



한국언론정보 기술 협회 웨비나 15 - 16 DEC, 2022

AI 기반의 콘텐츠 서비스 혁신 AWS AI 서비스 기반으로 빠르게 고도화하기

Jonathan Kim

Solutions Architect, M&E Vertical
AWS Korea

Agenda

AWS for Publishing

AI Services on AWS for Publishing

Summary



Publishing industry is going through a transformation



Value shift to digital platforms and rise of subscription revenues



Personalisation



Ubiquitous connectivity



Content consumption "on the go"



Disruption means changes to publishing business and operation models

Business Drivers

How AWS Helps

1 Monetization

구독자 확보 및 유지
새로운 서비스 출시

Customer Data Platform

Digital Engagement Tools
구독자 맞춤형 Personalization

2 Optimization

편집 워크플로우 최적화
비용 절감

AI/ML


반복적인 작업 자동화





AI Services on AWS

The Opportunity: Get more value from your **Data**



DATA



ANALYTICS



MACHINE
LEARNING

Becoming a data-driven organization

DATA IS AN
ORGANIZATIONAL ASSET

데이터는 더 이상
특정 부서만의 자산이
아님

DATA IS
DEMOCRATIZED

데이터는 필요로 하는
사람에게 안전하게
제공되어야 함

DATA IS PUT
TO WORK

데이터 기반의 분석과
학습을 통해 의사
결정과 워크플로우를
효율화

Data-Driven "...” 을 하기 위해서는?

1

데이터 보관
인프라의 개선

2

데이터 통합

3

Innovate with
ML

AI? ML? DL?



Artificial intelligence (AI)

Any technique that enables computers to mimic human intelligence using logic, if-then statements, and ML (including deep learning)



Machine learning (ML)

Subset of AI that uses machines to search for patterns in data to build logic models automatically



Deep learning

Subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition



Our mission at AWS


Put machine learning in the
hands of every developer

The AWS AI/ML Stack

AI SERVICES



ML SERVICES



ML FRAMEWORKS & INFRASTRUCTURE

PyTorch, Apache MXNet, TensorFlow

Amazon EC2

CPUs

GPUs

AWS Inferentia

AWS Trainium

Habana Gaudi

FPGA

Elastic inference





Amazon Rekognition

Deep-Learning Based Image & Video Analysis

Amazon Rekognition



Amazon Rekognition Image & Video



OBJECT, SCENE, AND ACTIVITY



CUSTOM LABELS



CONTENT MODERATION



TEXT



PPE DETECTION



FACE DETECTION AND ANALYSIS




FACE COMPARE & SEARCH



CELEBRITY RECOGNITION



VIDEO SEGMENTS



LIVE STREAM VIDEO



PATHING

Common Applications of Amazon Rekognition

Automated image and video analysis at scale



WORKPLACE SAFETY

PPE detection
Vehicle or person detection
Person of interest



IDENTITY VERIFICATION

Test taker fraud detection
Know your customer (KYC)
Service access verification



MEDIA ANALYSIS

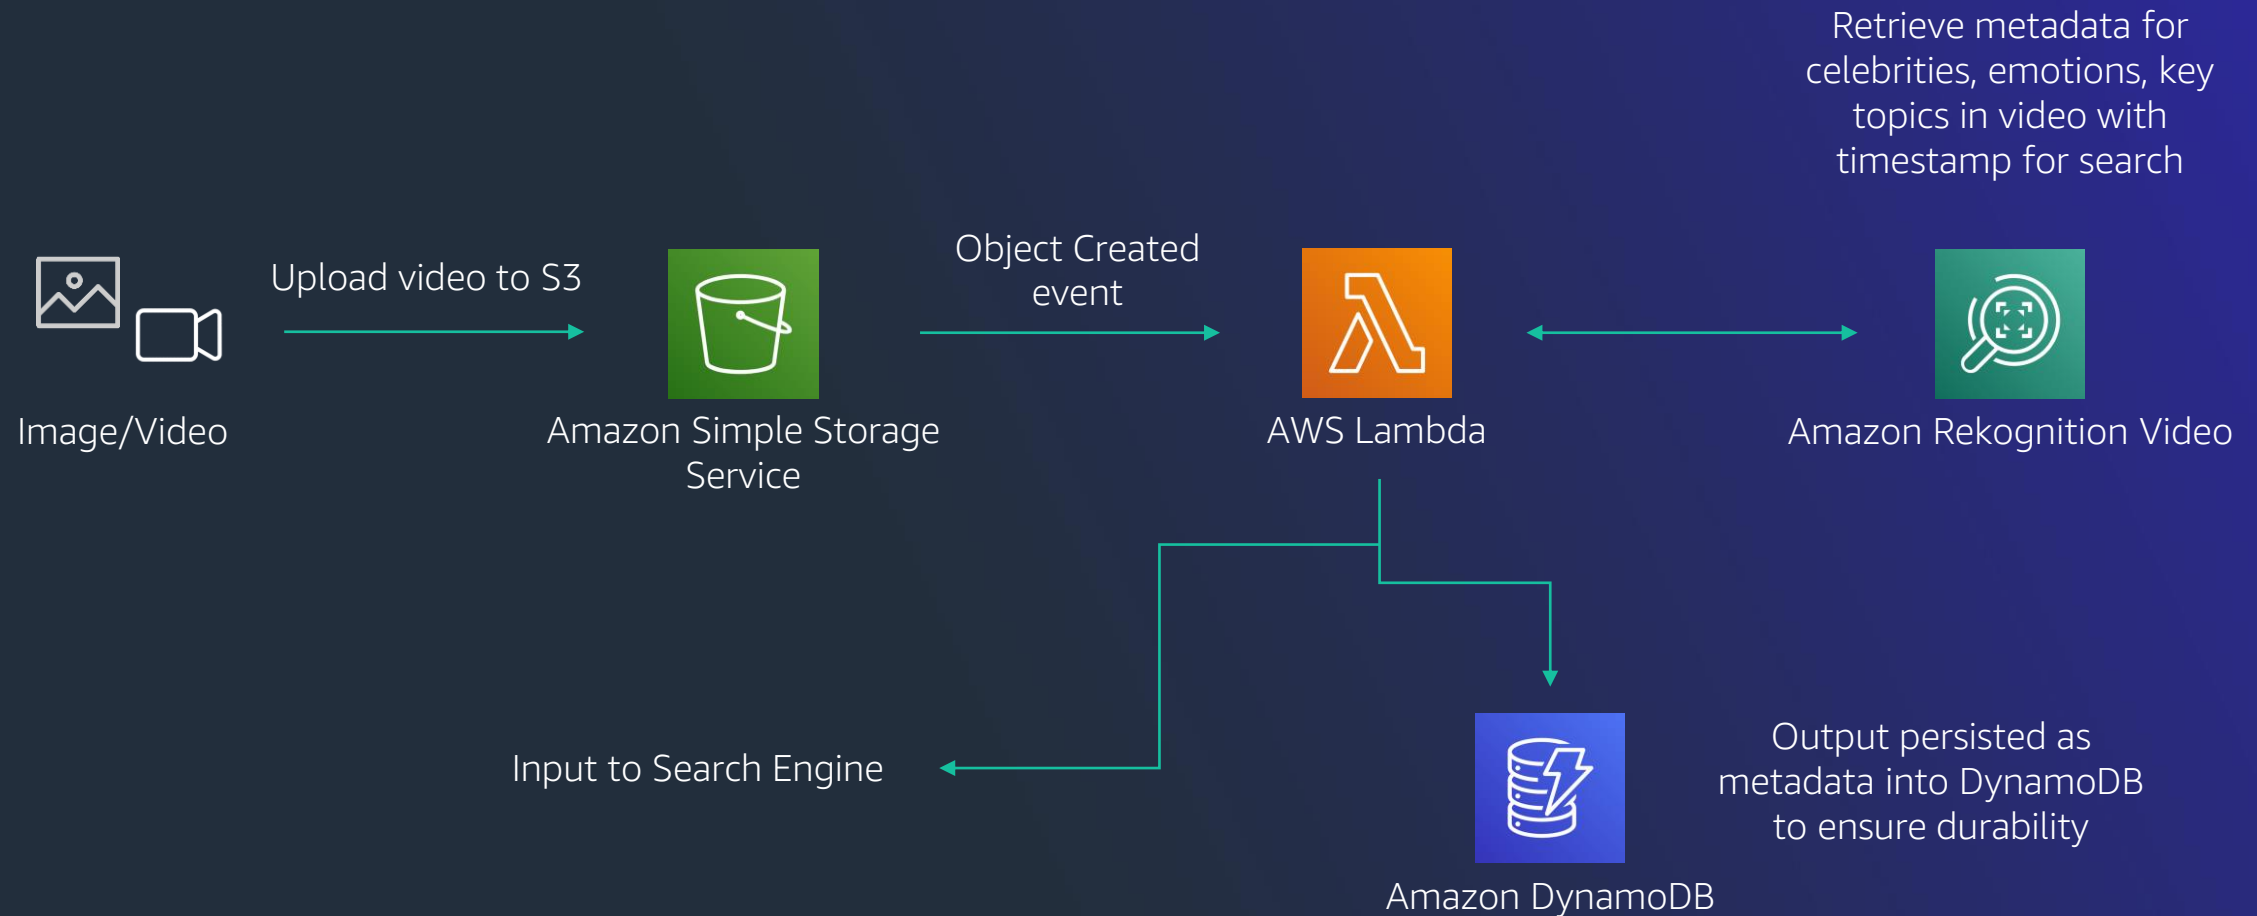
Search & discovery tagging
Content operations
Inappropriate content



