AWS re:Invent

ADM303

# Implementing Amazon Personalize across marketing channels

**James Jory**
Senior Partner Solutions Architect, DCX
Amazon Web Services

**Dylan Tong**
Senior Partner Solutions Architect, AI/ML
Amazon Web Services



© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

# Related breakouts

| | |
|---|---|
| ADM402 | Testing techniques for measuring personalization |
| AIM304 | Build a content-recommendation engine with Amazon Personalize |
| AIM346 | Personalized user engagement with machine learning |
| AIM349 | Build a hotel recommender with Amazon Personalize—no ML skills needed |
| AIM424 | Accelerate experimentation with personalization models |
| RET310 | Create higher-quality recommendations in your e-commerce platform |

# Chalk talk format refresher

- Intended to be interactive
- Share your experiences
- Ask questions
- Focus of this session is on Amazon Personalize

# AWS ML stack

## Broadest and deepest set of capabilities

### AI Services

| VISION | | | SPEECH | | LANGUAGE | | CHATBOTS | FORECASTING | RECOMMENDATIONS |
|---|---|---|---|---|---|---|---|---|---|
| REKOGNITION IMAGE | REKOGNITION VIDEO | TEXTRACT | POLLY | TRANSCRIBE | TRANSLATE | COMPREHEND & COMPREHEND MEDICAL | LEX | FORECAST | PERSONALIZE |

### ML Services

**Amazon SageMaker**

| Ground Truth | Notebooks | Algorithms + Marketplace | Reinforcement Learning | Training | Optimization | Deployment | Hosting |
|---|---|---|---|---|---|---|---|

### ML Frameworks + Infrastructure

| FRAMEWORKS | INTERFACES |
|---|---|
| TensorFlow, mxnet, PYTORCH | GLUON, Keras |

| INFRASTRUCTURE | | | | | | | |
|---|---|---|---|---|---|---|---|
| EC2 P3 & P3DN | EC2 G4 | EC2 C5 | FPGAS | DL CONTAINERS & AMIs | GREENGRASS | ELASTIC INFERENCE | INFERENTIA |



