service forwards the user's prompt to the PDF Actions Service, which extracts relevant content from its knowledge base.³

Step 8: The PDF Actions Service sends the extracted knowledge base content, along with the user's query, to Azure OpenAI.

³ The knowledge base in the PDF Actions Service is populated with data ingested from relevant Adobe HelpX pages and is refreshed whenever content changes on those pages.

Step 9: Azure OpenAI uses the information forwarded in Step 8 to determine the user's intent. Based on that, it selects the appropriate tool to execute the requested action or generates contextual guidance to assist the user in completing the task.

Step 10: Azure OpenAI returns the response to the PDF Actions Service. In the case of the user wanting to execute an action on a PDF, Azure OpenAI returns the name of the tool (or tools, in the case of multi-step actions) that can complete the user's request. In the case of providing guidance to the user, Azure OpenAI returns the steps that guide the user on how to complete the task.⁴

Step 11: The PDF Actions Service forwards the response to AI Assistant, which either executes the user's requested action or displays the steps to complete a task, depending on the user's prompt.

PDF Spaces

The following steps describe the flow and security of data when a user creates a new PDF Space, shares or updates an existing PDF Space, or views or interacts with a PDF Space that has been shared with them. The steps in the narrative map to the numbers in the diagram below:




Figure 4: PDF Spaces data flow diagram

Create a PDF Space

Step 1: After signing into Acrobat Studio, the user selects PDF Spaces in the left-hand navigation panel. The user is then prompted to select PDF documents, other supported file types, webpage links, and written text to add to their PDF Space in the dialog box. When the user has finished selecting their files, they click on "Create a PDF Space."

Step 2: Acrobat sends an authentication token to [Adobe Identity Management Services](#), which authenticates the user and their entitlements.

⁴ For example, if the user enters "Rotate page 5" in the prompt box, Acrobat will rotate page 5 and display a message to the user in AI Assistant indicating that the action has been completed. If the user enters "How do I rotate a page?" in the prompt box, AI Assistant will display the steps to the user that describe how to rotate the page. If the user enters a prompt that asks Acrobat to perform an unsupported action in the prompt box, then AI Assistant will display an error message.

Step 3: If the user is entitled to use Acrobat Studio Acrobat creates a PDF Space in the user's provisioned Adobe cloud storage location and uploads the user's selected files to the PDF Space. For information about file restrictions in PDF Spaces, see [HelpX](#).

Step 4: The Acrobat Generative AI Service extracts the content from the documents in the PDF Space and stores the extracted content in encrypted cloud storage within the service.

Step 5: The Acrobat Generative AI Service sends the extracted content to Microsoft Azure OpenAI for processing.

Step 6: Azure OpenAI generates insights and suggested prompts for the uploaded files in the PDF Space and an overview page for the PDF Space and then returns this information to the Acrobat Generative AI Service.

Step 7a: The Acrobat Generative AI Service sends the AI-generated overview page, insights, and suggested prompts to the user in AI Assistant; **Step 7b:** automatically generates two (2) podcasts⁵ and stores them in the user's provisioned cloud storage; and **Step 7c:** stores the generated overview page, insights, and suggested prompts in the user's provisioned cloud storage for future reference and use.

Step 8a: If the user wants to chat with the contents of their PDF Space, they enter a question in AI Assistant.

Step 8b: The Acrobat Generative AI Service retrieves the extracted content from encrypted storage, prepares it for processing and sends it to Azure OpenAI.

Step 9: Azure OpenAI generates the answer, attributes the answer in the source documents, suggests next steps, and returns this information to the Acrobat Generative AI Service.

Step 10a: The Acrobat Generative AI Service returns the answer, attributions, and suggested next steps to the user in AI Assistant and **Step 10b:** stores this information as chat history in the user's provisioned Adobe cloud storage, enabling the user to access their chat history whenever they relaunch the PDF Space.

Note: If a PDF Space is inactive for 30 days, the extracted content and search indexes maintained by the Acrobat Generative AI Service are archived. If the user later returns to an archived PDF Space and initiates a question, the Acrobat Generative AI Service rehydrates the PDF Space by re-extracting the content from the original documents and restoring it to active storage, after which the normal Q&A flow (Steps 8, 9, and 10) resumes.