INDUSTRIAL & RETAIL






Parts classification
Inventory management
Damage detection

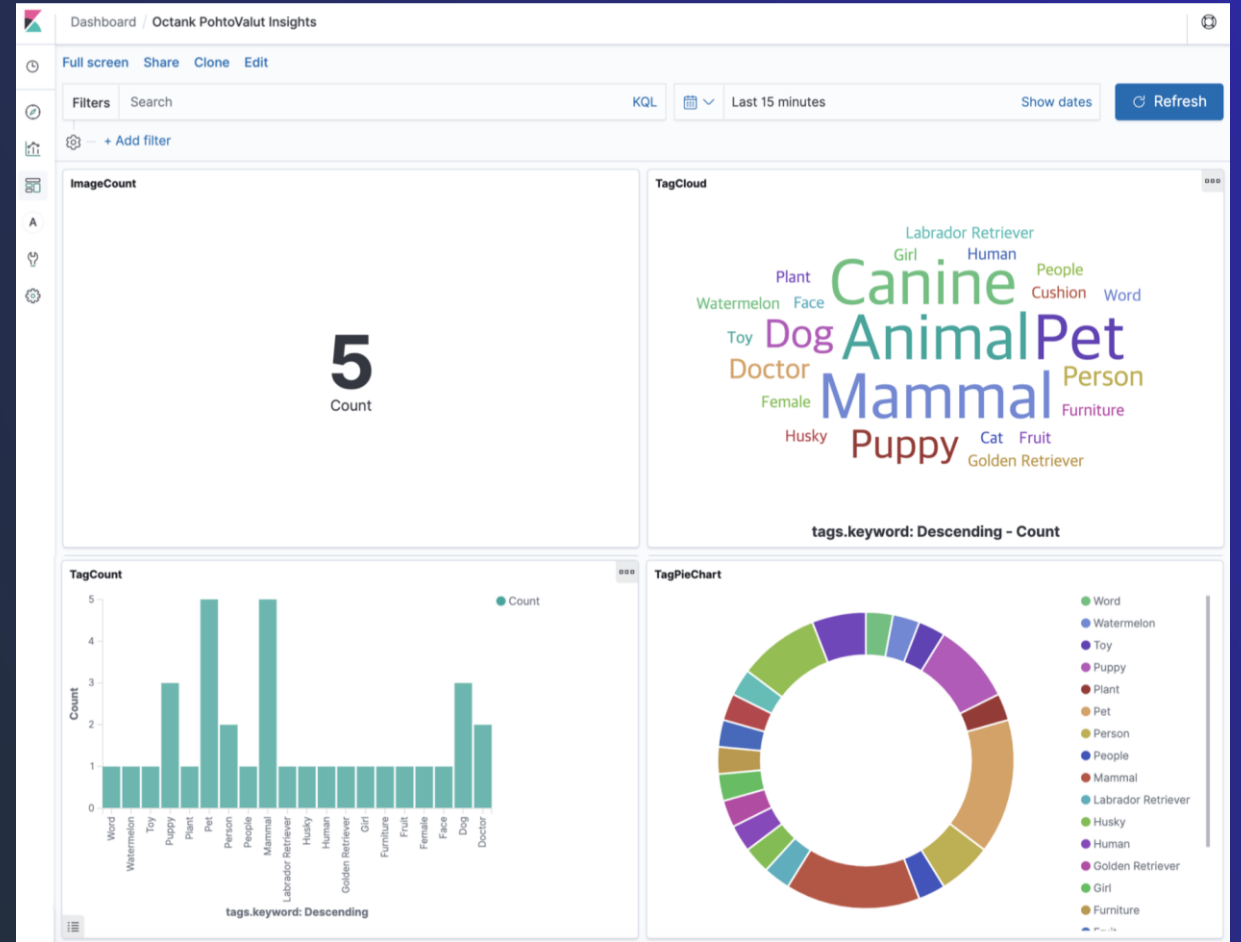
Video/Image Search Index using Amazon Rekognition



Try This – Build Image Insights Dashboard

Upload a file

	20191115_051237.jpg	49.9kB
	20191122_103157.jpg	8.3kB
	20191122_103254.jpg	13.0kB
	20191122_103805.jpg	5.4kB
	20191122_103903.jpg	6.1kB



<https://github.com/ksmin23/image-insights>



Amazon Polly & Transcribe

Text-to-Speech / Speech-to-Text

Amazon Polly



Text-to-speech



Speech Synthesis
Markup Language
(SSML)



Lexicons



Speech
marks



Brand voice

Amazon Polly – Language Support

EMEA

- Danish
- Dutch
- English (UK)
- English (SA)
- French
- German
- Icelandic
- Italian
- Norwegian
- Polish
- Portuguese
- Romanian
- Russian
- Spanish
- Swedish
- Turkish
- Welsh
- Welsh English

Americas

- Brazilian Portuguese
- Canadian French
- English (US)
- Spanish (Mexican)
- Spanish (US)

APAC

- Arabic
- English (AU)
- English (IN)
- Japanese
- Hindi
- **Korean**
- Mandarin
- English (NZ)

Try it – Amazon Polly

Amazon Polly

[Amazon Polly](#) is a cloud service that converts text into lifelike speech. You can use Amazon Polly to develop applications that increase engagement and accessibility. Amazon Polly supports multiple languages and includes a variety of lifelike voices, so you can build speech-enabled applications that work in multiple locations and use the ideal voice for your customers. With Amazon Polly, you only pay for the text you synthesize. You can also cache and replay Amazon Polly's generated speech at no additional cost.

For a full breakdown of the available voices and their respective locales, [view the docs here](#).

In this example, we're going to show how easy it is to send text to [Amazon Polly](#) to synthesize audio.

Methods:

[getSynthesizeSpeechUrl\(\)](#)

Send audio data to Polly, and return the URL where the result will be.

