Adobe

SECURITY FACT SHEET

Adobe Firefly Services

March 2025

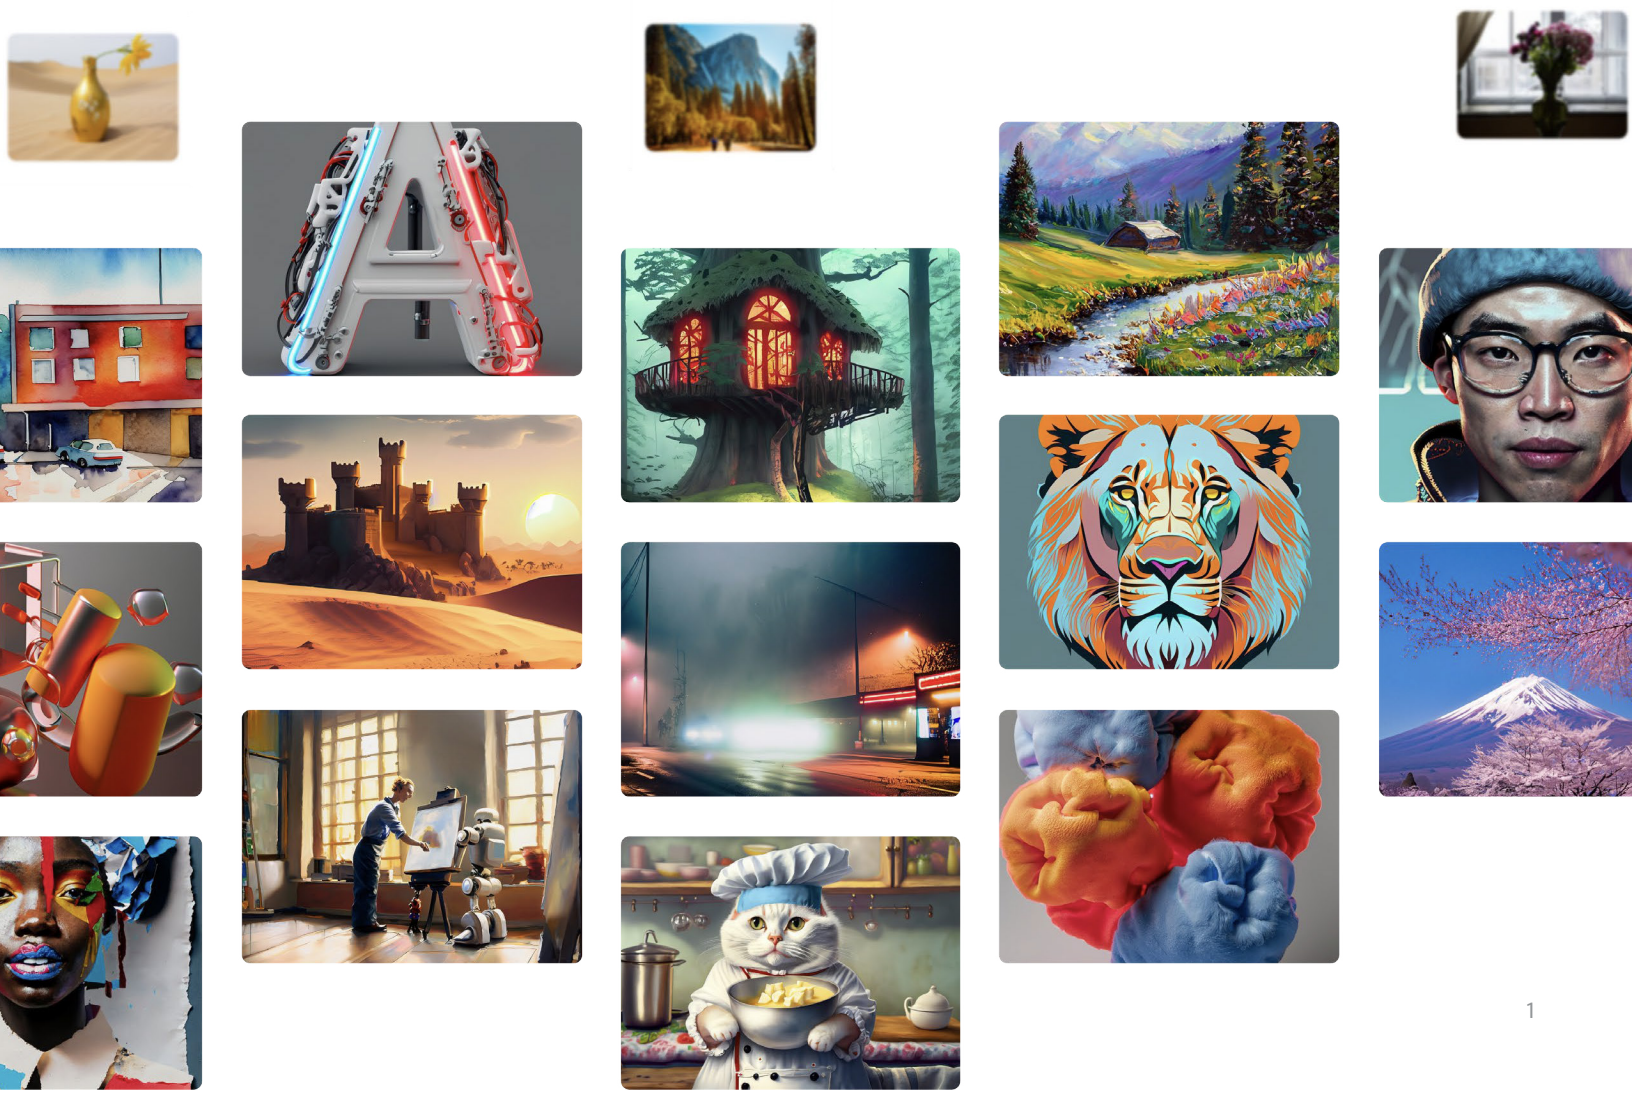

About Firefly Services

Firefly Services is a comprehensive set of non-generative AI, generative AI, and creative APIs that streamline workflows. By embedding these capabilities into any production or experience workflow, organizations can scale the generation, editing, and assembly of content and streamline its production while maintaining quality and control.

API Families

Firefly Services includes a range of APIs that are grouped into API families. More information on each API family is available on HelpX.

- Adobe Firefly API – Generative AI services that automate text-to-image generation for entire images and masked in-painting or expanded out-painting areas. For more information, please see Firefly API.
- Adobe Photoshop API – Non-generative AI and creative services that edit and process PSD, JPEG, PNG, and TIFF images using Adobe Photoshop technology. For more information, please see Photoshop API.
- Adobe Lightroom API – Non-generative AI and creative services that edit JPEG, PNG, RAW, and other supported file formats using Adobe Lightroom technology. For more information, please see Lightroom API.
- Adobe InDesign APIs – Creative services that merge data into and render JPEG, PNG, or PDF outputs from an InDesign document. For more information, please see InDesign APIs.
- Audio/Video APIs – Generative and non-generative AI services that create and modify audio and video through APIs that support a range of actions, such as reframing for different aspect ratios, generating speech- and actor-based video avatars from text, and translating and lip syncing videos. For more information, please see Audio/Video APIs and TLS APIs.
- Content Tagging API – Non-Generative AI services that extract intelligent features from content, organize and streamline content flow, and deliver more impactful, personalized experiences. For more information, please see Content Tagging API.

How to Access Firefly Services

To access Firefly Services APIs, the customer's Adobe Admin must assign the user to the *Developer* role in the organization's Adobe Admin Console, enabling the user to access the Developer Console. The Adobe Admin must also assign the user a Firefly Services entitlement in the Admin Console, which allows the user to create a project using the Firefly Services APIs and develop applications in the Developer Console.

In the Developer Console, the user can generate an API key (client ID) and client secret. Using these credentials, the user can have any application they create that leverages Firefly Services APIs generate an access token, which is passed along with every API request to validate the application's credentials.

Adobe Identity Management Services (IMS) is used to manage access, including authentication and authorization, to Firefly Services APIs. For more information, please see the Adobe Identity Management Services Security Overview.

API Authentication

Firefly Services support OAuth Server-to-Server credential authentication. For more information, please see the Server to Server authentication documentation.

Firefly Generative AI Services Security Architecture and Data Flow

Figure 1 below illustrates the data flow between an application built to leverage the Firefly Services generative AI APIs to generate content in Firefly using either the Firefly core model or a customer-defined Firefly Custom Model.

