



SECURITY OVERVIEW

Knowledge Base in Acrobat

August 2026



About Knowledge Base in Acrobat

Knowledge Base in Acrobat is an enterprise-grade generative AI solution that enables organizations to expose the knowledge contained in large, curated collections of documents and web content using natural language prompts in a conversational interface. Unlike traditional search-and-retrieve methodologies, where users can spend significant time searching through documents for the information they need, Knowledge Base in Acrobat helps users quickly find answers and receive responses grounded in the source documents with citations back to the source.

Knowledge Base in Acrobat components

Knowledge Base in Acrobat includes the following components:

- **Knowledge Base web application** – Enables users to create and manage curated collections by uploading documents, files, and images in a web application. Users can also add web links to dynamically retrieve pre-indexed content in their collections via the Knowledge Base web app UI.
- **Knowledge Base Indexer** – Ingests web sites/pages that require authentication, such as Wikis and other internal sites, into a Knowledge Base collection through a separate desktop application.¹
- **SharePoint and Google Drive connectors** – Include and continuously sync documents stored in these source repositories with their Knowledge Base collections.
- **Knowledge Base AI Assistant** – Enables users to ask natural language questions of their Knowledge Base collections and receive AI-generated answers, summaries, and insights grounded in the collection's content in a conversational interface.
- **Knowledge Base chatbot for Slack and Microsoft Teams** – Query a Knowledge Base collection from a Slack or Teams channel using an Adobe-developed Knowledge Base-specific chatbot. For more information, see HelpX.
- **Knowledge Base Generative AI Service** – Ingests documents, manages file storage, and divides and indexes files within collections using an advanced RAG platform optimized for large document collections. Leveraging Microsoft Azure OpenAI, the Knowledge Base Generative AI Service also enables users and developers to invoke a range of inference queries against the files in the target collections, including multi-turn Q&A, search, summarization, and entity extraction.
- **Knowledge Base data storage** – Stores documents, files, and other data ingested into the platform in a vector database as well as retains the original documents in file storage for fast retrieval when viewing citations.
- **Knowledge Base APIs** – Provide API access to build custom integrations and AI-powered workflow solutions.

¹ If a linked Wiki or other source changes, the user must refresh the connection to view the most recent content.

Accessing Knowledge Base

An organization's Adobe Administrator enables Knowledge Base in Acrobat using a dedicated toggle inside the Acrobat Studio for Enterprise product profile in the Adobe Admin Console.² For new customers, this toggle is automatically enabled. For more information, see [HelpX](#).

Users licensed for Acrobat Studio can create, manage, and query Knowledge Base collections. Any user with an Adobe ID but not licensed to use Acrobat Studio can query Knowledge Base collections that have been shared with them (either individually or within the same organization) but cannot create new collections or ingest documents into an existing collection.

Users are authenticated by Adobe Identity Management Services (IMS) using their Adobe credentials, granting access to the Knowledge Base web application. For more information about Adobe IMS, see the [Adobe Identity Management Services Security Overview](#).

Integration with PDF Spaces³

Knowledge Base includes an organization-wide setting⁴ that allows customers' Adobe Administrators to manage whether PDF Spaces in Acrobat Studio can access the organization's Knowledge Base collections. This integration is enabled by default.

Disabling the integration does not turn off PDF Spaces itself; it simply removes Knowledge Base content from the scope of what PDF Spaces can access. When the integration is disabled, no user in the organization can link a Knowledge Base collection to a PDF Space or access the contents of a Knowledge Base collection from one.⁵

Customers with regulatory compliance requirements should disable the integration before ingesting regulated data into a Knowledge Base collection.⁶

Data encryption

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

² The organization must have Acrobat Studio for Enterprise and Acrobat Shared Credits to control access and usage.

³ For more information about PDF Spaces, see the [Adobe Acrobat Studio security overview](#).

⁴ This setting can only be changed by users assigned to the product administrator role in the Adobe Admin Console.

⁵ All other Knowledge Base access surfaces, including Knowledge Base AI Assistant, Slack and Teams chatbots, and custom integrations that use the Knowledge Base APIs will continue to work as intended.

⁶ Disabling this integration does not certify or guarantee compliance with any regulatory framework. Customers are advised to consult their compliance attorney to determine compliance. Customers with regulated use cases (for example, Covered Entities under HIPAA) may have additional obligations, such as executing a Business Associate Agreement (BAA) with Adobe before using Knowledge Base for regulated data. Customers interested in HIPAA Ready for Knowledge Base should contact their Adobe account team to begin the BAA process.

Security architecture and data flow — creating a collection




Figure 1: Create a Knowledge Base collection data flow diagram

Data flow narrative

Step 1: The user logs into Acrobat Studio with their Adobe credentials and selects the Knowledge Base icon.

