Customer Data Platform for Airlines

Airline initiatives to build operations data stores and related services often don't adapt to change due to rigid schemas, long implementation times, siloed operations and analytics systems, and on-premises scaling limitations to add new domains. Use this data platform architecture to relieve or replace your on-premises data platform load, increasing development agility and cost savings; it uses the Implementing Travel & Hospitality Data Mesh reference architecture as its foundation for a domain-owned design approach, open data standards, purpose-built databases, and extensible serverless architecture.

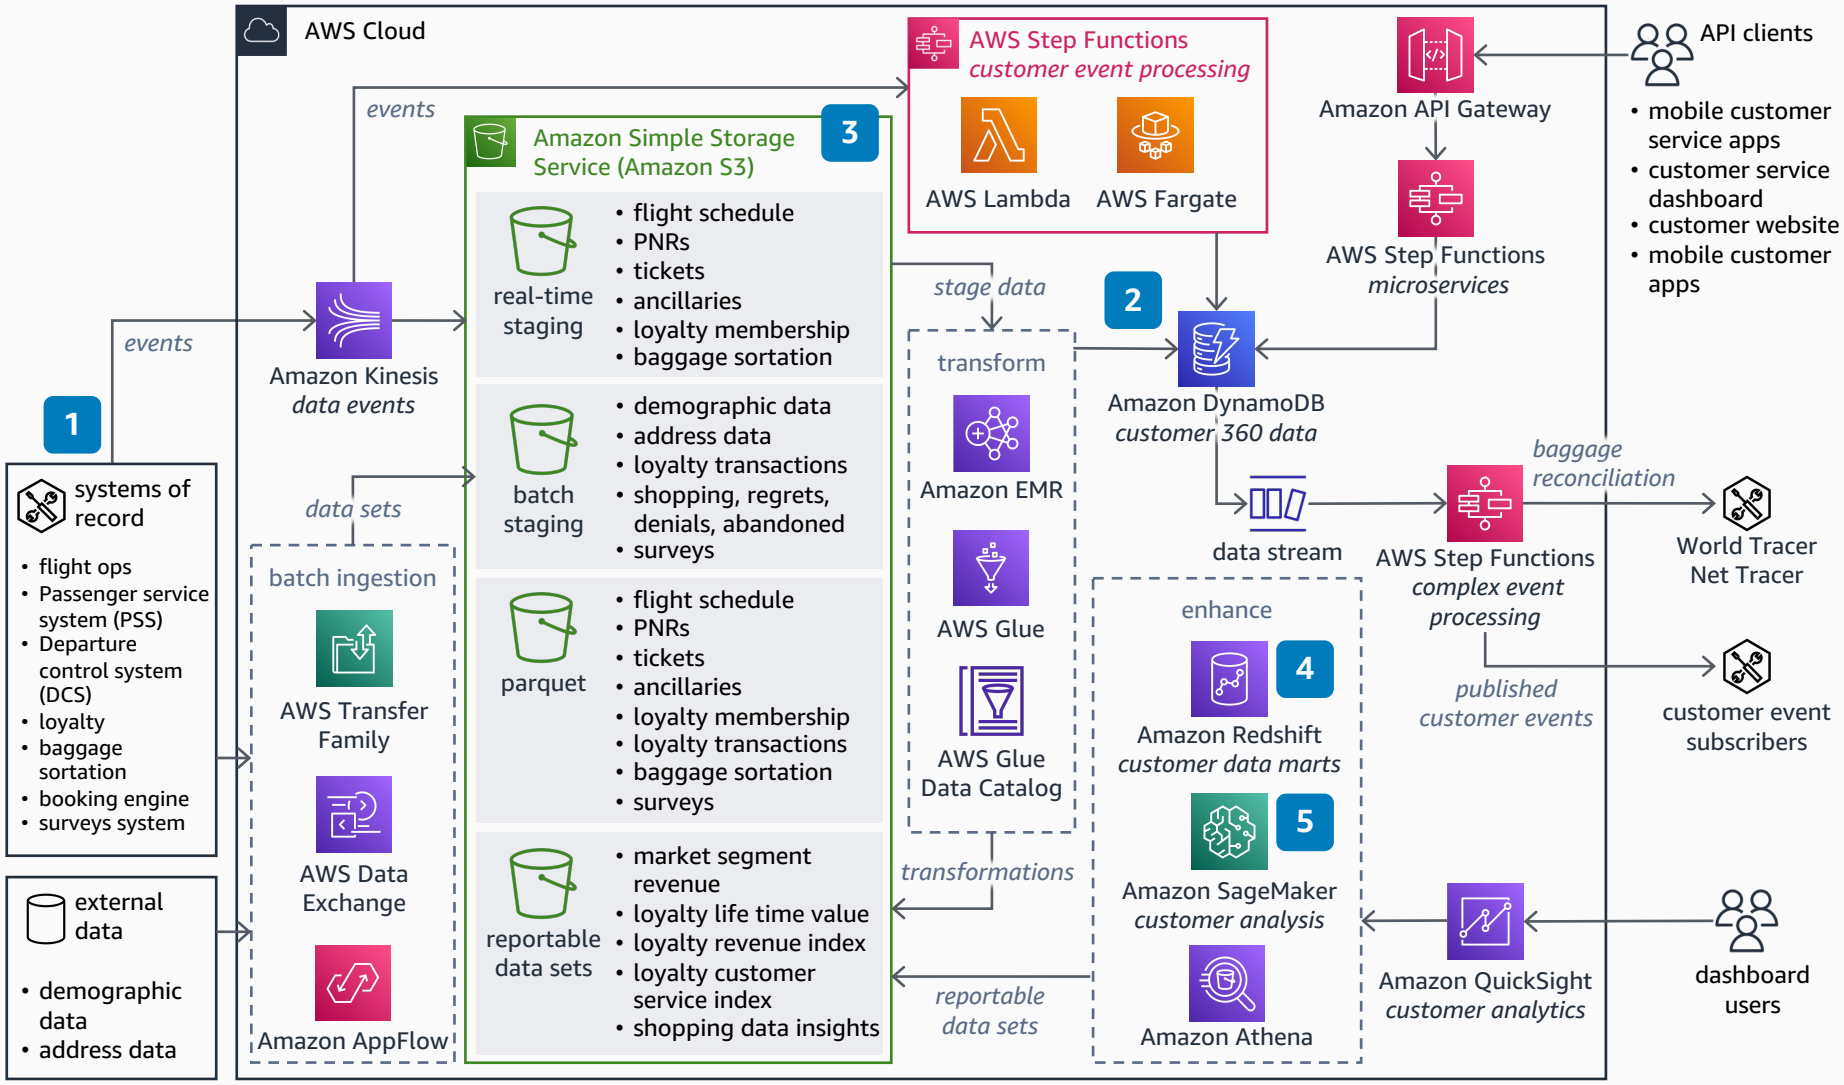

1. Build data products for relevant domains (like flight, passenger, and loyalty), separating storage from compute.
2. In the operational data store, use managed services and purpose-built databases with microservice and event-driven patterns. This enables you to deprecate expensive on-premises infrastructure (like operational databases, service-oriented architecture (SOA) infrastructure, and message-oriented middleware), replacing it with architecture components that allow you to scale based on client adoption.
3. Use open standards to build a data lake using the same data as in the operational data platform. Use a read-pattern schema to make the raw data and curated data readily available for all user roles.
4. For well-known query patterns, build standard enterprise-data-warehouse schemas and data marts in **Amazon Redshift**. For ad-hoc query requirements, publish the data catalog in **AWS Glue** and use **Amazon Athena** for querying the data lake directly. Extend the data warehouse based on your specific needs.
5. Use **Amazon SageMaker** to provide standard artificial intelligence (AI) and machine learning (ML) models for customer segmentations and lifetime value. Also, use **SageMaker** to build your own models on top of the data lake.

Reviewed for technical accuracy August 23, 2022

AWS Reference Architecture