Guidance for Real-Time Casino Player Analytics on AWS

This architecture diagram shows how you can use a real-time analytics pipeline and AI to build and deliver effective marketing offers during casino game sessions.

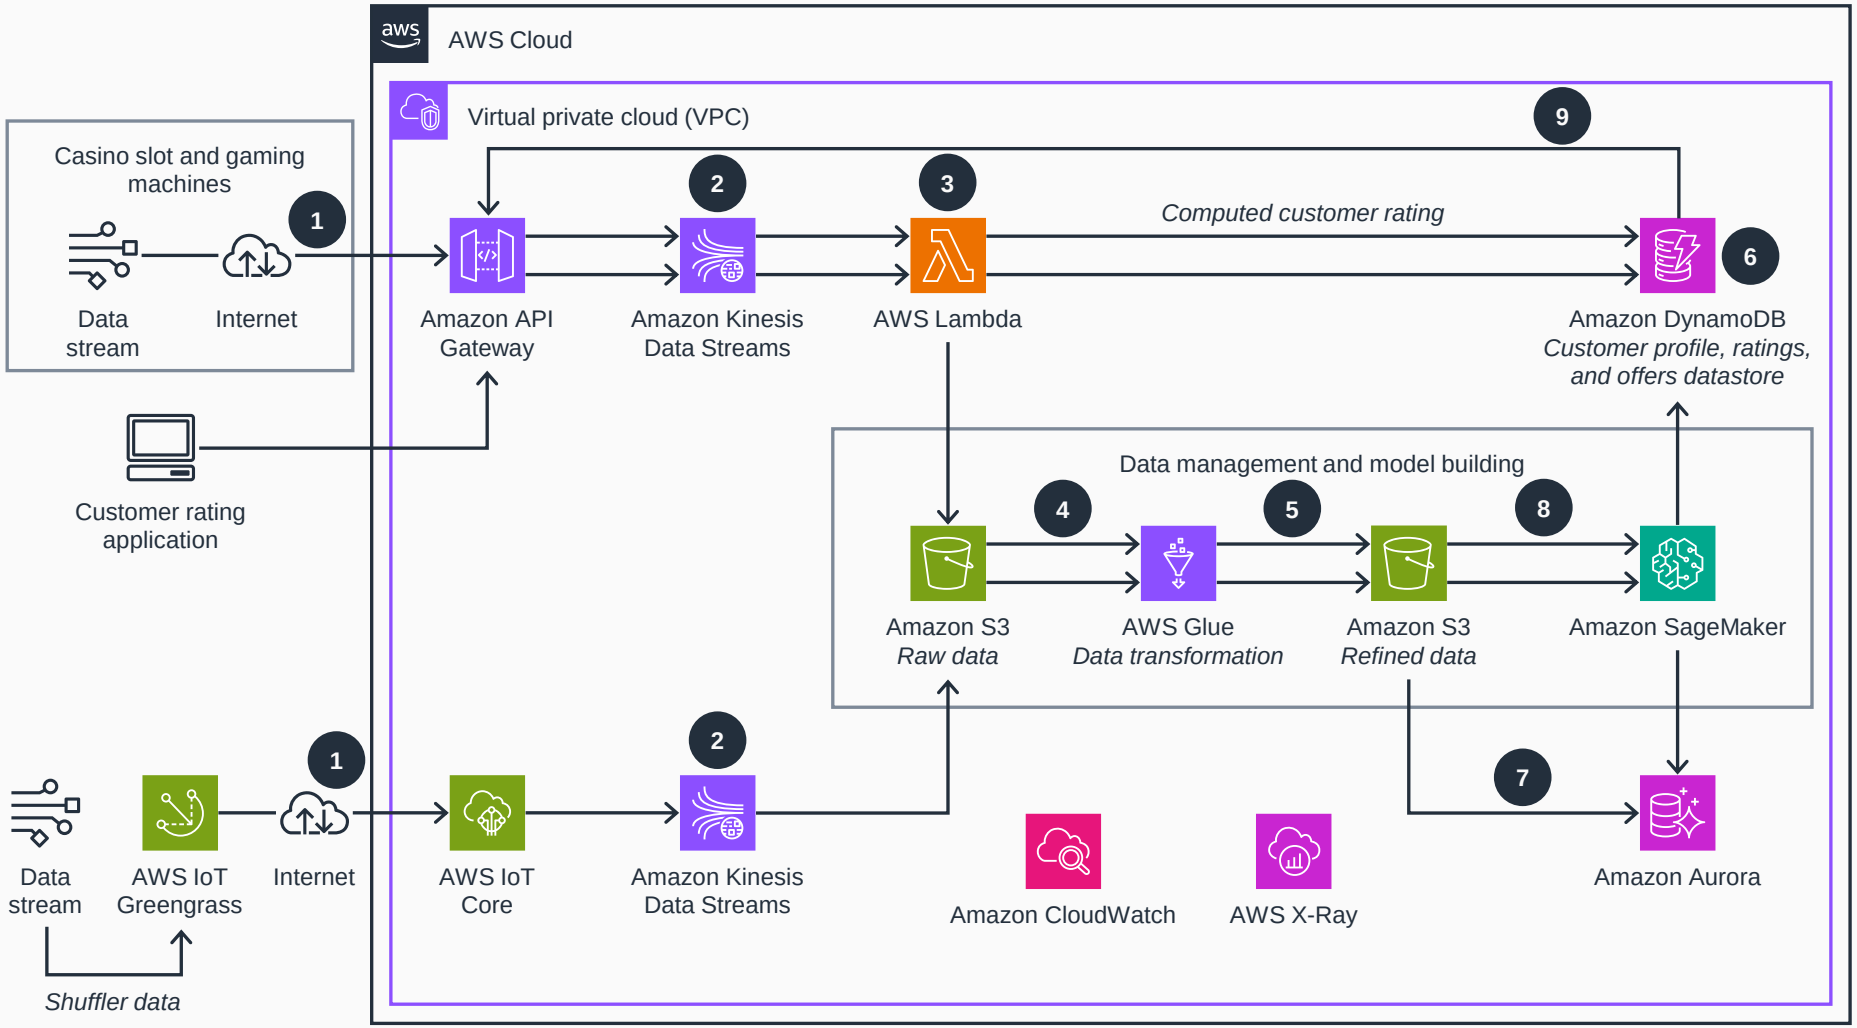

1. Casino slot machine data, as well as shuffler data integrated with **AWS IoT Greengrass**, are streamed from the casino floor through a private network into **Amazon API Gateway** and **AWS IoT Core**, respectively. **AWS X-Ray** can be used to help analyze any issues.
2. Data is then streamed into **Amazon Kinesis Data Streams**.
3. Slot data from **Kinesis Data Streams** is processed by **AWS Lambda** to calculate the customer rating and store a raw copy in **Amazon Simple Storage Service (Amazon S3)** for ML training.
4. Raw data from slot machines and shufflers is transformed to identify unique records through **AWS Glue** data transformation.
5. The refined data after transformation is stored in an **Amazon S3** bucket.
6. The customer profile, ratings, and offers are updated in **Amazon DynamoDB** for fast retrieval by slot machines or a customer rating application.
7. Refined shuffler data is stored for aggregation and retrieval in **Amazon Aurora**.
8. Refined slot data is used to train and update the ML model on **Amazon SageMaker**, which then predicts the best offers for the individual customer. Refined shuffler data is used to extract metrics to feed an ML model to predict failures.
9. The customer profile, ratings, and offers are made available to be consumed by games and applications to promote within the game or session.

Reviewed for technical accuracy August 5, 2024

AWS Reference Architecture