# Domino’s

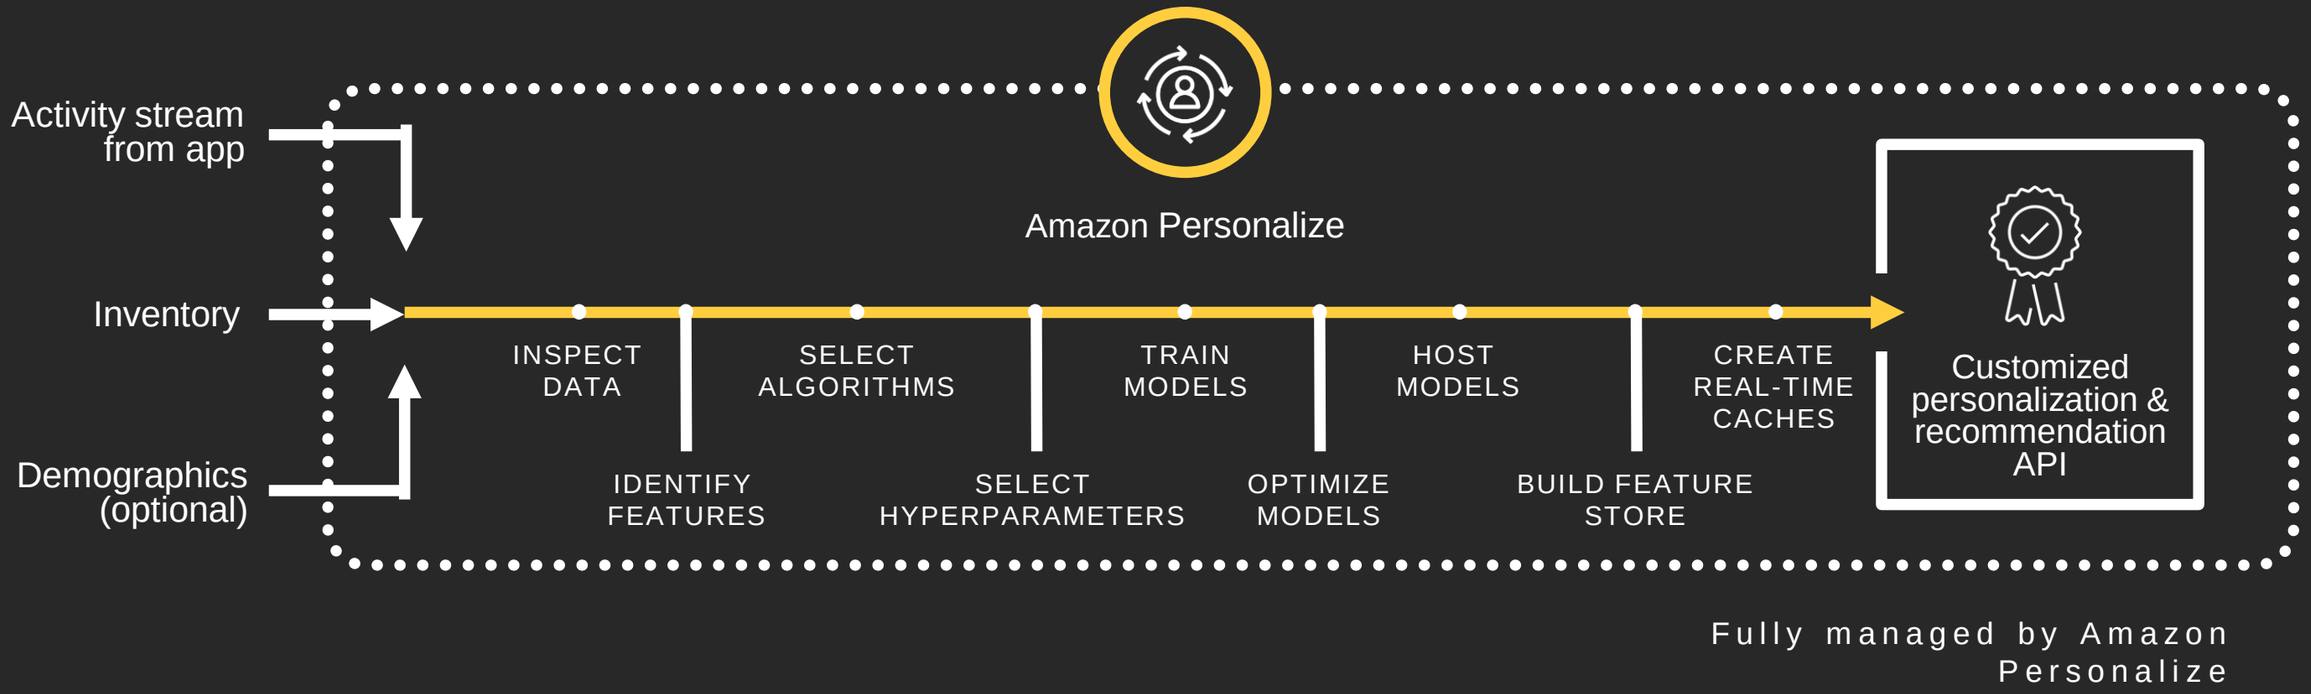
## Personalizing customer experiences

Domino’s uses Amazon Personalize to customize and scale relevant marketing communications to customers based on time, context, and content, thereby improving and enhancing their experience with the Domino’s brand.

# Amazon Personalize: Machine Learning personalization and recommendations



# Amazon Personalize: Machine Learning personalization and recommendations



# Why Amazon Personalize?

- Existing recommendation system not meeting expectations
- Expand personalization capabilities beyond recommendations
- Reduce complexity of existing recommendation system
- Reduce TCO of recommendation system
- Lack of machine learning expertise
- Time to market

# Marketing use cases for personalization

- Item recommendations
- Related item recommendations
- In-app messaging and digital assets (banners, product images, video, etc.)
- Outbound messaging (email, push, SMS, social)
- Optimizing customer journey/funnel
- Landing pages (ABM)

# Best practices integrating Amazon Personalize

- Define your business objectives
- Event identification and instrumentation
- Combine with other AWS services (Amazon Elasticsearch Service, Amazon Rekognition, Amazon Comprehend, and others)
- Streaming real-time events using Amazon Kinesis and AWS Lambda
- Handling anonymous users
- Recommending users for an item
- Incorporate experimentation/testing
- Automate solution building and campaign deployment
- Controlling costs