Step 1: Choose a voice


Standard ▾
Seoyeon ▾

Step 2: Write text

Enter the text you would like Polly to say

Voice My Message Using Polly

Step 3: Get Result

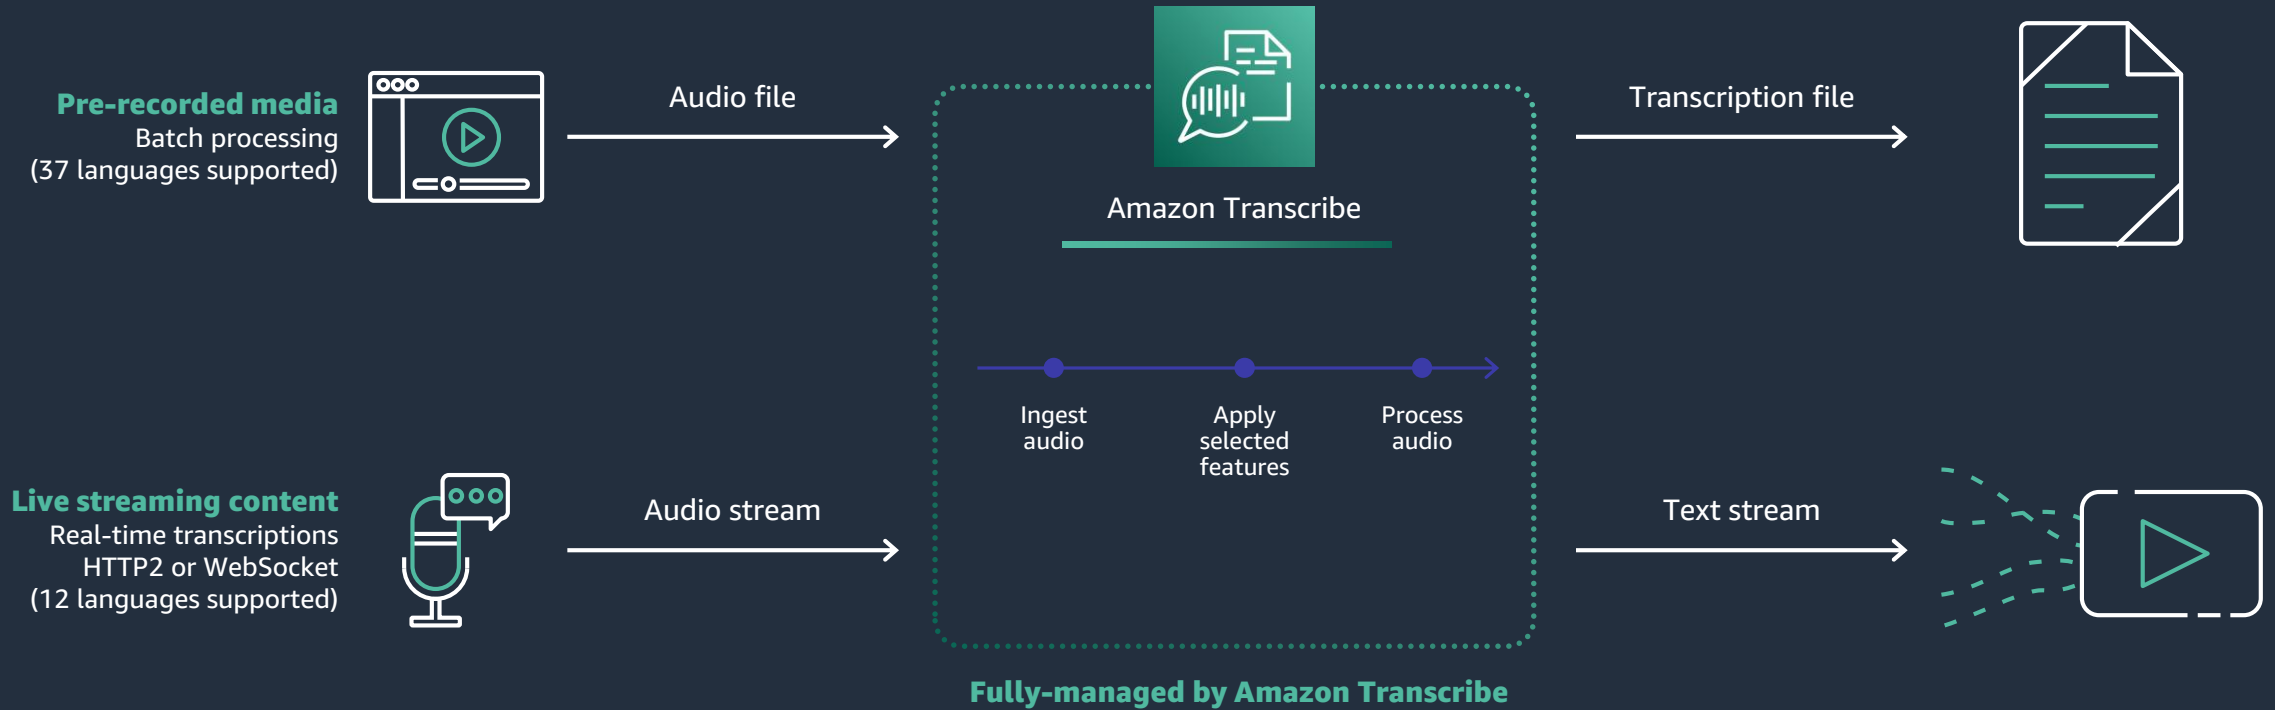


<https://ai-service-demos.go-aws.com/polly>



Amazon Transcribe

A fully-managed and continuously trained automatic speech recognition (ASR) service



