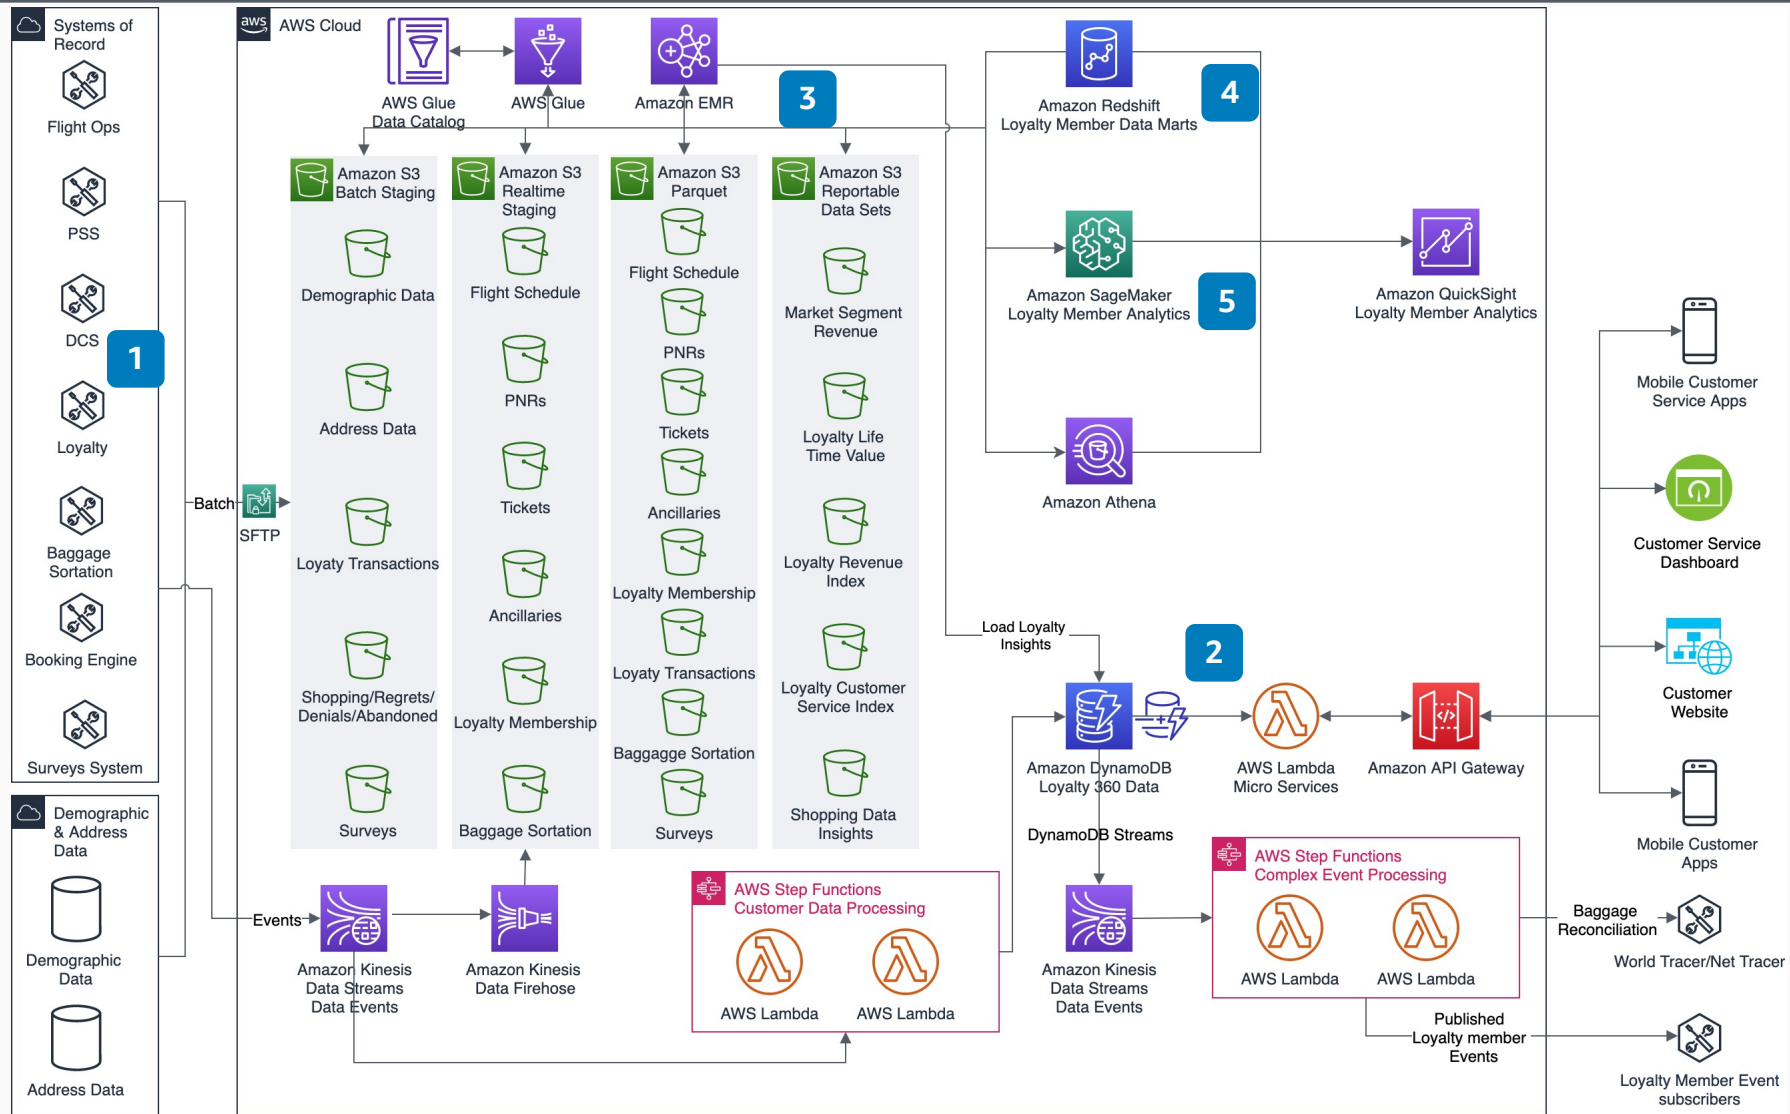


Airline Data Platform

Build a data platform to serve both operational and analytics needs using open data standards, purpose-built databases, and extensible serverless architecture.



Airline initiatives to build operations data stores and related services typically do not adapt to change because of rigid schemas and long implementation times. In addition, siloed systems for operations and analytics do not work well together, and building these systems on-premises prevents scaling to add new domains. This data platform architecture helps relieve and eventually replace the on-premises data platform load leading to cost savings and an agile environment.

- 1** Build basic data as a service with the most important domains: Flight, Passenger, and Loyalty with the key tenet of separating storage from compute.
- 2** Use purpose-built databases and serverless architecture to deliver microservices and events for operational data store. These serverless and fully managed services allow customers to scale based on adoption. This architecture enables customers to replace expensive on-premises operational databases, service-oriented architecture (SOA) infrastructure, and message-oriented middleware.
- 3** Use open standards to build the data lake using the same data as the operational data platform. Use a read pattern schema to make the **raw data** and **curated data** readily available for all user roles seamlessly.
- 4** Build standard enterprise data warehouse schemas and data marts in **Amazon Redshift** for known and well used usage patterns. For ad-hoc requirements, publish the data catalog and use **Amazon Athena** for analysis directly using the data lake. Customers can also extend the data warehouse based on their specific needs.
- 5** Use **Amazon SageMaker** to provide standard AI/ML models for customer segmentations and lifetime value. Customers can also use **Amazon SageMaker** to build their own models on top of the data.

