SECURITY FACT SHEET

Adobe Firefly Custom Models

January 2025

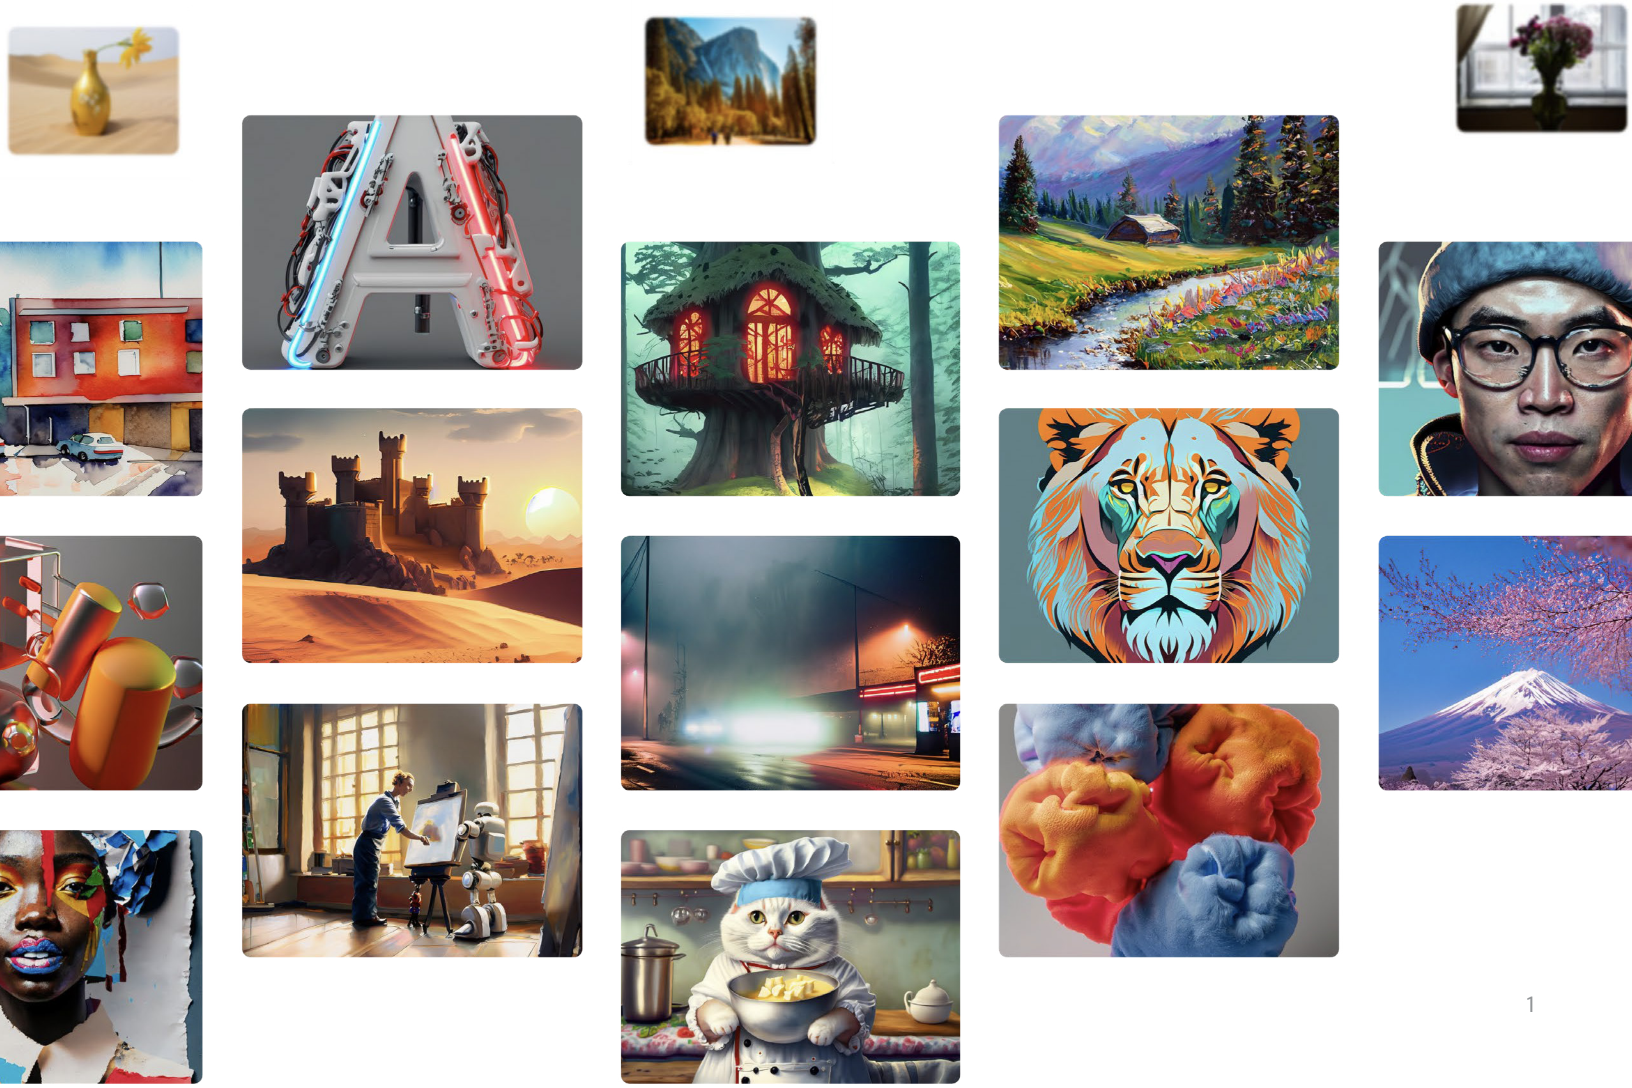

About Adobe Firefly Custom Models

Firefly Custom Models 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.

Enterprise Access and Controls

Customers who have purchased Custom Models can manage user access via the Adobe Admin Console. The customer's Adobe Admin assigns users to the **Trainer** role in the console, which allows a user to train or fine-tune Firefly models with their brand assets on the Firefly web app.

Trainers may share Custom Models with other users using the enterprise-grade granular governance features of Custom Models. By setting permissions for collaborators to edit, review, or use, other members of the organization can collaborate in training the model, reviewing its performance, and generating images.

