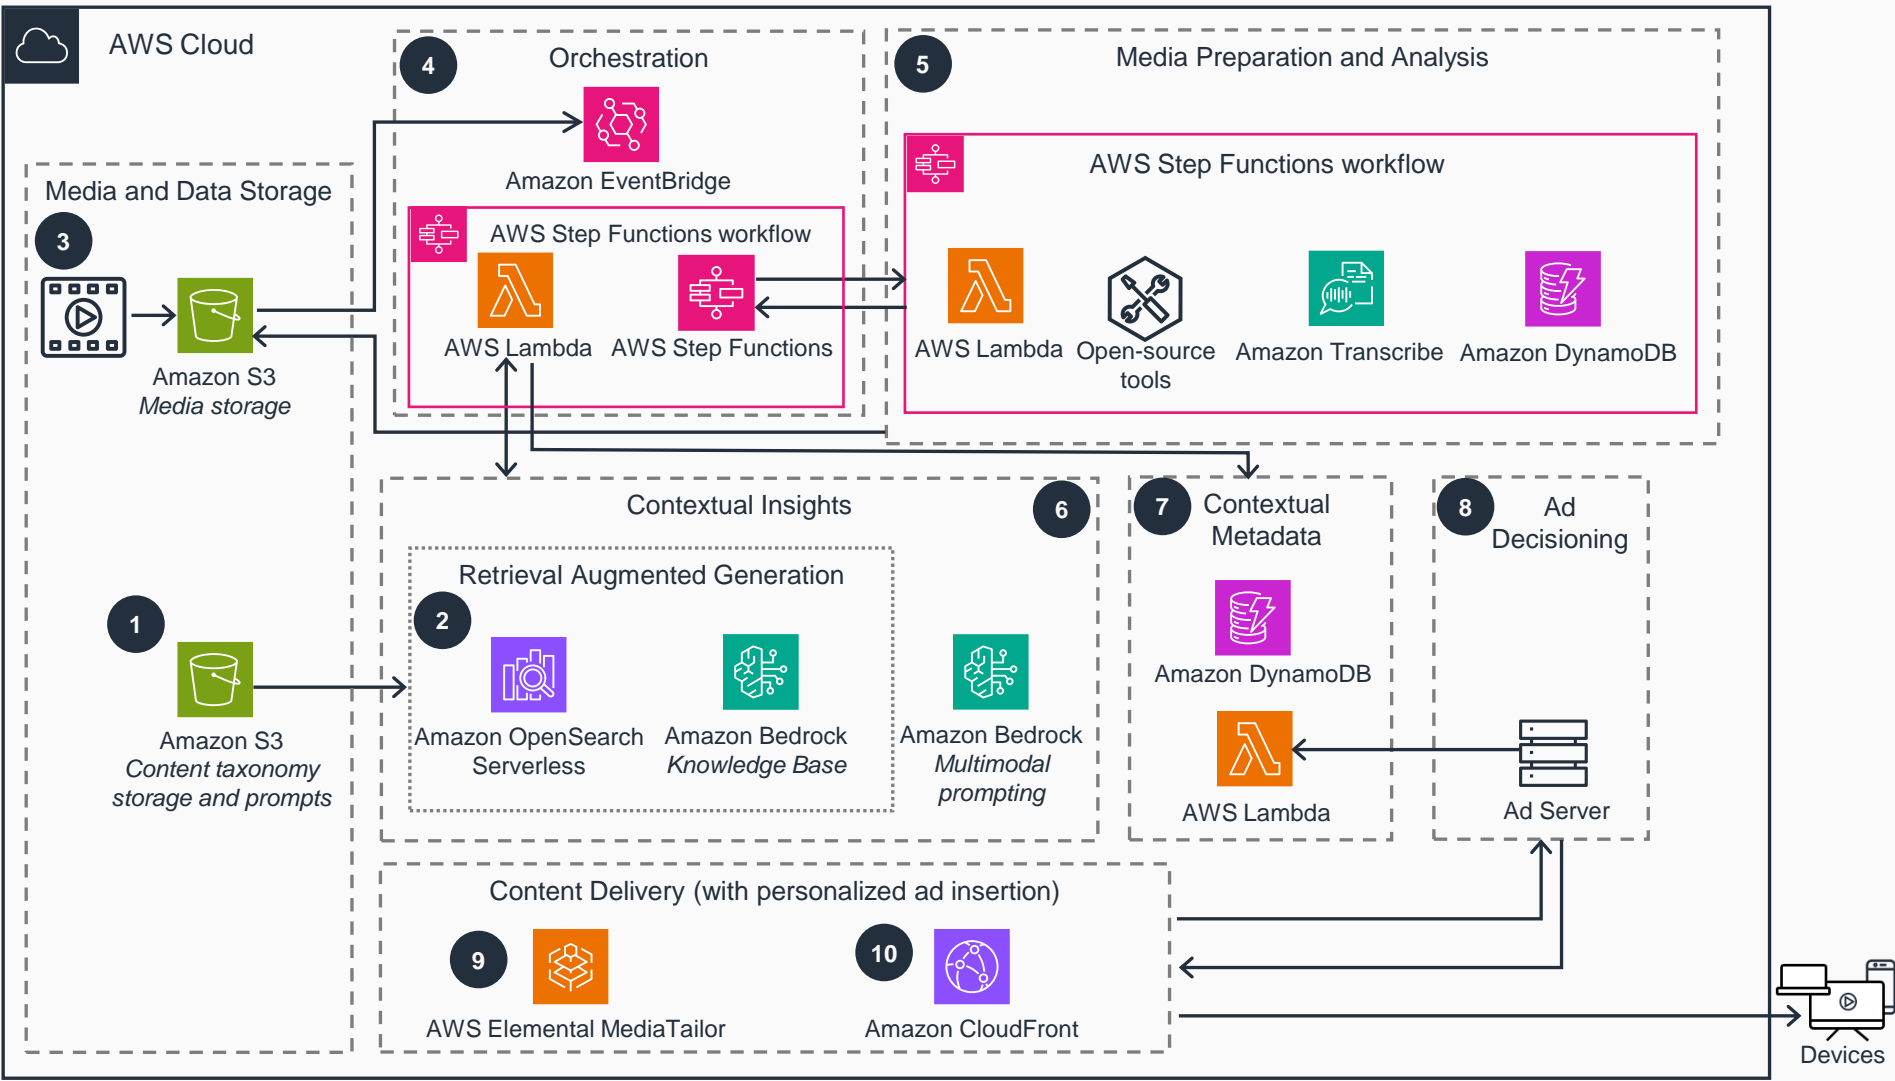


Guidance for Contextual Intelligence Advertising Using Generative AI on AWS

This architecture diagram shows how to use generative AI to derive contextual insights from multimedia assets and identify relevant audience segments for ad placements. It uses multimodal large language models (LLMs) to extract insights from visuals and transcripts, and a vector database with industry taxonomy to monetize content. The diagram includes 10 steps, with 1-5 shown here and steps 6-10 shown on the next slide.