# Thank you!

**James Jory**
jjory@amazon.com

**Dylan Tong**
dylatong@amazon.com

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

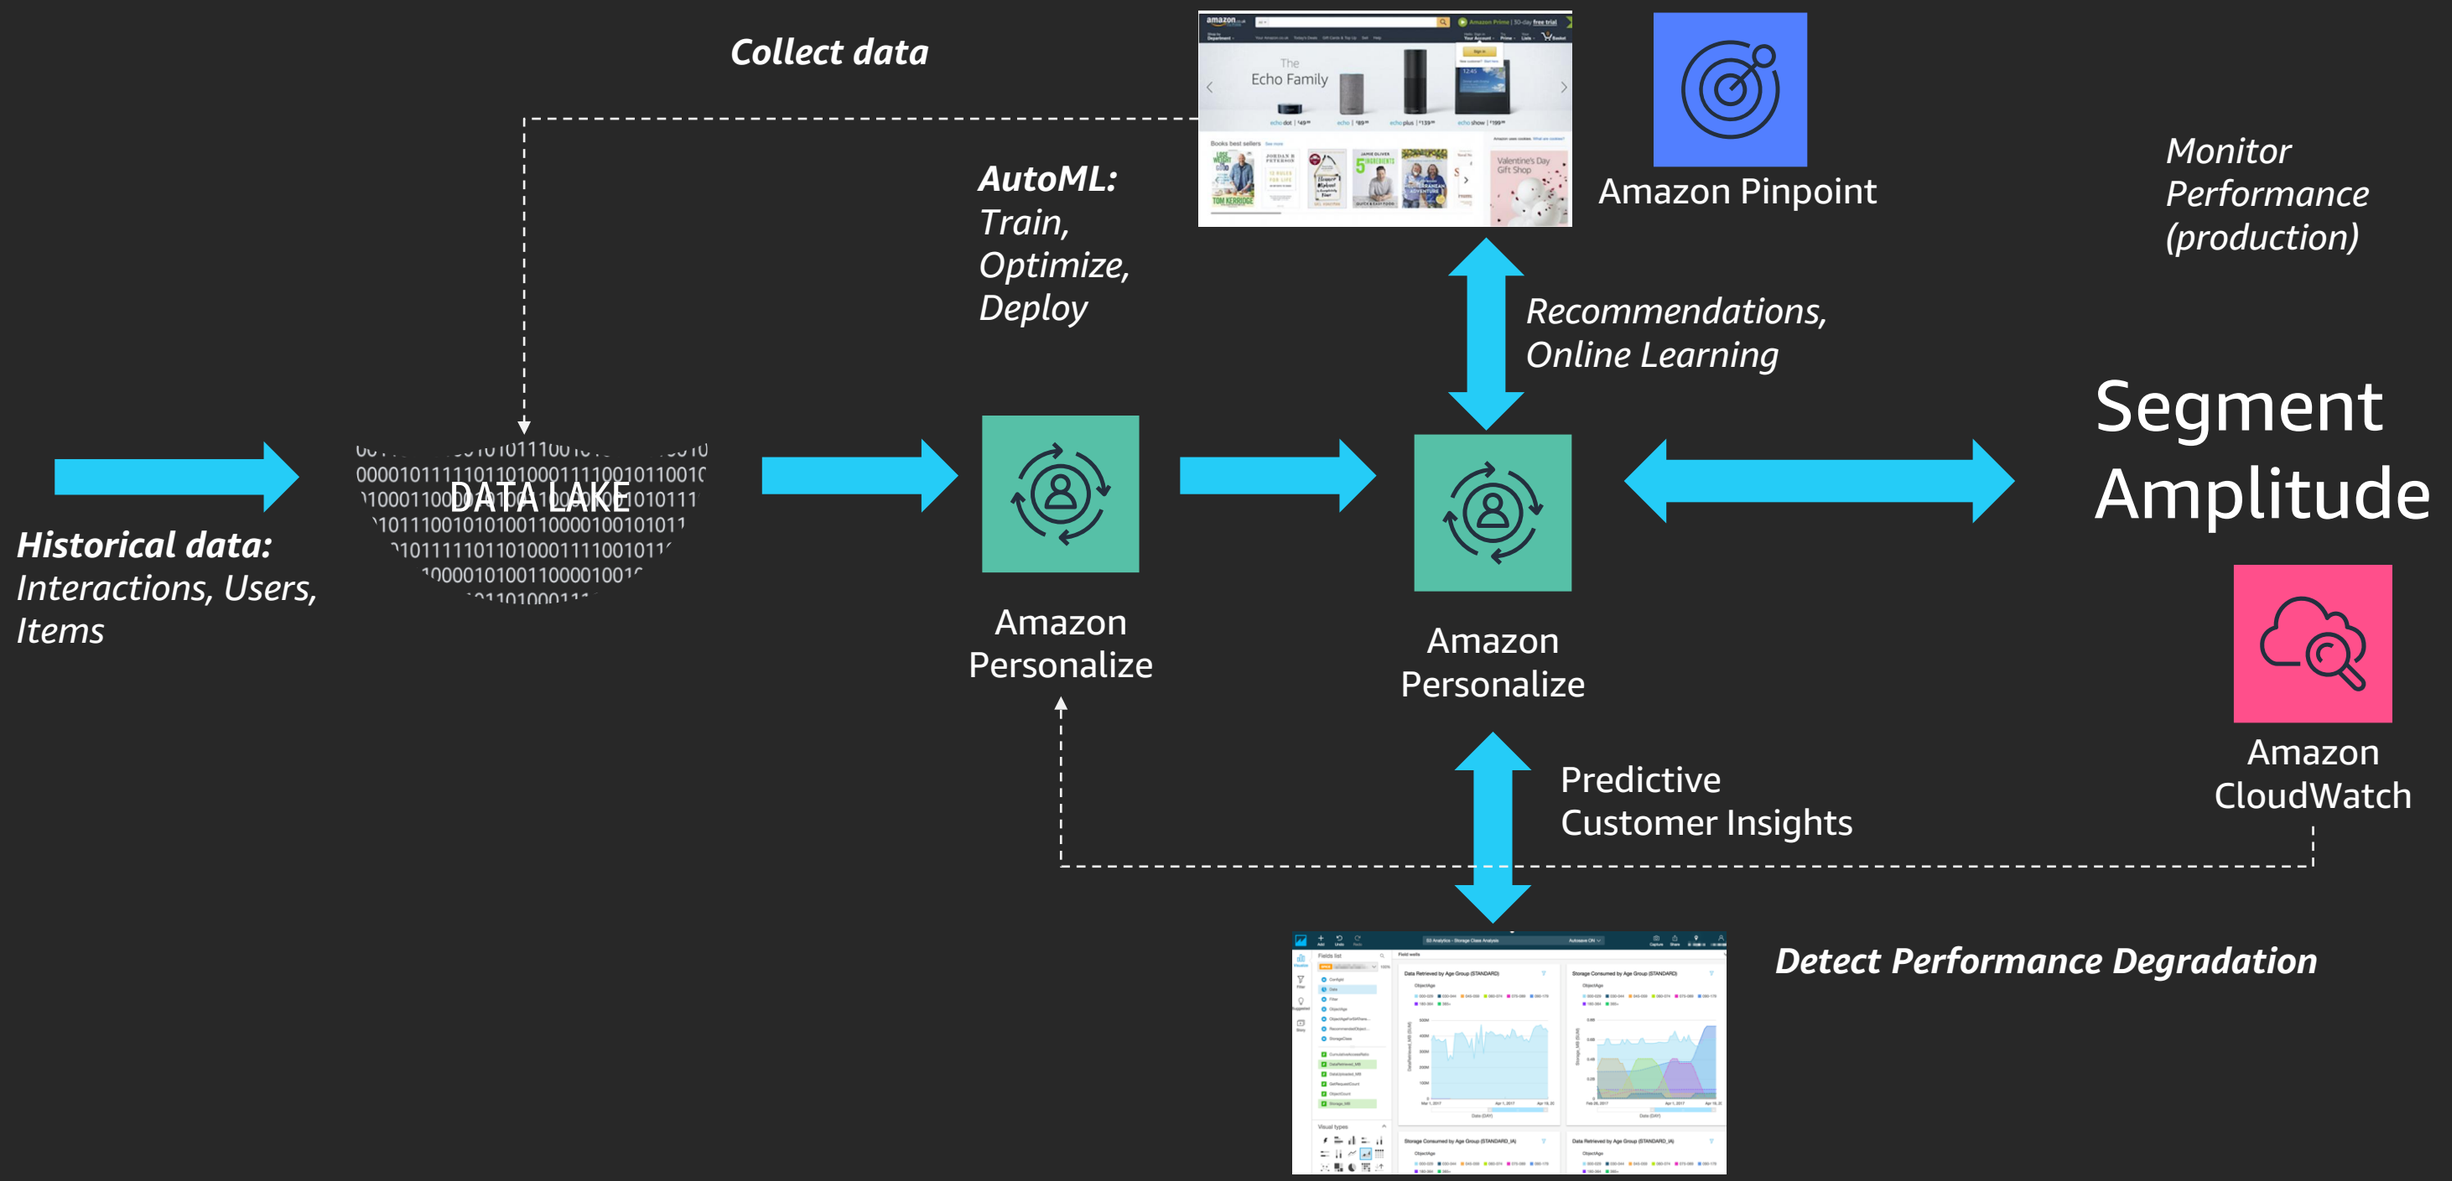