Users can only generate assets or images using Custom Models if they've been explicitly invited to do so by an existing collaborator via the Share menu or if the Custom Model belongs to a Project to which they have been given access. For more information, please see Custom Model Projects.

Security Architecture and Data Flow – Model Training

Figure 1 (below) illustrates the data flow when a user with a Trainer entitlement initiates a model training activity: All data is encrypted in transit over HTTPS using TLS 1.2 and a minimum of AES 128-bit GCM encryption.

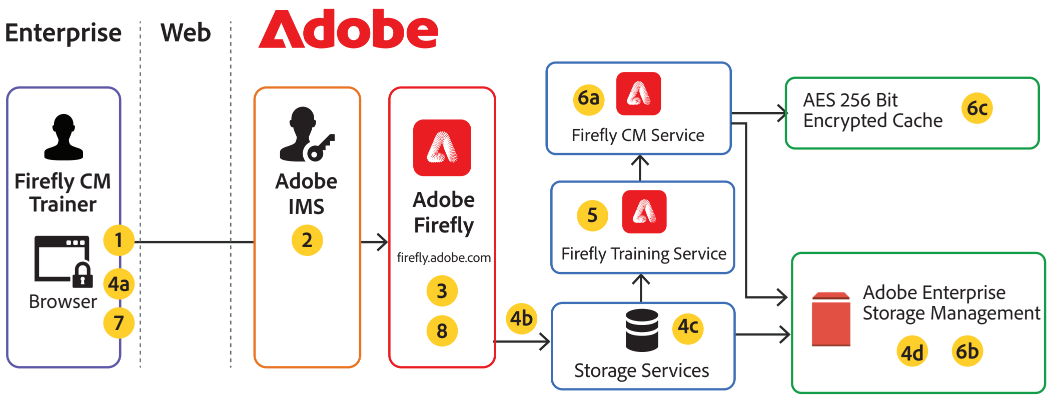

Figure 1: Firefly Custom Models data flow diagram

Step 1: In their web browser, the user signs into the Firefly web app (firefly.adobe.com) using their Trainer credentials and selects "Custom models."

