

AWS
re:Invent

AIM303

Stop guessing: Use AI to understand customer conversations

Dirk Fröhner

Senior Solutions Architect
Amazon Web Services

Boaz Ziniman

Principal Technical Evangelist
Amazon Web Services

Agenda

- Introduction Amazon Web Services (AWS) AI/ML offering
- Introduction Amazon Connect
- Architecture of the labs
- Work on the labs
- Wrap-up


Related breakouts

- AIM211 - AI document processing for business automation
- AIM212 - ML in retail: Solutions that add intelligence to your business
- AIM222 - Monetizing text-to-speech AI
- AIM302 - Create a Q&A bot with Amazon Lex and Amazon Alexa

Intro: AWS AI/ML Services

Natural language processing (NLP)

- Automatic speech recognition (ASR)
- Natural language understanding (NLU)
- Text to speech
- Translation





amazon



UBER

Tide

AMEX



THIS DAY IN
HISTORY



eipi

ROYAL OPERA HOUSE

MR. ROBOT
DAILY FIVE / NINE

StubHub!



TED

WebMD

SHOWTIME



100,000+



Use cases for NLP



Voice of customer applications



Knowledge management




Education



Customer service/
call centers



Semantic search



Accessibility



Enterprise
digital assistant



Captioning workflows



Information bots




Personalization




Localization

The Amazon ML stack: Broadest & deepest set of capabilities

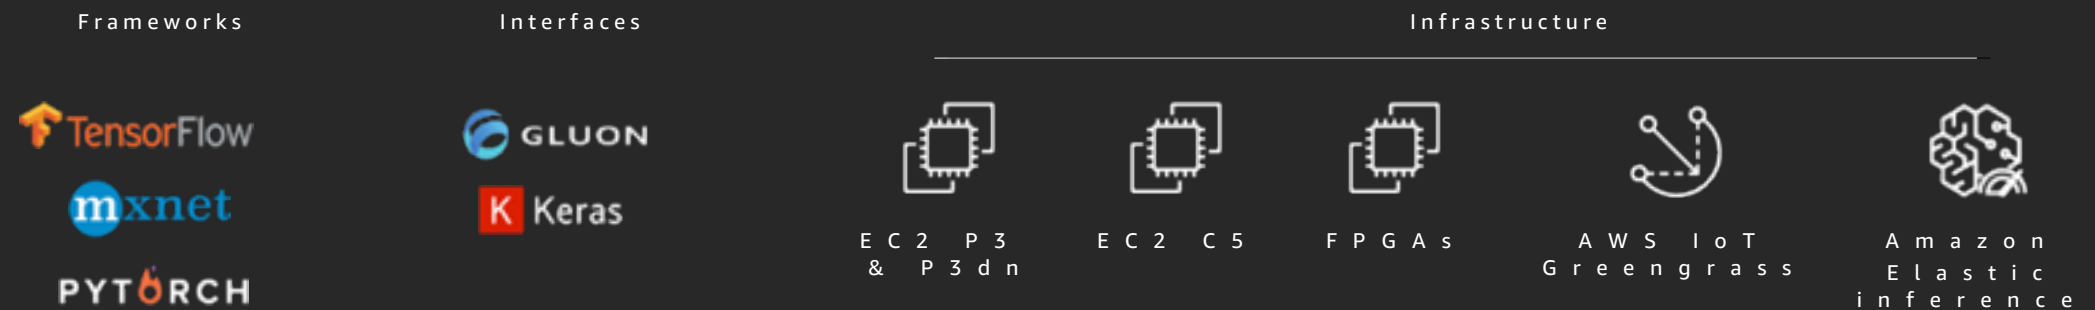
AI services



ML services





ML frameworks & infrastructure




AI services

AI services



Pretrained AI services that require no ML skills or training





Ability to easily add intelligence to your existing apps and workflows




Quality and accuracy from continuously learning APIs

AI services

AI services



Pretrained AI services that require no ML skills or training



Ability to easily add intelligence to your existing apps and workflows



Quality and accuracy from continuously learning APIs

Amazon Polly

Turn text into lifelike speech using deep learning

Amazon Polly: Use cases




Navigation



Amazon Polly

Podcasting,
voiced blogs
& news articles



Special needs
AI assistant



Language
learning



Voiced videos
& presentations




Contact
centers



Amazon Polly: Text in, lifelike speech out

"Today in Seattle, WA,
it's 11°F"



Amazon Polly



"Today in Seattle, Washington,
it's 11 degrees Fahrenheit"




29 languages



Amazon Polly: Text in, lifelike speech out

“Genießt ihr den Rest
des Tages”



Amazon Polly



“Genießt ihr den Rest
des Tages”



29 languages

Add semantic meaning to text

<https://www.w3.org/TR/speech-synthesis/>

< speak >

The spelling of my name is

< prosody rate='x-slow' >

< say-as interpret-as="characters" >Boaz< /say-as >

< /prosody >



< /speak >

Standard vs. neural TTS

Text

Sentence to synthesize.



'sɛntəns tə 'sɪnθə ,saɪz.

Standard vs. neural TTS

Text

Sentence to synthesize.

Phonetic transcription

'sɛntəns tə 'sɪnθə ,saɪz.

Concatenative TTS



Neural TTS

'sɛnt

sɛntəns tə

'sɪnθ

ə ,saɪz.



Neural TTS with Amazon Polly

Standard vs. neural

“Hi. My name is Boaz, and I will be your host today during this Innovate online session.”



Standard

Neural TTS with Amazon Polly

Standard vs. neural

“Hi. My name is Boaz, and I will be your host today during this Innovate online session.”



Standard



Neural TTS

Neural TTS with Amazon Polly

Standard vs. neural

“Hi. My name is Boaz, and I will be your host today during this Innovate online session.”



Standard



Neural TTS

Styles with neural TTS

“Hi. My name is Boaz, and I will be your host today during this Innovate online session.”



Newscaster NTTS

A person's hands are shown holding a white smartphone. The background is a blurred office environment with other people working at desks. The overall tone is professional and modern.

The Washington Post

Add voice to your app



“This is a new technology that can give users more choice and better accessibility to our content, so we wanted to create an experiment to dive deeper into the user experience. After a month, we’ll take what we’ve learned about how users engage with this feature to develop our first iteration of a product with **Amazon Polly**.”

Joseph Price, Product Manager
The Washington Post

Amazon Transcribe

Automatic speech recognition

Automatic speech recognition service



Amazon Transcribe



“Hello. This is Allan speaking.”

Amazon Transcribe: Key features



Custom
vocabulary



Speaker
identification



Channel
identification



Punctuation and
capitalization



Word-level
time stamps



Word-level
confidence scores

Amazon Transcribe: Streaming transcription

Bidirectional stream over:

HTTP/2 protocol


or

WebSocket (WSS) protocol



Source file or
microphone

Audio stream



Amazon Transcribe

Text stream





Speech to text

RingDNA is an end-to-end communications platform for sales teams.

Hundreds of enterprise organizations use RingDNA to increase productivity, engage in smarter sales conversations, gain predictive sales insights, and improve their win rate.

*"A critical component of RingDNA's Conversation AI requires best-of-breed speech-to-text to deliver transcriptions of every phone call. RingDNA is excited about **Amazon Transcribe** since it provides high-quality speech recognition at scale, helping us to better transcribe every call to text."*

Howard Brown, CEO & Founder
RingDNA

Amazon Translate

Natural and accurate language translation

Key features



54 languages

2,804 combinations



\$15/1M characters

Or \$0.000075 per word; pay as you go, 2M characters monthly free tier



Real-time

< 500ms / sentence on average
< 150ms / conversational / short form



Tag handling

XML tag placement maintains styling and formatting through translation



Data security

Data ownership
Encryption
Access management



Ease of use

Simple API calls and partner solutions

HIPAA eligible

Amazon Translate

Natural and fluent language translation

"Hello, what's up? Do you want to go see a movie tonight?"



Amazon Translate



"Hallo, wat is er? Wil je vanavond naar de film gaan?"

Translate API example

```
import boto3
translate = boto3.client("translate")
lang_flag_pairs = [
    ("fr", "🇫🇷"), ("de", "🇩🇪"), ("es", "🇪🇸"),
    ("pt", "🇵🇹"), ("zh", "🇨🇳"), ("ja", "🇯🇵"),
    ("ru", "🇷🇺"), ("it", "🇮🇹"), ("zh-TW", "🇹🇼"),
    ("tr", "🇹🇷"), ("cs", "🇨🇪"), ("he", "🇮🇱")]
for lang, flag in lang_flag_pairs:
    print(flag)
    print(translate.translate_text(
        Text="Hello, World",
        SourceLanguageCode="en",
        TargetLanguageCode=lang
    )['TranslatedText'])
```

Translate API example



<https://github.com/ziniman/aws-translate-demo>



Bonjour, Monde



Hallo, Welt



Hola, Mundo



Olá, Mundo



您好, 世界



ハローワールド



Привет, Мир



Ciao, Mondo



大家好, 世界



Merhaba, Dünya.



Ahoj, světe.



שלום, עולם.










Amazon Comprehend

Discover insights and relationships in text

Amazon Comprehend

Discover insights and relationships in text



-  Entities
-  Key phrases
-  Language
-  Sentiment
-  Syntax
-  Grouping

Run Amazon Comprehend on an S3 bucket

```
import boto3
import json

s3 = boto3.resource('s3')
bucket_name = 'my_bucket'
region_name = 'us-east-1'
bucket = s3.Bucket(bucket_name)

comprehend = boto3.client(service_name='comprehend', region_name=region)

for obj in bucket.objects.all():
    body = obj.get()['Body'].read()
    text = body

    sentiment_response = comprehend.detect_sentiment(Text=text, LanguageCode='en')
    print(json.dumps(sentiment_response, sort_keys=True, indent=4))
```

"The Babel fish is small, yellow, leech-like—and probably the oddest thing in the universe. It feeds on brain wave energy, absorbing all unconscious frequencies and then excreting telepathically a matrix formed from the conscious frequencies and nerve signals picked up from the speech centres of the brain, the practical upshot of which is that if **you stick one in your ear, you can instantly understand anything said to you in any form of language**: the speech you hear decodes the brain wave matrix."