Step 2: Acrobat Studio forwards the request to [Adobe Identity Management Services \(IMS\)](#), which authenticates the user and checks whether the user is entitled to use Knowledge Base in Acrobat Studio.

Step 3: Once authenticated, the user is redirected to the Knowledge Base web application.

Step 4: The user creates a Knowledge Base collection by uploading documents from their desktop or Adobe Cloud Storage, adding web links to public websites, and connecting SharePoint or Google Drive repositories in the Knowledge Base web application.^{7 8}

Step 5: The Knowledge Base Generative AI Service receives the uploaded documents and content from the web app and stores them as-is in Knowledge Base data storage. Concurrently, the Knowledge Base Generative AI Service parses the document content into chunks and stores the chunks in a vector database in Knowledge Base data storage.

Step 6: The Knowledge Base Generative AI Service sends the content chunks to Azure OpenAI to generate semantic embeddings of the document, which enables Knowledge Base to retrieve document content based on the semantic meaning of the input text.

Step 7: Azure OpenAI returns the generated semantic embeddings to the Knowledge Base Generative AI Service.

⁷ When a user syncs documents from SharePoint or Google Drive with Knowledge Base, copies of those documents are stored in Knowledge Base data storage and are continuously synced for quick citation retrieval.

⁸ To ingest content from sites requiring authentication, the user leverages the separate Knowledge Base Web Indexer desktop app.

Step 8: The Knowledge Base Generative AI Service stores the semantic embeddings and other document metadata in Knowledge Base data storage.

Step 9: Once all the documents in the collection are successfully ingested, the Knowledge Base Generative AI Service notifies the user with a message in the Knowledge Base user interface.

Security architecture and data flow — querying a collection




Figure 2: Query a Knowledge Base collection data flow diagram

Data flow narrative — from the Knowledge Base UI

Step 1: The user selects a file or a collection in the Knowledge Base UI, clicks on the AI Assistant button, and submits a question in the chat window.

Step 2: The Knowledge Base Generative AI Service receives the user's question and identity and authenticates the user's identity with Adobe IMS.

Step 3: Based on the user's prompt, the Knowledge Base Generative AI Service retrieves the relevant document content from the Knowledge Base data storage and leverages Azure OpenAI for optimizations during the retrieval flow.

Step 4: The Knowledge Base Generative AI Service sends the user's prompt and document metadata retrieved from the data store to Azure OpenAI to determine the relevant document content to answer the customer's prompt.

Step 5: Azure OpenAI returns the generated response to the Knowledge Base Generative AI Service.

Step 6: The Knowledge Base Generative AI Service sends the retrieved document content and the user's prompt to Azure OpenAI, which generates an answer based on the provided context.

Step 7: Azure OpenAI returns the generated response to the Knowledge Base Generative AI Service.

Step 8: The Knowledge Base Generative AI Service stores the generated response and the chat history in Knowledge Base data storage.

Step 9: Simultaneously, the Knowledge Base Generative AI Service streams the response back to the user to view in the Knowledge Base user interface.

Data flow narrative — from Slack and Teams chatbots

Step 1: The user asks a question in a Slack or Microsoft Teams channel connected to a Knowledge Base collection and configured with a Knowledge Base chatbot. For information on managing a chatbot for a collection, see [HelpX](#).

Step 2a: The Knowledge Base Generative AI Service receives the user's question and identity from Slack or Microsoft Teams and **Step 2b:** validates the user's identity with Adobe IMS.

Steps 3-8: The same as in the prior data flow.

Step 9a: Simultaneously, the Knowledge Base Generative AI Service streams the response back to the user to view in the Slack or Microsoft Teams channel.

Security architecture and data flow — querying from PDF Spaces⁹




Figure 3: Query a Knowledge Base collection from PDF Spaces data flow diagram

Data flow narrative

Note: Before the user can query a Knowledge Base collection from PDF Spaces, they must link that collection to their PDF Space. For more information, see [HelpX](#).

Step 1: In a PDF Space, the user asks a question and **Step 1a:** the query is sent to the Acrobat Generative AI Service.

Step 2: The Acrobat Generative AI Service calls the Knowledge Base Generative AI Service, which receives the user's request and identity and **Step 2a:** authenticates the user's identity with Adobe IMS.

Steps 3-7: The same as the prior two data flows.

Step 8: The Knowledge Base Generative AI Service caches the generated response and chat history.

Step 9: The Knowledge Base Generative AI Service sends the generated response to the Acrobat Generative AI Service.

⁹ Available only when PDF Spaces integration is enabled by the organization's Adobe Administrator.

Step 10: The Acrobat Generative AI Service synthesizes the response with content from local PDF Spaces documents using Azure OpenAI.¹⁰ **Step 10a:** Azure OpenAI returns the generated response to the Acrobat Generative AI Service.

