



» Whitepaper

How to harness the power of AWS Personalize

to deliver content recommendations quickly,
at scale, and with ease



Lead. Don't follow

Author: Chris Birkinshaw
Principal Technologist
Merapar



How to harness the power of AWS Personalize

to deliver content recommendations quickly, at scale, and with ease.

Introduction

Most OTT (Over The Top Video) video businesses are interested in providing content recommendations to their users. The primary reason is usually to increase content viewership and breadth of catalogue consumption, so that the user feels more satisfied with the value they get from the service, and is therefore less likely to churn voluntarily. It is arguably also the case that recommendations have become a standard part of the UX (User experience), and viewers simply expect them to be there to assist them in content discovery.

Many COTS (Commercial Off The Shelf) software solutions for personalized recommendations are available in the marketplace, however the high capital costs and per subscriber cost of these solutions does not generally deliver a return on investment for a typical subscription OTT video streaming service. However, building a custom recommendations solution using machine learning technology has previously been impractical for most organisations due to the specialised skill sets required.

AWS Personalize

AWS (Amazon Web Services) have disrupted many technology areas, combining their technical experience gained running amazon.com with the vast scale of AWS, thereby commoditizing many of the services that comprise modern software architectures.


Many AWS products were key parts of their internal infrastructure for years before being offered to the public, e.g. DynamoDB, Kinesis Streams and now Personalize. With Personalize, AWS have released the technology used to power customer recommendations in amazon.com and made it simple to use without any knowledge of machine learning (ML).



AWS Personalize Benefits

- Cost effective to run.
- Usable without any knowledge of machine learning.
- Algorithms have been developed by Amazon based on their extensive experience:
 - Data is private and isolated and only used for your recommendations.
 - Serverless hosting model, so AWS takes care of security and scalability.

Solution Overview



Recommendation use cases

User Recommendations	This is the classic recommendations use case, where the client app sends a request with the user ID and gets back a list of items recommended for that user. It might be shown in a “Recommended for you” carousel on a home page.
Related Items	In this use case the user interaction data is used to infer relationships between items. The client app can then query for related items by passing an item ID. It could be used to power a “More like this” or “Other users also watched” carousel on the page for a particular item (e.g. TV Show).
Reranking	In the final use cases, Personalize is used to re-rank or reorder an existing list of items for a specific user ID. This could be used to re-rank search results or even to re-rank lists of editorially curated content recommendations. This use case is very interesting, as it allows for a hybrid of editorial and algorithmic recommendations and ensures each user sees something unique with minimal input from the editor.
Promotions	The platform operator may want to boost a particular content item(s) in the recommendations results. The recommender could be configured to show a % of content from a promotions category. This feature can also be used to enforce business rules, like ensuring 50% of the items are movies, or live programmes.
Contextual Recommendations	A user’s viewing habits are likely to vary based on their viewing context, for instance the kind of device they are using (TV, laptop, mobile, etc), the time of day, or their physical location. This information can be passed into the recommendation requests, and into the data feed of user interactions, and Personalize will use them to improve the accuracy of recommendations.

Systems Integration

Personalize supports 3 types of data, but only the Interactions dataset is required.

	Required Fields	Description
Users	User ID	Data about the end users. May also include demographic data, which can be important in increasing the accuracy of recommendations.
Items	Item ID (e.g. Show ID)	The IDs of the items that will be recommended, e.g. Show, Season, Episode, Product
Interactions	User ID Item ID Timestamp	Any interaction between the user and the items, e.g. a video playback, adding to a favourite list, sharing with a friend.

Data can be imported both as historical batch data, and real-time events. Note that:

- A batch data import for a dataset replaces the previous data import, so for the interactions data, the real-time integration will be required to append to the data set.
- For the users and items datasets, this might not be such a big issue but for larger catalogs and user bases the event based approach will be significantly more efficient.

