SECURITY FACT SHEET

Adobe Firefly Custom Models

April 2024

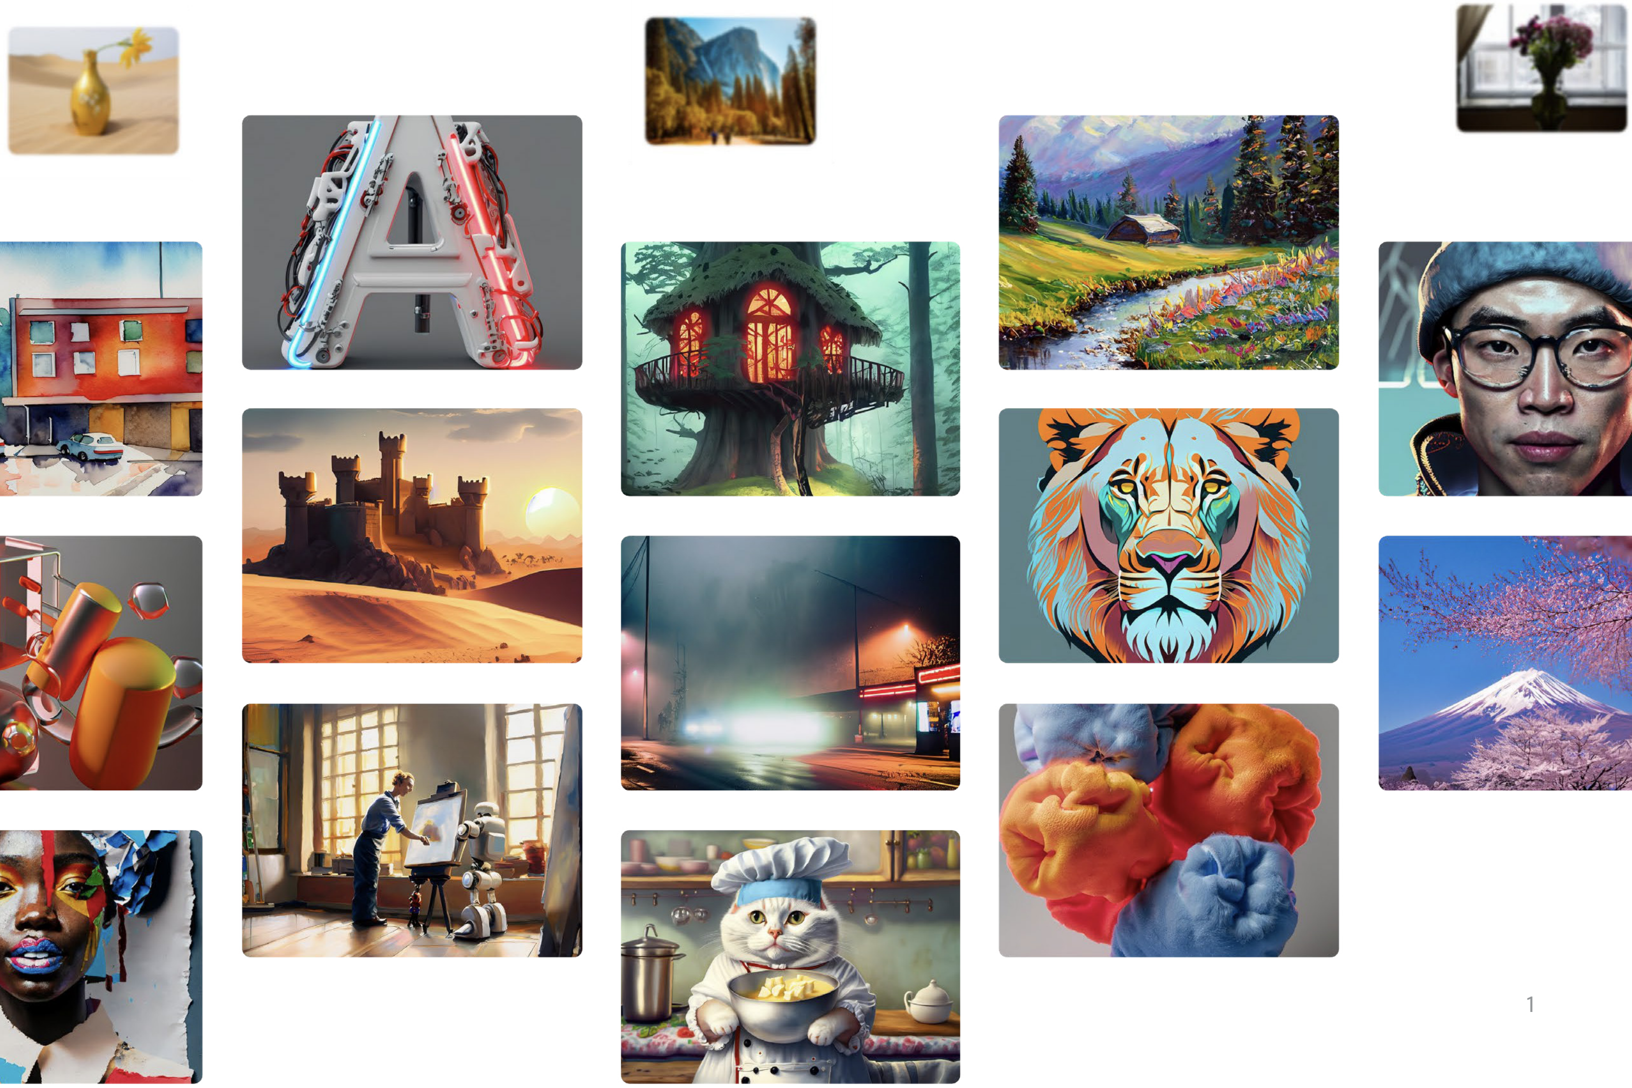

About Adobe Firefly Custom Models

Firefly Custom Models is an enterprise add-on offer that enables brands to fine tune the foundation Firefly generative AI model by training on their signature brand style, campaign style, character, or object. Using Custom Models, organizations can consistently create on-brand creative assets at scale, transforming their style or subject to explore new ideas, visualize different surroundings, generate innovative content, and tailor content to specific segments. This fact sheet covers the security posture and capabilities of Custom Models.

Custom Models Product Profiles