Step 2: Adobe Identity Management Servies (IMS) validates the user and their Trainer entitlements. IMS returns a user token that authenticates the Trainer with both the Firefly Custom Models Service and the Firefly Training Service.

Step 3: The Trainer selects "Train a New Model" in the Firefly web app along with one of two training modes:

- **Style** trains the model on the colors, shapes, and background aesthetic.

- **Subject** trains the model on an object or character.

The Trainer also selects an existing or creates a new Project folder in which to store the Custom Model.

Note: The Trainer can only store a Custom Model in a Project to which they have appropriate permissions. For more information, please see Manage Custom Model Project Access.

Step 4: The Trainer (a) selects at least 10 images they want to use to train the Custom Model. These images are uploaded to Storage Services, which (b) scans the images for viruses, (c) encrypts the images using a customer-managed encryption key (CMK), and (d) stores the images in the organization's AES 256-bit encrypted S3 bucket.

Step 5: The Firefly Training Service uses the uploaded images to train a Custom Model. Once the training is complete, the model will appear on the "Your models" page in firefly.adobe.com with a *Ready* status.

Step 6: The Firefly Custom Models Service (a) creates the Custom Model and stores it, along with the delta weights and other metadata, in both (b) the organizational storage for long-term storage (as noted in Step 4) and (c) in an encrypted cache to speed review and use of the model.

Note: The cache is kept for as long as users are actively accessing or fine-tuning the Custom Model and is deleted after 14 days of inactivity.

Step 7: The Trainer can preview and test the Custom Model to confirm it matches their intention before publishing it. The Trainer may also share the Custom Model with collaborators. Users with "Review" access can also preview and test the Custom Model to validate that the output matches the Trainer's intention.

Step 8: When the Custom Model is ready to use, the Trainer can publish it and share it with other collaborators.

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). More information on input and output storage and processing can be found in the [Adobe Firefly for enterprise security fact sheet.]()

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

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

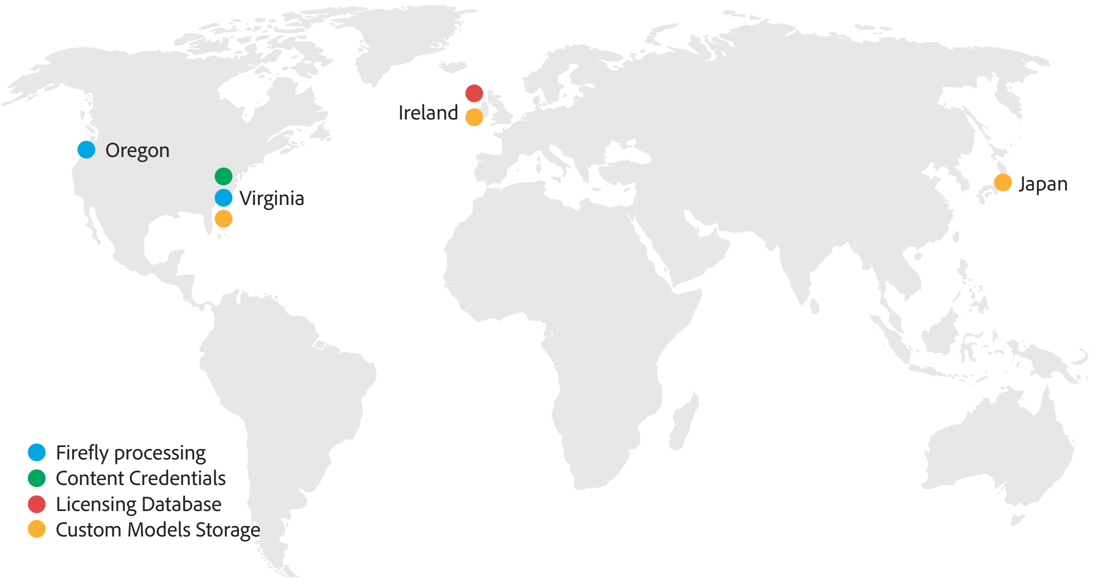

Figure 2: Firefly Custom Models data storage locations

Reference images uploaded for Custom Models, as well as corresponding delta weights and other metadata, 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.

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)
- 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)

Testing

Adobe teams rigorously test our generative AI products to reduce the potential for biased and harmful outcomes. 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).

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.