

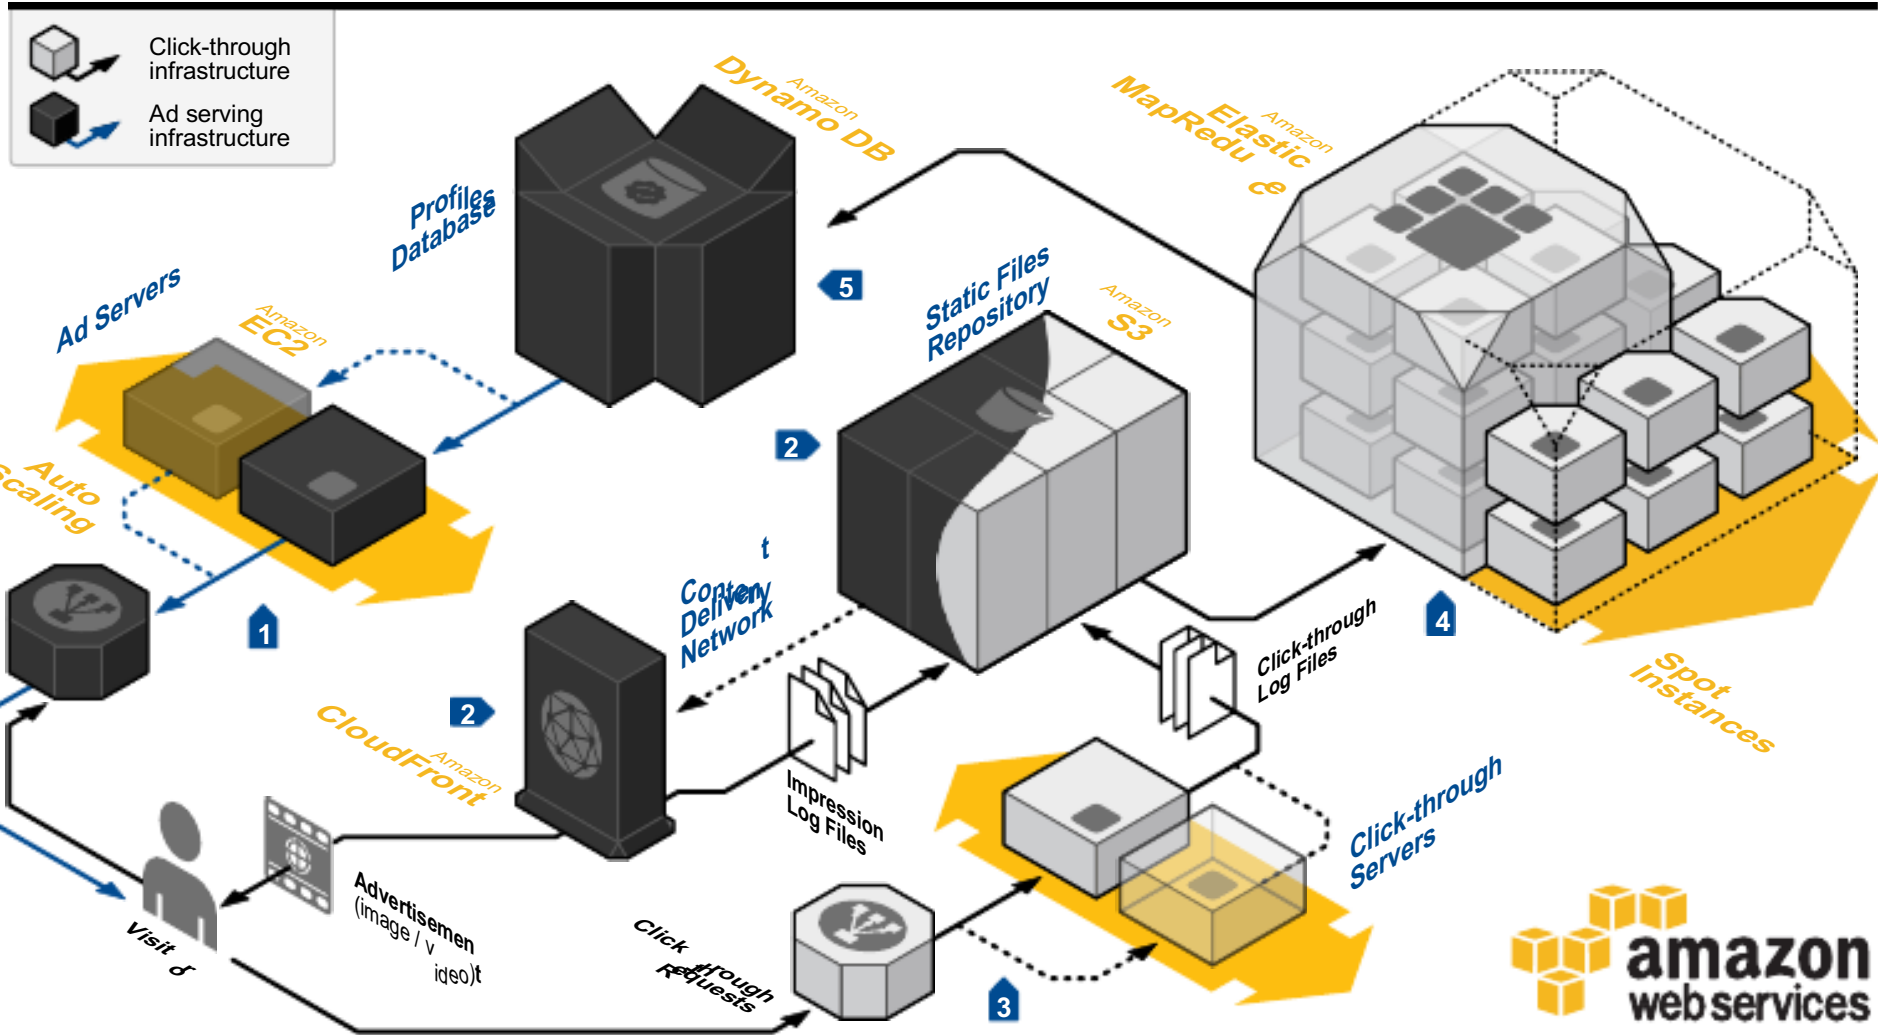
ADVERTISEMENT SERVING

Internet advertising services need to serve targeted advertising and must do so under limited time. These are just two of multiple technical challenges they face.

Amazon Web Services provides services and infrastructure to build reliable, fault-tolerant, and highly available ad serving platforms in the cloud. In this document, we describe the two main parts of such a system: ad serving infrastructure and click-through collection featuring a data analysis cluster.

AWS Reference Architectures

- Amazon EC2
- Auto Scaling
- Elastic Load Balancing
- Amazon CloudFront
- Amazon EC2 Spot
- Amazon EMR
- Amazon S3
- Amazon Import/Export
- Amazon DynamoDB



System Overview

1 When visitors load a web page, ad servers return a pointer to the ad resource to be displayed. These servers are running on **Amazon Elastic Compute Cloud** (Amazon EC2) instances. They query a data set stored in an **Amazon DynamoDB** table to find relevant ads depending on the user's profile.

2 Ad files are downloaded from **Amazon CloudFront**, a content delivery service with low latency, high data-transfer speeds, and no commitments. Log information from displayed ads is stored on **Amazon Simple Storage Service** (Amazon S3), a highly available data store.

3 The click-through servers are a group of **Amazon EC2** instances dedicated to collecting click-through data. This information is contained in the log files of the click-through web servers, which are periodically uploaded to **Amazon S3**.

4 Ad impression and click-through data are retrieved and processed by an **Amazon Elastic MapReduce** cluster using a hosted Hadoop framework to process the data in a parallel job flow. The cluster's capacity can be dynamically extended using **Spot Instances** to reduce the processing time and the cost of running the job flow.

5 Data processing results are pushed back into **Amazon DynamoDB**, a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. **Amazon DynamoDB** tables can store and retrieve any amount of data, and serve any level of request traffic, both of which are specific requirements for storing and quickly retrieving visitors' profile information.

The high availability and fast performance of **Amazon DynamoDB** enable ad server front-ends to serve requests with predictable response time, even with high traffic volumes or large profile's data sets.