Douglas Adams

The Hitchhiker's Guide to the Galaxy

Build your own Babel fish




Person speaks in
English

Build your own Babel fish



Person speaks in
English



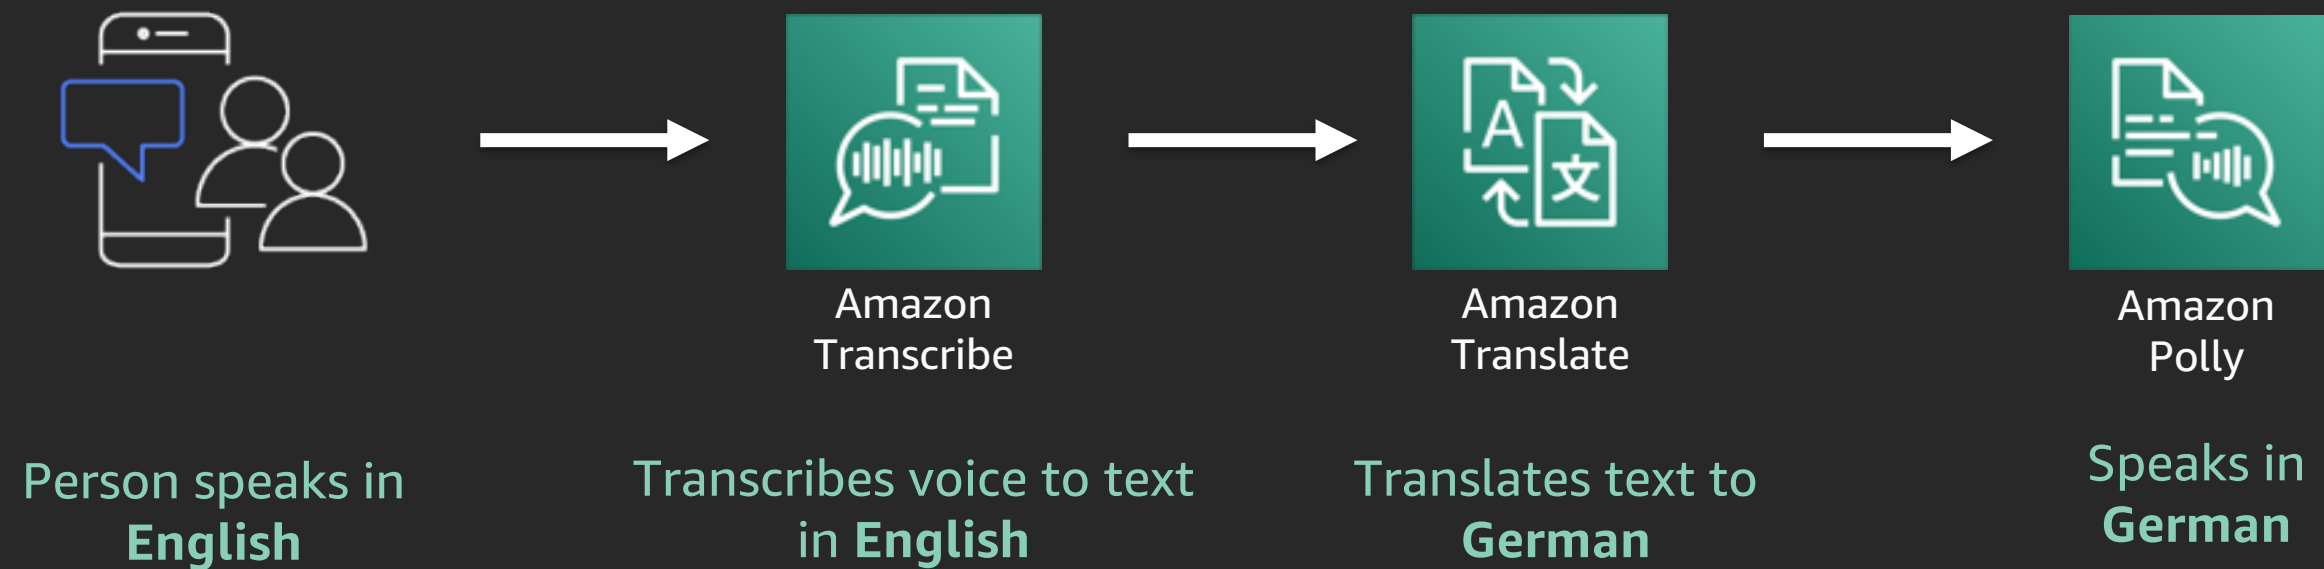
Amazon
Transcribe

Transcribes voice to text
in **English**

Build your own Babel fish



Build your own Babel fish



Demo: Real-time audio transcription & translation

<https://voice.boaz.cloud/>




Intro: Amazon Connect

Amazon Connect



Easy-to-use, cloud-based contact center solution
that scales to support businesses of any size

With tools that grow with your needs



Skills-based routing
(Automatic call distribution [ACD])