aws



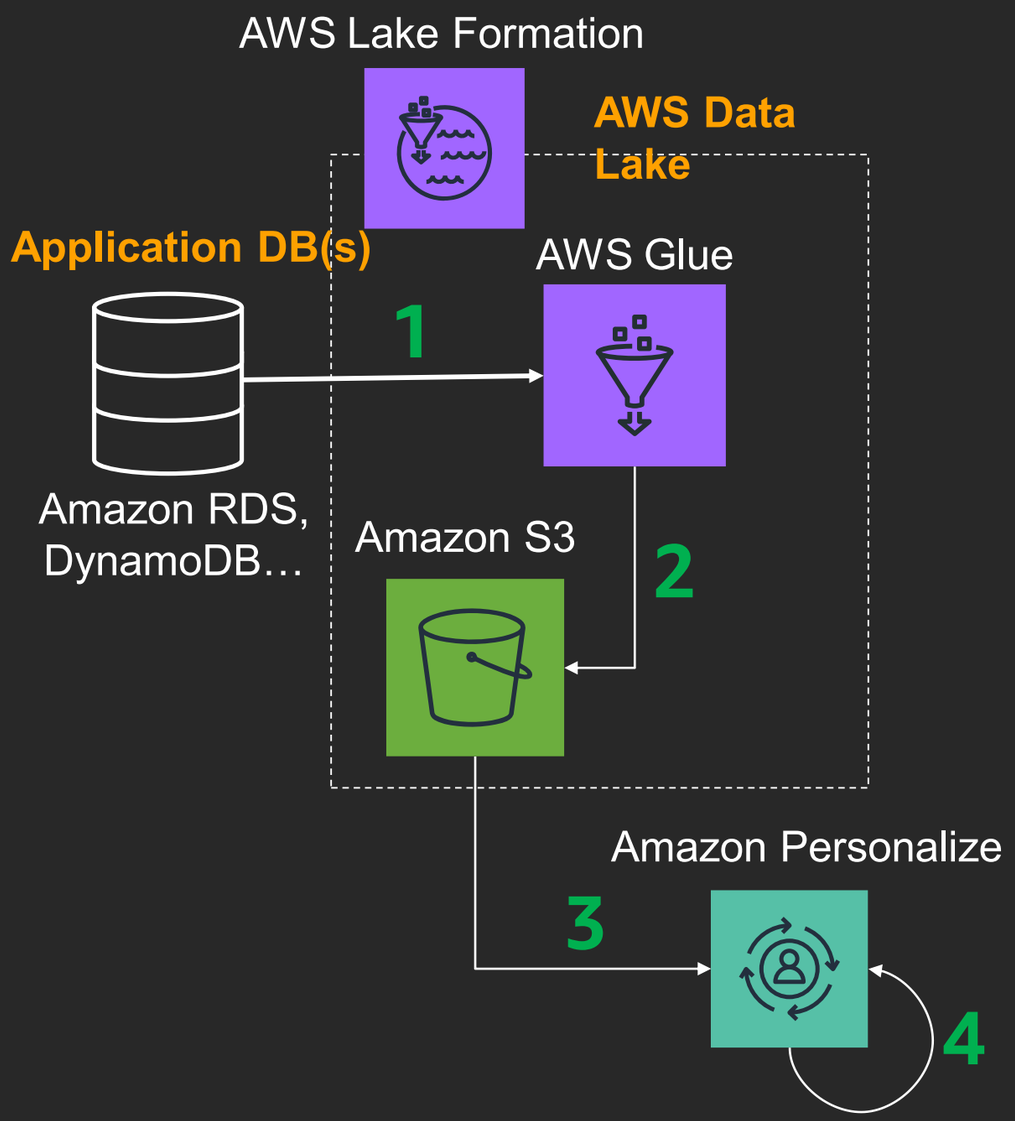
# Please complete the session survey in the mobile app.



© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.