Customers who have purchased Custom Models can manage user access through the Product Profile in their Adobe Admin Console. Administrators can assign a user to one of the following roles:

- **Trainer** – Allows a user to train or fine-tune Firefly models with their brand assets on the Firefly web app.
- **Generator** – Allows a user to generate content with the trained Custom Model/s through the Firefly web app.

Custom Models Security Architecture and Data Flow – Trainer

When a user with the Trainer entitlement initiates a model training activity, the data flows as shown in Figure 1:

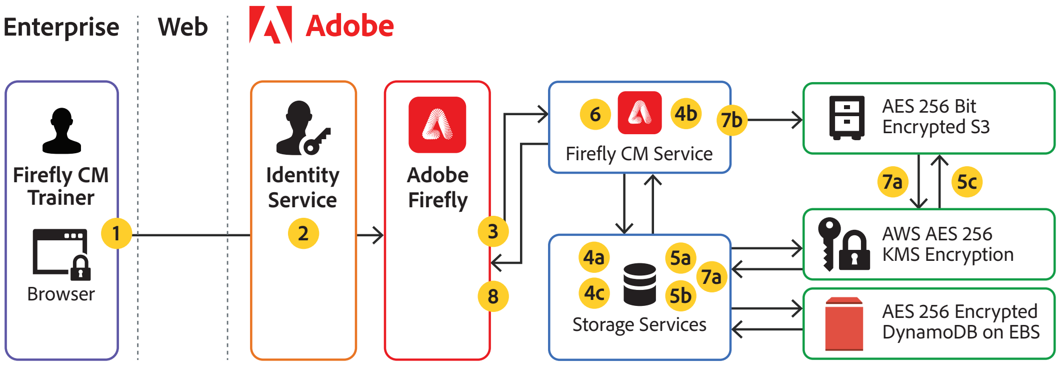

Figure 1: Firefly Custom Models model training security architecture and data flow

