Creating Multi-Account IoT Pipelines On-The-Fly

Automatically deploy AWS IoT Analytics pipelines into different AWS Organizations accounts by leveraging Amazon Kinesis Data Analytics to buffer and route data from new devices and groups.

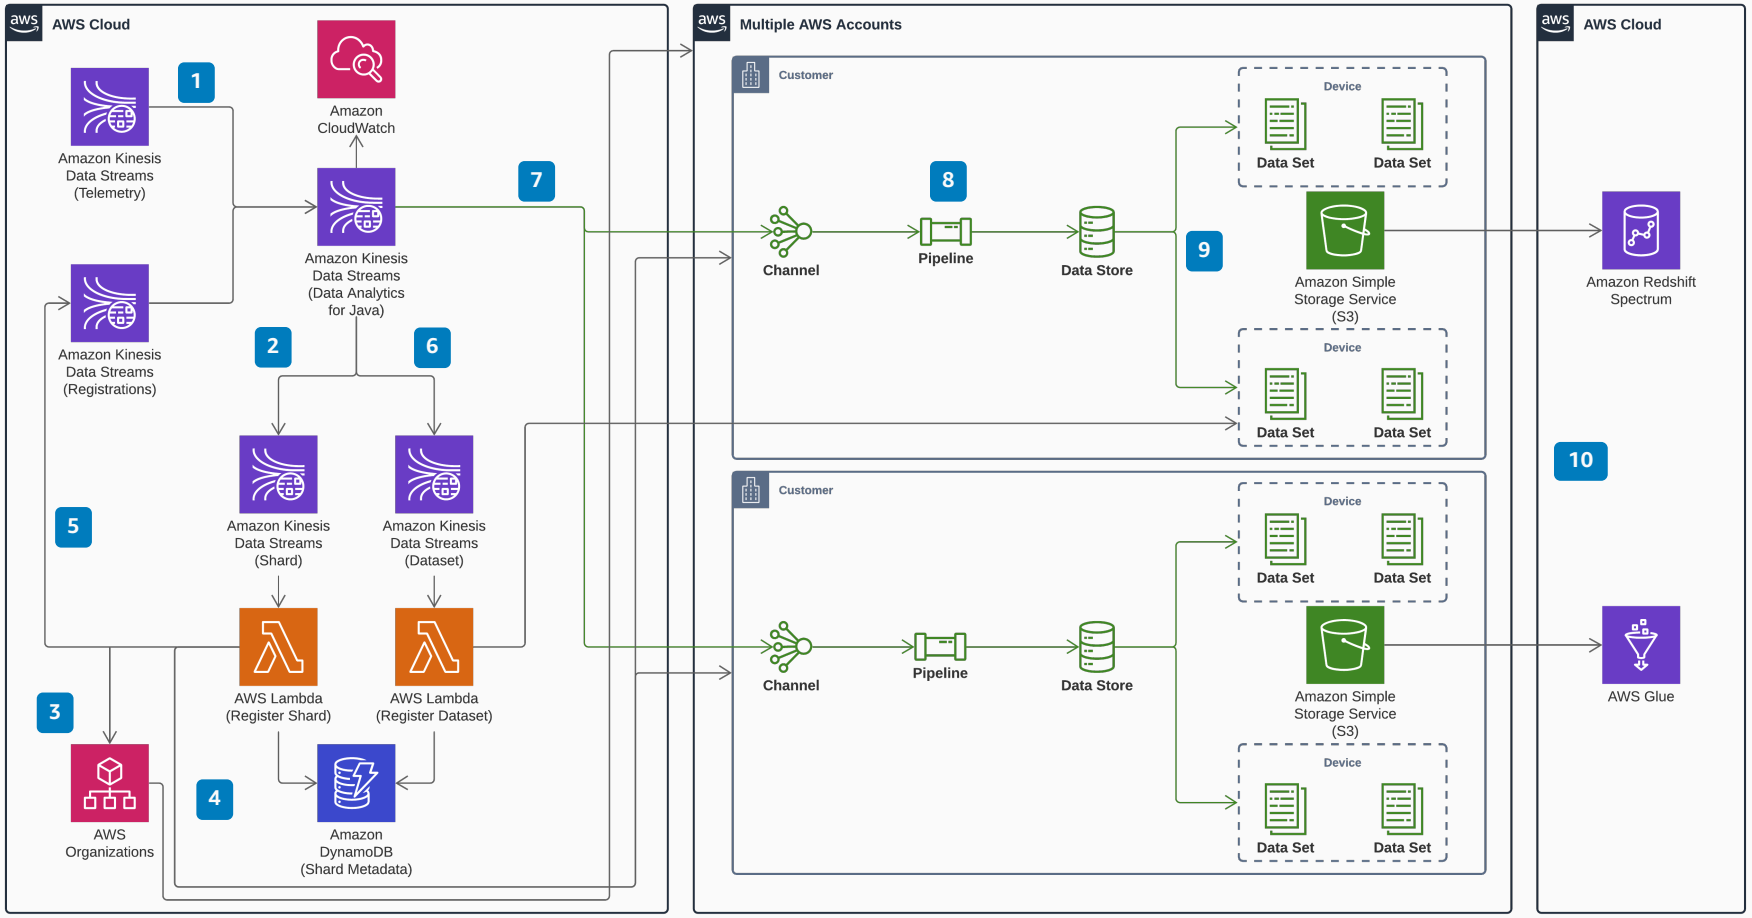

1. Ingest device telemetry data through the telemetry **Amazon Kinesis** stream.
2. Detect when an unregistered device group (shard) is ingested into **Kinesis Data Analytics (KDA)**, buffer incoming messages in a custom window, and emit an event.
3. The shard **AWS Lambda** function creates a new account and **Amazon Simple Storage Service (Amazon S3)** bucket if an account doesn't have enough capacity.
4. Utilize the AWS SDK to create an **AWS IoT Analytics** channel, Pipeline, and data store for the new shard and register them in **Amazon DynamoDB**.
5. Emit a registration event that contains the **AWS IoT Analytics** channel information and trigger the KDA window to continue processing buffered messages for this shard.
6. Detect new devices in the group and emit an asynchronous event to the register dataset Lambda function to create an **AWS IoT Analytics** dataset while continuing to ingest the telemetry.
7. KDA utilizes the **AWS IoT Analytics** channel registration for each device group to upload the group's device data to the correct **AWS IoT Analytics** channel in the registered account.
8. **AWS IoT Analytics** stores device group messages in the channel, transforms them in the Pipeline, and stores the transformed data in the data store and is backed by **Amazon S3**.
9. Each device's **AWS IoT Analytics** dataset uses a delta time query on the datastore to filter the data down to a specific device for the most recent 5 minutes for easier query.
10. Query device data with **AWS Glue** and/or **Amazon Redshift Spectrum**.

AWS Reference Architecture