Call
recording

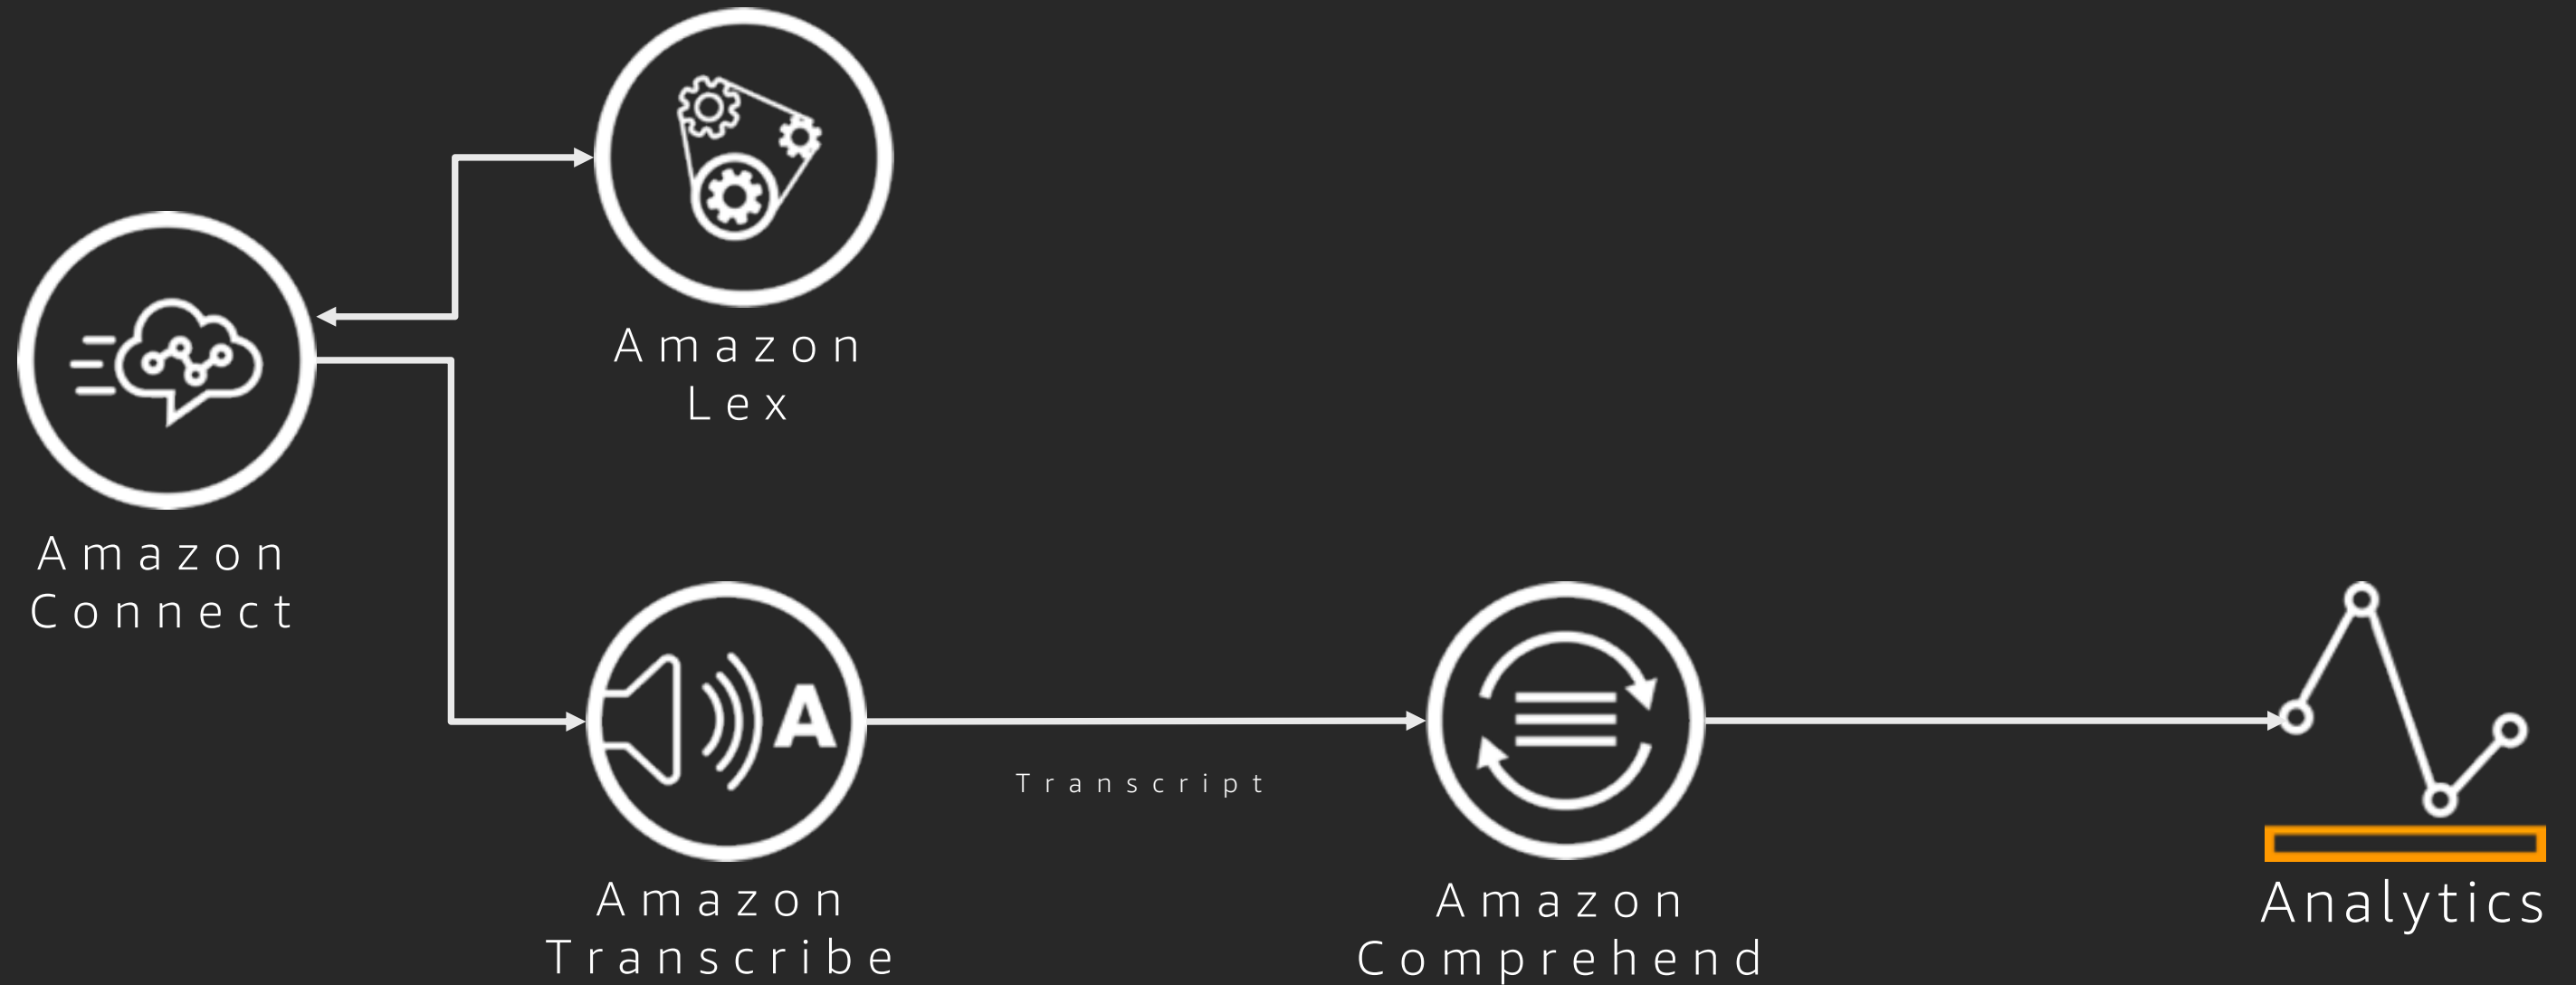


Real-time and
historical analytics



High-quality
voice capability

Improving contact centers with artificial intelligence



What we will build today

Our use-case today

Processing of contact center
call recordings

Continual analyses for
sentiment and key words

Actionable insights from data
as strategic asset

Cloud-first & serverless



Choose your path

Labs for different experience levels to respect a heterogeneous audience

Foundation lab to generate or obtain sample audio recordings

Exploration path for less advanced participants

Deep-dive path for more advanced participants



Choose your path

Exploration


labs 1–3

Deep-dive

labs 4–6

Foundation

lab 0



Foundation


Foundation pillar for both paths

Create or obtain your sample call recordings



Foundation: Lab 0


Create or obtain your sample call recordings, three options available



Foundation: Lab 0

Create or obtain your sample call recordings, three options available


Amazon
Connect




Foundation: Lab 0

Create or obtain your sample call recordings, three options available

Amazon
Connect




Amazon
Polly




Foundation: Lab 0

Create or obtain your sample call recordings, three options available


Amazon
Connect



Amazon
Polly




Amazon
S3




Foundation: Lab 0

Create or obtain your sample call recordings, three options available



Amazon
Connect




Amazon
Polly



Amazon
S3



Amazon
S3



Exploration path

Manual soldering in the AWS management console

Learning objectives:

- Explore speech and language AI services
- Understand how these services can help you without the need to train your own ML models



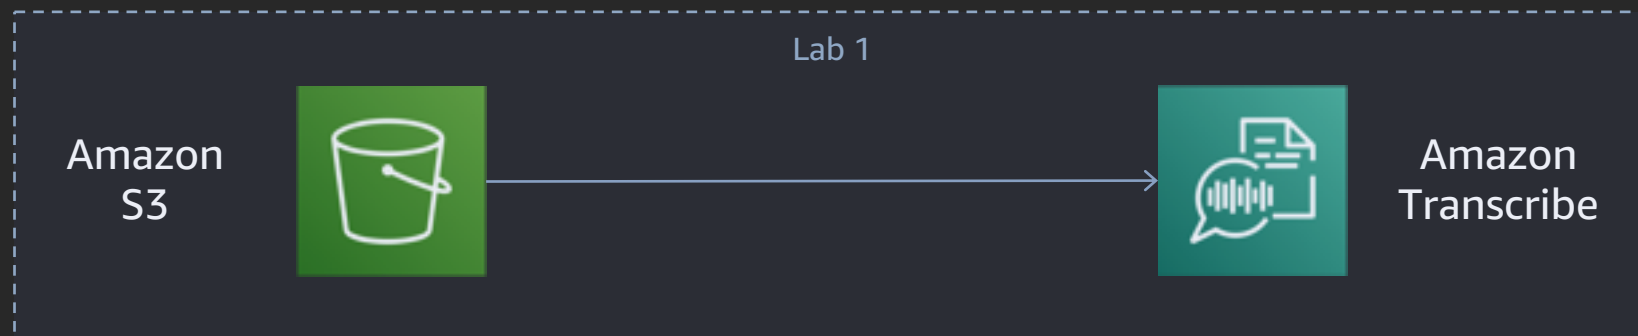
Exploration: Labs 1–3

Manually run through ASR and NLP analyses, look into language translation as well



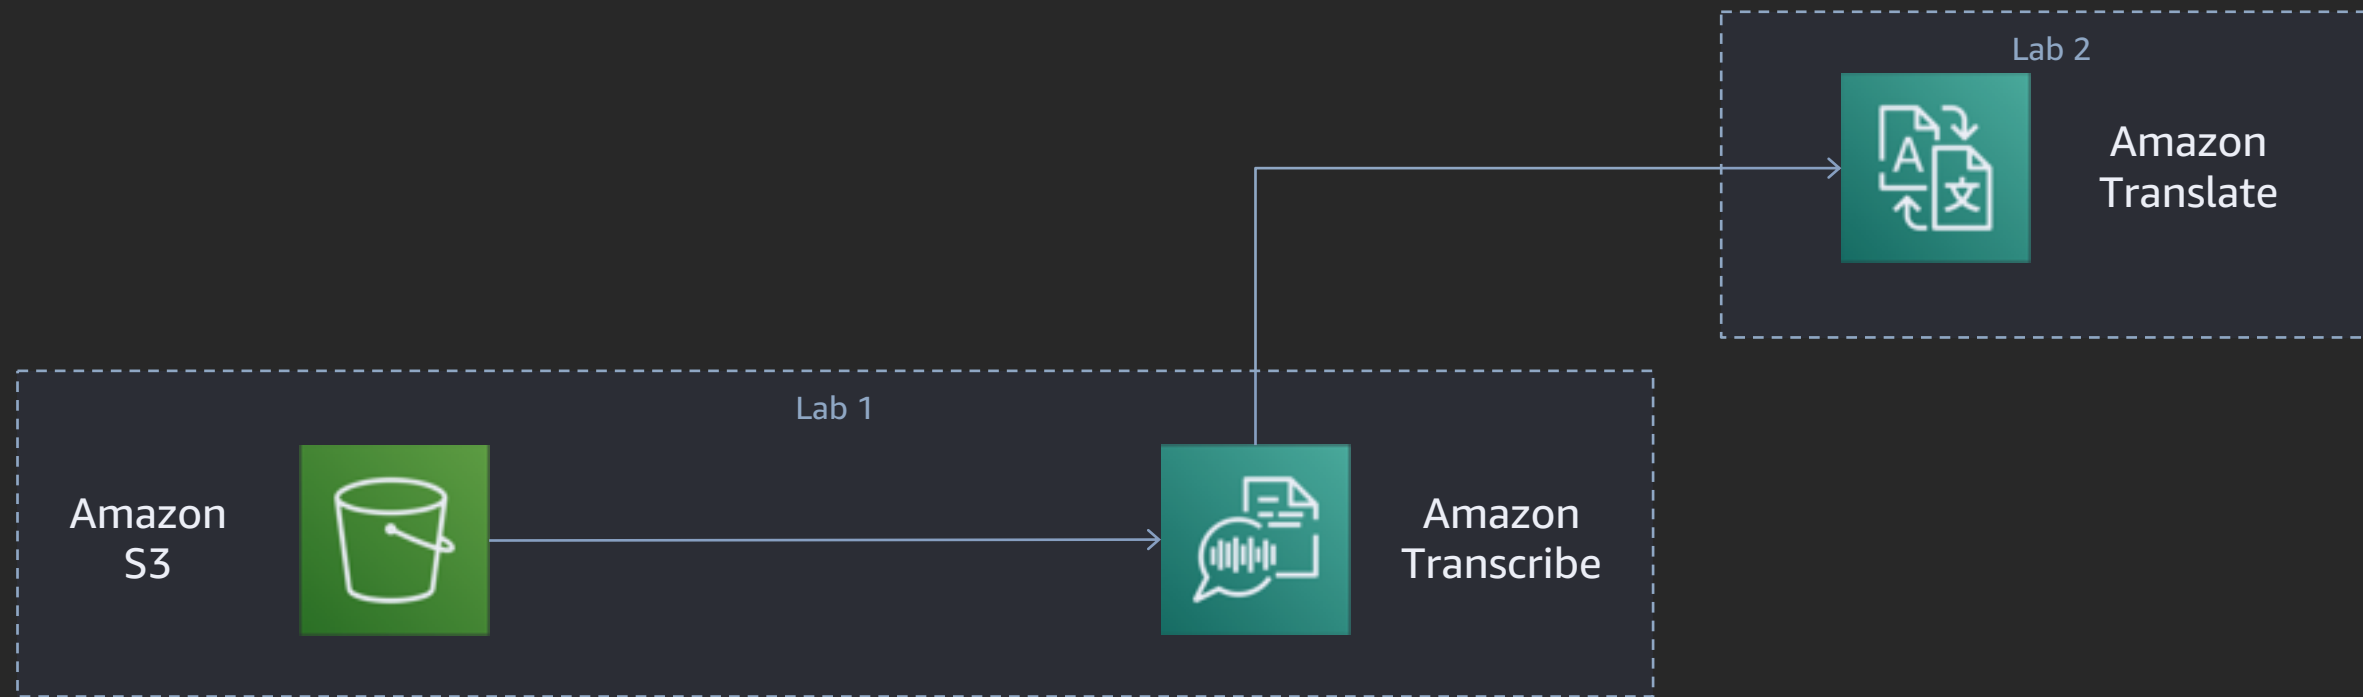
Exploration: Labs 1–3

Manually run through ASR and NLP analyses, look into language translation as well