⁵ The two podcast formats available in PDF Spaces are Highlights, which delivers a focused summary of the key points, and Deep dive, which explores the content in greater detail for deeper understanding. Once created, the podcasts live in the Creations tab and, if the podcast is tied to a PDF Space rather than a single file, it updates automatically as the user adds or removes files from that space. See Generate a podcast in a PDF Space below for details about the data flow for this action, including the generative AI models used for script and voice generation.

Share a PDF Space

When the user shares the PDF Space with others, Acrobat Studio instructs the Acrobat Generative AI Service to (a) share the AI-generated insights, suggested prompts, overview heading, and audio overview (if any) with each recipient in AI Assistant and (b) store this information in the recipient's provisioned Adobe cloud storage for future reference and usage. This process repeats for each user with whom the PDF Space is shared.

For more information on initiating a sharing transaction and sharing options, including access and permission levels, see [HelpX](#).

Note: The owner of a PDF Space can revoke sharing at any time, making the PDF Space immediately inaccessible to others.

Update a PDF Space

If the owner of a PDF Space edits any overview-related, text summary, or other overview artifacts, Acrobat Studio saves the updated artifacts in the owner's provisioned Adobe cloud storage location.

If the owner of the PDF Space adds images, text blocks, or logos in the overview, they are stored in the owner's provisioned Adobe cloud storage location. The owner can also add other artifacts, such as a generated podcast, notes, or other assets, in the PDF Space as separate blocks in the overview pane.

View or interact with a shared PDF Space

Recipients with whom a PDF Space is shared receive a prompt to log into Acrobat (if they are not already logged in) before they can open the PDF Space. After Adobe IMS authenticates the recipient, they can interact with the PDF Space with the privileges of the role granted to them by the owner.⁶

A user with whom a PDF Space has been shared can chat with the documents in the PDF Space in a new private chat, regardless of the role granted to them by the owner. Chat history is stored in the recipient's own provisioned Adobe cloud storage location.

Note: If the owner has created a custom AI Assistant for the PDF Space, this customized assistant is used for all chats initiated by recipients with whom the PDF Space was shared.

If a user with whom a PDF Space is shared with View or Add Comments permission wants to create their own notes or add new source content to the PDF Space, they must copy the PDF Space to their own Adobe cloud storage location.⁷ This action copies the original source documents and notes to the new PDF Space but does not copy the PDF Space owner's chat history, which remains private.

A user with whom a PDF Space is shared can report any suspicious content inside the PDF Space. See the Content Storage and Processing section below for more information.

⁶ Provided the organization's Adobe Admin has turned PDF Services "On" for that user. If the recipient has PDF Services turned "OFF," they can view the PDF Space but cannot interact with its contents, including AI Assistant features.

⁷ Provided the organization's Adobe Admin has turned PDF Services "On" for that user.

Generate a podcast within a PDF Space

The following steps describe the flow of data when a user generates podcast, including a transcript and accompanying audio, from the files in a PDF Space. The steps in the narrative map to the numbers in the diagram below:




Figure 5: Generate podcast data flow diagram

Note: This data flow assumes the user is already signed into Acrobat Studio.

Step 1: In a PDF Space, the user selects the Generate Podcast option.

Step 2: Adobe Identity Management Services (IMS) confirms that the user has the required role-based entitlement to generate a podcast.⁸

Step 3a: The Acrobat Generative AI Service extracts the content from the PDF (or other supported document formats) and **Step 3b:** stores the extracted content in encrypted cloud storage.

Step 4: The Acrobat Generative AI Service sends the content to Azure OpenAI.

Step 5: Azure OpenAI generates a transcript based on the extracted content and returns it to the Acrobat Generative AI Service.

Step 6: The Acrobat Generative AI Service forwards the transcript to the user in the PDF Space in Acrobat Studio.

Step 7: PDF Spaces in Acrobat Studio breaks the transcript into logical chunks and sends them to the Acrobat Generative AI Service.

Step 8: The Acrobat Generative AI Service sends each chunk of transcript to Google Chirp to synthesize the audio for that chunk. *Note: On subsequent requests for the same chunk of audio, the pre-generated audio is served directly from the Acrobat Generative AI Service encrypted storage without calling Google Chirp.*

⁸ Only Owners and Contributors (Editors) of the PDF Space can generate podcasts or other artifacts. Reviewers and Viewers cannot.