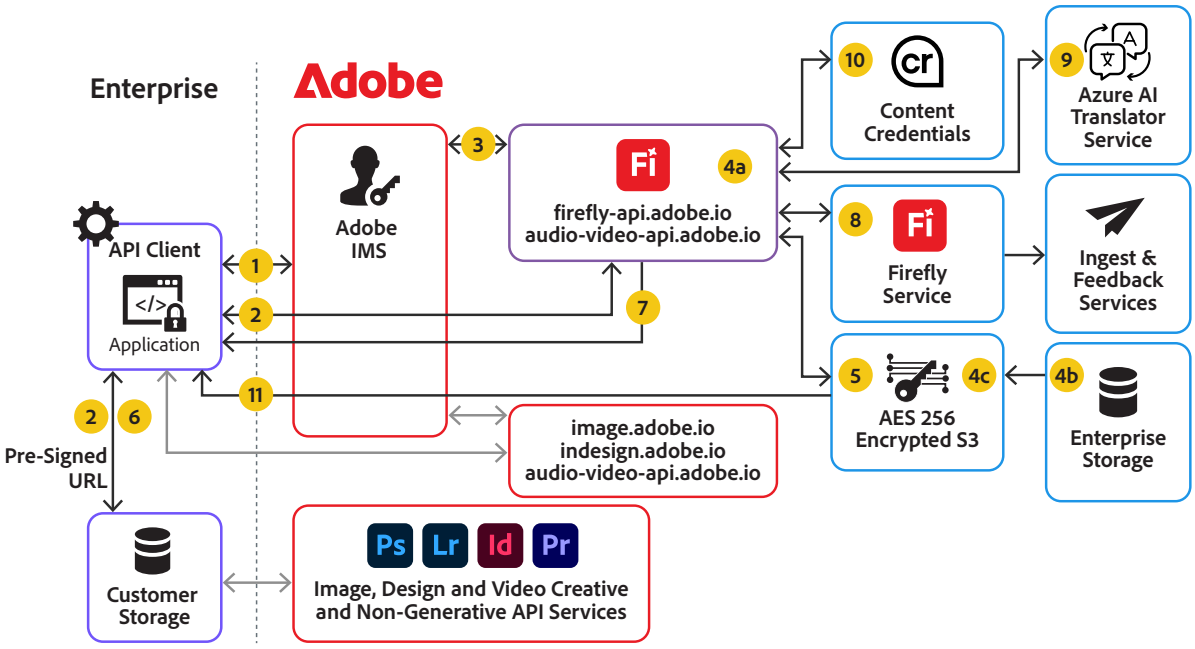

Figure 1: Adobe Firefly Generative AI Services security architecture and data flow

Data Flow Narrative

Step 1: The client application requests an access token from Adobe Identity Management Services (IMS) using an API Key (client ID) and client secret.

Step 2: The client application submits an API request to the Firefly API containing the prompt, the configuration settings, and optionally, the pre-signed URL to the user's defined storage location.[1]

Note: The Firefly API leverages the Firefly Service for inferencing and generating content. The Adobe Firefly for enterprise (core and web) security fact sheet contains more information about the security posture of Firefly.

Step 3: The Firefly API validates the included API key (client ID) and the access token with Adobe IMS.

Step 4: If the inference process utilizes a Firefly Custom Model, the client application (4a) retrieves the Model ID from the Firefly Service (which is used to identify the model in Step 8 to generate the content) and (4b) retrieves the custom model delta weights associated with that Model ID from enterprise storage. Both the Model ID and the delta weights are cached in an AES256-bit encrypted AWS S3 storage location (4c) for use during inferencing.

Note: Delta weights are cached for two weeks.

Step 5: If the inference process for Firefly requires reference content (e.g., Generative Match, Generative Fill, or Generative Expand), the user's uploaded content is stored in an AES 256-bit encrypted AWS S3 storage location and referenced by a pre-signed URL.

Note: Reference content is deleted 24 hours after initial creation.

Step 6: If the developer uses the generative AI audio/video APIs (i.e., Translate or Lipsync), the Firefly Service retrieves the relevant files from the customer's pre-defined storage via the pre-signed URL.

Step 7: If the developer calls the Firefly Service asynchronously, the Firefly Service returns a status URL that the developer can query to determine when the process has finished.

Step 8: The Firefly Service generates content based on the prompt text, configuration settings and, if utilized, the Firefly Custom Model and/or the uploaded reference content. The Firefly Service retrieves any reference content from the pre-signed URL noted in Step 5.

1 As of March 2025, supported locations include Adobe Enterprise Storage for business, Amazon Web Services, Azure, Dropbox, and Google Drive.

Step 9: If the developer uses the Translate API (one of the audio/video APIs), then the Firefly Service sends the content to the Azure AI Translator Service to translate the text.