Collect data
AutoML:
Train,
Optimize,
Deploy
Amazon Pinpoint
Monitor
Performance
(production)
Recommendations,
Online Learning
Historical data:
Interactions, Users,
Items
DATALAKE
Amazon
Personalize
Amazon
Personalize
Segment
Amplitude
Amazon
CloudWatch
Predictive
Customer Insights
Detect Performance Degradation

# Recommender system training and deployment



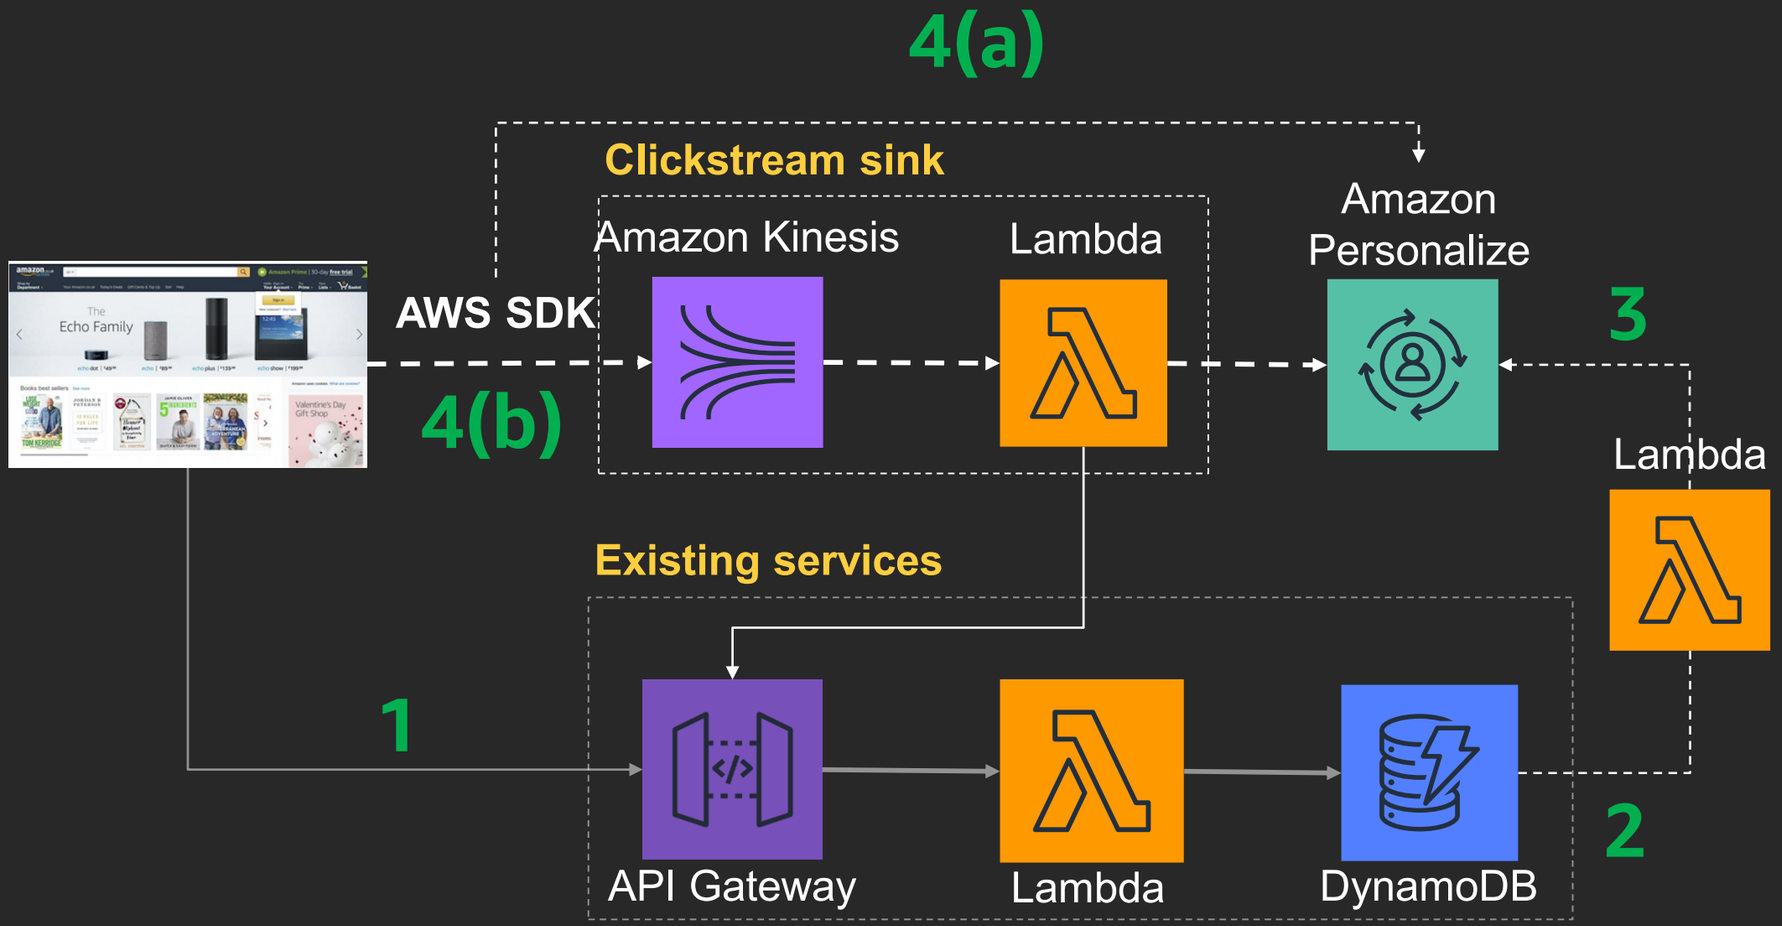

1. **AWS Lake Formation** provides access to a variety of options to extract data from sources: Glue jobs, crawlers and blueprints