Data Flow Narrative

All data is encrypted in transit over HTTPS using TLS 1.2 and a minimum of AES 128-bit GCM encryption.

Step 1: In their web browser, the user signs into the Firefly web app (firefly.adobe.com) using their Trainer-entitled credentials.

Step 2: Adobe Identity Management Services (IMS) validates the user and their entitlement as a Trainer.

Step 3: The user selects "Train a New Model" and a training mode:

- **Style** trains the model on the colors, shapes, and background aesthetic.
- **Subject** (tech preview) trains the model on an object or character. Include images of a single subject with the same traits in different backgrounds and poses.

The user also must add a sample prompt that reflects a practical example of effectively using the model; the training mode (e.g., subject or style) should be included in the prompt text.

The user then uploads reference images on which to train the Custom Model.

Step 4: The Custom Models Service (a) checks the user's permissions with Storage Services, (b) creates the new Custom Model, and (c) sends both the Custom Model and the uploaded images to Storage Services.

Step 5: Storage Services (a) scans the uploaded images for viruses, (b) sends them for encryption with AWS KMS, which encrypts them using a customer-managed encryption key, and (c) stores the images in the organization's AES 256-bit encrypted S3 bucket.

Note: Access to uploaded images is only available through the Custom Models interface, firefly. adobe.com/cme/train

Step 6: The Custom Models Service trains the new model on the uploaded reference images.

Step 7: The Custom Models Service stores the delta weights for the newly trained Custom Model in both (a) the organizational storage (as noted in Step 5) and (b) in an encrypted cache.

Step 8: The user tests the model and then publishes the Custom Model to the Custom Models Service for use by Generator users.

Custom Models Security Architecture and Data Flow – Generator

When a user with the Generator entitlement initiates an image generation activity using the previously trained Custom Model, the data flows as shown in Figure 2:

