Adobe

SECURITY OVERVIEW

Adobe Brand Intelligence

September 2026

About Adobe Brand Intelligence

Adobe Brand Intelligence is an enterprise-grade, continuous learning system designed to operationalize brand governance at scale. It ingests brand assets from multiple sources — including brand guidelines, design systems, approved assets, and review annotations — and structures them into a machine-readable brand ontology.

Adobe Brand Intelligence skills can be integrated into other Adobe solutions for which the customer holds a license, such as Adobe Experience Manager, Adobe Workfront, and Adobe Frame.io, as well as third-party and custom enterprise applications via a headless API-first architecture.

Deployment of Adobe Brand Intelligence, as well as any integrations with Adobe solutions or third-party applications, requires engagement with the Adobe forward-deployed engineering (FDE) team.[1]

Adobe Brand Intelligence components

Adobe Brand Intelligence consists of the following components:

- **Brand Control Center** – Enables Adobe forward-deployed engineers (FDEs) to configure customer instances of Adobe Brand Intelligence and acts as the central hub for managing the brand ontology.
- **Brand Context Layer** – Stores and serves the structured brand ontology, including explicit brand rules and implicit organizational knowledge, and functions as the intelligence backbone during validation and assembly workflows.
- **Intelligence Engine** – Powers brand reasoning via specialized models:
  - **Computer vision models** – Validate visual elements, such as logos, color, typography, and layout
  - **Reasoning models** – Compare asset elements against relevant brand guidelines to determine brand conformance
- **Validate Skill** – Submits assets for automated brand review, returning structured findings with bounding boxes, guideline evidence, and violation categories.
- **Instruct to Assemble Skill** – Creates on-brand asset variants from campaign briefs by intelligently applying brand-approved content to templates.
- **Adobe Brand Intelligence API** – Allows integration with other Adobe solutions, third-party applications, and custom enterprise systems.

[1] Forward-deployed engineering services purchased separately. Additional fees apply.

Provisioning and configuration

A white-glove managed service, Adobe's forward-deployed engineering team provisions the customer's Adobe Brand Intelligence instance using the Brand Control Center. Once provisioned, the FDE configures the customer's Adobe Brand Intelligence instance by uploading brand guidelines, design kits, approved assets, and annotations from connected sources.

Customer data is stored in the Brand Context Layer, which is hosted in Adobe cloud storage. The configured brand instance is then available for integration into customer workflows using the Adobe Brand Intelligence API.

Accessing Brand Intelligence

All users in the customer organization may access Adobe Brand Intelligence. User requests are authenticated using Adobe Identity Management Services (IMS), and authorizations are enforced by domain name validation.

Data encryption

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

Security architecture and data flow

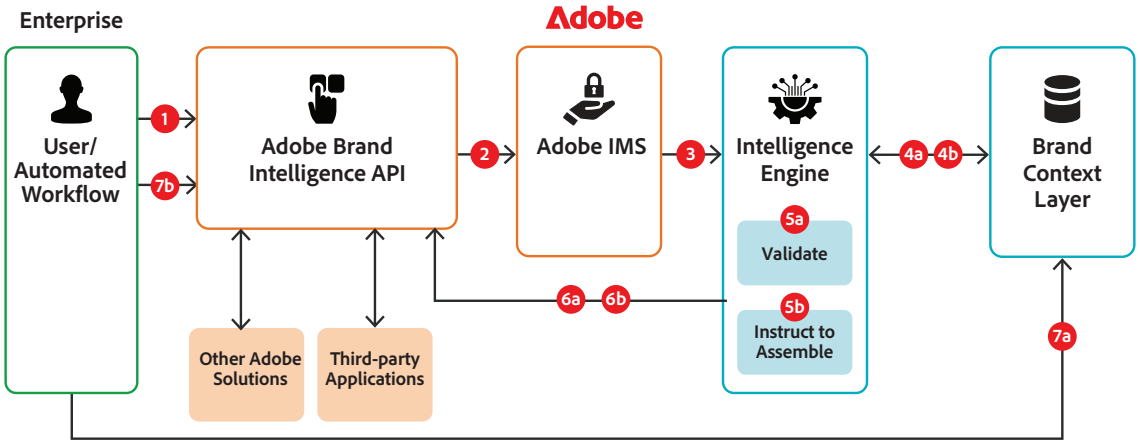

Figure 1: Adobe Brand Intelligence data flow diagram

Data flow narrative – Validate

Step 1: The user or customer-defined automated workflow submits an asset for validation via the Adobe Brand Intelligence API.

