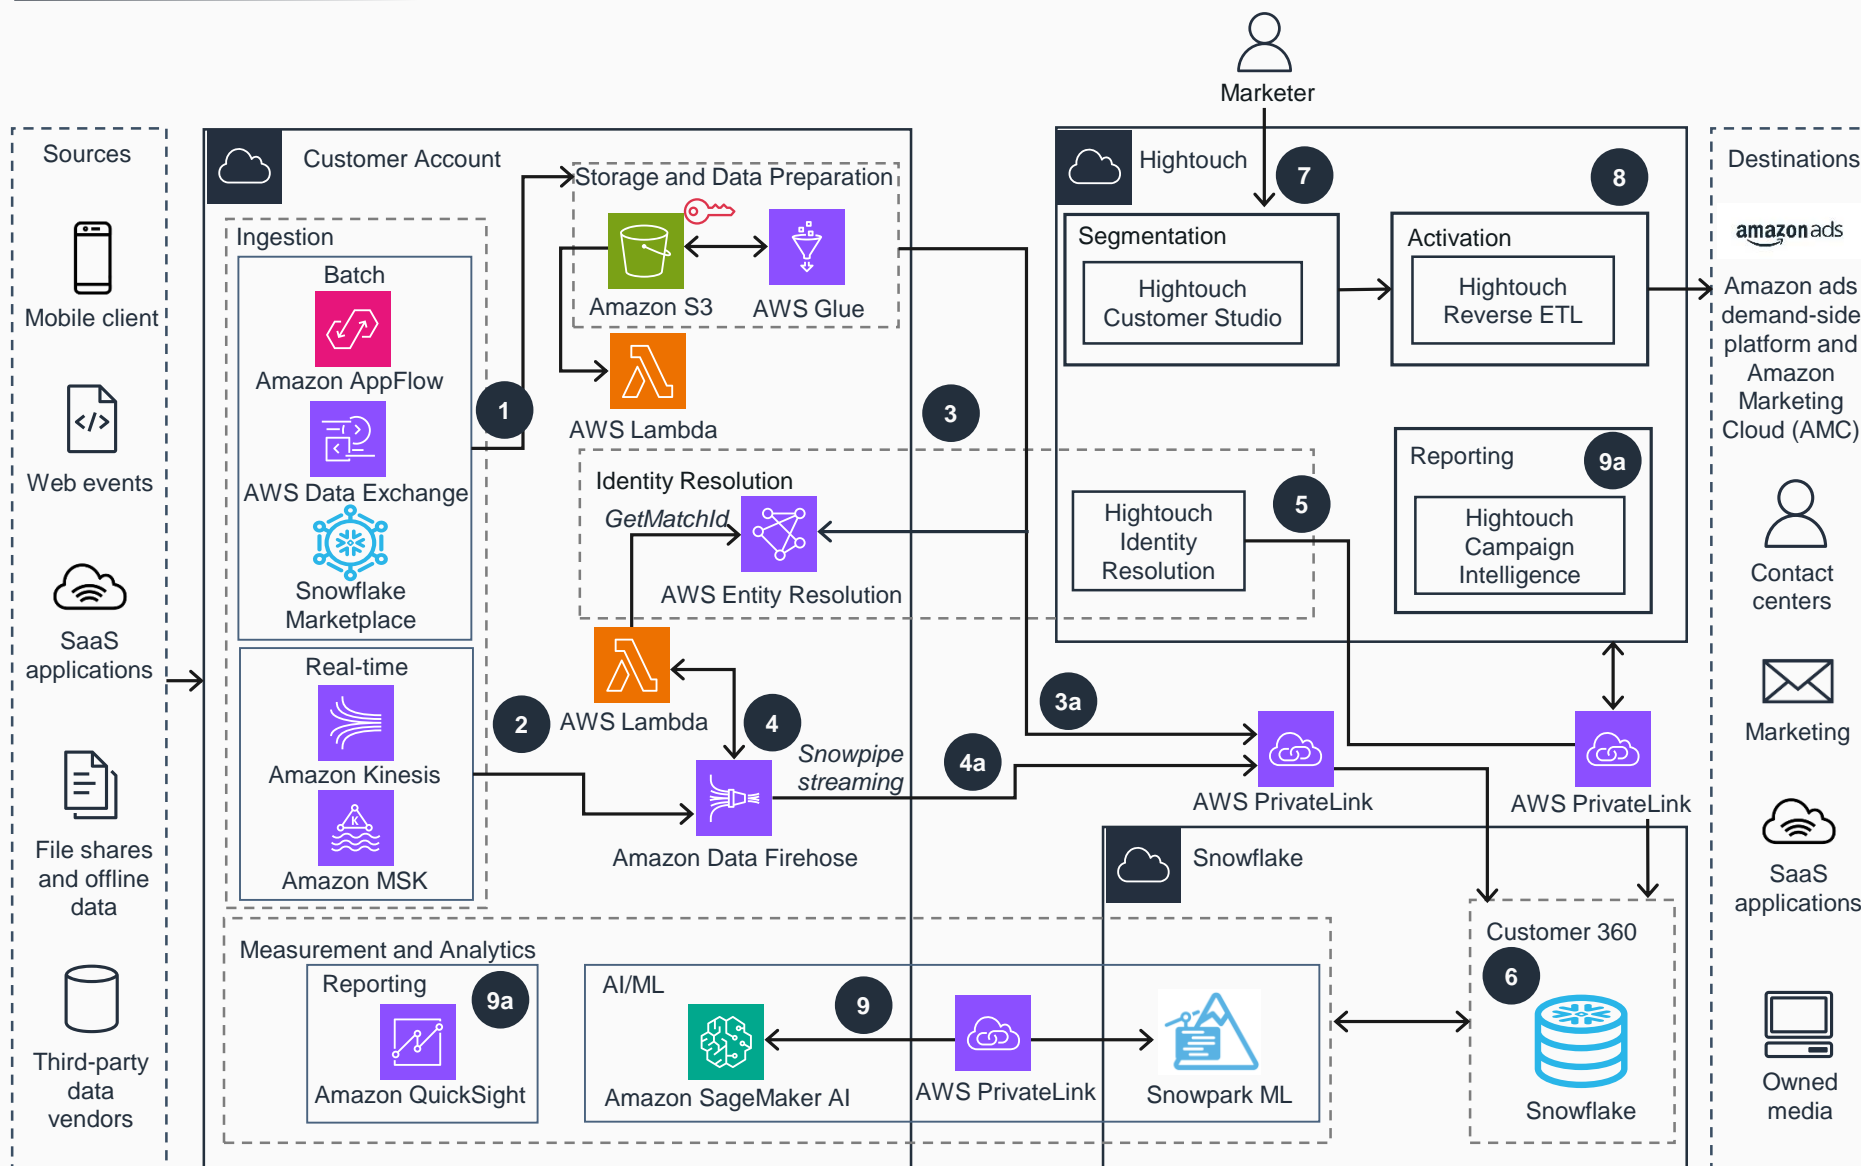# Guidance for Building a Composable Customer Data Platform on AWS