Testing and Monitoring

Accuracy of recommendations

AWS Personalize splits the historical data that is ingested to train the machine learning models and retains some of the data to test the model with. By testing the model with a subset of the real data, Personalize can output metrics on the accuracy of the model. Personalize can be configured to train multiple models, adjust the parameters automatically and then pick the best performer.

Business value of recommendations

Even if we conclude that the recommendation model is sufficiently accurate, we would still like to know whether users engage with the suggestions. It is vital to ensure that clickstream data from the OTT apps is sent to a suitable solution, with enough event metadata to attribute clicks through to the content pages, back to the recommendations. This setup allows the business to continuously monitor KPIs like:

- **Click through rate (CTR)** on the recommendations
- **Conversion rate** (e.g. to a play or a subscription) of journeys that started with a recommendation click
- **Uplift** in overall engagement of customers that click recommendations or watch recommended content


The analytics solution could be composed of AWS services or utilise SaaS platforms like the Google or Adobe analytics suites.

Tuning the Solution

AWS Personalize's own metrics on model accuracy, and the KPIs extracted from the analytics solution should be used to constantly evaluate and optimise the solution.

Imagine the CRM team may have some demographic data on the users that can be added to AWS Personalize to improve accuracy, or videos have been processed by AWS Recognition to extract a cast list. Once this data has been added to the data-sets, the change in recommendation accuracy and user engagement can be measured to evaluate the impact.

Merapar Development Kit (MDK) for AWS Personalize



Deploying AWS Personalize is more complex than most other AWS managed services. The architecture is provisioned in layers and depends on long running batch processes. For example, it is not possible to create an endpoint for recommendations without first importing interactions data and training a machine learning model. These operations are time consuming, because the datasets needed are very large.

Although Personalize is a very comprehensive solution, significant effort can be spent creating tools to manage streaming and batch data manipulation and ingest. The recommendations are only ever as good as the quality of the data. For some use cases, timeliness of the data is also important, for example adding new content to the catalogue or filtering recommendations that users have already watched.

‘At Merapar, we have developed a toolkit that takes the hard work out of provisioning AWS Personalize, and the supporting infrastructure needed to transform and ingest data at scale.’

About Merapar

Merapar is a technology consulting and software engineering firm with deep domain expertise in cloud and media.

Our consultants and engineers are highly trained in cloud-native technologies, combining this with advanced software development techniques and agile methodologies. We enable our customers to transform their business to the latest Infrastructure as a Service (IaaS) technologies, iteratively, and at speed.

We develop and deliver strategic infrastructures, workflows and consumer facing applications which are scalable, robust, repeatable, and always with cost optimisation in mind. Our core expertise is in cloud engineering, digital transformation, and video platforms for the media and entertainment industry.

Once you are up and running, our managed service ensures that your mission-critical systems run with broadcast-grade availability (>99,99%). Our Continuous Integration and Continuous Delivery (CI/CD) platforms guarantee a scalable and future-proof solution.

As an Advanced Consulting Partner of Amazon Web Services, we're able to provide our clients with the reassurance of the world's #1 cloud-hosting provider.

AWS Advanced Consulting Partner

Merapar is an official Advanced Consulting Partner within the Amazon Web Services Partner Network. We have a team of AWS certified developers and architects on board. With this partnership, Merapar offers our customers with the best of both worlds: the assurance of using Amazon Web Services, the nr. 1 Cloud Technology Provider, and Merapar's proven track record in building scalable, flexible and cost effective solutions. reassurance of the world's #1 cloud-hosting provider.





Merapar is exhibiting at IBC 2023

15th - 18th September

RAI Amsterdam

Hall 5, Stand B36

 [Click here](#)

to book a 1-2-1 demo of the Merapar Development Kit (MDK)
at this years IBC show.

See first hand how our MDK takes the complications out of
integrating AWS Personalize into your platforms, with:

- Accelerated time to market
- Trusted Software Components
- Zero licence fee access to code

How to harness the power of AWS Personalize

to deliver content recommendation quickly, at scale, and with ease