Step 9: Google Chirp returns the synthesized audio to the Acrobat Generative AI Service.

Step 10: The Acrobat Generative AI Service stores each audio chunk in Adobe encrypted cloud storage.

Step 11: The Acrobat Generative AI Service sends each chunk of audio to the user in the PDF Space in Acrobat Studio and combines it with the transcript.

Step 12: The transcript and accompanying audio are stored in the user's PDF Space in their provisioned Adobe cloud storage.

Note: If the user manually updates a previously generated transcript, the Acrobat Generative AI Service sends the updated transcript to Microsoft Azure OpenAI for processing and regenerates an updated audio summary. Acrobat Studio then saves the updated transcript and accompanying audio in the owner's provisioned Adobe cloud storage location.

Generate a presentation within a PDF Space

The following steps describe the flow of data when a user wants to generate a presentation from files stored in a PDF Space.⁹ The steps in the narrative map to the numbers in the diagram below:

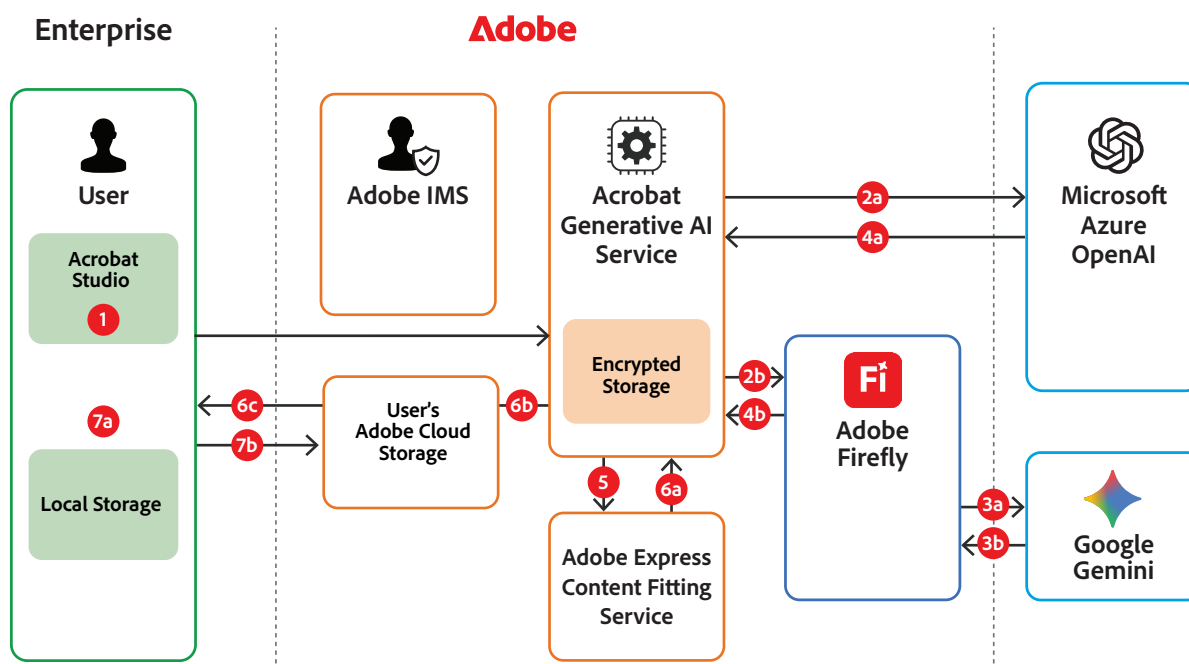


Figure 6: Generate presentation data flow diagram

Note: This data flow assumes the user is already signed into Acrobat Studio.

Step 1: In the chat window within the selected PDF Space, the user prompts Acrobat Studio to generate a presentation from the files in the PDF Space.

⁹ The generative presentations feature is currently limited to 20 files within a PDF Space. If the PDF Space contains more than 20 files, only the first 20 will be used in presentation generation. For more information, see [HelpX](#).