This architecture diagram uses customer data on Snowflake and Hightouch capabilities to solve your marketing needs, helping you adapt to marketing changes, integrate necessary components, and secure data. This slide shows Steps 1-4a.


AWS Reference Architecture

1 Batch data is ingested from various SaaS platforms using available **Amazon Appflow** connectors. Third party data insights are ingested using **AWS Data Exchange**. Ingested batch data is stored in **Amazon Simple Storage Service (Amazon S3)** and transformed using an **Amazon S3** trigger that invokes **AWS Lambda** function. **AWS Glue** can help with data preparation and quality checks. You can also use Snowflake Marketplace to bring third party insights directly into Snowflake tables.

2 Real-time data is ingested using **Amazon Kinesis Data Streams** or **Amazon Managed Streaming for Apache Kafka** (Amazon MSK). Data streams are then sent to **Amazon Data Firehose** which can do a Direct PUT to Snowflake in seconds.

3 To resolve identities of ingested data, you can choose between **AWS Entity Resolution** or Hightouch Identity Resolution based on your use case. Use Guidance for Preparing and Validating Records for Entity Resolution on AWS to resolve identities of batched transformed data in **Amazon S3** using AWS Entity Resolution matching workflows.

3a Transformed and resolved data is loaded from Amazon S3 into Snowflake over **AWS PrivateLink**.