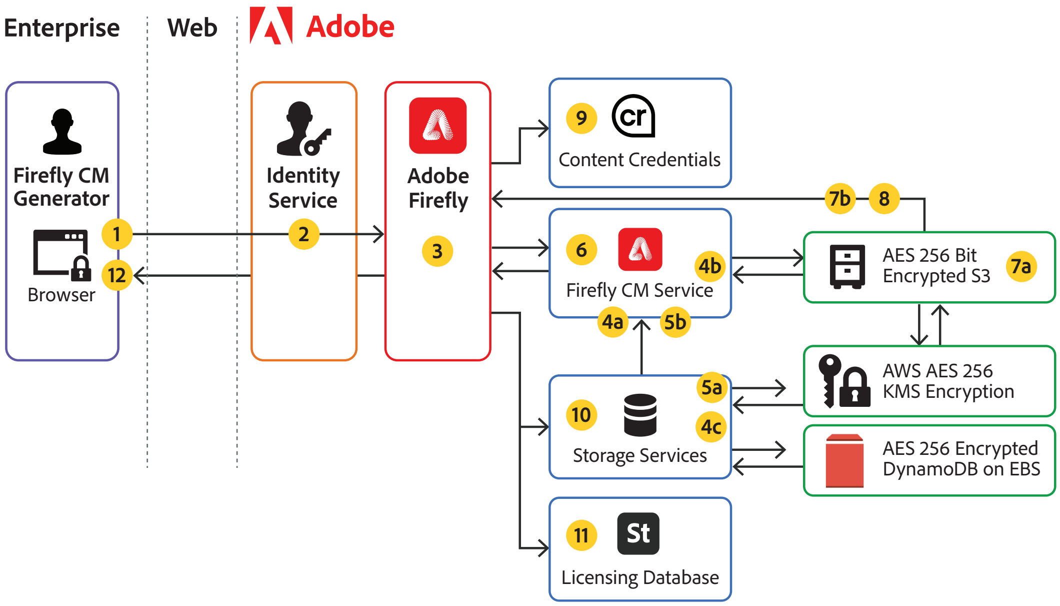

Figure 2: Adobe Firefly Custom Models generator security architecture and data flow

Data Flow Narrative

All data is encrypted in transit over HTTPS using TLS 1.2 and a minimum of AES 128-bit GCM encryption.

Step 1: In their web browser, the user signs into the Firefly web app (firefly.adobe.com) using their Generator- or Trainer-entitled credentials.

Step 2: Adobe Identity Management Services (IMS) validates the user and their entitlement for Custom Models as well as their permissions to access certain Custom Models.

Step 3: The user initiates a “Text-to-Image” workflow and selects a Custom Model from the “Your Models” drop-down list which also populates the prompt field with Trainer supplied sample prompt.

Step 4: The Custom Models Service (a) checks user permissions with Storage Services and (b) requests the delta weights from the encrypted cache or (c) from the Storage Services if the cache has expired.

Step 5: If required, the Storage Services (a) retrieves and decrypts the delta weights using the customer's managed encryption key (CMK) and (b) sends the content and metadata to the Custom Model Service.

Step 6: The Custom Models Service generates images by combining the delta weights with the base Firefly model for inference.

Step 7: The generated images are (a) temporarily stored in an application-managed cache storage and (b) a pre-signed URL for the cached images is returned to firefly.adobe.com.

Step 8: If the user performs an action on the generated images such as "Download," "Save to library," "Copy Image," or "Edit in Adobe Express." The firefly.adobe.com application loads the generated output from that pre-signed URL.

Step 9: Firefly.adobe.com attaches a Content Credentials manifest to the downloaded, saved, or edited image and saves this manifest to the Content Credentials cloud.

For more information on Content Credentials, please see the "Content Credentials" section below.

Step 10: If the user chooses "Save to library," or "Edit in Adobe Express," firefly.adobe.com will send the image and the attached Content Credentials to the respective application, which will store it in user-managed storage through Storage Services.

Step 11: If the purchased enterprise offer includes output indemnification, then for certain text-to-image workflows, firefly.adobe.com sends a full resolution copy of the generated image with embedded Content Credentials to the Licensing Database.

Step 12: If the user chooses to "Download" the image, firefly.adobe.com will send the image and the attached Content Credentials to user's desktop.

Content Credentials

Adobe automatically generates Content Credentials for every Firefly-generated asset 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

- Summaries of the type of actions taken in Firefly (such as use of a reference file, edit activity, etc.)

- 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.

Content Storage and Processing

Uploaded images used to train the Custom Models and the associated delta weights are stored in Adobe storage for business, which is a secure cloud storage hosted in Amazon Web Services (AWS) data centers. (see "Storage Services" in data flow narrative above).

Adobe does not train our foundation Firefly generative AI models on any Creative Cloud subscriber's personal content.

Enterprise Access and Control

Enterprises can control access to Custom Models through the Adobe Admin Console. Only users entitled with a Trainer product profile can train a Firefly Custom Model, upload images for training, or manage those uploaded images. Only users entitled with either a Trainer or a Generator product profile can generate new images using a Custom Model.

Users not provisioned with a Trainer or Generator entitlement in their product profile cannot access Custom Models nor can they generate images using a Custom Model.