Step 11: The Acrobat Generative AI Service stores the generated responses and the chat history in the user's provisioned encrypted Adobe Cloud Storage.

Step 12: Simultaneously, the Acrobat Generative AI Service streams the response back for the user to view in PDF Spaces.

Knowledge Base and Azure OpenAI

Knowledge Base currently leverages Azure OpenAI and enhances it with Adobe proprietary technology to glean relevant information and insights that the user is interested in extracting from the uploaded files.

To generate answers or extract values, Adobe may pass the following data to Azure OpenAI:

- Customer queries
- The relevant user document data needed to answer the user's query
- The instructions for how to answer the user's question

Adobe has disabled logging in Azure OpenAI, so that Microsoft cannot collect, review, or store any data sent to Azure OpenAI for processing. More information is available at [Azure OpenAI data privacy and security](#).

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

Content filtering

Adobe leverages the [Microsoft Azure AI Content Safety](#) content filtering service to moderate hate, sexual, violent, self-harm, and protected material content. Adobe has configured content filtering to "Medium" blocking for both input and output that may violate Azure OpenAI User Guidelines.

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 Adobe's generative AI solutions, see the [Generative AI Built for Business solution brief](#).

Knowledge Base in Acrobat is built on the Acrobat Sign infrastructure, which operates under Adobe's Common Controls Framework (CCF). This provides a standardized set of security and compliance controls designed to support external audits. Adobe plans to include Knowledge Base in Acrobat in applicable certification, attestation, and audit reporting scopes in the upcoming assessment cycle.

¹⁰ See Step 8a through Step 10 in the "Create a PDF Space" data flow narrative in the [Adobe Acrobat Studio security overview](#) for more information.

Data retention and deletion

AI Assistant chat history

All chat interactions in Knowledge Base are recorded in the user's history and retained in Knowledge Base data storage until the relevant document or collection of documents is deleted.

User-uploaded documents

Knowledge Base stores user-uploaded documents, along with the content indexes and citations generated from these files, in Knowledge Base data storage until the user deletes the file(s). If a user deletes a document in the Knowledge Base web application, the document and its corresponding metadata are immediately deleted, no longer accessible to any user, and not recoverable through the Knowledge Base application.

Knowledge Base does not currently support document version history, so previous versions cannot be restored. Document changes are recorded in internal application logs, but customer-accessible audit logs are not currently available.

Knowledge Base encrypts user-uploaded documents on a per-document basis. User-uploaded documents are only visible to the user unless they choose to share the collections in which the documents are ingested with other authorized users within the same organization by clicking the "Share" button on a collection.

Syncing with SharePoint and/or Google Drive

If the user has chosen to sync documents from their SharePoint/Google Drive repository, Knowledge Base mirrors the retention policy of the SharePoint/Google Drive folder(s) respectively.

Users may leverage version history and audit logging of content synchronized via the SharePoint or Google Drive connector from the source repository, provided the user's SharePoint or Google Drive credentials remain active. If that authentication becomes invalid, synchronization is interrupted and the associated documents will remain in Knowledge Base data storage rather than following the source repository's version, audit, or deletion state.

Disabling/removing a user account

If the individual assigned to the "product administrator" role in the organization's Adobe Admin Console disables or removes a user account from Knowledge Base, the user's personal data (such as chat history, configured connectors and chatbots, feedback, and usage records) is deleted from the platform.¹¹ However, content uploaded by that user, including documents in collections and content associated with connectors or chatbots, is retained unless the organization's Adobe Administrator explicitly deletes it during the offboarding process.

Ownership of the content automatically transfers to the organization's Adobe Admin to enable continued access. At that time, the Adobe Admin may retain ownership or transfer ownership to another user. Ownership transfer is supported at the collection level rather than for individual documents.

¹¹ Adobe does not currently log these events.

Terminating a Knowledge Base contract

Upon termination of the organization's contract with Adobe, the customer should export or delete any content they wish to retain or remove prior to account closure. The customer may also notify Adobe at the time of contract termination and request deletion of their data. Adobe does not automatically delete customer data upon contract termination or expiration.

Data processing and storage locations

Knowledge Base

The Knowledge Base web application, Knowledge Base Generative AI Service, and corresponding data and file storage are hosted in AWS data centers in Virginia, USA and Frankfurt, Germany.

Although Knowledge Base in Acrobat is built on the Acrobat Sign infrastructure, data remains separate and is not transferred or shared between the Knowledge Base in Acrobat and Acrobat Sign solutions.

PDF Spaces

For information about data storage and processing in PDF Spaces, see the [Adobe Acrobat Studio security overview](#).

Azure OpenAI

Microsoft hosts the Azure OpenAI Service in Microsoft Azure data centers. Adobe restricts processing of Knowledge Base data to the customer's provisioned data sovereignty zone (either North America or EU).

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](#).