Amazon Transcribe Features



**Punctuation,
capitalization**




**Speaker
labelling/channel
identification**



**Word-level
time stamps**



**Word-level
confidence scores**



Language ID



**Custom
vocabulary**



**Custom Language
Models (CLM)**



**Alternative
Transcriptions**

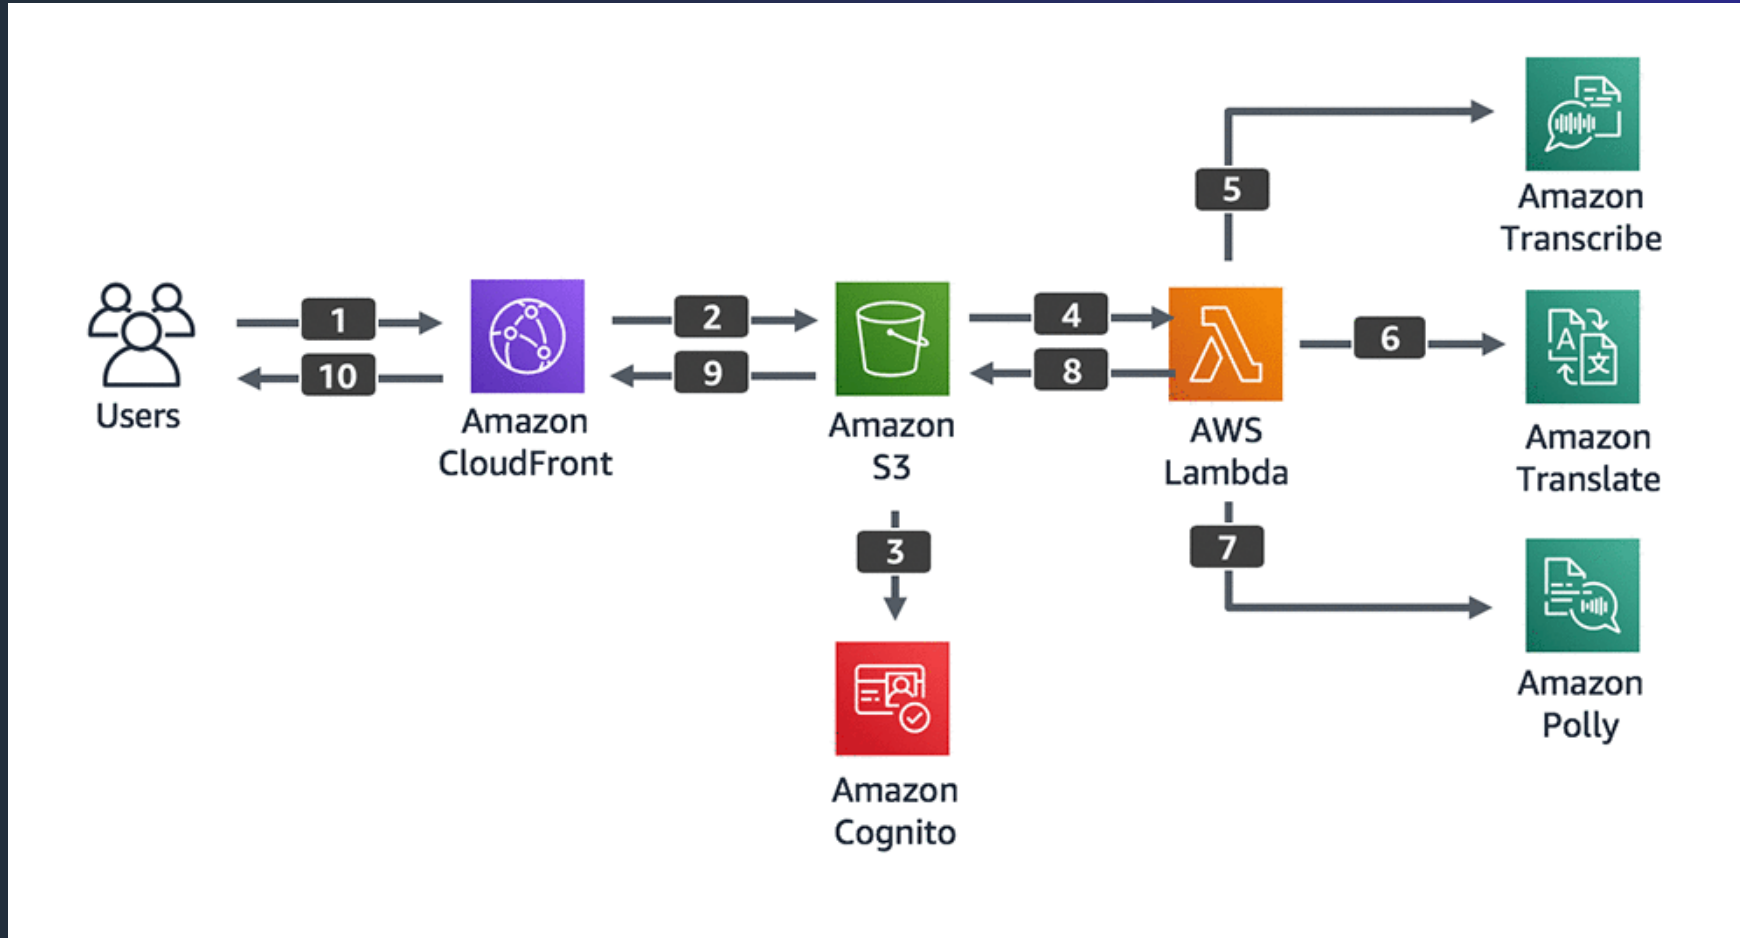


**Vocabulary
Filtering**



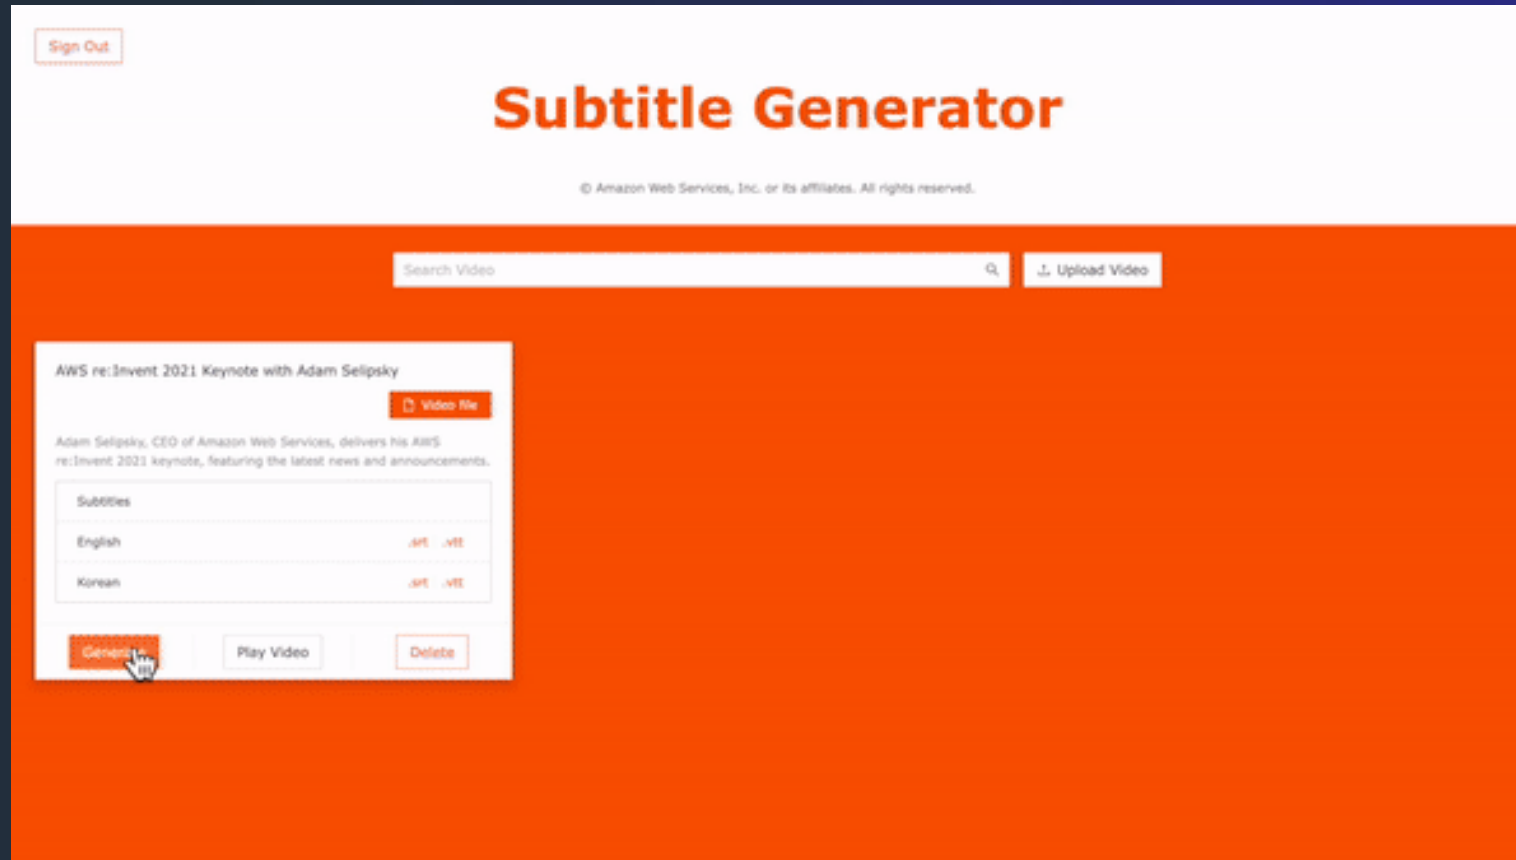
**PCI/PII content
redaction**

원하는 형태로 **AWS** 서비스를 조합



<https://aws.amazon.com/ko/blogs/machine-learning/build-your-own-real-time-voice-translator-application-with-aws-services/>

Try It – Subtitle Generator with Amazon Transcribe



<https://github.com/aws-samples/subtitle-generator-cdk-app>



Amazon Personalize

Improve your customer experience

Amazon Personalize

Create personalized user experiences, no machine learning expertise required



Delivers personalized user experiences faster



Responds quickly to changing user intent using real-time recommendations



Managed ML service reduces time to market



Automates steps to deploy recommendations



Easily integrates with existing systems

Amazon Personalize

How it works

User events/interactions
(views, signups,
conversions, etc.)



Item metadata
(details of articles,
products, videos, etc.)



User metadata
(age, location, etc.)