2. **AWS Glue** transforms source data into the schemas expected by **Amazon Personalize**

3. Import interaction, users, and inventory data into **Amazon Personalize**

4. **Amazon Personalize** provides AutoML to automate the process of training and evaluating multiple recipes and delivering a custom recommender as a managed API endpoint

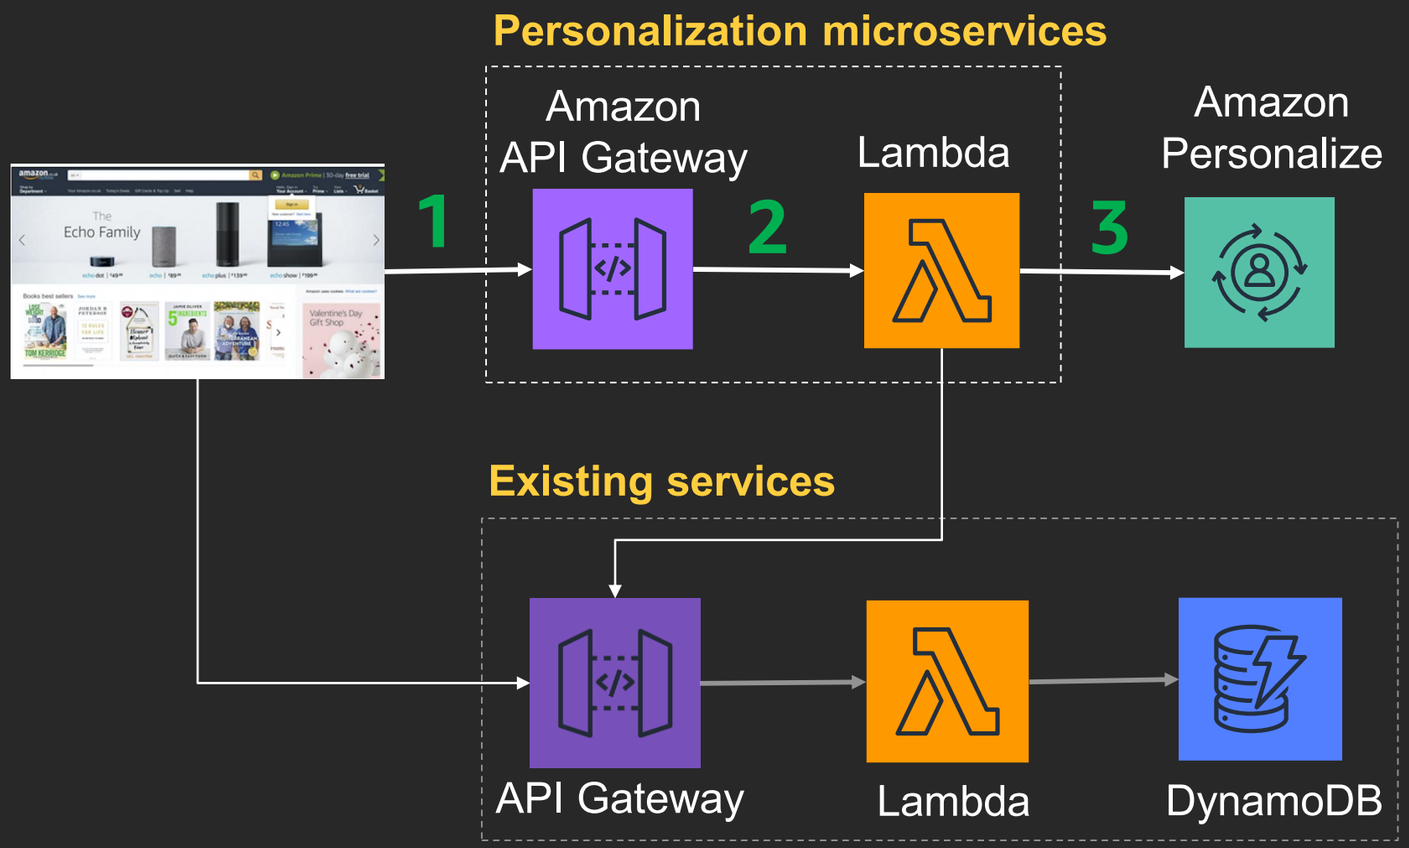
# Event recording for online learning




1. Application services continue to collect relevant data for training
2. AWS Lambda is triggered by DynamoDB streams as new events are triggered
3. Events are written to Amazon Personalize through the event ingestion SDK
4. **(a)** Some events might not be stored in the app db. The client can write events directly to **Amazon Personalize** through the SDK. **(b)** The event recorder has limits. For high throughput cases, write the events to **Amazon Kinesis**. Lambda is triggered to micro-batch events to **Amazon Personalize**