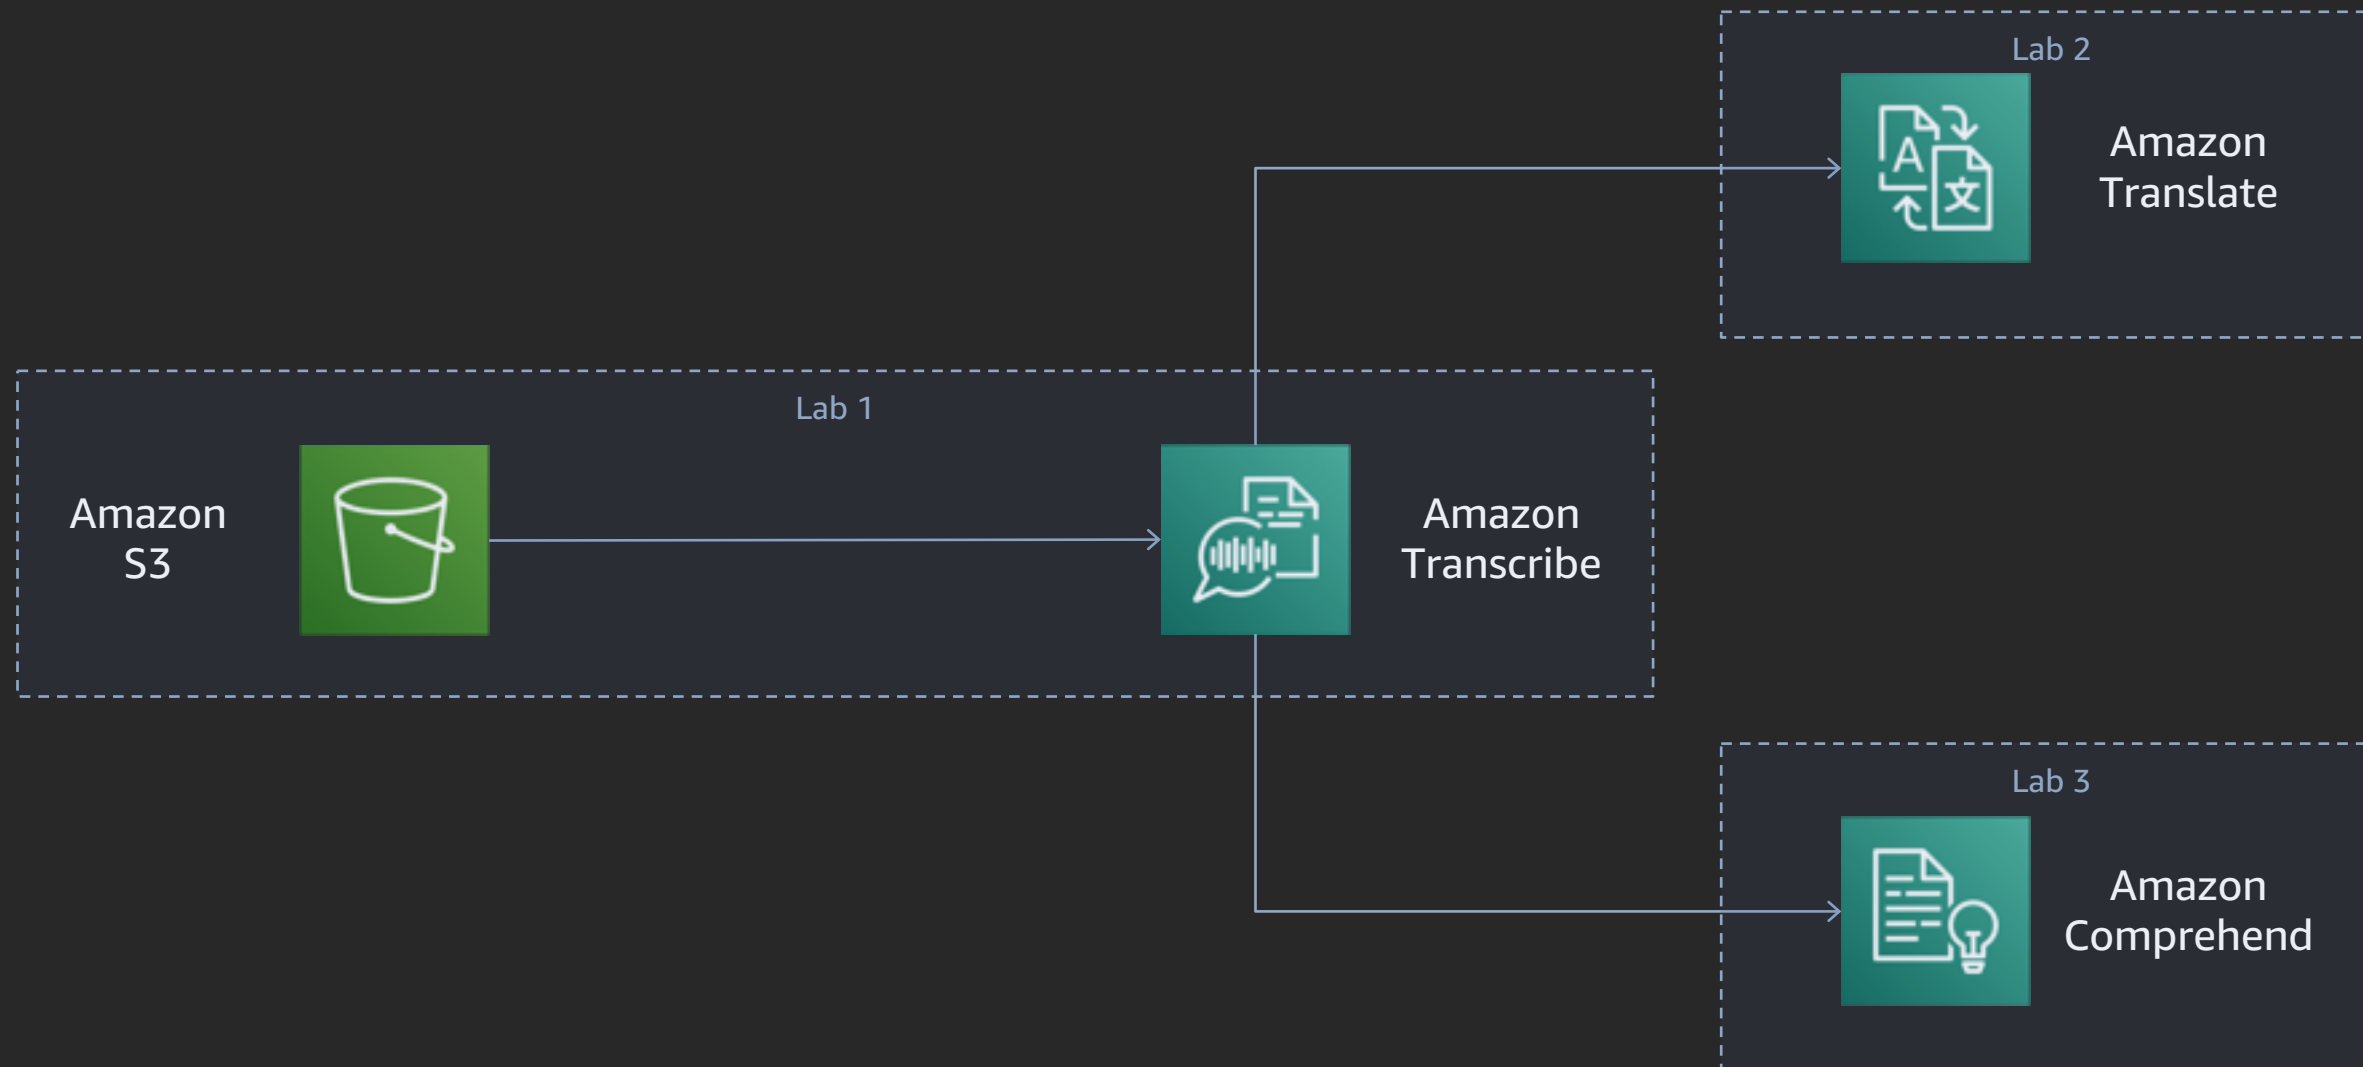
Exploration: Labs 1-3

Manually run through ASR and NLP analyses, look into language translation as well



Exploration: Labs 1–3

Manually run through ASR and NLP analyses, look into language translation as well




Deep dive path

Launching a processing pipeline to automatically run ASR + NLP analyses


Learning objectives:

- Explore how to use speech and language AI services with their APIs programmatically
- Understand how to further analyze and visualize these results with AWS analytics services



Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics

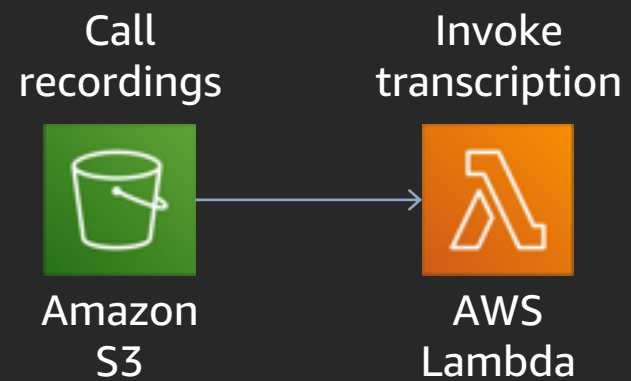
Call
recordings



Amazon
S3

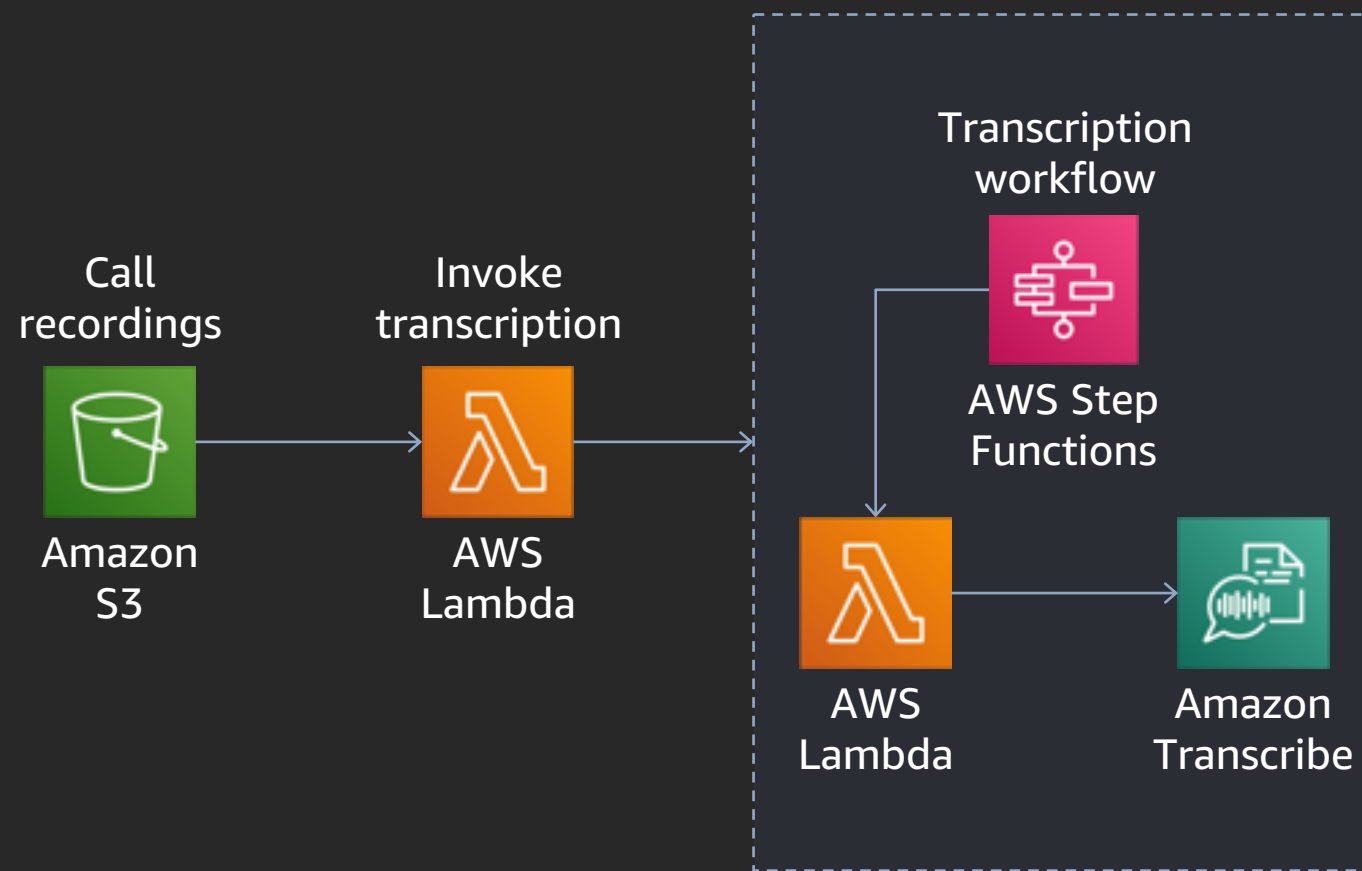
Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



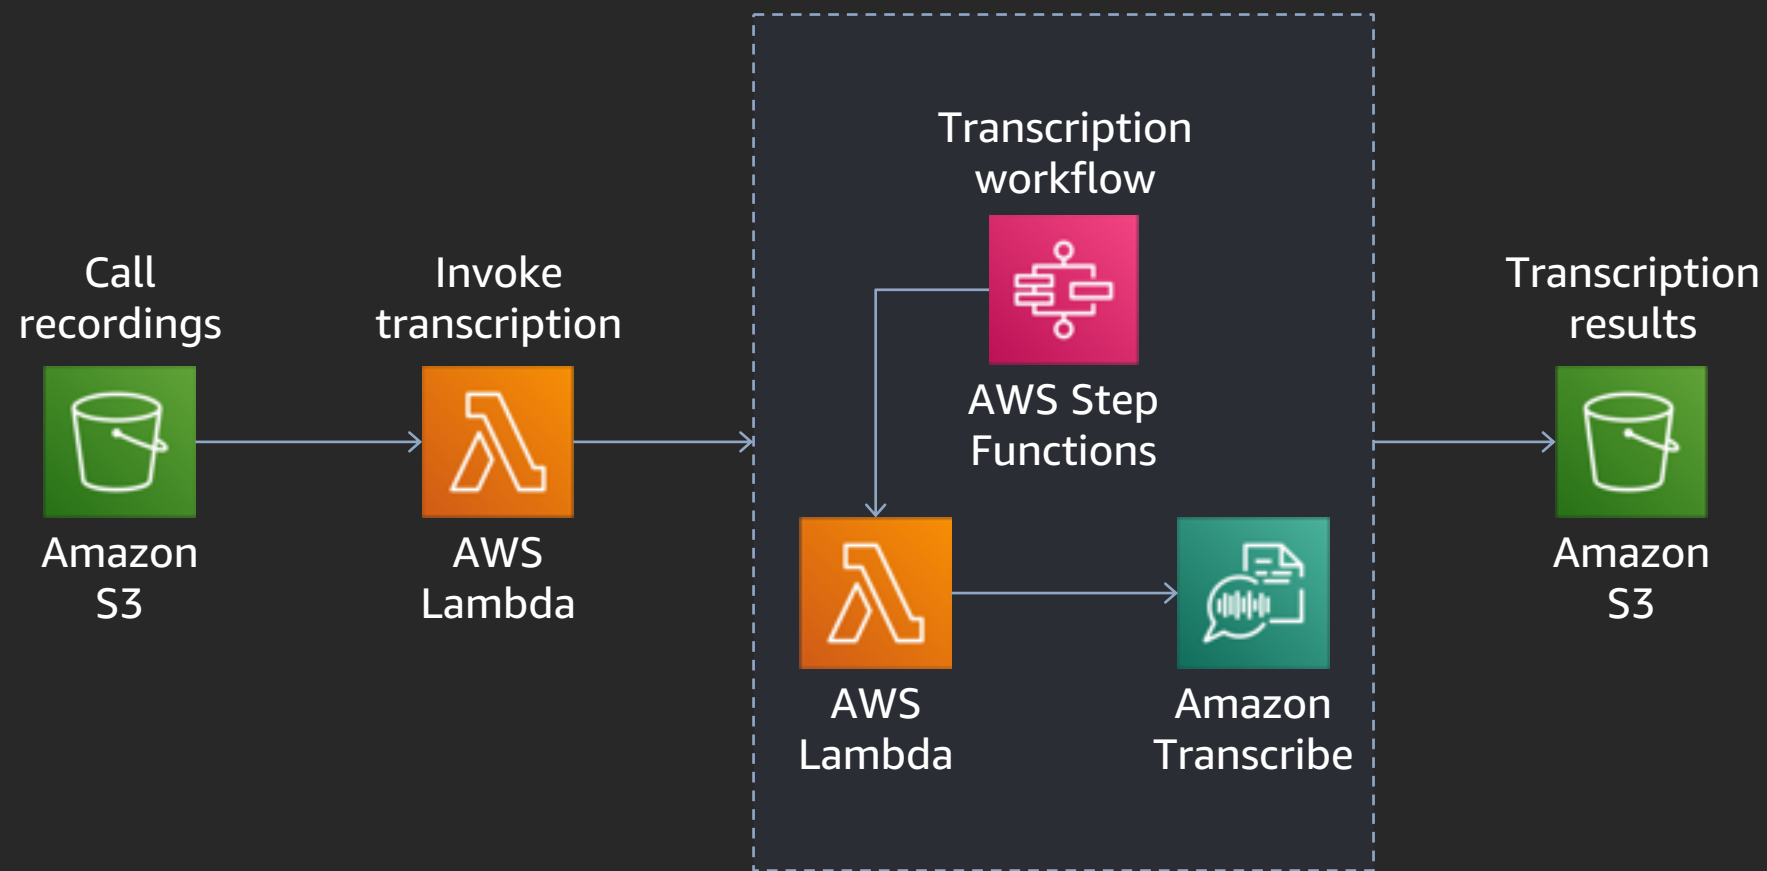
Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



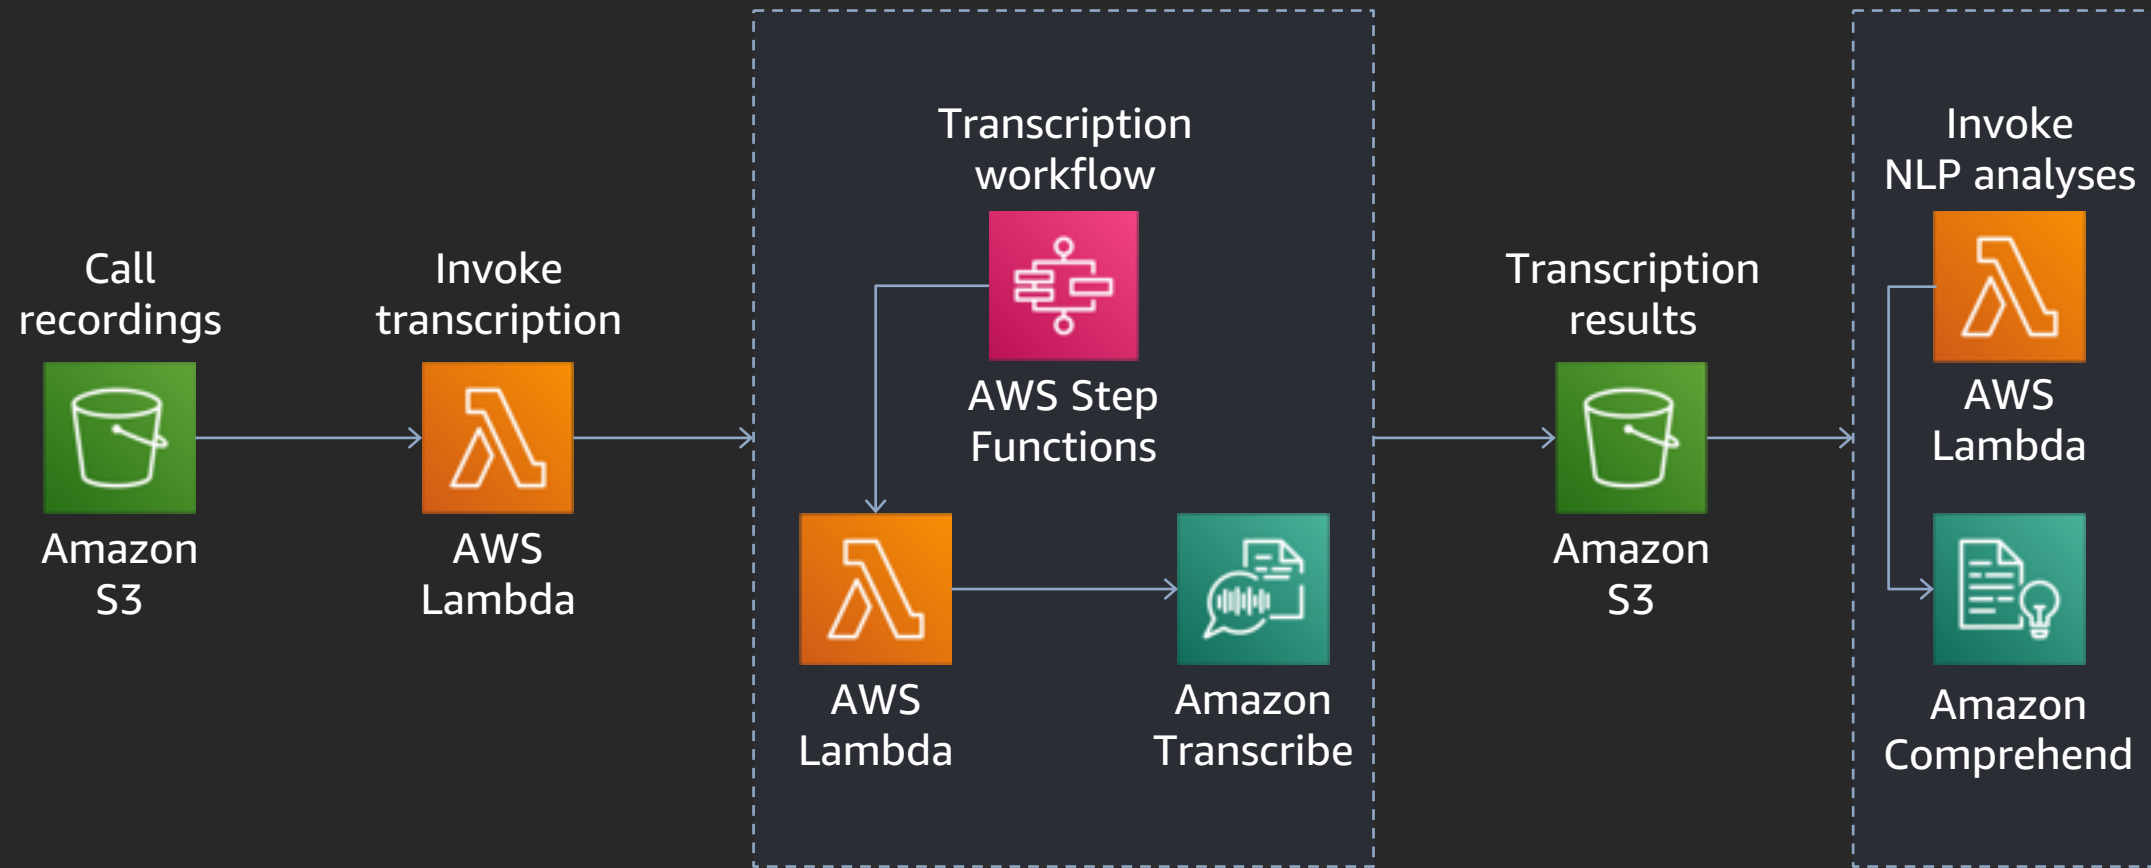
Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