4 Stream data through **Firehose** to invoke a **Lambda** function and transform data. Additionally, invoke **AWS Entity Resolution** GetMatchId API to find matches for the given data record.

4a The transformed and matched data is then sent to Snowflake using Snowpipe Streaming to build Customer 360 over an Amazon private network using **PrivateLink**.

Guidance for Building a Composable Customer Data Platform on AWS

This architecture diagram uses customer data on Snowflake and Hightouch capabilities to solve your marketing needs, helping you adapt to marketing changes, integrate necessary components, and secure data. This slide shows Steps 5-9a.
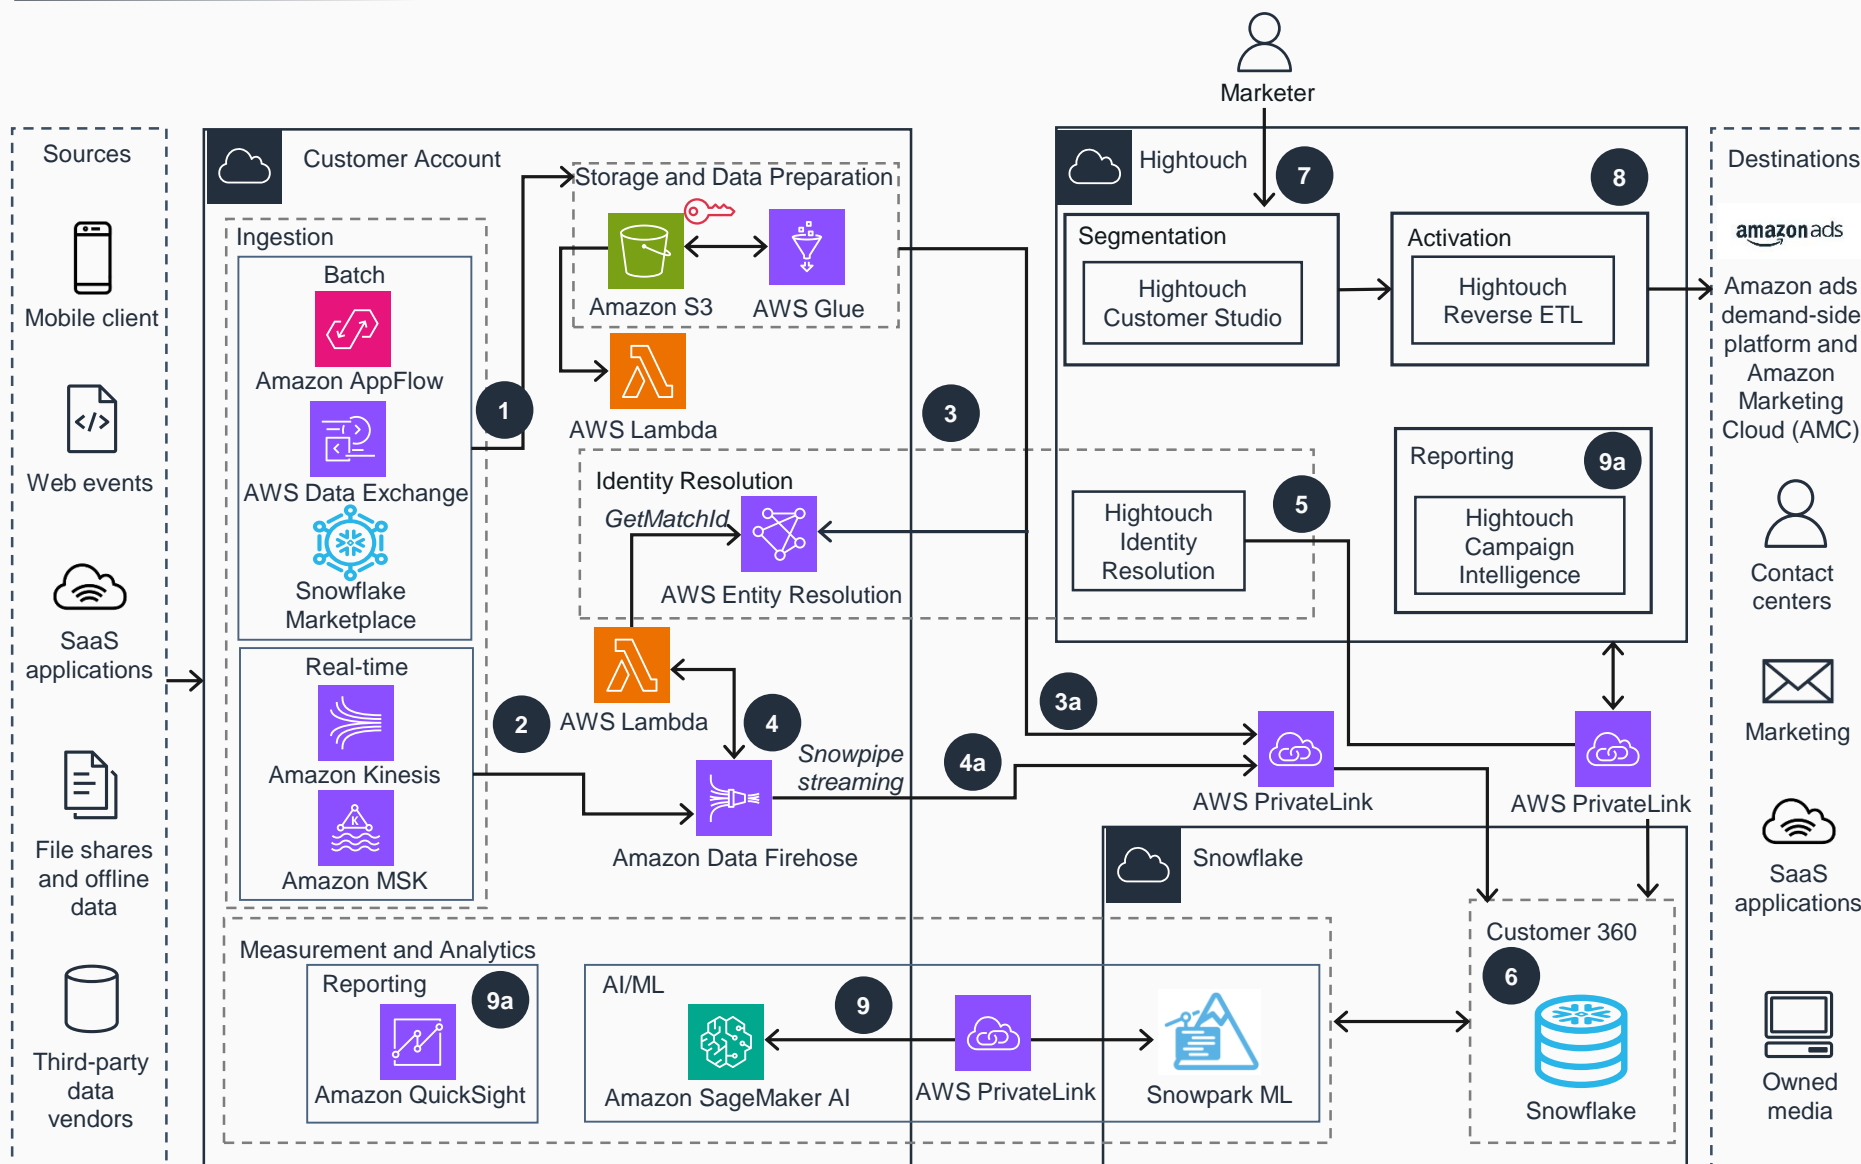
5 Alternatively, use Hightouch Identity Resolution to stitch your data to build customer profiles directly from data in Snowflake tables. The identity graph table is stored back in Snowflake. All traffic between Hightouch and Snowflake is securely managed using **PrivateLink**.

6 Build a Customer 360 view once you centralize all your customer data in Snowflake tables to launch personalized customer experiences across various channels.

7 Marketers can use Hightouch Customer Studio to build, manage, and analyze audiences with no SQL knowledge required. All traffic between Hightouch and Snowflake goes over **PrivateLink**.

8 Create sync pipelines using Hightouch Reverse ETL to activate your audiences to various destinations.

9 For any AI/ML needs, you can train and deploy models using **Amazon SageMaker AI** and Snowpark ML.

9a Use **Amazon QuickSight** or Hightouch Campaign Intelligence to analyze data from various marketing campaigns.

AWS Reference Architecture