https://docs.aws.amazon.com/personalize/latest/dg/recording-events.html

# Real-time personalization

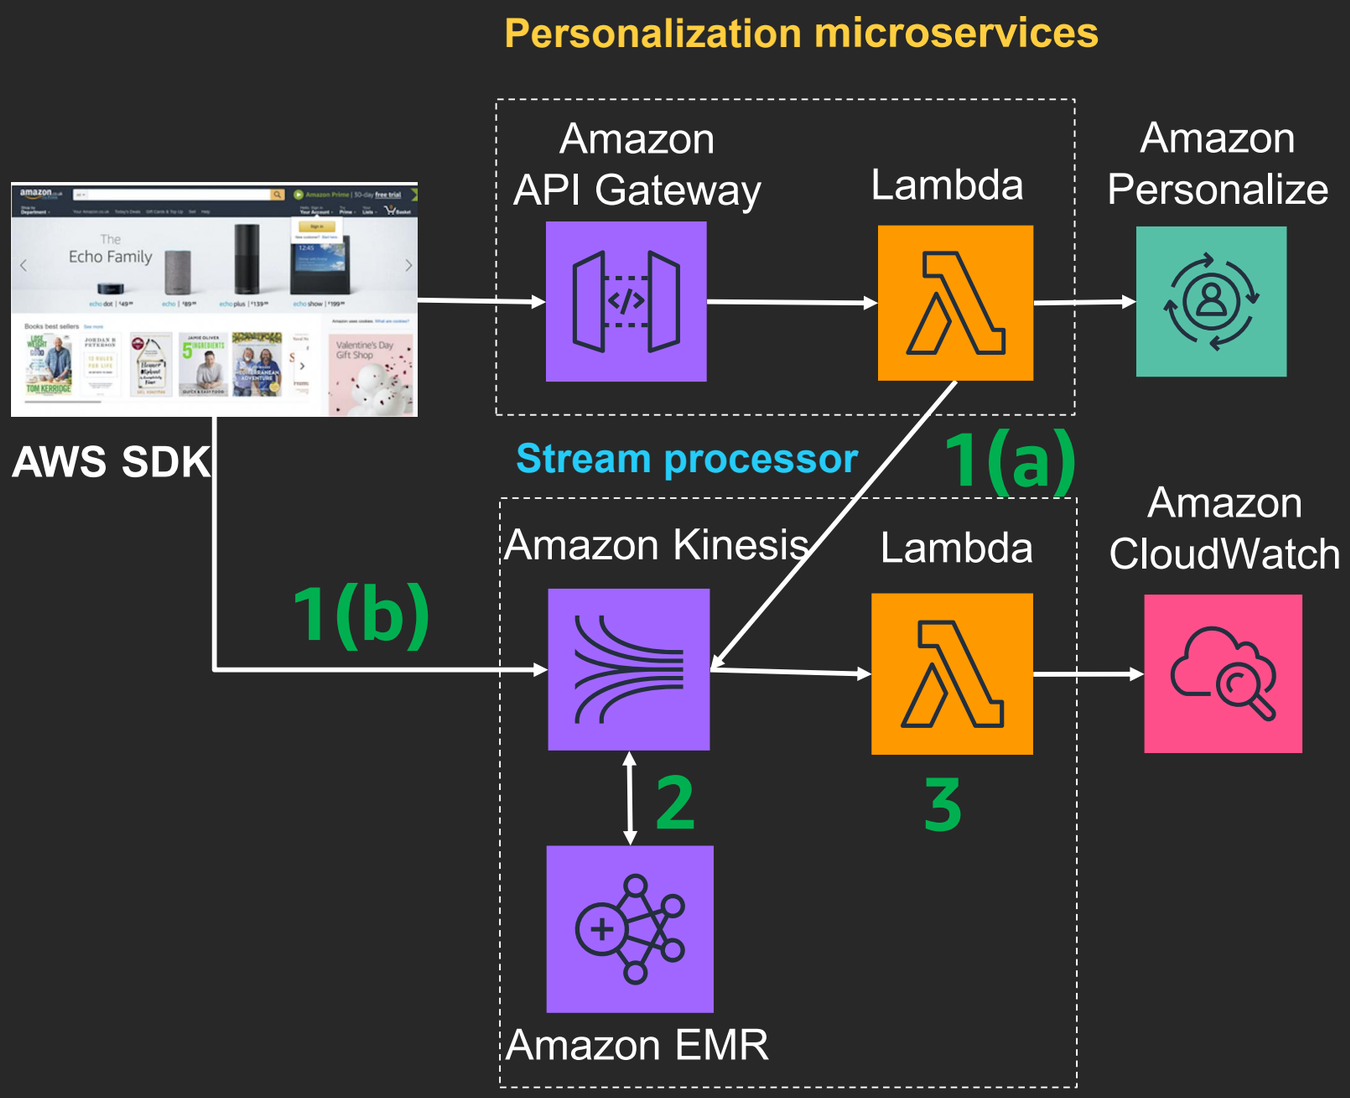



1. Application communicates with a microservice fronted by **Amazon API Gateway**. This provides API management functionality.

2. AWS Lambda serves the business logic for the microservice. Additional user and product information will need to be retrieved to serve the request. Personalize only returns Item IDs.

3. **Amazon Personalize** serves real-time requests for recommendations and ranked lists.

# Production drift detection [build]

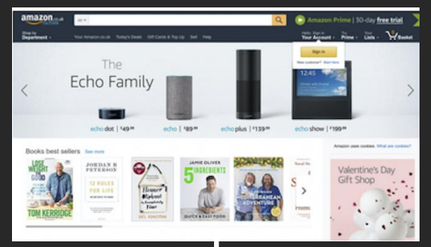



1. (a) As your personalization microservice responds to requests for recommendations, record the recommendations in Kinesis. (b) In parallel, the clients should record clicks or conversion events into Kinesis.

2. Use Spark streaming on **Amazon EMR** to reconcile successful “hits” on recommendations using time window analysis. Write the results back into Kinesis as a “successful hits” data stream.

3. Trigger AWS Lambda to write metrics into Amazon CloudWatch to track metrics like hit ratio. Trigger alerts or automatic processes to reconcile performance degradation.