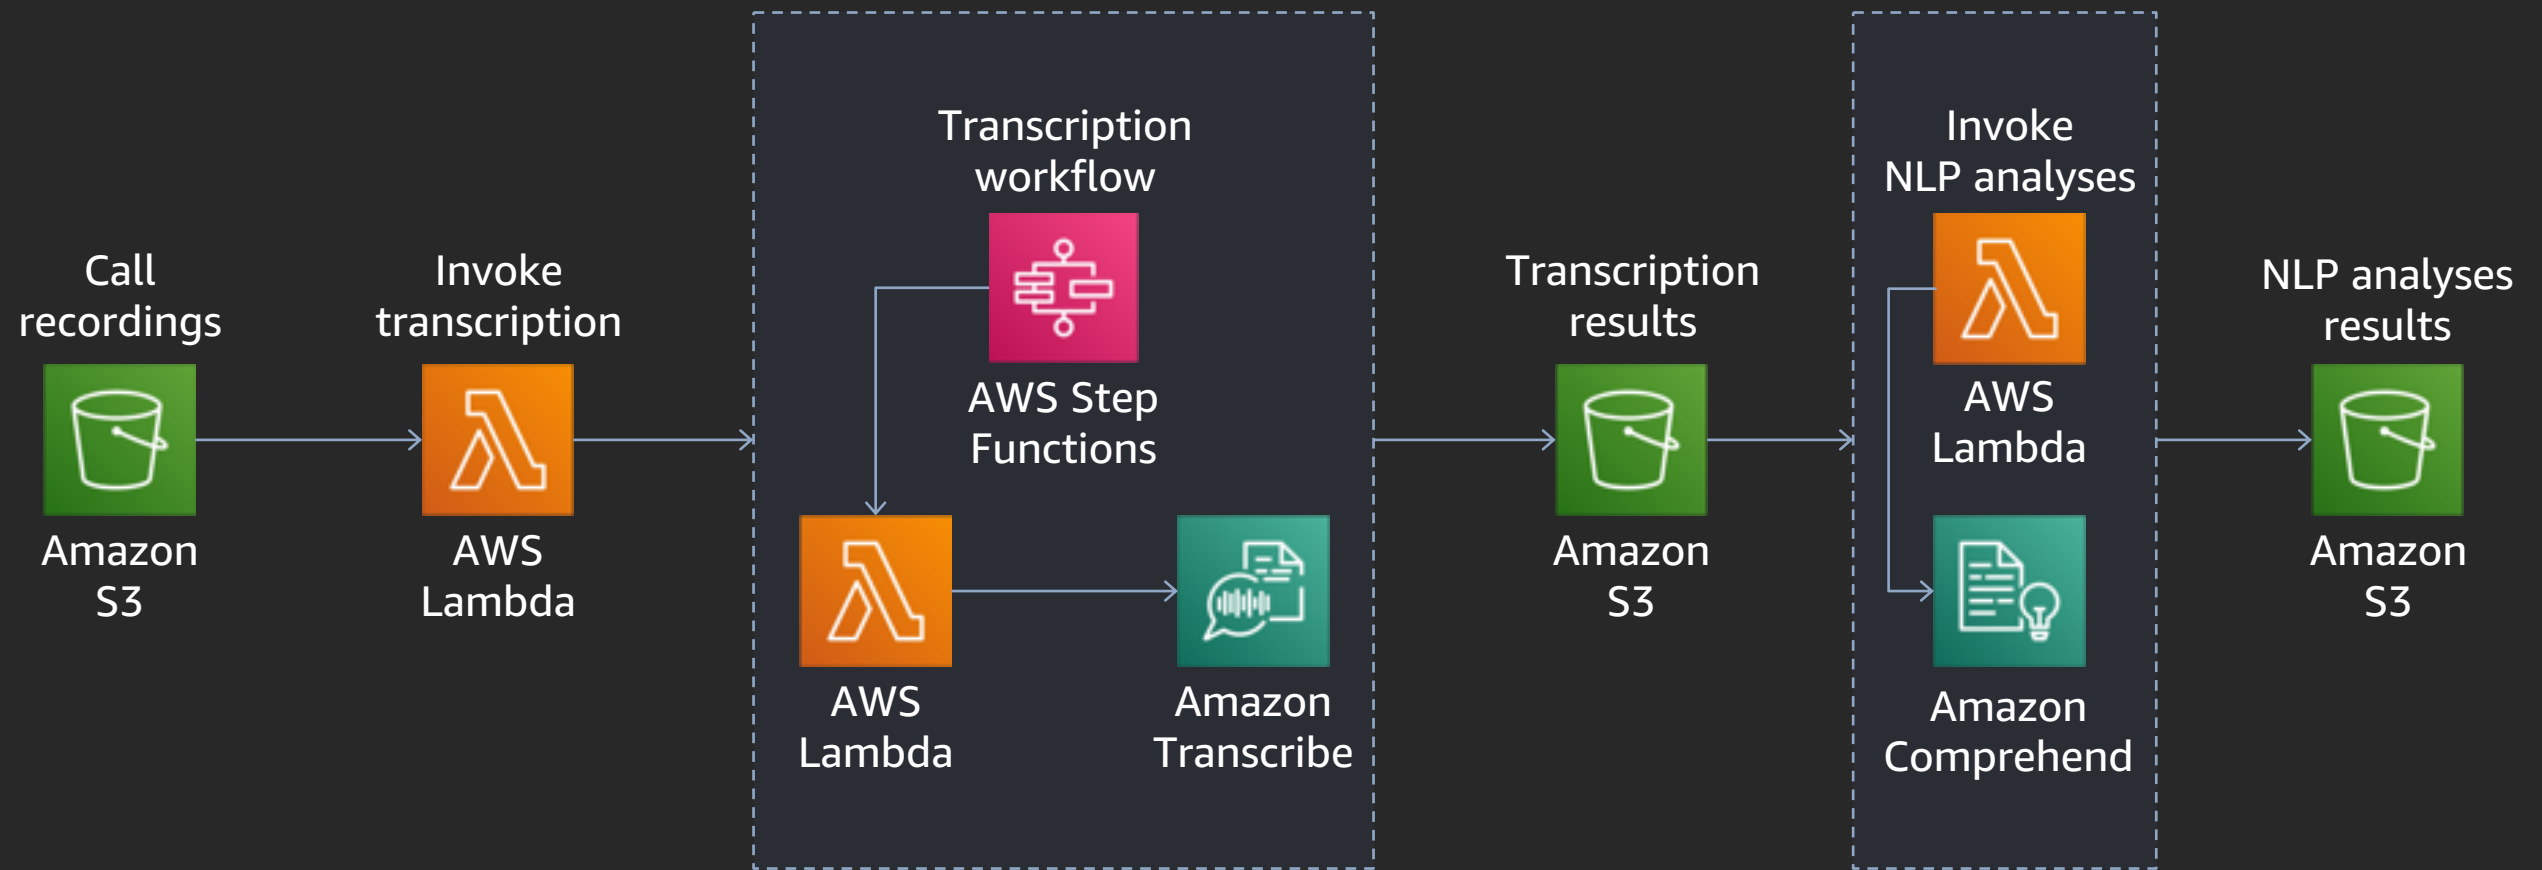
Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics




Deep dive: Lab 4

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Lab 5

Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics

Lab 5

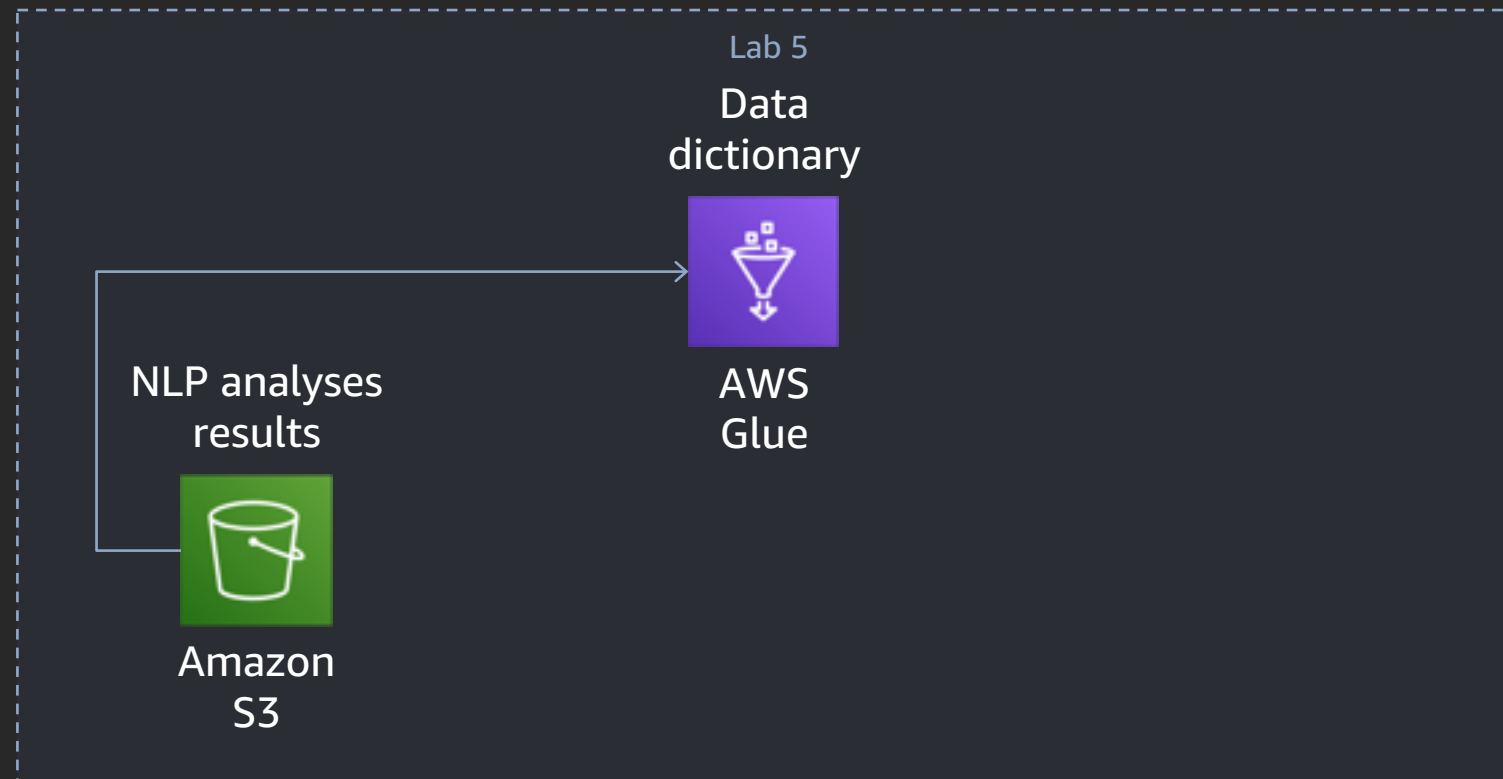
NLP analyses
results



Amazon
S3

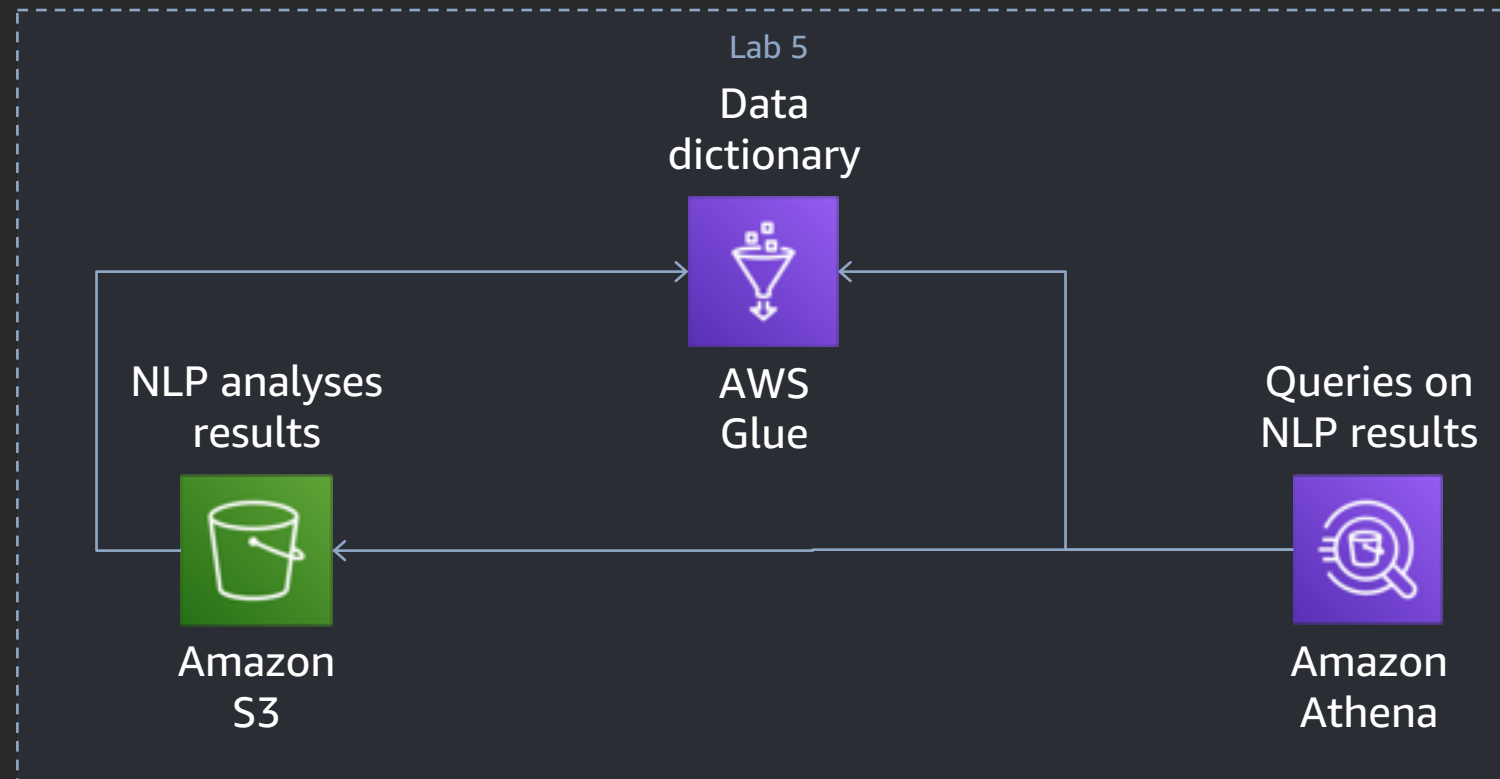
Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



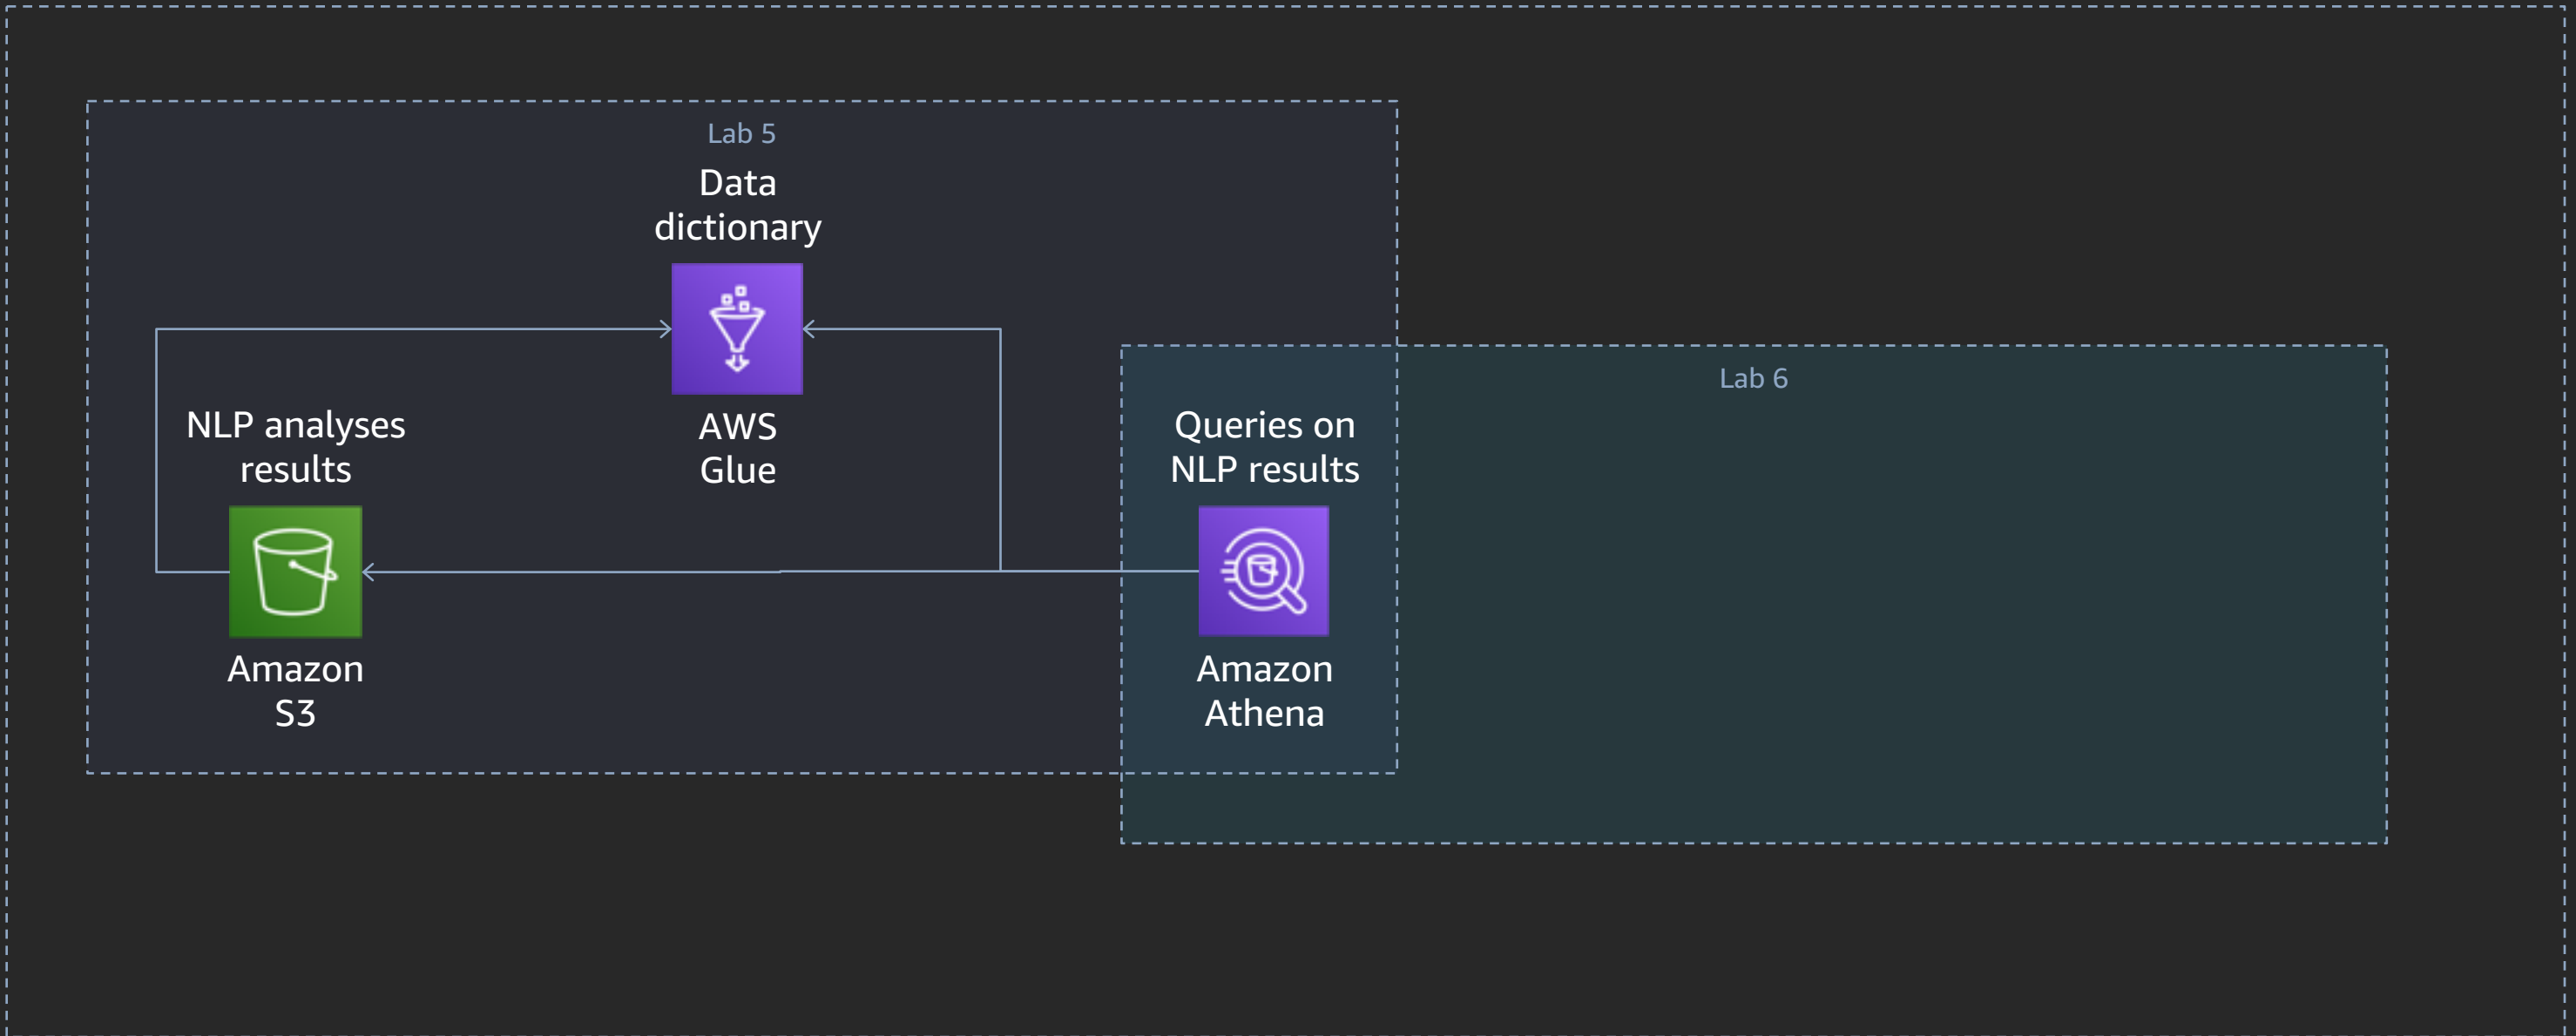
Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



