Adobe

SECURITY OVERVIEW

Adobe LLM Optimizer

January 2026

About Adobe LLM Optimizer

Adobe LLM Optimizer is a generative AI-first application for Generative Engine Optimization (GEO), designed to help brands enhance their visibility, accuracy, and influence in AI-driven search environments. Going beyond simply tracking mentions, Adobe LLM Optimizer provides insights into brand presence in AI-generated answers, offers prescriptive content recommendations, and automates optimization fixes.

LLM Optimizer combines what AI systems say about the customer's brand (via LLM responses), how AI agents and users access the customer's site (via CDN logs, analytics platforms, and operational telemetry data), how accessible and optimized the customer's site content is (via SEO and crawler data) and provides recommendations to improve engagement (for both on-site and off-site opportunities). For a complete list of data sources used by LLM Optimizer, see the *Data sources* section below.

LLM Optimizer components

Adobe LLM Optimizer includes the following components:

- **LLM Optimizer User Interface** – Enables customers to configure prompt settings, determine data sources for optimization, and monitor a variety of factors through a range of dashboards, including:
  - **Customer Configuration** – Enables customers to set up categories, topics, and prompts that enable LLM Optimizer to tailor insights using accurate visibility, traffic, and opportunity analysis. Customers also set up CDN log forwarding in this dashboard.
  - **Brand Presence** – Helps customers understand how their brand is appearing in AI-generated responses, including metrics for brand mentions, citations, and sentiment, filterable by period, category, and platform.
  - **Agentic Traffic Insights** – Tracks bot visits to the customer's website, signaling LLM visibility, including metrics on hits and success rates segmented by market and page type.
  - **Referral Traffic Insights** – Provides visibility into how often users click on the links in AI-generated answers, connecting LLM visibility to real engagement, including metrics on bounce rates and referral sources and regions.
  - **URL Inspector** – Helps customers understand how website pages are performing, revealing how URLs are cited across prompts and categories and showing trending citations to help identify which pages are working well and which can be improved.
  - **Optimization Opportunities** – Automatically identifies visibility gaps, suggests improvements, and provides prescriptive guidance tailored to the customer's brand, assigning a projected traffic value for each opportunity to prioritize changes for optimal impact. Customers can implement any recommended onsite optimization with a single click.

- **LLM Optimizer Core Agent** – Processes and analyzes data retrieved from the configured data sources for brand visibility metrics (including mentions, citations, and sentiment in LLM responses), optimization opportunities across owned and external properties, AI agent-driven traffic metrics, and referral traffic data and patterns.

- **LLM Optimizer Tools** – Support data collection, analysis, and recommendations through the following three (3) services:

  - **Data Retrieval Service** – Monitors and collects responses to configured prompts using the integrated and customer-configured data sources

  - **Data Analysis Service** – Processes and analyzes collected data using Azure OpenAI to make recommendations

  - **Opportunity Detection Service** – Identifies actionable insights to improve site and external presence to increase brand visibility in AI search

Accessing LLM Optimizer

To enable users in an organization to use LLM Optimizer, the customer's Adobe Administrator must add the product profile to the Adobe Admin Console. All users in the organization are automatically entitled to access LLM Optimizer; no role-based or user-group-based permissions apply. More information is available on Experience League.

LLMO Optimizer leverages Adobe Identity Management Services (IMS) to manage user access and authentication.

Data encryption

- **In Transit** – All data is encrypted in transit over HTTPS using TLS 1.2 or greater.

- **At Rest** – Any data stored by Adobe LLM Optimizer is encrypted at rest using AES 256-bit encryption. LLM Optimizer does not collect or store any personally identifiable information (PII) data.

Adobe LLM Optimizer security architecture and data flow narrative

The following example illustrates how data typically flows in Adobe LLM Optimizer:

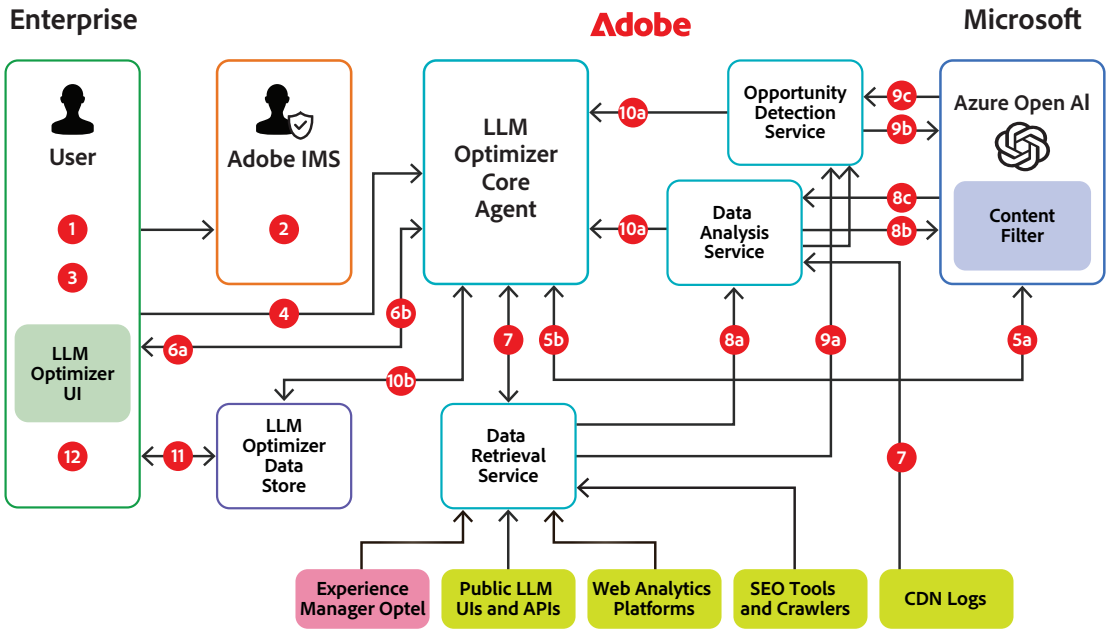

Figure 1: Adobe LLM Optimizer data flow diagram

Data flow narrative

Step 1: The user opens LLM Optimizer in the AEM Experience Hub or via direct access using https://llmo.now and logs in using their Adobe credentials.

Step 2: LLM Optimizer authenticates the user via Adobe IMS (Identity Management Services) and verifies their entitlement to use the service.

Step 3: The user configures their brand in the LLM Optimizer user interface, adds prompts (i.e., questions they want LLM Optimizer to answer about their website and its visibility in LLMs), selects the LLM-based assistants to query, and connects the CDN used by their website to LLM Optimizer for data ingestion.

Step 4: The LLM Optimizer user interface sends configuration information to the LLM Optimizer Core Agent, which analyzes the information and generates additional suggested prompts.

Step 5: The LLM Optimizer sends the additional suggested prompts to Azure OpenAI, which returns category and topic information for the prompts.

Step 6: The user reviews and accepts or rejects the suggested additional prompts in the LLM Optimizer user interface. Once they are satisfied with the configuration, they submit the configuration back to the LLM Optimizer Core Agent.

Step 7: Based on the approved configuration from Step 6, the LLM Optimizer Core Agent begins collecting data from the selected LLM-based assistants for each configured prompt using the Data Retrieval Service. CDN log data is ingested directly into the Data Analysis Service, which processes the logs inside the Agentic Traffic pipeline to compute aggregated metrics (agentic traffic, referral traffic, crawl failures, etc.).

Step 8: The Data Retrieval Service forwards the collected data (answers from the AI chat bots) to the Data Analysis Service, which uses Azure OpenAI to determine if the brand was mentioned. Azure OpenAI returns the information to the Data Analysis Service.

Note: LLM Optimizer does not process or store Personally Identifiable Information (PII) that a customer has included in text prompts.

Step 9: In parallel, the Opportunity Detection Service sends historical prompt data where brand visibility is low or has recently declined to Azure OpenAI, which proposes content changes to the relevant pages that help improve or restore visibility. Additionally, the service sends information about technical errors resulting from LLM-based assistants' inability to access the customer's site to Azure OpenAI, which analyzes the errors and recommends fixes to resolve the errors.

Step 10: The Data Analysis Service and the Opportunity Detection Service send the results generated by Azure OpenAI to the LLM Optimizer Core Agent, which stores the data in the LLM Optimizer data store.

Step 11: The LLM Optimizer user interface retrieves the stored data from the LLM Optimizer data store and organizes it into the appropriate dashboards.

Step 12: The user reviews the alerts and recommendations in the Optimization Opportunities dashboard and accepts or rejects each recommendation.

Adobe LLM Optimizer and Azure OpenAI

Adobe LLM Optimizer currently leverages Azure OpenAI to organize suggested prompts by category and topic, detect opportunities for optimization, and generate suggestions for website improvements. For these activities, LLM Optimizer may pass the following data to Azure OpenAI:

- Content from customer's public website
- Configuration data, including brand name and prompts, as well as answers from the LLM-based assistants received by the Data Retrieval Service
- Website performance metrics
- Data from third-party services related to website performance and ranking
- Fragments of website markup and scripts

Adobe never transmits raw CDN logs—or any log-level data—to Azure OpenAI. Only public website content and de-personalized page-level data is used by Azure OpenAI for prompt/topic categorization.

Adobe has disabled logging in Azure OpenAI, helping ensure that Microsoft does not collect or store any data sent for processing to Azure OpenAI. More information is available at Azure OpenAI data privacy and security.)

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

Content filtering

Adobe leverages Azure OpenAI's content filtering service to moderate hate, sexual, violent, and self-harm content. The service uses Microsoft's collection of proprietary models for content filtering that has both contextual and semantic understanding of text. Adobe has configured the content filter to filter "medium" and "high" severity outputs from the model but not to filter any input.

