Guidance for Personalized Engagement Using Online & Mobile User Behaviors on AWS

This architecture diagram shows how RudderStack Open Source and AWS services can help you collect and store user behavior data from various sources, such as mobile apps and websites, and make it available for personalized engagement.

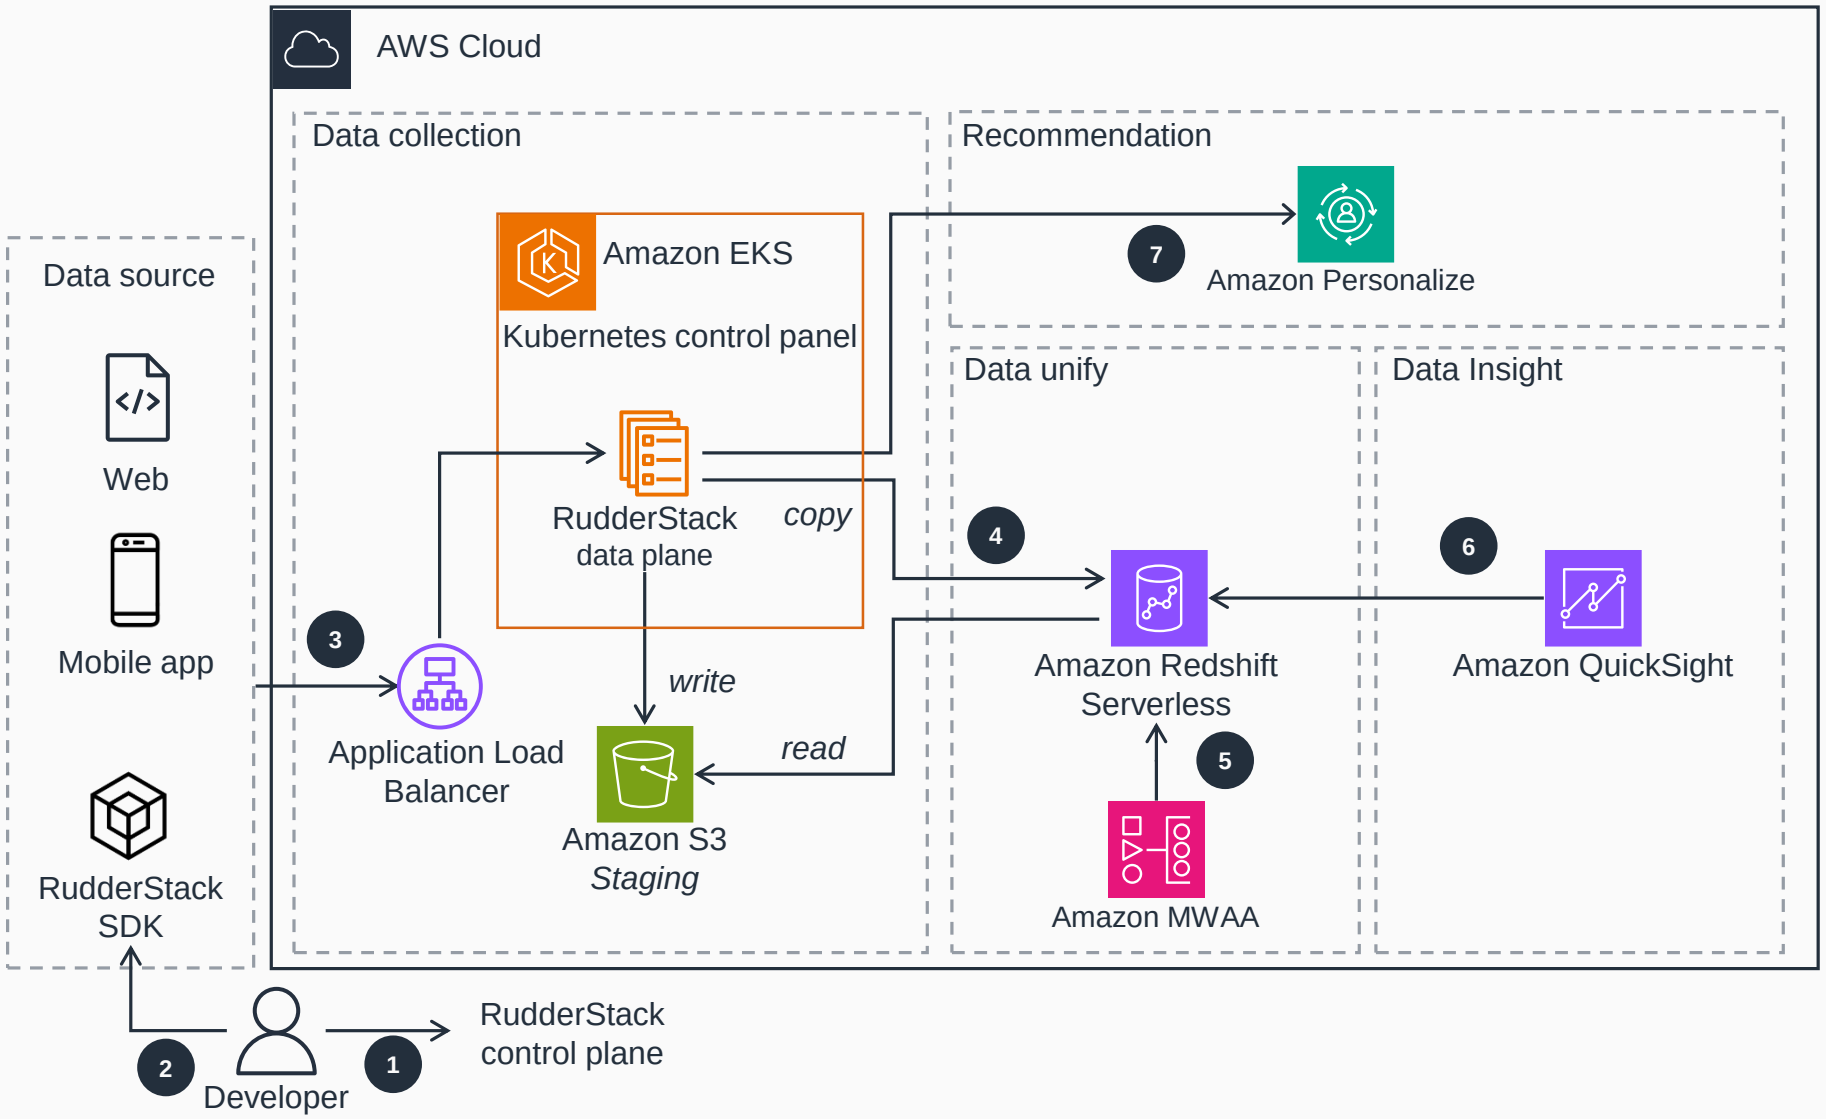

1. Developers configure data sources (such as websites and mobile applications), data destinations, like **Amazon Redshift**, and connections in the RudderStack control plane.

2. Developers use the software development kit (SDK) provided by RudderStack's event stream to develop tracking for data sources.

3. The SDK sends tracking events to **Application Load Balancer**, which then enters the RudderStack data plane deployed on **Amazon Elastic Kubernetes Service (Amazon EKS).** The data plane writes the events to the **Amazon Simple Storage Service (Amazon S3)** staging bucket.

4. The RudderStack data plane periodically sends copy commands, data merge SQL, and data definition language (DDL) to **Amazon Redshift Serverless**, importing the event data files from the **Amazon S3** staging bucket into **Amazon Redshift** tables.

5. Using **Amazon Redshift Serverless**, the event table is processed according to analysis requirements to create user behavior analysis detail tables, summary tables, and user profile tables. Use **Amazon Managed Workflows for Apache Airflow (Amazon MWAA)** for task scheduling.

6. Use **Amazon QuickSight** to create dashboards like user behavior analysis, web attribution reports, and funnel analysis, with the data source being the summary level tables read through **Amazon Redshift Serverless**.

7. The interaction data between users and items are sent in real-time as events from the RudderStack data plane to **Amazon Personalize**. Based on different recommendation algorithms, corresponding recommendation results are generated.

Reviewed for technical accuracy February 2, 2024

AWS Reference Architecture