# Production drift detection [partner solutions]

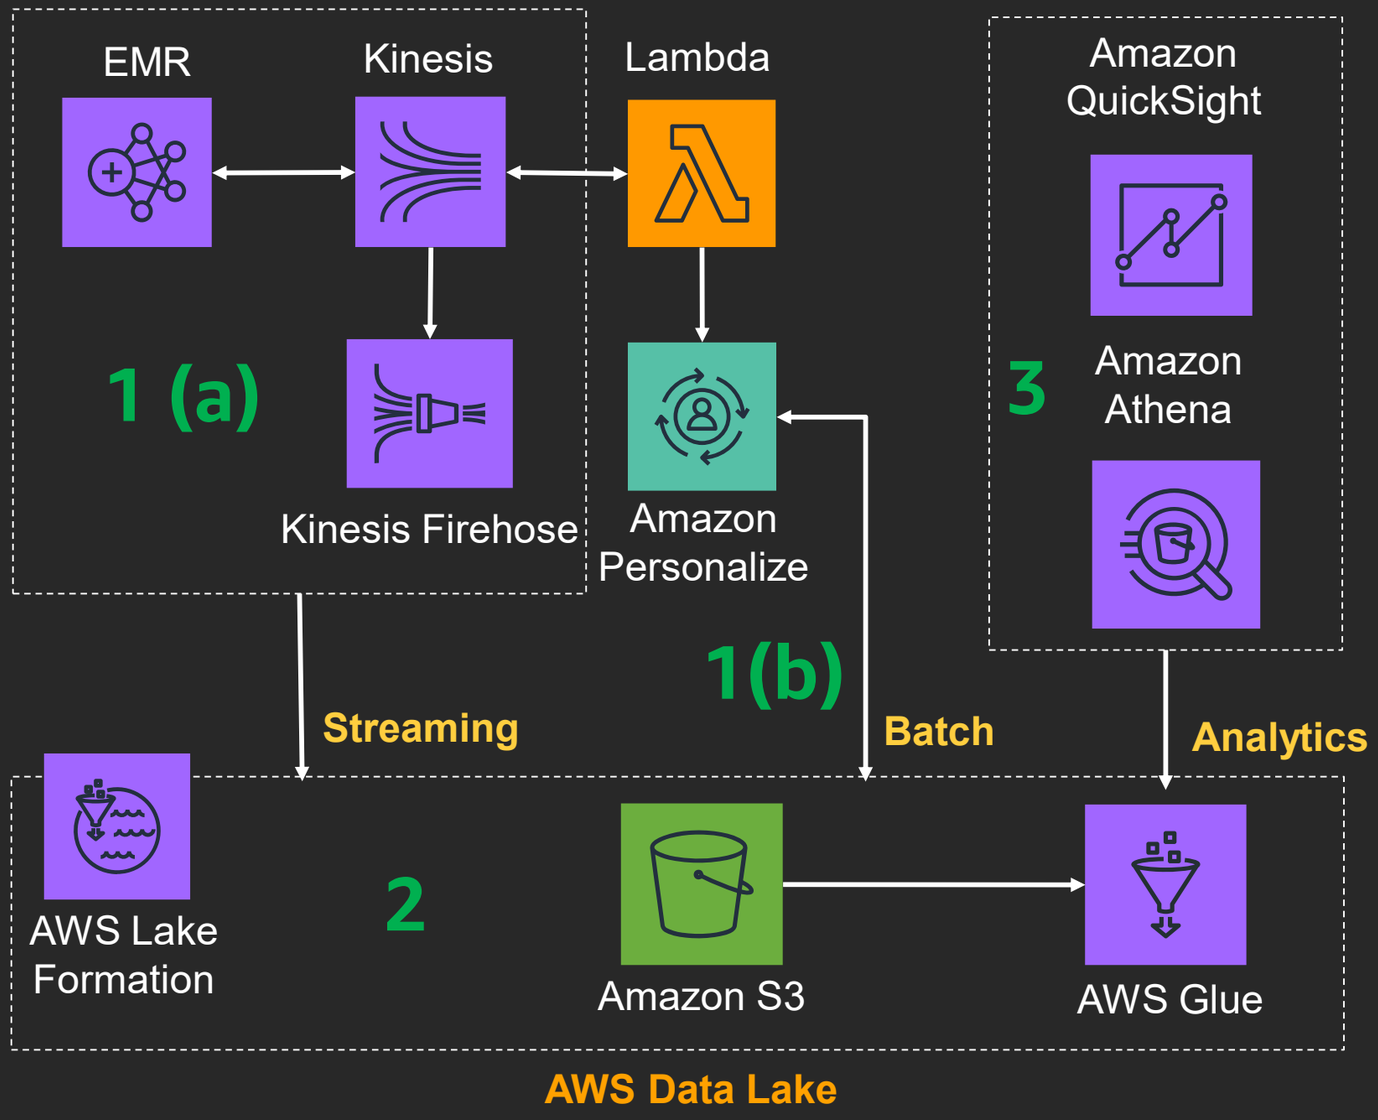



Amazon
Personalize

Segment

Amplitude

# Predictive Customer Insights




1. (a) Data streams are ingested into **Amazon Kinesis**. **AWS Lambda** uses **Amazon Personalize** to created streams enriched with recommendations. **Amazon EMR** provides access to a variety of Hadoop tools to perform streaming analysis. Kinesis Firehose provides a simple way to write data back into the data lake where incremental updates provide a means for intra-day offline analysis. (b) **Amazon Personalize** provides the ability to batch predictions.

2. Data is governed, processed, and optimized through AWS data services.

3. Recommendations provide predictive insights into customer product affinity. These insights facilitate activities ranging from campaign targeting to guiding cross and up-sell opportunities. **Amazon Athena** provides a serverless SQL engine to query this data. **Amazon Quicksight** turns these queries into visualizations and dashboards.