Note: The Azure AI Translator Service does not store, retain, train, or use the processed content in any manner other than to translate the specific text.

Step 10: The Firefly API attaches a Content Credentials manifest to the generated content and saves this manifest to the Content Credentials cloud.

Step 11: The Firefly API stores the generated content in an AES 256-bit encrypted AWS S3 storage location, generates a pre-signed Firefly URL, and returns the URL to the client application either directly or via status URL, if the user choses to work asynchronously.

Firefly Non-Generative AI and Image, Design, and Audio/Video Services Security Architecture and Data Flow

Figure 2 below illustrates the data flow between an application built to leverage the Firefly Services non-generative AI and creative APIs to manipulate content. While the underlying services differ,[2] the general data flow is the same.

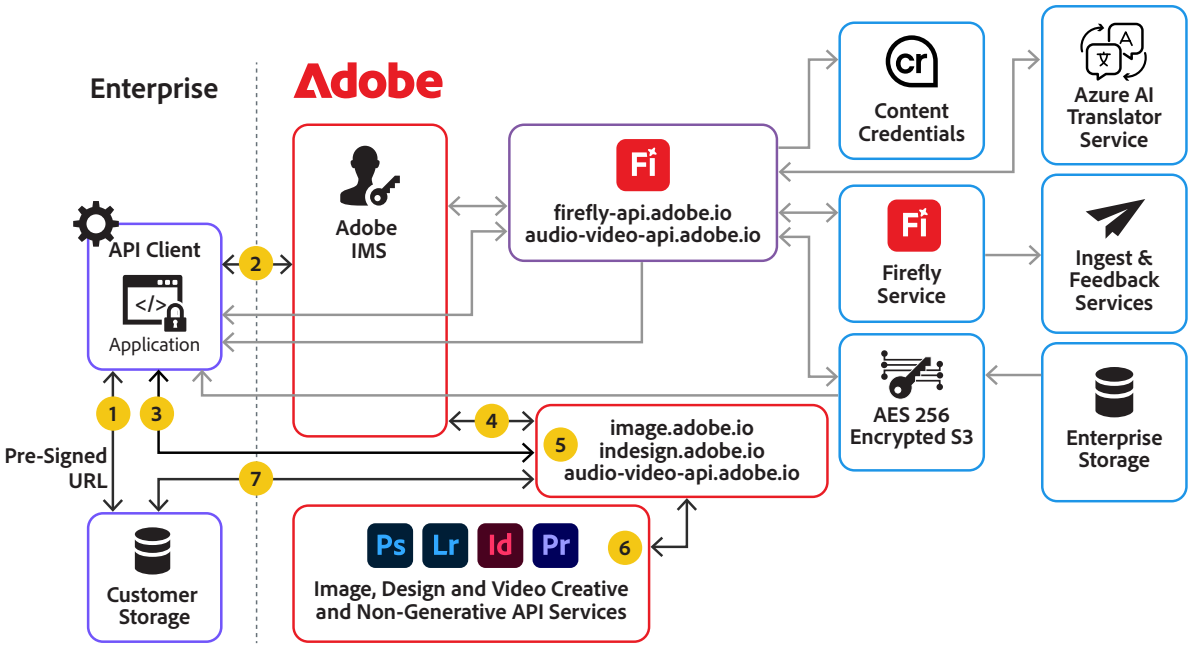

Figure 2: Adobe Firefly Services Non-generative AI and Image, Design, and Audio/Video Services security architecture and data flow

2 For example, using the Photoshop API to remove the background of an image or the InDesign API to merge data and export to PDF, or using the Video API to change the aspect ratio a video to export out to multiple platforms.

Data Flow Narrative

Step 1: The client application ensures that relevant content is available via a pre-signed URL to the customer-defined storage location.

Step 2: The client application requests an access token from Adobe Identity Management Services (IMS) using an API Key (client ID) and client secret.

Step 3: The client application submits an API request to the relevant service referencing the pre-signed URL from Step 1.

Step 4: The API validates the included API key (client ID) and the access token with Adobe IMS.

Step 5: The API service returns a status URL to the client application.

Step 6: The API service asynchronously performs the requested tasks using the relevant creative services (e.g., Photoshop, Premiere Pro, InDesign, or Lightroom). In-process images and content are cached for the duration of the API transaction and then immediately deleted.

Step 7: The client application queries the status URL provided in Step 5. When the API service completes the required processing, it stores the content in the cusomter-defined storage location, and the client application downloads the completed asset using the pre-signed URL.