Amazon
Personalize

Inspect
data

Identify
features


Select
hyper-
parameters

Train
models

Optimize
models

Host
models

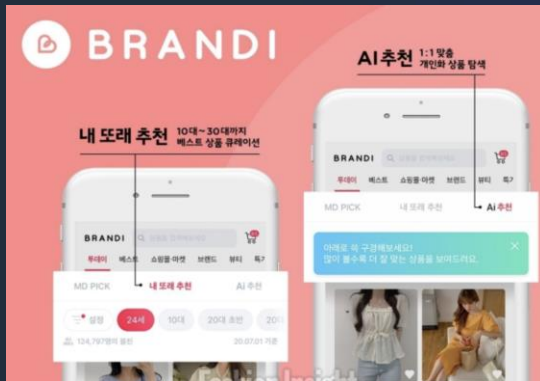
Real-time
feature
store



Customized
personalization
API

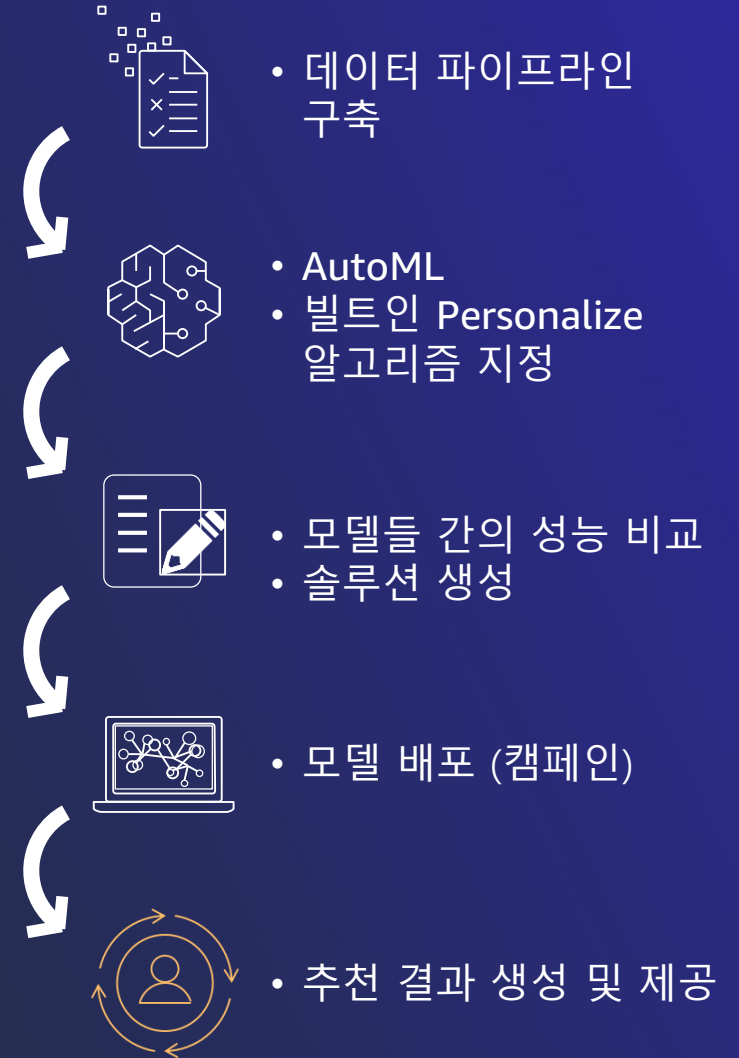
Amazon Personalize – Customer Cases

다양한 고객사례: <https://aws.amazon.com/ko/personalize/customers/>



Amazon Personalize 워크플로우

- Dataset 및 Dataset Group 생성
- 트레이닝 데이터 준비
 - ✓ Dataset group으로 트레이닝 데이터 Import
 - ✓ Dataset group에 사용자 이벤트를 기록(record) - 실시간
- 레시피Recipe를 이용하여 개인화 & 추천 모델(솔루션Solution)생성
- 성능 평가 지표metrics를 이용한 솔루션 성능 분석
- 캠페인 생성 (:= 솔루션 배포)
- 사용자에게 추천 결과 제공
- 사용자 액티비티를 기반으로 한 솔루션의 지속적인 업데이트



Try It – Amazon Personalize A to Z



데이터 준비 및 가공



Amazon Personalize



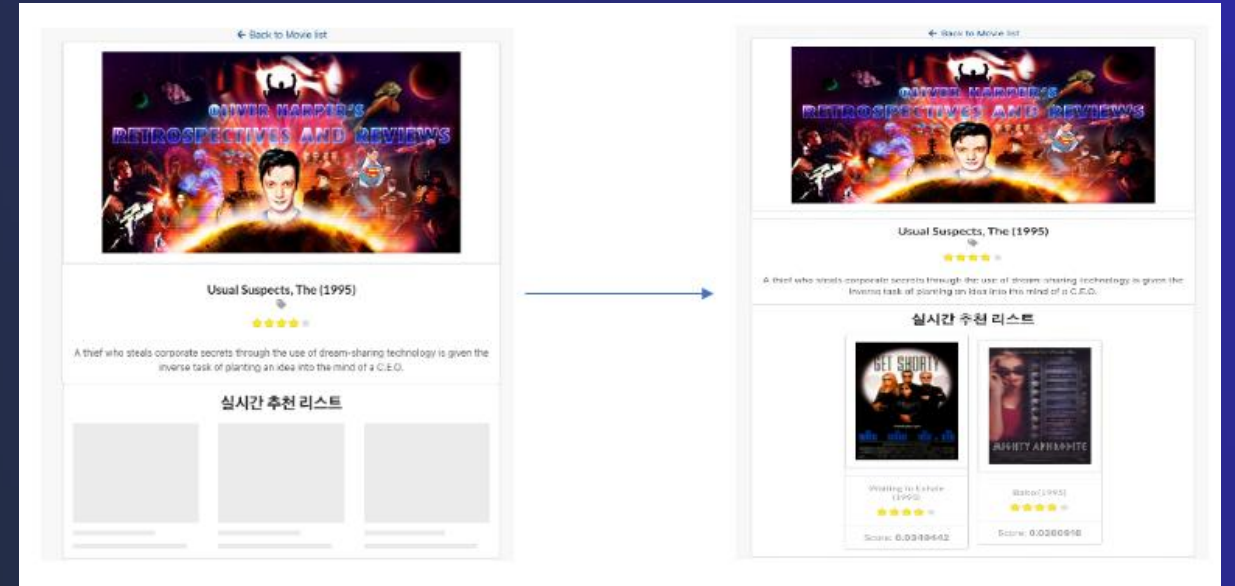
모델 훈련 및 배포



실시간 사용자 행동기반 추천



배치 추천



추천 서비스 결과값 확인하기

<https://catalog.us-east-1.prod.workshops.aws/workshops/ed82a5d4-6630-41f0-a6a1-9345898fa6ec/ko-KR>



Amazon Kendra

Highly accurate intelligent search service powered by ML

Find more accurate answers with Amazon Kendra

EASY TO FIND WHAT YOU ARE LOOKING FOR



Natural language queries

Intuitive search within unstructured content



NLU and machine learning (ML) core

Instant answers
FAQ matching
Document ranking



Broad domain expertise

Pre-trained for 14 domains and industries



Continuous improvement

Incrementally learns from user feedback

QUICK SETUP



Accelerators

Native and partner connectors
Automate ingestion with custom metadata enrichment
No-/low-code UI builder

SECURE



Secure search

Encrypted in transit and rest
Token-based access control
AWS SSO integration

Employees and customers want answers quickly

Traditional search

The image shows a side-by-side comparison of two search interfaces. On the left is a 'Traditional search' interface titled 'Intranet Search'. It features a search bar with the text 'it support desk', a 'Your recent searches' dropdown, and a 'Not finding relevant r' message. Below the search bar are sections for 'SEARCH IN:' (Everything (21), Wiki (17), Email List Archive (3), Show more...), 'REFINE:', 'CATEGORIES' (Service (1), Team (1)), and 'CREATOR' (admin (1), abcde (1), it (1), corp (1)). The search results display a list of items with titles like 'IT_Support_Training_Program.W', 'Com_Support_Wiki.Web', and 'OperationalBestPractices.EventM', each with a snippet and a URL. On the right is an 'Intelligent search' interface titled 'RESULTS PAGE'. It has a search bar with the question 'Where is the it support desk?'. Below the search bar is a 'Kendra's suggested answer' section with the heading '1st floor' and a paragraph of text. Below this is a 'Frequently asked questions' section with three questions: 'Where do I get IT help?', 'What are the IT support hours?', and 'Where can I get IT help corporate campus?'. A white arrow points from the search results in the traditional interface to the intelligent search results.

Intelligent search

Amazon Kendra



Natural language queries

“Where is the IT support desk in Kumo?” → “1st floor”

Keyword queries

“Health benefits” or “parking allowance”

The screenshot shows a search interface with a search bar containing the query "what are less common COVID-19 symptoms?". Below the search bar, it indicates "1-10 of about 94,600" results and a "Sort by Best match" dropdown menu. The top suggested answer is titled "Optimal Control Measures to Combat COVID19 Spread in Sri Lanka: A Mathematical Model Considering the Heterogeneity of Cases". The text of this answer includes a highlighted sentence: "There are some less common symptoms including sputum production, headache, hemophiliacs, and diarrhea [3]". Below this, there is a section for "More suggested answers (2)". The first suggested answer is titled "Long-Term Neurological Threats of COVID-19: A Call to Update the Thinking About the Outcomes of the Coronavirus Pandemic". Its text includes a highlighted sentence: "headache, abdominal pain, diarrhea, nausea, and vomiting (10)".

Amazon Kendra

NLU and ML core



Natural language queries



NLU and ML core

Better answers
Reading comprehension
FAQ matching
Document ranking

Where is the it support desk in kumo?

Kendra's suggested answer

1st floor

... our IT help desk, deskside, which are located all around our buildings and open for support at most hours. The one in Kumo is on the 1st floor and is open from 12:30 to 5 p.m. daily.

WEB | <https://w.amazon.com/bin/view/ILDRTEST/>

Frequently asked questions

Where do I get IT help in Kumo?

What are the IT support hours in Dopler?

Where can I get IT help Seattle campus?

1-10 of 486 Recommended documents

AWS Kumo

...this is how we're doing it. Access controls Access to this host is governed by the permissions associated with the AWS-SUPPORT-KUMO hostclass. Permissions to the aws-support-kumo group... (More)

WEB | <https://w.amazon.com/bin/view/kumo-support>

Amazon Kendra

Domain expertise



Natural language queries



NLU and ML core



Domain expertise

14개 주요 도메인에 최적화

IT, Financial Services, Insurance, Pharma, Industrial, Energy, Legal, Media and Entertainment, Travel and Hospitality, Health, HR, News, Telecommunications, and Automotive

what is the production increase in lithium carbonate?