Step 2: The Brand Intelligence API authenticates the user with Adobe IMS.

Step 3: Adobe IMS routes the request to the Intelligence Engine within the customer's provisioned instance.

Step 4a: The Intelligence Engine invokes the Validate skill, which identifies key elements (for example, logos, typography, color, layout, and copy) and retrieves the relevant brand guideline chunks for the detected asset elements from the Brand Context Layer, limiting context to only what is necessary for accurate validation.

Step 5a: The Intelligence Engine performs a multi-modal comparison of asset elements against brand standards and compiles a structured findings report, including bounding boxes that note the specific elements of the asset that are not brand-compliant and references to specific guideline pages.

Step 6a: The Intelligence Engine returns the structured findings report to the user or automated workflow in the application in which they are working. A human reviewer confirms, rejects, or annotates the findings in the report.

Step 7a: Humanly validated findings are logged and fed back into the Brand Context Layer to improve future outputs.

Data flow narrative – Instruct to Assemble

Note: Steps 1-3 are the same as the above data flow except the user submits a creative brief via the Adobe Brand Intelligence API.

Step 4b: The Intelligence Engine invokes the Instruct to Assemble skill, which selects the key creative elements (for example, hero images, headlines, and legal text) from the Brand Context Layer and assembles on-brand content variants that match the brief.

Step 5b: The Intelligence Engine returns the on-brand content variants to the user or automated workflow in the application in which they are working.

Step 6b: The user reviews the content variants and makes any final changes (for example, resizing images or editing text).

Step 7b: The user publishes the approved content in the relevant application via the Adobe Brand Intelligence API.

Adobe Brand Intelligence and third-party AI models

Adobe Brand Intelligence uses a variety of models under the hood to serve different purposes:

- **Computer Vision Models** – A combination of Adobe first-party models and third-party models used to validate visual elements, such as logos, color, typography, and layout.
- **Reasoning Models** – Third-party models used to evaluate asset compliance against brand guidelines.

The following data may be passed to third-party models during processing:

- Visual representations of submitted assets, such as screenshots and rendered pages
- Relevant excerpts from the customer's brand guidelines
- Structured brand ontology data used to ground model reasoning

Adobe has disabled logging in all third-party LLM services, helping ensure that they do not collect or review any data sent for processing to the LLM.

Adobe does not use customer data to train or fine-tune third-party AI models used by Adobe Brand Intelligence.

Human-in-the-loop feedback

Adobe Brand Intelligence includes a feedback mechanism that allows users to review and annotate validation findings:

- Users may confirm a finding as accurate, reject a finding for a specific reason, edit a finding, or add new comments.
- Validation results are not legally binding, do not constitute regulatory advice, and are not the opinion of Adobe; the customer remains solely responsible for human review to validate both content accuracy and whether it meets the customer's requirements.
- All feedback is logged and used exclusively to improve model performance for the customer's Adobe Brand Intelligence instance.
- Feedback can be collected automatically via direct integrations with other Adobe solutions or manually through Adobe FDE-supported structured inputs.

Adobe does not use customer feedback data to train models for other customers or for general model improvement outside the customer's instance.

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 our generative AI solutions, see the Generative AI Built for Business solution brief.

Data retention

Brand assets and guidelines

Ingested brand data (guidelines, assets, annotations) is retained for the duration of the customer's contract and used solely to power that customer's Adobe Brand Intelligence instance.

Validation logs and feedback

Validation results and human-in-the-loop feedback are stored and used for model improvement within the customer's tenant. Reporting and audit logs are available upon request.

Conversation and session data (API)

Request and response payloads are not retained beyond the processing window unless explicitly stored by the customer or required to log feedback.

Data processing and storage locations

Adobe Brand Intelligence

All Adobe Brand Intelligence components are hosted in Adobe-managed data centers. Specific regional data residency options are confirmed at provisioning.

Third-Party model providers

Requests routed to third-party AI model providers are subject to those providers' hosting locations.

Questions?

If you have any additional questions about the security posture and capabilities of Adobe Brand Intelligence, please contact your Adobe account manager. For detailed information about Adobe's security programs, processes, compliance certifications, see the Adobe Trust Center.