Data Encryption

In Transit: All data is encrypted in transit over HTTPS using TLS v1.2 and a minimum of AES 128-bit GCM encryption.

At Rest: If the user chooses Adobe-managed storage as their customer-defined storage location, which is determined on a per-transaction basis, Adobe encrypts the content at rest using AES 256-bit encryption. Generated content is cached in AES 256-bit encrypted storage.

Input and Output Content Storage and Processing

Adobe stores input and output information in connection with the services it provides.

Input

By default, Firefly Services store or process the following inputs:

- **Generative AI APIs (firefly-api.adobe.io, audio-video-api.adobe.io)**
    - Prompt text input and configuration settings, such as seed values, image styles, tone, depth of field, aspect ratio, etc. (retained for 90 days)
    - A pseudonymous user ID (e.g., i001ad83a-d41f-4afb-9f5c- 7b72c88ae873a) (logged and stored for 90 days with other data elements as noted)
    - Reference content, if uploaded by the API (cached for 24 hours, as described above)
- **Photoshop, Lightroom, InDesign, and audio/video non-generative AI and creative APIs**
    - Configuration settings (logged for 30 days)
    - Uploaded images (processed and cached for 24 hours)
    - Audio/video input assets are transactional and are deleted immediately after use

Note: Adobe does not include enterprise user content (including Firefly inputs and outputs) in datasets used to train Firefly foundation models. This does not apply to (1) the use of Firefly as part of any feedback or improvement programs in which the customer/user can control the use of their content for training through an opt-in/opt-out menu or (2) the customer's use of any Adobe product or feature that utilizes AI training to provide the customer with a service.

Output

By default, Adobe does not permanently store Firefly outputs; instead, they are temporarily cached in application-managed cache storage for a maximum of seven (7) days and can be accessed using a pre-signed URL.

- **Firefly API (firefly-api.adobe.io)**
  - Generated image (stored for 24 hours and accessible via pre-signed URL as described above)
  - A pseudonymous user ID (e.g., i001ad83a-d41f-4afb-9f5c- 7b72c88ae873a) (logged and stored with the other data elements as noted)
  - A cryptographic hash of the image for Content Credentials and stored in the Content Credentials Cloud with the manifest
- **Audio/Video APIs (audio-video-api.adobe.io)**
  - Processed content is stored in AES 256-bit encrypted cache for seven (7) days and is accessible via an API generated pre-signed URL. Additionally, processed content may be stored in customer-defined storage as described above.
  - For generative AI outputs, a cryptographic hash of the generated content for Content Credentials is stored in the Content Credentials cloud repository with the manifest.
- **Photoshop, Lightroom, and InDesign APIs**
  - The processed content (stored in customer-defined storage as described above)

Content Credentials

Adobe automatically generates Content Credentials for certain Firefly-generated assets to help provide transparency that the asset was created using generative AI.

Content Credentials typically contain the following metadata:

- In certain cases, a thumbnail of the generated image
- The tool/tools used to generate the asset
- Whether the asset was completely generated by Firefly or combined with other content

- A cryptographic hash of the image and its metadata in a verifiable, tamper-evident signature that provides proof that the image and metadata have not been altered. The cryptographic hash is irreversible.

Content Credentials are attached to the exported asset file and stored in the Content Credentials cloud repository, which allows recovery of the Content Credentials in the event it is stripped from the exported asset.

Note: Text prompts are never included in any automatically generated Content Credentials.

User-Generated Content

Firefly Services accept and process user-generated content (UGC). This content is downloaded, secured, and temporarily cached as part of normal service operations as described above.

If the Firefly API workflow includes reference content, it will upload and store that reference content in an AES 256-bit encrypted AWS S3 storage location referenced by a pre-signed URL.

The image, design, and audio/video API workflows require a customer-defined storage location for user-generated content. The API will upload from that location and store processed and edited content into that location via a pre-signed URL.

Hosting Locations

Except as noted, the server-side components of the Firefly Services APIs are hosted on Amazon AWS data centers in the US-East and US-West regions.

Audio/Video APIs utilize edge servers in EMEA (Ireland), but all processing occurs in the US-East and US-West regions.

InDesign APIs are processed in Azure and AWS data centers in the US-East region.

Azure AI Translator Service processing occurs in the US-East and US-West regions.

Testing

Adobe teams conduct rigorous 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. For the annual Security Testing Report for Adobe Firefly, please see the Adobe Firefly Security Testing Report (NDA required).

Questions?

If you have any additional questions about the security posture and capabilities of Adobe Firefly, 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.