1-6 of 6 results

Amazon Kendra suggested answers

Can Tesla Really Produce A \$25,000 Self-Driving Electric Car?

It's just like widely available." Related: The World's Most Expensive Crudes Get Expensive Again

There are 63kgs of **lithium carbonate** in a Tesla 70kWh battery. So that's 18.9 million tonnes if everyone wants to trade in their clunker for a Model Y. Global **lithium carbonate** equivalent **production** last year was **365.000 tonnes, an increase of 32% from the year before.** Just 32%?

Clearly, the **lithium** mining industry simply hasn't even been trying. Salting away Musk solves difficult problems from a "first principle standpoint" which he frequently referred to on stage.

<https://s3.us-west-2.amazonaws.com/.../Can Tesla Really Produce A...?...>



What are Amazon Kendra suggested answers? [Info](#)

Amazon Kendra

Incremental learning



Natural language queries



NLU and ML core



Domain expertise



Continuous improvement

Amazon Kendra improves automatically over time

Captures click-through and user feedback

Retrains models periodically

how much amazon refund for postage?

1-10 of about 3,100 results

Amazon Kendra suggested answers

[Amazon.com Help: About Our Returns Policies](#)

up to \$20

Amazon will automatically **refund up to \$20** for return **postage** costs on receipt of the item. If your return **postage** costs exceed \$20, you may contact our Customer Service Center to request a **refund** of the remaining **postage** cost. If you return a defective, damaged or incorrect item, youâ??ll be **refunded** the full **postage** cost, as well as your Import Fees Deposit, once your return is processed. **Amazon** Alexa Voice Shopping Non-digital products purchased using Alexa are eligible for free returns. To return a product purchased using Alexa, process your return as normal in the Returns Center.

<https://s3.us-east-1.amazonaws.com/amazon-service.../201819200.html>

What are Amazon Kendra suggested answers? [Info](#)

Sort: Relevance ▼

[Amazon.com Help: About Our Returns Policies](#)

...Customer Service Center to request a **refund** of the remaining **postage** cost. If you return a defective, damaged or incorrect item, youâ??ll be **refunded** the full **postage** cost, as well as your Import Fees Deposit, once your return is processed. **Amazon** Alexa Voice Shopping Non-digital...

<https://s3.us-east-1.amazonaws.com/amazon-service.../201819200.html>

Amazon Kendra

Native connectors



Natural language queries




NLU and ML core



Native and partner connectors

Popular data sources including Microsoft SharePoint, Atlassian Confluence and Jira, Adobe Experience Manager, ServiceNow, and Salesforce