User Identity Information

Adobe uses named user licensing to uniquely identify users of any Adobe product, including Custom Models. Custom Models is fully integrated with Creative Cloud for Enterprise identity access and management using Adobe Identity Management Services (IMS), allowing multi-factor authentication (MFA) to any SAML2-compliant provider.

More information on named user licensing can be found in the Adobe Identity Management Services Security Overview.

Data Storage Locations

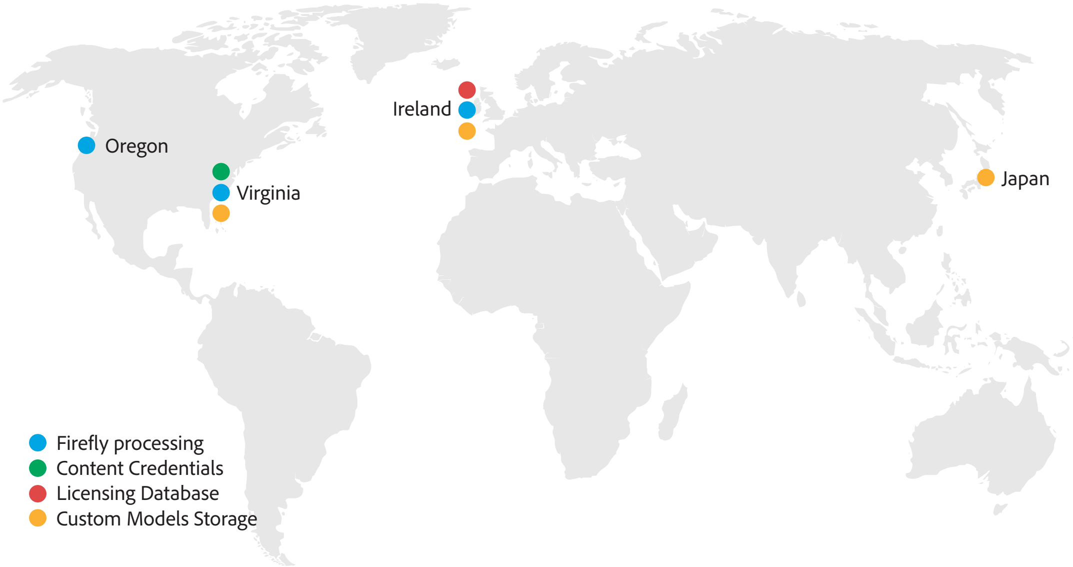

Reference images uploaded for Custom Models as well as the delta weights are stored in the customer's assigned regional data center — US-East (Virginia), EMEA West (Ireland), or APAC (Japan.)

Adobe currently processes, caches, and stores additional Firefly input content (such as Generative Match reference images) in Amazon Web Services (AWS) data centers in the US-East (Virginia) and US-West (Oregon) regions, regardless of the user's location. Adobe offloads some prompt pre- and post-processing to AWS data centers in the EMEA West (Ireland) region.

Adobe currently stores Content Credentials in AWS data centers in the US-East (Virginia) region, regardless of the user's location.

Data Types and Retention

Adobe retains customer data in accordance with the customer's Enterprise Term License Agreement (ETLA) contract and Adobe's product-specific licensing terms.

The following types of data are potentially stored by Adobe, depending on the user actions as described above:

- Uploaded reference content
- Any prompt text
- Data to identify the model(s) used in inferencing

- Configuration settings (such as aspect ratio, content type, styles, tone, etc.)
- A timestamp (based on multiple NTP Stratum 1 satellite-connected and atomic reference clocks)
- A cryptographic hash of the image for Content Credentials
- Binary data from the Firefly-generated content (if the user initiates specific actions as noted above)
- User identity data in the form of a pseudonymous ID (e.g., GUID i001ad83a-d41f-4afb-9f5c- 7b72c88ae873a)
- Data required for IP Indemnification as noted above (if purchased)

Conclusion

If you have any additional questions about the security posture and capabilities of Custom Models, 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.

Adobe