Step 2a: The Acrobat Generative AI Service retrieves the extracted content from encrypted cloud storage and sends it to Azure OpenAI to generate text for the presentation. **Step 2b:** Concurrently, the Acrobat Generative AI Service sends the extracted content to Adobe Firefly to generate images for the presentation.

Step 3a: Adobe Firefly¹⁰ sends the extracted content to Google Gemini to generate images and **Step 3b:** Google Gemini returns the images to Adobe Firefly.

Step 4a: Azure OpenAI returns the generated text to the Acrobat Generative AI Service and **Step 4b:** Adobe Firefly returns the generated images to the Acrobat Generative AI Service.

Step 5: The Acrobat Generative AI Service sends the generated text and images to the Adobe Express Content Fitting Service, which assembles them into a presentation layout in Adobe Express file format.

Step 6a: The Adobe Express Content Fitting Service returns the complete presentation as an Adobe Express file to the Acrobat Generative AI Service, which **Step 6b:** stores it in the user's private chat history in their provisioned Adobe Cloud Storage, and **Step 6c:** forwards it to Acrobat Studio.

Step 7a: The user can download the presentation in PDF, PNG, JPG, or PPTX format to local storage or open it in Adobe Express (by double-clicking the presentation icon) to view or edit. **Step 7b:** The user can also save the presentation to their PDF Space (in their provisioned Adobe Cloud Storage location) to preserve it beyond the private chat history and to share it with others who have permissions to access the PDF Space.

¹⁰ For generative presentations, Adobe Firefly uses Google Gemini, a supported third-party model hosted by Google, for image generation. For more information, see the Partner Models in Adobe Creative Cloud security fact sheet. If Google Gemini is unable to complete a generative image request, Acrobat Studio will instead insert a representative Adobe Stock image.

Other Acrobat Studio data flows

The following data flows begin when the user is on the Acrobat Studio home page.

Stylize or translate a PDF

The following steps describe the flow of data when a user wants to stylize or translate a PDF. The steps in the narrative map to the numbers in the diagram below:

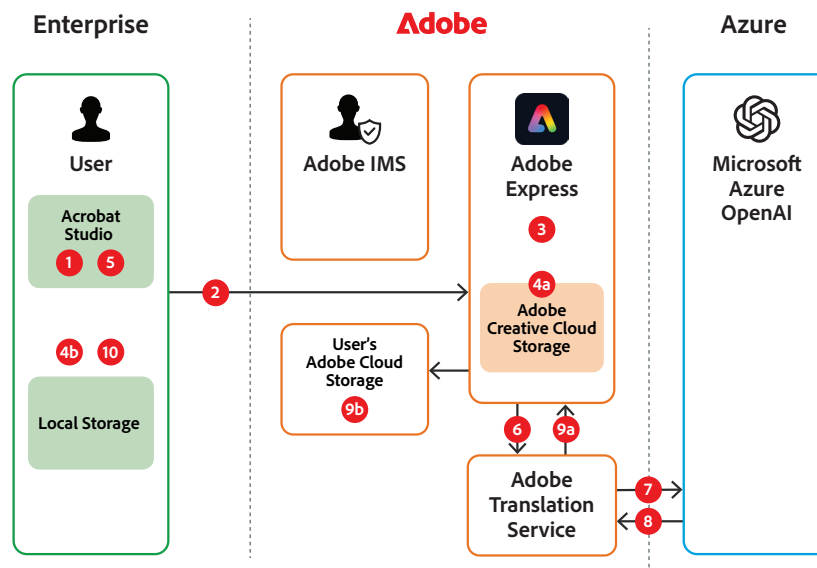


Figure 7: Stylize or translate PDF data flow diagram

Note: This data flow assumes the user is already signed into Acrobat Studio.

Step 1: The user opens a PDF and selects “Stylize this PDF” or “Translate this PDF” in the left-hand navigation pane.

Step 2: Acrobat passes the PDF to the Adobe Express PDF import service, which converts it into an Express-compatible file. The original PDF remains in its original storage location (either on their local system or in the user’s PDF Space in their provisioned Adobe cloud storage location).

Stylize

Step 3: If the user wants to stylize the PDF, the file opens in Adobe Express, and the user selects a style to apply or makes other design edits (for example, colors and fonts) to the file. Each change is automatically saved in the Adobe Express file in Creative Cloud storage.