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, please see the Generative AI Built for Business solution brief. A Security Testing Report for LLM Optimizer is expected to be available in the first half of 2026.

Data sources

LLM Optimizer uses several types of data sources as core inputs, depending on the specific customer configuration:

- **LLM response data (input to Brand Presence and URL Inspector dashboards)** – Calls public LLM APIs and crawls public LLM UIs, runs customer-defined prompts as well as autogenerated prompts from SEO signals, and analyzes the responses. Configured by the customer via categories, topics, and prompts in the LLM Optimizer UI.

- **Customer CDN logs (input to Agentic Traffic and Referral Traffic dashboards)** – Aggregates HTTP request logs from the customer's CDN[1] (e.g., Akamai, Fastly, and CloudFront). The logs enable LLM Optimizer to detect how often AI crawlers and browsers fetch the customer's pages and measure the traffic health (e.g., success rates, failed URLs, and failure location). Log forwarding is configured by the customer during onboarding. LLM Optimizer does not process or store PII at any stage of CDN log ingestion or subsequent data handling. See *CDN log details* below for the subset of fields collected by LLM Optimizer.

- **SEO and crawler insights (input to Optimization Opportunities dashboard)** – Uses a combination of SEO tools and proprietary crawlers in LLM Optimizer to understand how the customer's web content is structured and exposed to AI systems. Generates smart prompts and topics for brand presence analysis, detects "hidden" content that AI crawlers might not see, and sends technical gaps (including blocked AI bots, broken pages, and missing structured data). Configured by the customer by defining domains in scope.

[1] If the customer stores their CDN logs in an observability system, such as Datadog, Splunk, or Grafana, they must manually forward the logs to LLM Optimizer. LLM Optimizer cannot ingest logs directly from those platforms.

- **Web analytics platforms, including Adobe Analytics, Adobe Customer Journey Analytics, and Google Analytics (input to Referral Traffic dashboard)** – Ingests aggregated metrics per URL and referrer post-consent (non-anonymized), such as page views, entries, and bounces; orders, revenue, and conversion indicators; and device and geography breakdowns. This information shows how AI-driven visits perform compared to other sources, quantifies the business impact of LLM visibility, determines conversions and revenue attributed to LLM referrals as well as funnel performance from AI-driven landing pages. Configured by the customer via secured exports and APIs.

- **AEM operational telemetry[2] (input to Referral Traffic dashboard)** – Ingests limited samples of metrics per URL, such as visits, bounce rate, consent rate, device, region, channel, and explicit LLM channel classifications. By design, AEM optel preserves the privacy of end-users by sampling and anonymizing data, giving an approximation of overall traffic. This information helps LLM Optimizer estimate how referral traffic is engaging with the brand after landing on the website by checking if the visitor bounced and if a cookie consent banner was displayed. Automatically configured if the customer is using Adobe Experience Manager as a Cloud Service. For more information, see Experience League.

CDN log details

LLM Optimizer only ingests a restricted, non-PII subset of fields from CDN logs. Although raw log field names vary by CDN provider, they are normalized and presented as:

- URL (path only)
- User agent
- Status code
- Referrer header
- Host header
- Time to first byte (ttfb)
- Request method
- Timestamp
- Content type

These normalized fields are exposed through the agentic view. On the Referral Traffic dashboard, CDN logs are utilized to display page hit metrics.

CDN log analysis is done using Adobe's own ingestion and analytics services and only aggregated metrics are retained. Logs are deleted after processing.

Data usage

LLM Optimizer does not use any customer data to train or fine-tune LLMs.

Data refresh

By default, Brand Presence data is refreshed on a weekly basis. During onboarding, the customer may request daily data refresh instead.[3]

[2] If the customer deploys their site using Adobe Experience Manager as a Cloud Service or AEM Sites.
[3] Additional fees may apply.

Data is also refreshed when the customer makes the following configuration changes:

- Addition of new prompts – Refresh will occur in the next weekly (or daily) cycle
- Deletion of category or topic – Related data will be removed within 24 hours
- Addition/deletion of brand alias/aliases – Historical data will be reprocessed within 24 hours

Data processing and storage locations

Adobe hosts all components of LLM Optimizer in AWS US-East-1.

Microsoft hosts the Azure OpenAI Service in Microsoft Azure data centers. LLM Optimizer sends all data for processing by Azure OpenAI to the instance in the Azure US-East location. No data is stored in Azure OpenAI.

Data retention

Adobe stores the following customer-specific data generated by LLM Optimizer in the LLM Optimizer data store:

- Brand configuration and execution data, including mentions and citations, is stored until deletion by the customer or upon customer subscription cancellation
- Agentic traffic and referral traffic analysis is stored until customer subscription cancellation
- Identified opportunities and automated improvement suggestions are stored until deletion by the customer or upon customer subscription cancellation

Questions?

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