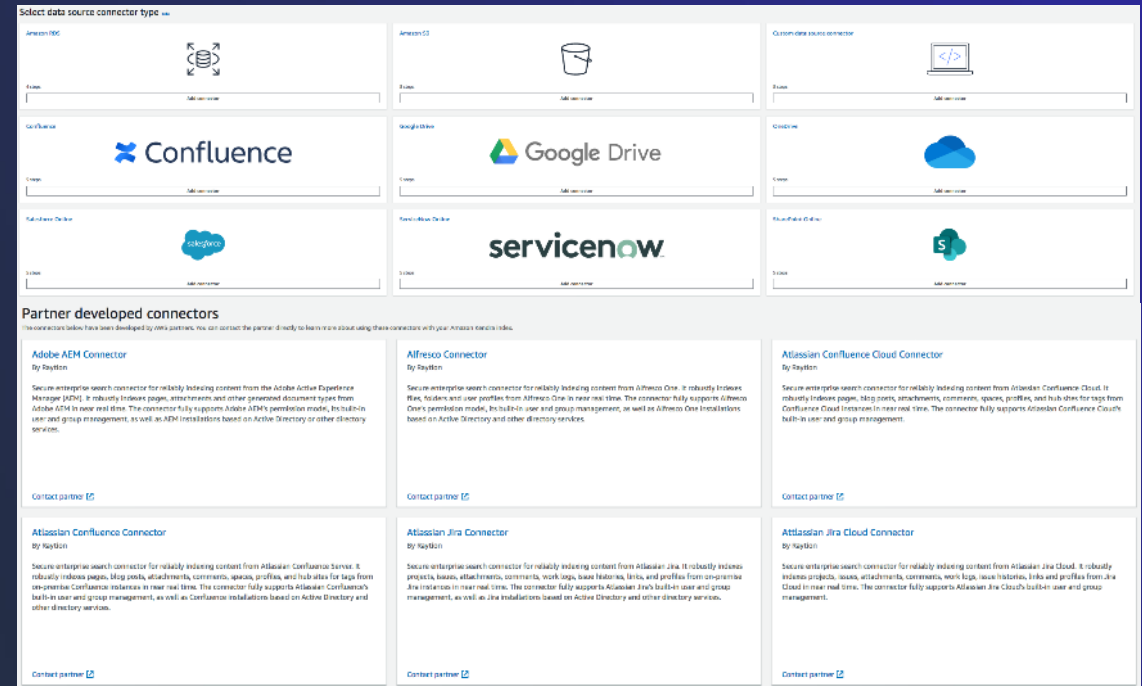
Add credentials, then point and shoot



Domain expertise



Continuous improvement



Amazon Kendra

Native connectors




Natural language queries



NLU and ML core



Native and partner connectors



Domain expertise

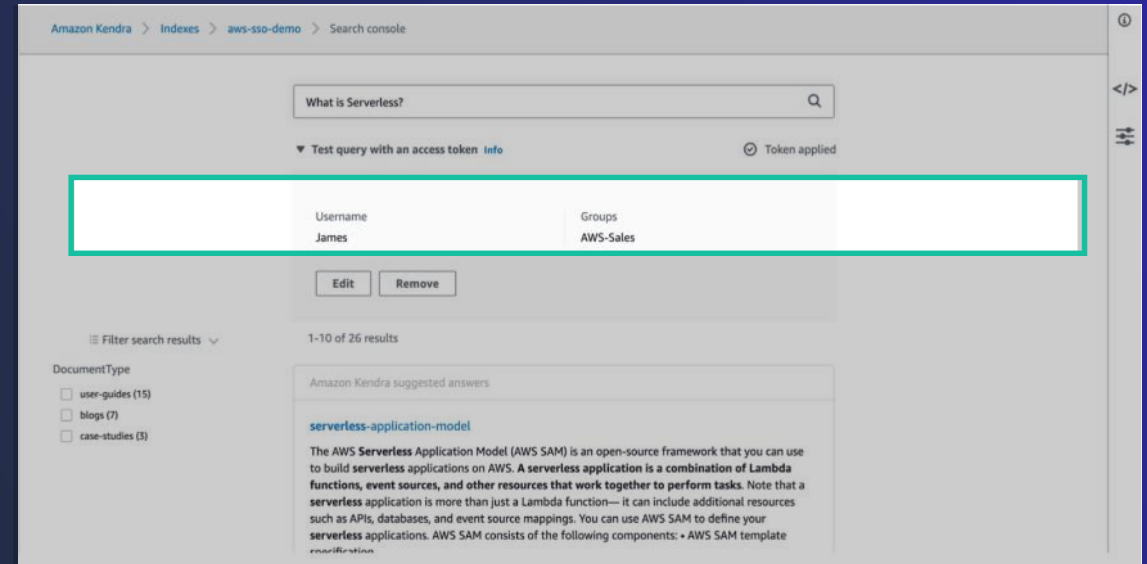


Continuous improvement



Secure search

Highly secure

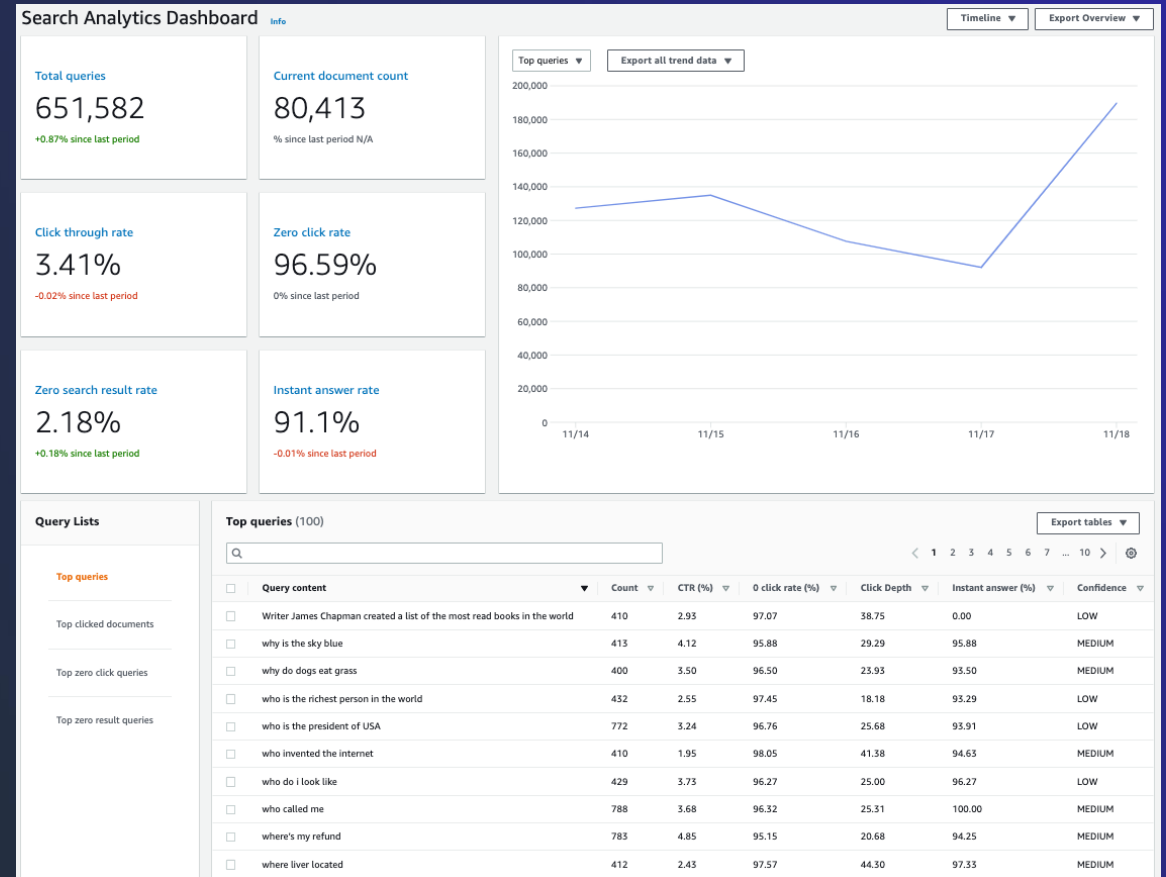


Data encrypted in transit and at rest

Token-based user access control

Amazon Kendra analytics dashboard

- Key quality and usability metrics
- Helps admins and content creators
 - Understand how easily users are finding relevant search results
 - Quality of the search results
 - Gaps in the content
 - Role of instant answer in providing answers to users' results
- Viewed in console dashboard or accessed through an API



Amazon Kendra Experience Builder – key features


- Deploy Amazon Kendra Search UI in a few clicks
- Drag and drop Search UI builder
- AWS SSO integration for secure access and search filtering
- Fully managed/hosted
- No ML or coding expertise needed

The AWS AI/ML Stack – Choose What you Need!

AI SERVICES



ML SERVICES



ML FRAMEWORKS & INFRASTRUCTURE

PyTorch, Apache MXNet, TensorFlow

Amazon EC2

CPUs

GPUs

AWS Inferentia

AWS Trainium

Habana Gaudi

FPGA

Elastic inference




Start your Journey with AWS



TRAINING



PROOF OF CONCEPT



PRODUCTION

Resources

AWS ML/AI website & documentation

<https://ml.aws>

AWS Artificial Intelligence Services website

<https://aws.amazon.com/machine-learning/ai-services/>

AWS Machine Learning Blog

<https://aws.amazon.com/blogs/machine-learning/>

AWS Training and Certification portal

<https://www.aws.training/LearningLibrary>



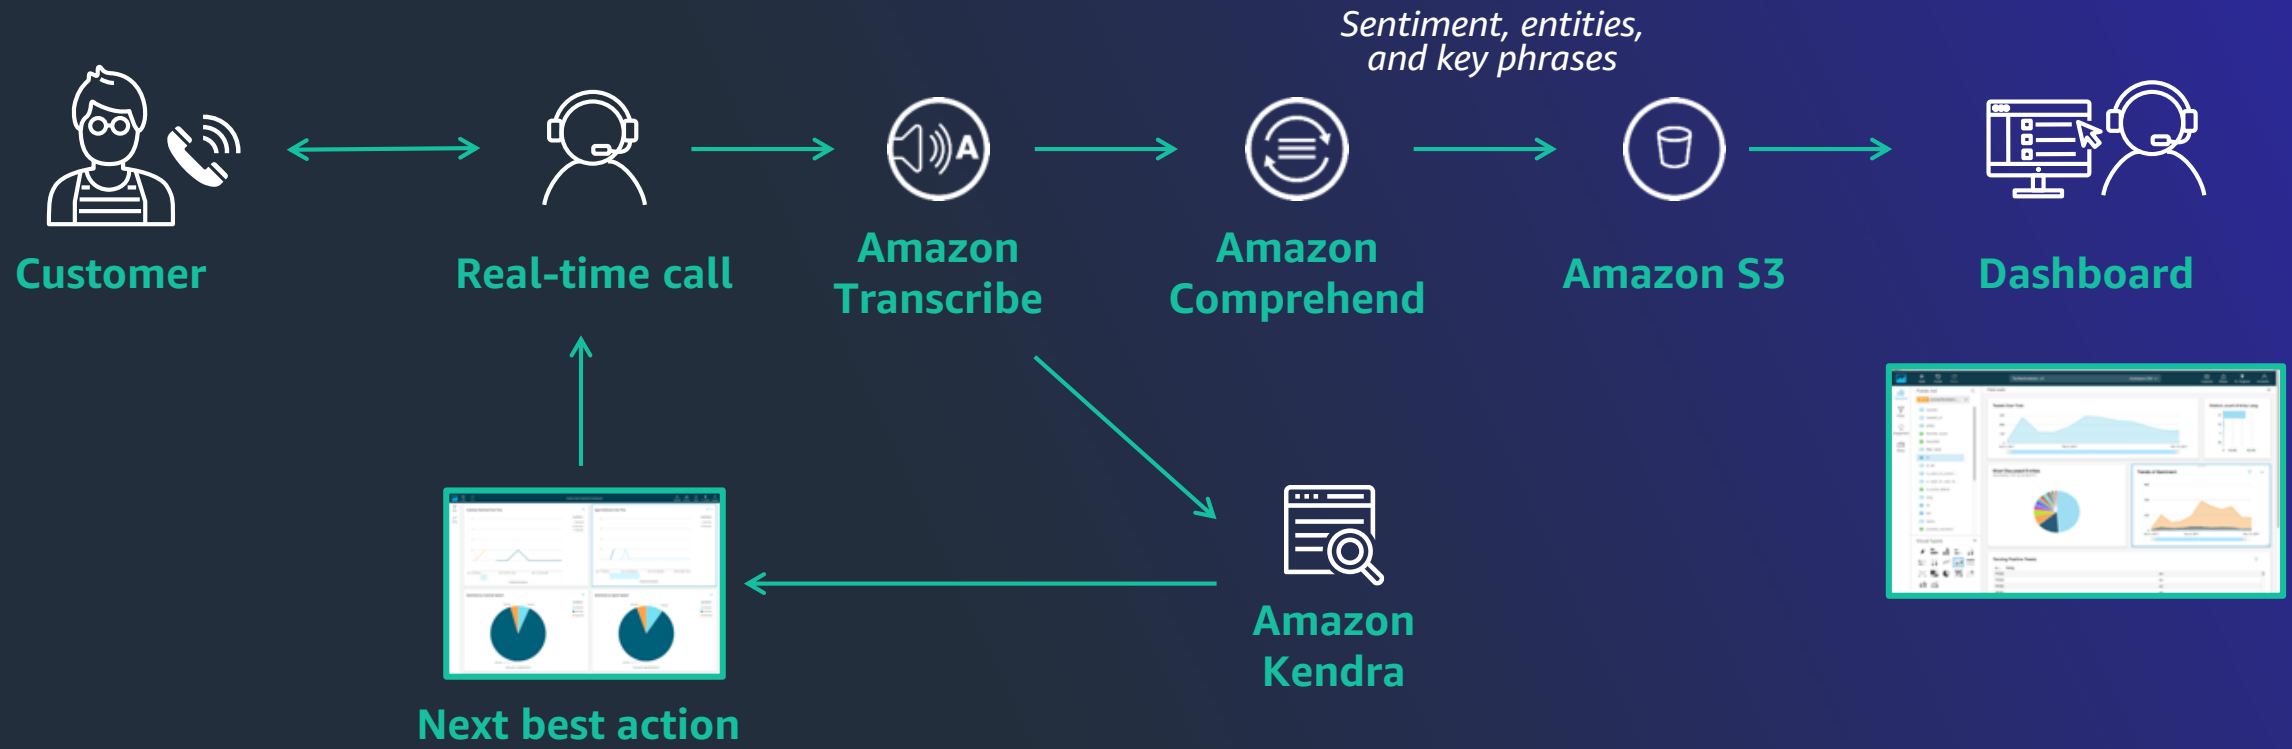
Thank you!