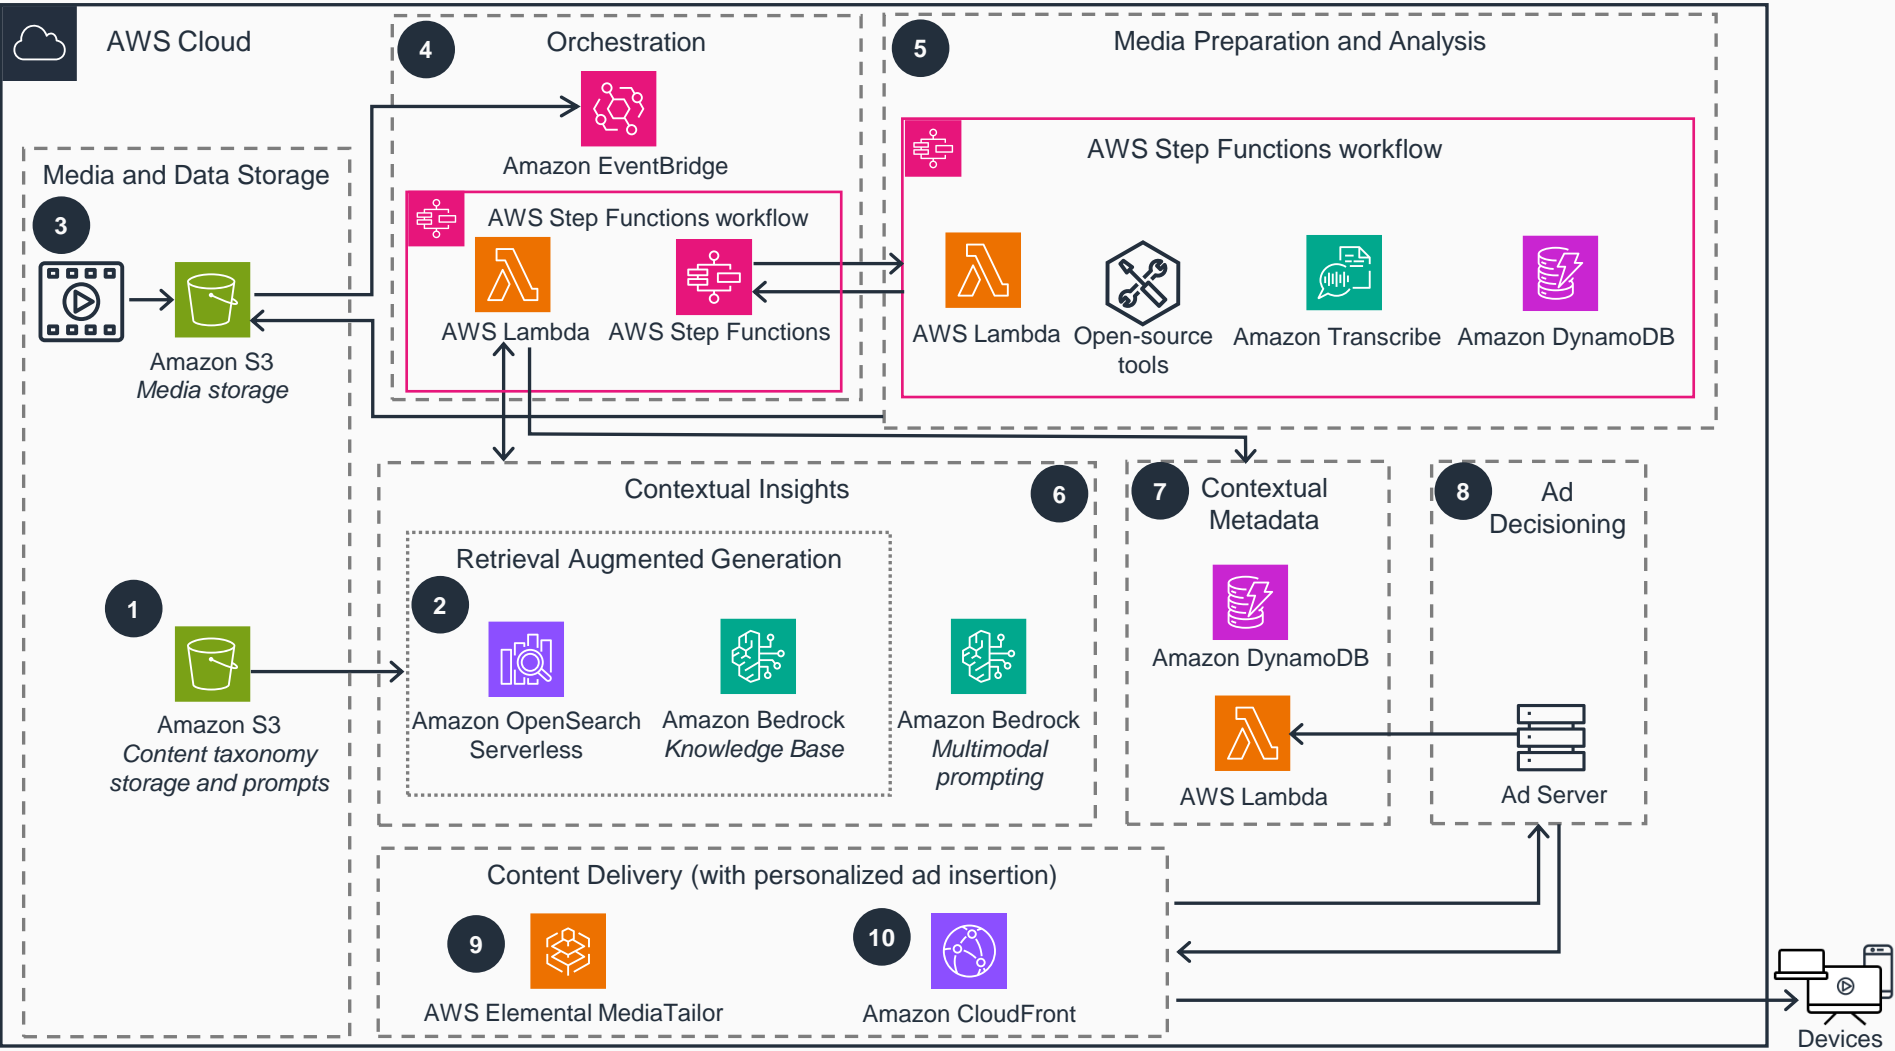


- 1 Upload the latest version of the IAB Content Taxonomy to **Amazon Simple Storage Service (Amazon S3)**.
- 2 Create a knowledge base in **Amazon Bedrock**, a fully managed service that offers a choice of high-performing foundation models (FMs). Select **Amazon OpenSearch Serverless** as the vector database and **Amazon Titan Text Embeddings v2 model** as the embeddings model.
- 3 Upload media content to an **Amazon S3** bucket.
- 4 **Amazon EventBridge** receives object creation notifications from **Amazon S3**. This triggers an orchestration workflow that executes **AWS Step Functions** for media preparation and analysis, as well as **AWS Lambda** functions to invoke **Amazon Bedrock** for contextual insights extraction.
- 5 The **Step Functions** workflow executes media preparation and analysis tasks, invoking **Lambda** functions that use open-source tools like ffmpeg and perceptual hashing, generating transcriptions using **Amazon Transcribe**. The workflow metadata is stored in **Amazon DynamoDB**, and the processed media files are persisted in **Amazon S3**.



Guidance for Contextual Intelligence Advertising Using Generative AI on AWS

Steps 6-10



- 6 The **Lambda** function invokes the multimodal large language models (LLMs) in **Amazon Bedrock** to extract contextual insights. It uses the Knowledge Bases for Bedrock to map the content to the IAB content taxonomy, utilizing a managed Retrieval Augmented Generation (RAG) search pattern.
- 7 The **Lambda** function persists the contextual insights extracted from **Amazon Bedrock** into a **DynamoDB** table, storing the contextual metadata.
- 8 When an ad request is received, the ad server invokes a **Lambda** function to fetch the IAB categories and other relevant contextual metadata from the **DynamoDB** store. The ad server then uses the retrieved contextual data to select the most relevant advertisement.
- 9 **AWS Elemental MediaTailor** retrieves the relevant advertisement from the ad server, performs ad stitching, and serves a digital Video Ad Serving Template (VAST) manifest file containing the video content and the associated advertisement.
- 10 **Amazon CloudFront** delivers the content, including the contextually relevant advertisements, to the end-user devices.