Step 4: The user can save a PDF version of the Adobe Express document to their local system using the “download as PDF” option.¹¹

¹¹ This action invokes the Adobe Express PDF export service and creates a new PDF that matches the content and styling of the Express-format document.

Translate

Step 3: If the user wants to translate the PDF, the file opens in Adobe Express and the user selects a language into which they want to translate the file (or selected subset of pages) and invokes the Translate feature.

Step 4: Adobe Express extracts the relevant text from the document and sends it to the AI-powered Adobe Translation Service.

Step 5: The Adobe Translation Service constructs a translation prompt and sends both the content and the prompt to Azure OpenAI.

Step 6: Azure OpenAI translates the selected text and returns it to the Adobe Translation Service, which formats the translated content into Adobe Express file format.

Step 7: Adobe Express stitches the translations into the original file, creating a new page for each translated page, along with the original source text. The file is automatically saved to the user's provisioned Adobe cloud storage location.

Step 8: The user can create a PDF version of the translated document and save it to their local system using the "download as PDF" option.¹²

Edit existing or generate new images in PDF

The following steps describe the flow of data when a user wants to edit an existing image in a PDF or generate a new image based on selected content in the PDF. The steps in the narrative map to the numbers in the diagram below:




Figure 8: Edit or generate image data flow diagram

Note: The following two (2) data flows assume the user is already signed into Acrobat Studio.

¹² This action invokes the Adobe Express PDF export service and creates a new PDF that matches the content and styling of the Express-format document.

Edit image

Step 1: The user opens a PDF. If the user wants to edit the local image, they select the image they wish to edit in the PDF and the data flow continues with **Step 2a**. If the user wants to generate a new image based on content they select in the PDF, they highlight the content in the PDF they wish to use to generate a new image and select *Generate Image*. The data flow continues with **Step 2b**.

Step 2a: Acrobat passes the selected image in PNG format to Adobe Express, which saves it as a temporary Express-formatted file.

Step 3a: The user edits the image in Express by selecting one or more of the supported editing options that appear after they choose “Edit image.” When satisfied, they apply the changes to the temporary Express file.

Step 4a: Express returns the edited image in PNG format to Acrobat Studio, which replaces the previously selected image in the PDF.

Step 5a: The user can save the PDF to their provisioned Adobe cloud storage or to their local system.

Generate new image

Step 2b: Acrobat launches Adobe Express using the selected text as the prompt.

Step 3b: Adobe Express sends the prompt to Adobe Firefly, which uses the prompt to generate four (4) image options based on the prompt text¹³.

Step 4b: Adobe Firefly returns the generated image options to the user in Adobe Express and the user selects one of the options for use in their PDF.

Step 5b: Adobe Express returns the selected PNG to Acrobat and the user places the image in their preferred location in the PDF.

Step 6: The user can save the updated PDF in their provisioned Adobe cloud storage or on their local system.

¹³ This feature uses Adobe Firefly first-party models for image generation.

Model training

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

Acrobat Studio and LLM Services

Azure OpenAI

Acrobat Studio leverages Microsoft Azure OpenAI and enhances it with Adobe proprietary technology to glean relevant information and insights from user-provided content. Adobe shares the provided content with Azure OpenAI to generate a response to the user's prompt.

Adobe has disabled the abuse monitoring service in Azure OpenAI, helping ensure that Microsoft does not access or use Adobe customer data to train Azure OpenAI.

Adobe has also disabled logging in Azure OpenAI, helping ensure that Microsoft does not collect, review, or store any data sent for processing to Azure OpenAI. For more information, see [Azure OpenAI data privacy and security](#).

Microsoft hosts the Azure OpenAI Service in Microsoft Azure data centers. For location information, see the [Adobe Trust Center](#). No Adobe customer data is stored by Azure OpenAI.

Google Chirp

For more information about Google Chirp Cloud Text-to-Speech, see [Chirp documentation](#). No Adobe customer data is stored by Google Chirp.

Content storage and processing

For Acrobat AI Assistant (including Generative Summary):

The Acrobat Generative AI Service extracts the content from the user-selected document and caches it with a corresponding session key for 12 hours. The session key uniquely identifies the PDF document of a specific user. If the user makes a subsequent request of the Acrobat Generative AI Service within 12 hours, the extracted content will be retrieved from the cache, improving performance.