Jonathan Kim

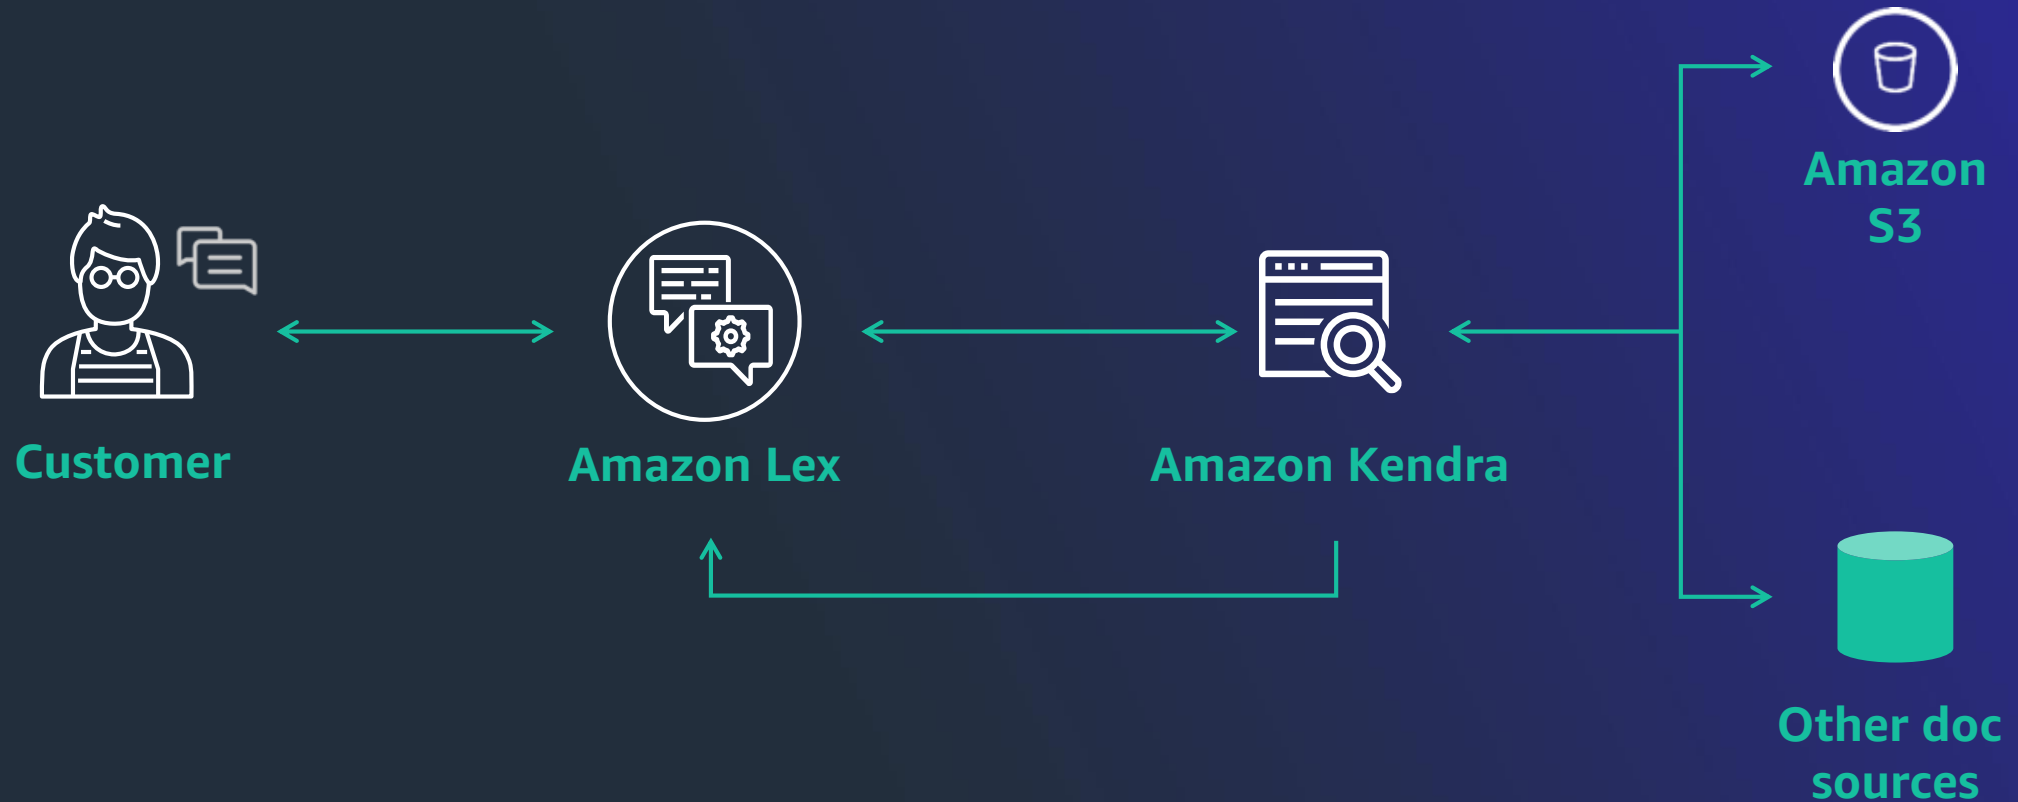
Solutions Architect, AWS Korea



Intelligent search for agent assist



Self-service chatbots with intelligent search



Should I use Amazon Kendra or Amazon OpenSearch Service?

Amazon Kendra	Amazon OpenSearch Service
<p>Why? No-code, ML-powered, fully managed search service with built-in connectors, WYSIWYG search experience builder, access control, and incremental learning</p>	<p>Why? Managed service delivering OpenSearch – the open-source search and analytics engine. Use for log-based monitoring. Build on its REST APIs for a highly customized search experience.</p>
<p>Use cases Internal search (R&D, compliance, enterprise-wide), contact center/support (agent assist, chatbot Q/A, website), ISV (Salesforce, Atlassian, DevFactory, and so on)</p>	<p>Use cases Logs – infrastructure, application, IoT, and security monitoring. Trace analytics for observability. ML features for anomaly detection. Search – eCommerce, ad serving.</p>
<p>Type of search High accuracy with keyword and natural language Q/A: <i>Ex: Q: What time is the IT desk open? A: 9:00 am – 5:00 pm</i> Rich keyword-based query language for complex queries: <i>ex: title:"history" AND "movies" NOT "fiction"</i></p>	<p>Type of search Rich keyword-based query language for complex queries Built-in, adjustable ranking Aggregations for faceted search and log analysis <i>Ex: color:"blue" AND category:"men's clothing" AND description:"wool"</i></p>
<p>Who is the ideal persona for this service? CIO/VP/Director responsible for employee productivity, R&D, compliance, knowledge management VP/Director responsible for customer service VP/Director of engineering at ISVs, building applications with intelligent search</p>	<p>Who is the ideal persona for this service? CIO/VP/Director of engineering managing IT departments, developers using DevOps CISO for enterprise security monitoring and compliance</p>
<p>How do customers get data to this service? Amazon Kendra has pre-built connectors with Amazon RDS, Amazon S3, Amazon WorkDocs, Atlassian Confluence + Cloud, Microsoft OneDrive, Microsoft SharePoint, ServiceNow, Salesforce, Google Drive, Web Crawler, and custom data through API. Total catalog is 100+ connectors.</p>	<p>How do customers get data to this service? Can use native integration with Amazon Kinesis Data Firehose or other AWS services to build and manage your own ingestion ecosystem</p>
<p>Pricing Discounted pricing is available for customers who are willing to make certain minimum volume commits. Please contact AWS for details. Credits available for POCs.</p>	<p>Pricing Charges for instances, storage, and networking. On-demand, and Reserved Instances. Tiered storage for cost-optimal, long-term retention of log data. Credits available for POCs.</p>