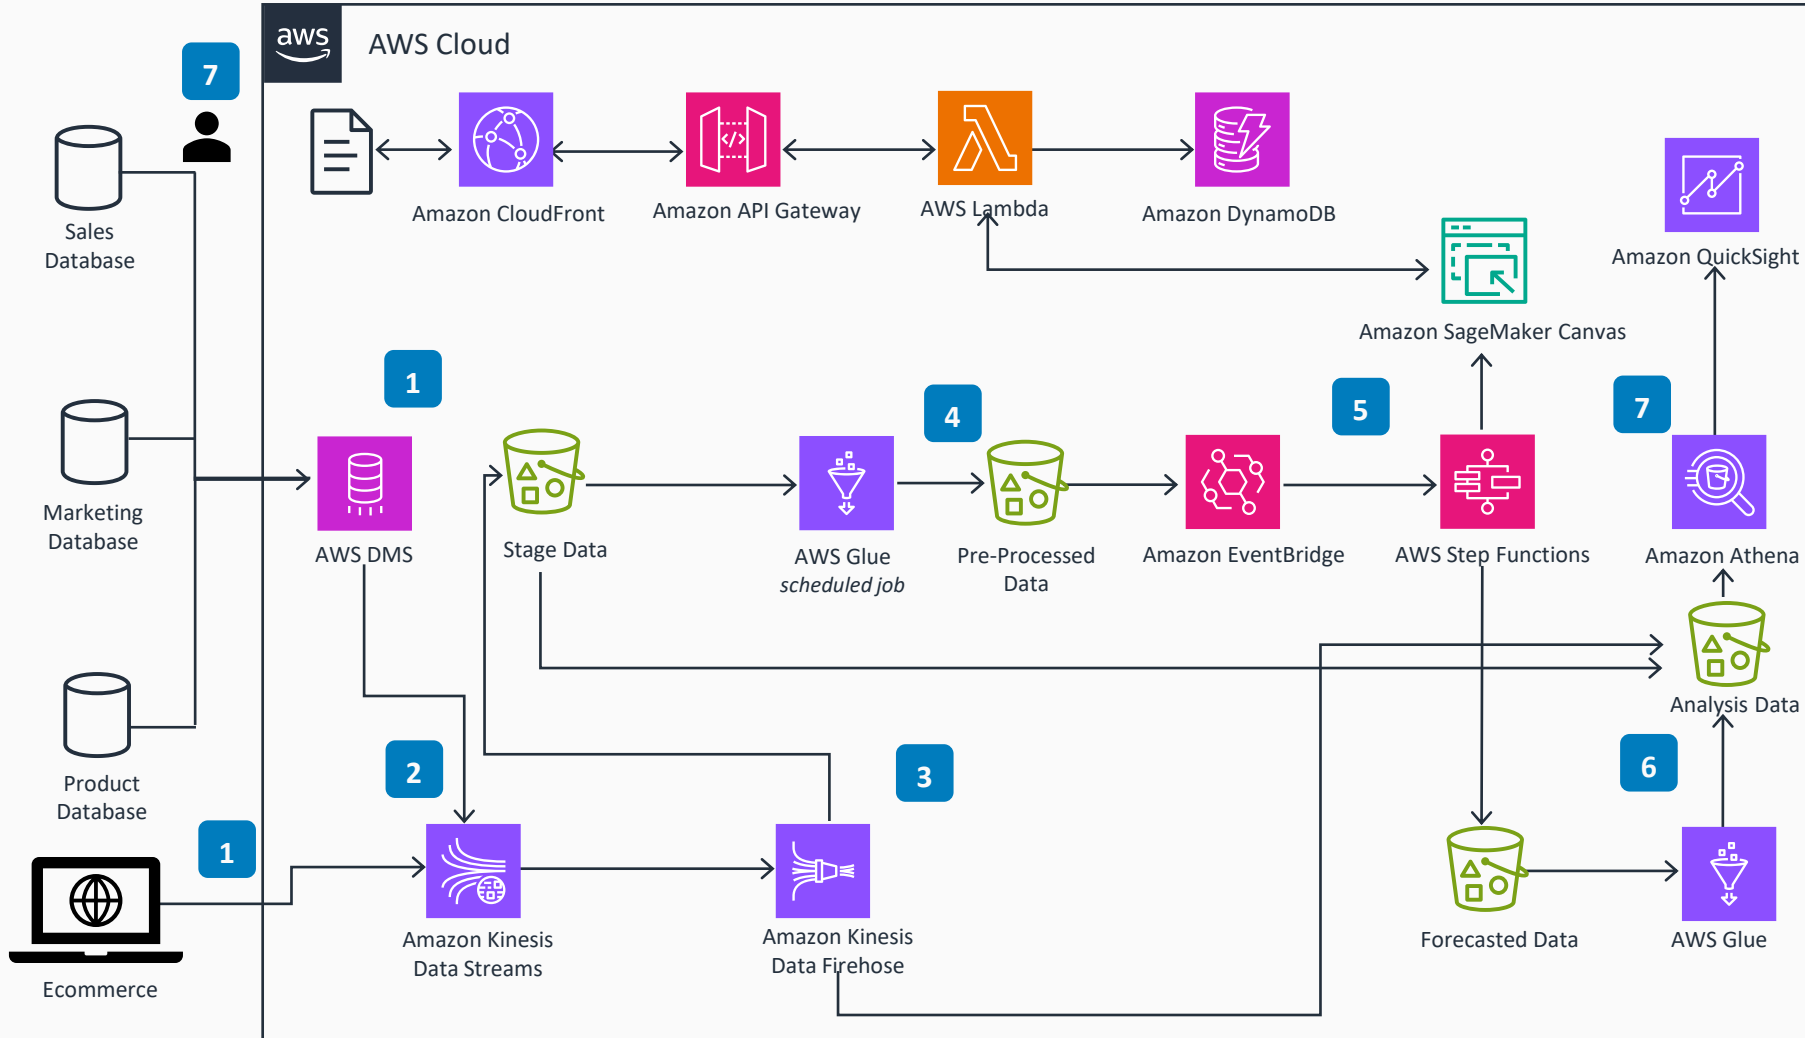


# Guidance for Demand Forecasting & Planning on AWS

This architecture shows how data is collected from your databases and ecommerce sites, processed, and exported to a dashboard for visibility and forecasting.



- 1** **AWS Database Migration Service (AWS DMS)** will connect with your databases using change data capture (CDC). Data about sales, products, and marketing campaigns is collected and sent to **Amazon Kinesis Data Streams**. Optionally, you can send data from your e-commerce site to **Kinesis Data Streams** directly.
- 2** **Amazon Managed Service for Apache Flink** identifies sales with anomalous behavior and sends to a data stream for outliers.
- 3** **Amazon Kinesis Data Firehose** sends data to an **Amazon Simple Storage Service (Amazon S3)** staged data bucket in parquet format. Items that were in the data stream as outliers are sent to this analysis bucket and are consumed by **Amazon QuickSight** with alerts configured.
- 4** A scheduled **AWS Glue** job runs Extract, Transform, and Load (ETL) to convert the staged data and the pre-processed data into the forecast format. It prepares the target time series, related time series, and item metadata datasets.
- 5** An **Amazon S3** event notification is sent to an **Amazon EventBridge**, invokes a rule, and calls **AWS Step Functions**. It creates datasets, a dataset group, and a predictor. It also analyzes metrics, creates a forecast, and exports results using **Amazon SageMaker Canvas**.
- 6** An **Amazon S3** event notification invokes an **AWS Glue** job. This enriches the exported forecast with data from the stage data bucket, and prepares it for **Amazon Athena**.
- 7** A dashboard is available in **QuickSight** with forecasted values. Optionally, a webpage will call forecasts for a product and store it in **Amazon DynamoDB**.