AWS Reference Architectures

- Amazon Kinesis
- AWS Lambda
- Amazon DynamoDB
- Amazon CloudWatch
- Amazon S3

AWS LAMBDA: REAL-TIME STREAM PROCESSING

Amazon Kinesis is a fully managed, cloud-based service for real-time data processing over large, distributed data streams. Amazon Kinesis can continuously capture and store terabytes of data per hour from hundreds of thousands of sources. AWS Lambda lets you run code without provisioning or managing servers. You can set up your code to be triggered automatically from other AWS services. You can use AWS Lambda and Amazon Kinesis to process real-time streaming data for application activity tracking, transaction order processing, click stream analysis, data cleaning, metrics generation, log filtering, indexing, social media analysis, and IoT device data telemetry.

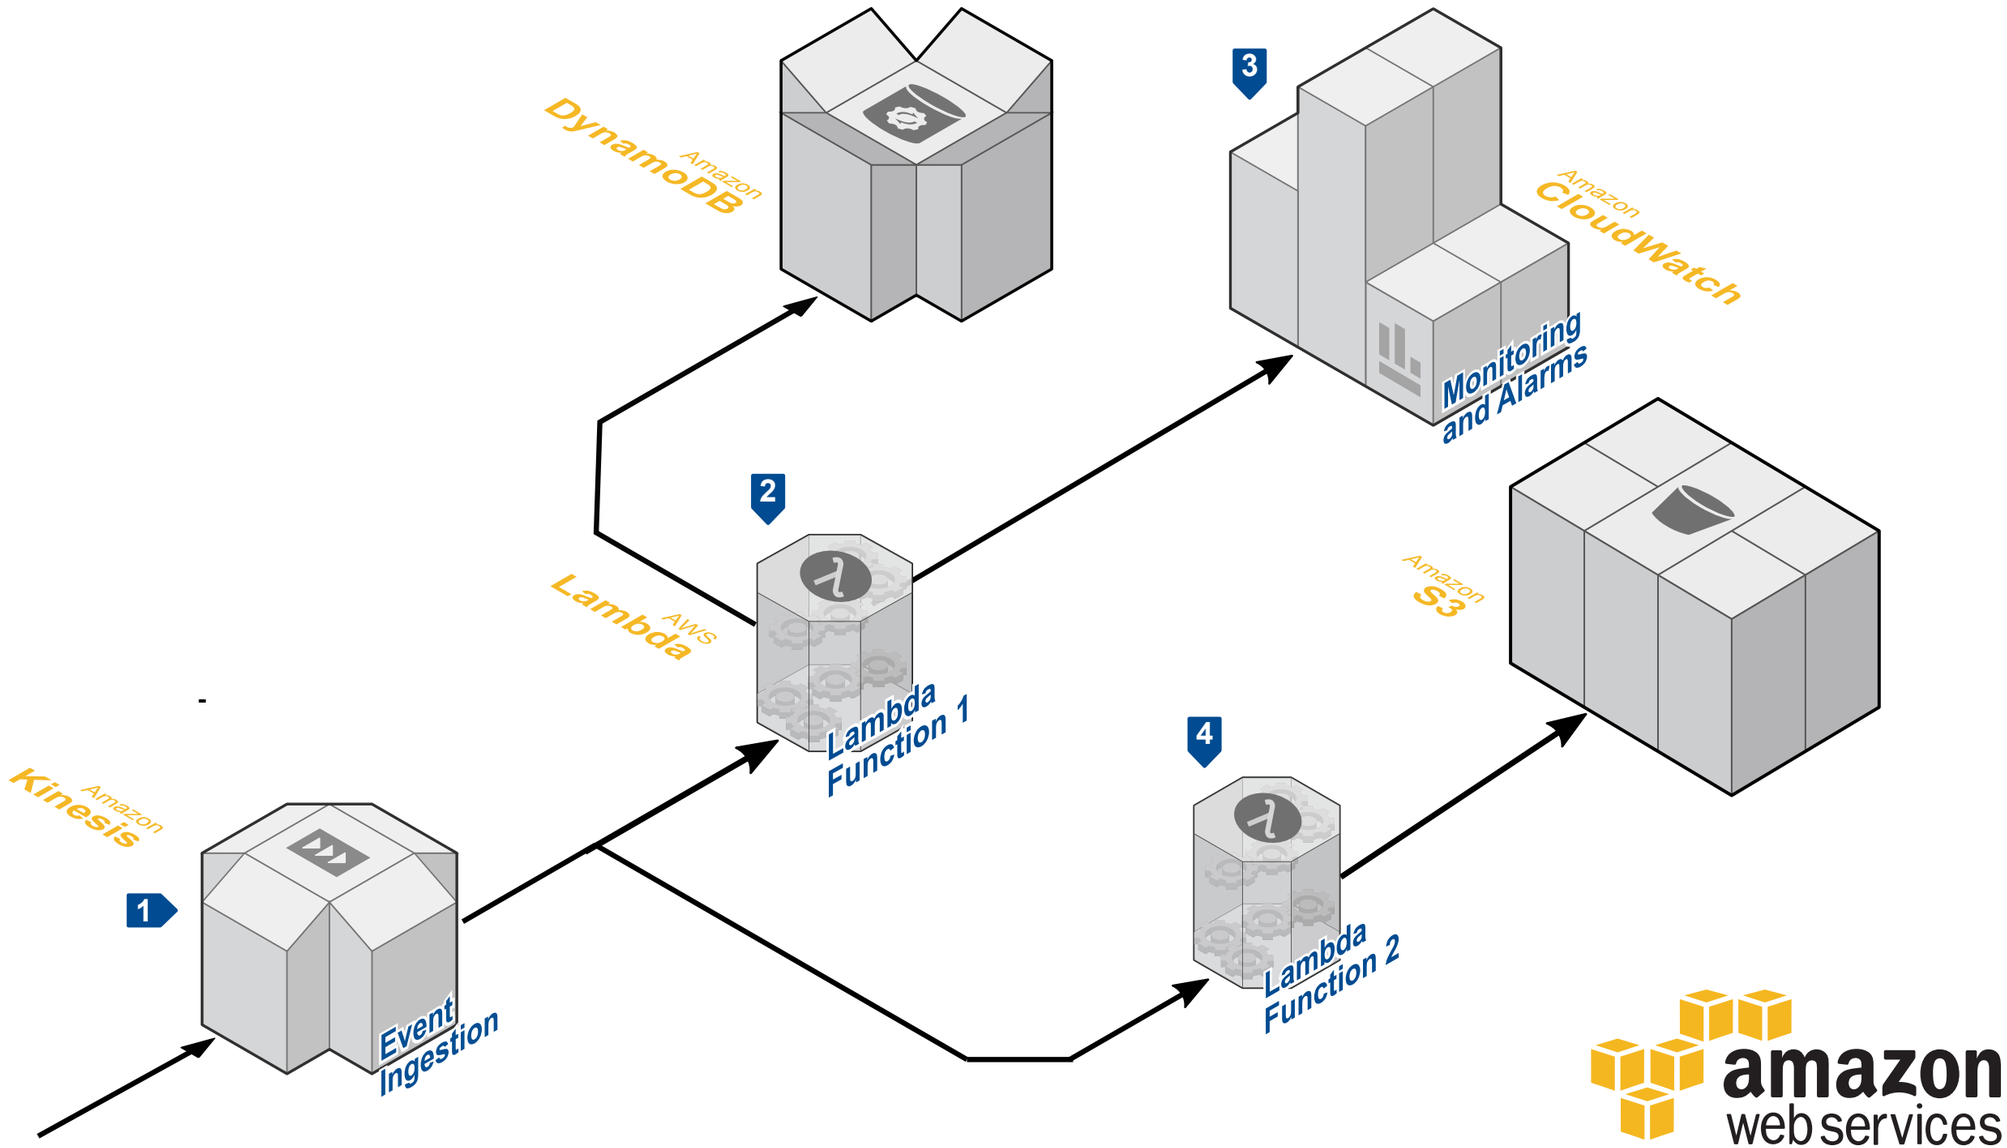

System Overview

1 Real-time event data is sent to **Amazon Kinesis**, which provides large-scale durable storage of the events for 24 hours and allows multiple **AWS Lambda** functions to process the same events.

2 In **AWS Lambda**, Lambda Function 1 processes the incoming events and stores the event data in a table in **Amazon DynamoDB** for low-latency access. You can provision the needed capacity of the DynamoDB table just by changing a configuration value.

3 The Lambda function also sends the values to **Amazon CloudWatch** for simple monitoring of aggregate metrics.

4 Another **AWS Lambda** function, Lambda Function 2, stores the incoming events in **Amazon S3** for durable, cost effective long-term storage. Storing data in Amazon S3 makes the data easily accessible for downstream processing and analytics.

amazon web services