Guidance for Demand Forecasting for Restaurants on AWS

This reference architecture showcases an end-to-end pipeline to deliver restaurant order demand forecasts in a data format that non-technical users can update and consume.

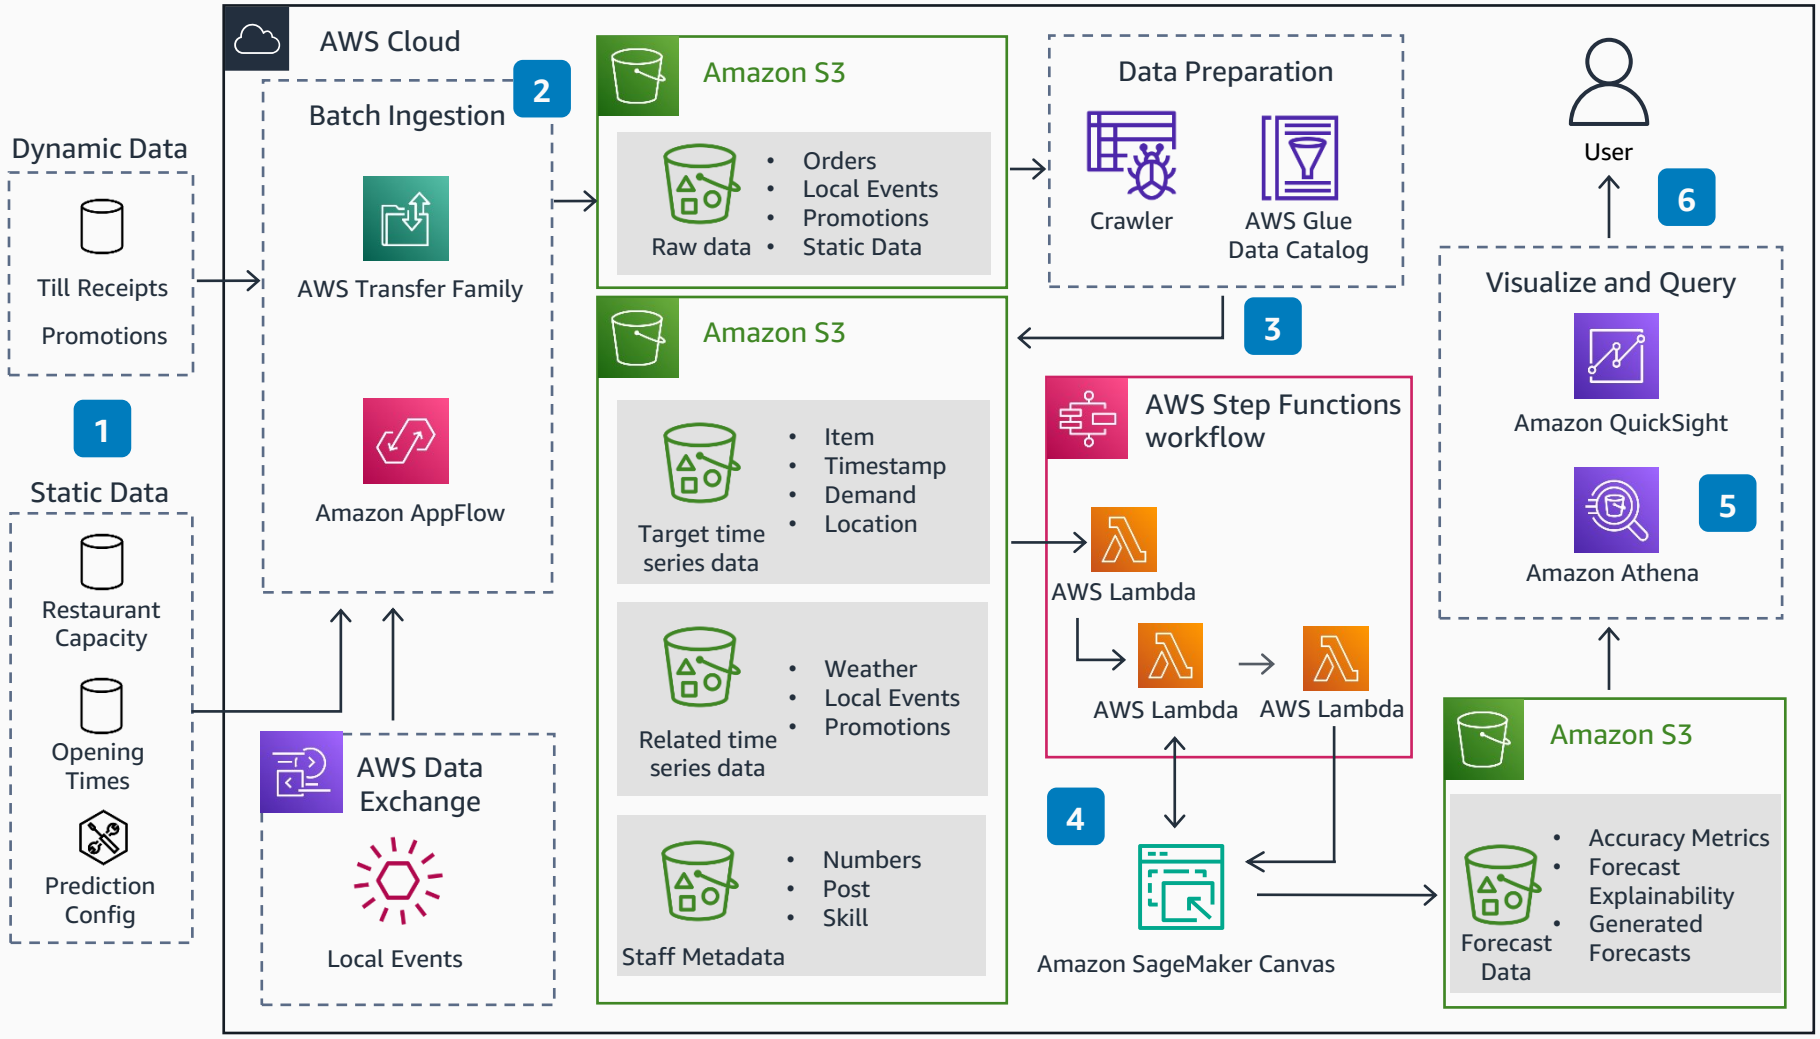

1. Export dynamic and static data from front of house systems in CSV or Excel format or to other restaurant management systems.

2. Upload structured files with pre-defined schemas to **Amazon Simple Storage Service (Amazon S3)** using **AWS Transfer Family**. You can access optional data through **AWS Data Exchange** (for local events) or **Amazon AppFlow** (for promotions).

3. **AWS Glue Crawler** creates an **AWS Glue Data Catalog** after transforming the raw data. The **AWS Glue Data Catalog** puts data in a format that **Amazon SageMaker Canvas** can consume.

4. Using **AWS Lambda**, **SageMaker Canvas** and **AWS Step Functions** build and train a **SageMaker Canvas** model known as a predictor:

   a) **Step Functions** orchestrates the ingestion of processed data for **SageMaker Canvas**.

   b) Processed data helps train the **SageMaker Canvas** predictor.

   c) The predictor then makes predictions about a pre-configured future period and publishes the outputs to **Amazon S3.**

5. **Amazon Athena** queries the final forecast data and creates reports that **Amazon QuickSight** accesses.

6. **QuickSight** uses the forecasted data to visualize and create analytics dashboards for non-technical users.

AWS Reference Architecture