If the user makes subsequent requests after the 12-hour period, the extracted content is no longer cached, and the Acrobat Generative AI Service begins the extraction and caching process again. The cache is stored in the same servers as the Acrobat Generative AI Service.

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.

For generated podcasts:

Acrobat Studio stores the generated podcast transcript and associated audio chunks in the user's provisioned Adobe Cloud Storage until the user proactively deletes them. If the user deletes the transcript, the associated audio chunks are also deleted; if the user regenerates the transcript the associated audio chunks are also regenerated.

For generated presentations:

When a user generates a presentation from within their PDF Space, the presentation is stored as an Adobe Express file in the user's PDF Space private chat history for that PDF Space. If or when the user clears the chat history, any presentation the user generated but did not explicitly save to the PDF Space is deleted. To preserve a presentation beyond the chat history, the user must save a copy of the presentation to the PDF Space. The source files remain stored with the PDF Space, and their extraction and storage follow the process described in *For PDF Spaces*.

For PDF Spaces:

Adobe stores all the original PDF files uploaded by user as well as the converted PDF files (from other supported formats) in Adobe Document Cloud storage. For the geographic locations of Adobe Document Cloud storage, see the [Adobe Trust Center](#).

Adobe also extracts, encrypts, and stores the content of PDF files in Acrobat Generative AI Service hosting locations to enable seamless use of Acrobat PDF Spaces.

Chat history remains in the user's control for review or query continuation. For all users, chat history is persistently stored in the user's provisioned Adobe Cloud Storage location until the user clears the chat history.

For Adobe Express (except generated presentations):

See the [Adobe Express for enterprise security overview](#) for content storage and processing information.

For content flagged as harmful, illegal, or offensive:

Adobe stores only the flagged content for 30 days (unless there are other legal reasons to retain it) in accordance with the [Adobe Privacy Policy](#). Enterprises may contact Adobe using the form in the [Adobe Privacy Policy](#) if they want Adobe to delete this data sooner.

Enterprises may also explicitly consent Adobe to access their document and its associated data for product improvement. If an enterprise provides this consent, Adobe stores the related session data (including the file, prompt, and response) for 30 days for product improvement purposes. This session data is stored in encrypted and secured storage.

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 described above and user-reported content, Adobe automatically deletes the uploaded documents, submitted prompts, and generated responses from Adobe Cloud Storage after 12 hours.

Testing

For the annual Security Testing Report for Adobe Acrobat, including AI Assistant and PDF Spaces, see the [Adobe Acrobat Security Testing Report](#) (NDA required). For the annual Security Testing Report for Adobe Express, including Express APIs used by Acrobat to communicate with Express in Acrobat Studio, see the [Adobe Express Security Testing Report](#) (NDA required).

Adobe teams also 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, see the [Generative AI Built for Business solution brief](#).

Data storage locations

User-generated content

Content generated using Acrobat (for example, PDF documents and PDF Spaces) is stored in the user's provisioned Document Cloud storage location in Adobe Cloud Storage unless the user explicitly stores it on their local system. For more information about Document Cloud hosting locations, see the [Adobe Trust Center](#). *Note: Documents within a PDF Space are stored in a separate folder only accessible through that PDF Space.*

Content generated using Adobe Express is stored in the user's provisioned Creative Cloud storage in the same Adobe Cloud Storage location associated with their Acrobat license unless the user explicitly stores it on their local system. For more information about Creative Cloud hosting locations, see the [Adobe Trust Center](#).

All user-generated content stored in Adobe Cloud Storage is encrypted at rest.

Acrobat Generative AI Service

Adobe hosts the Acrobat Generative AI Service in Amazon Web Services (AWS) data centers in North America, EMEA, and Asia-Pacific regions. Any content stored by the Acrobat Generative AI Service is co-located in the organization's provisioned Acrobat data center.

Questions?

If you have any additional questions about the security posture and capabilities of this or any Adobe solution, please contact your Adobe account manager. For detailed information about Adobe's security programs and processes and compliance certifications, see the [Adobe Trust Center](#).