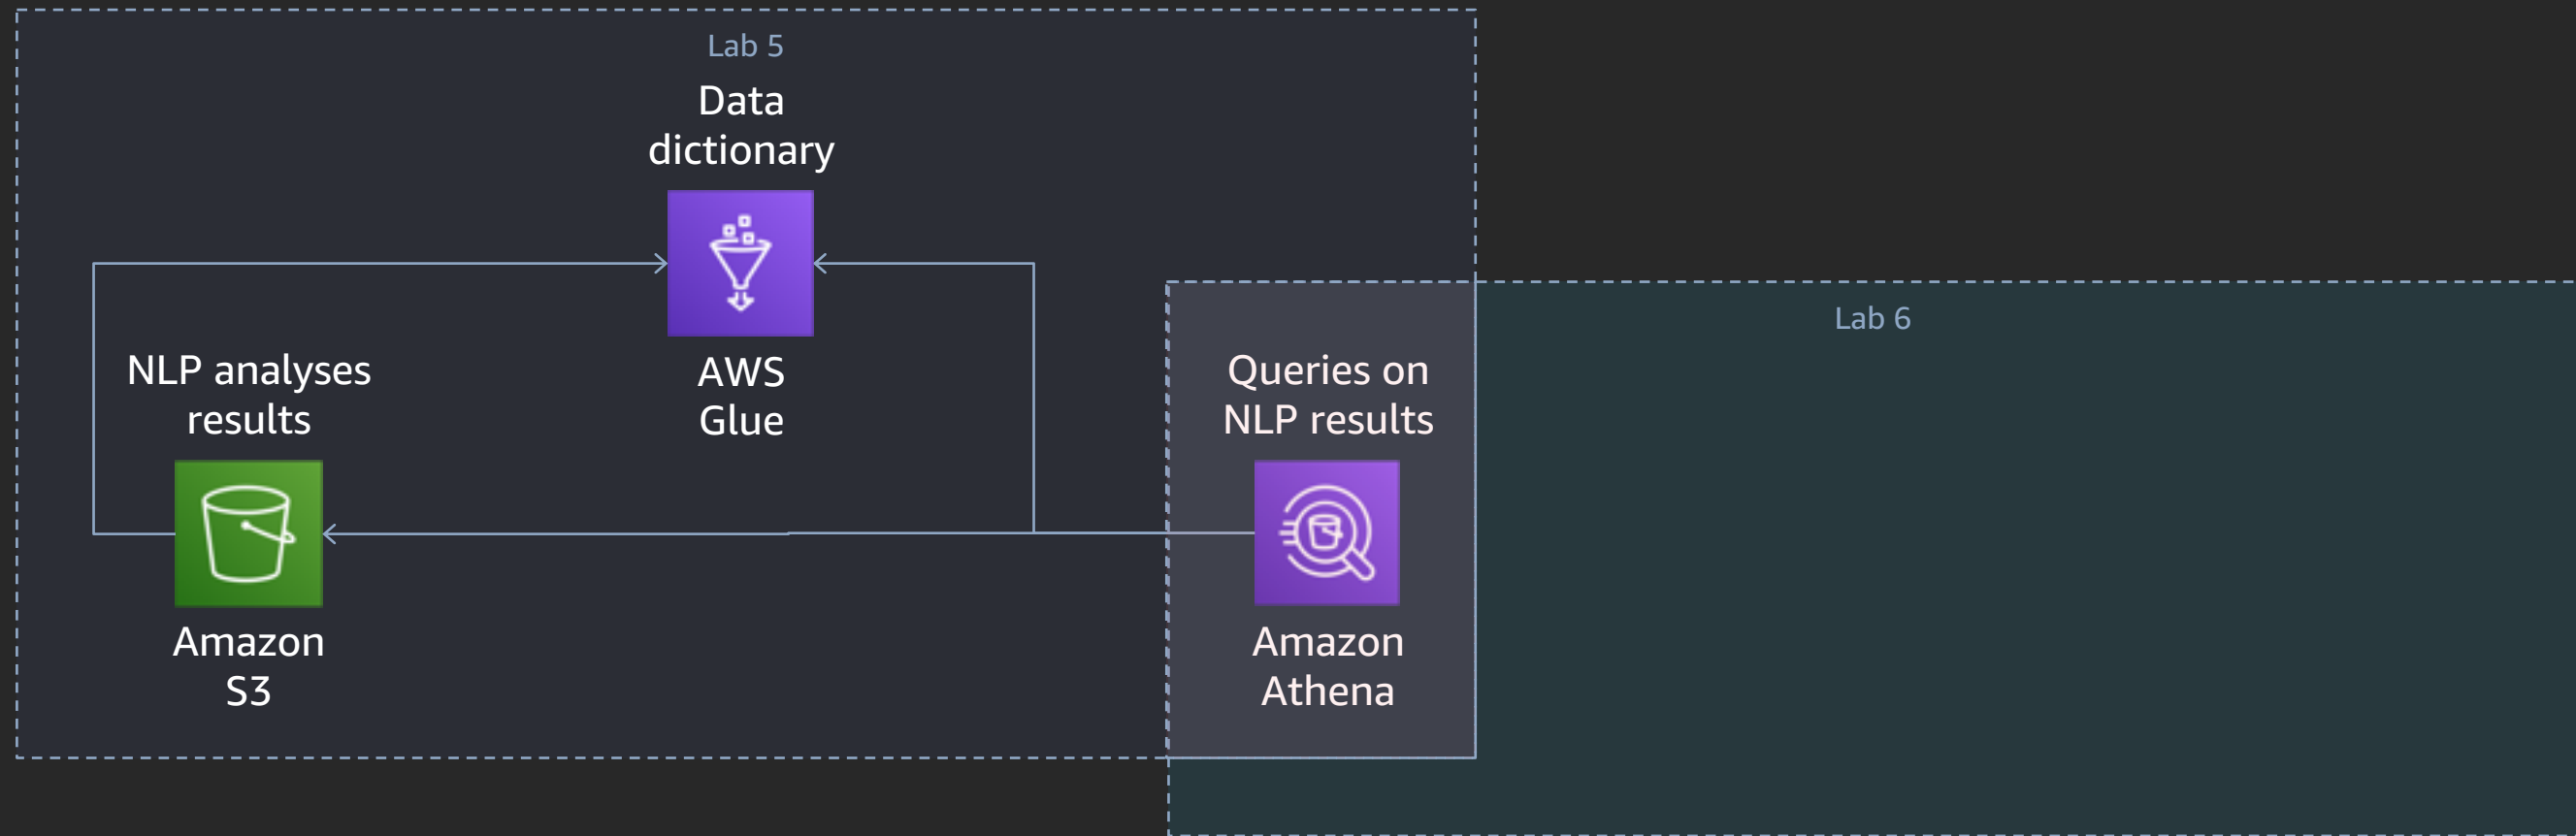
Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



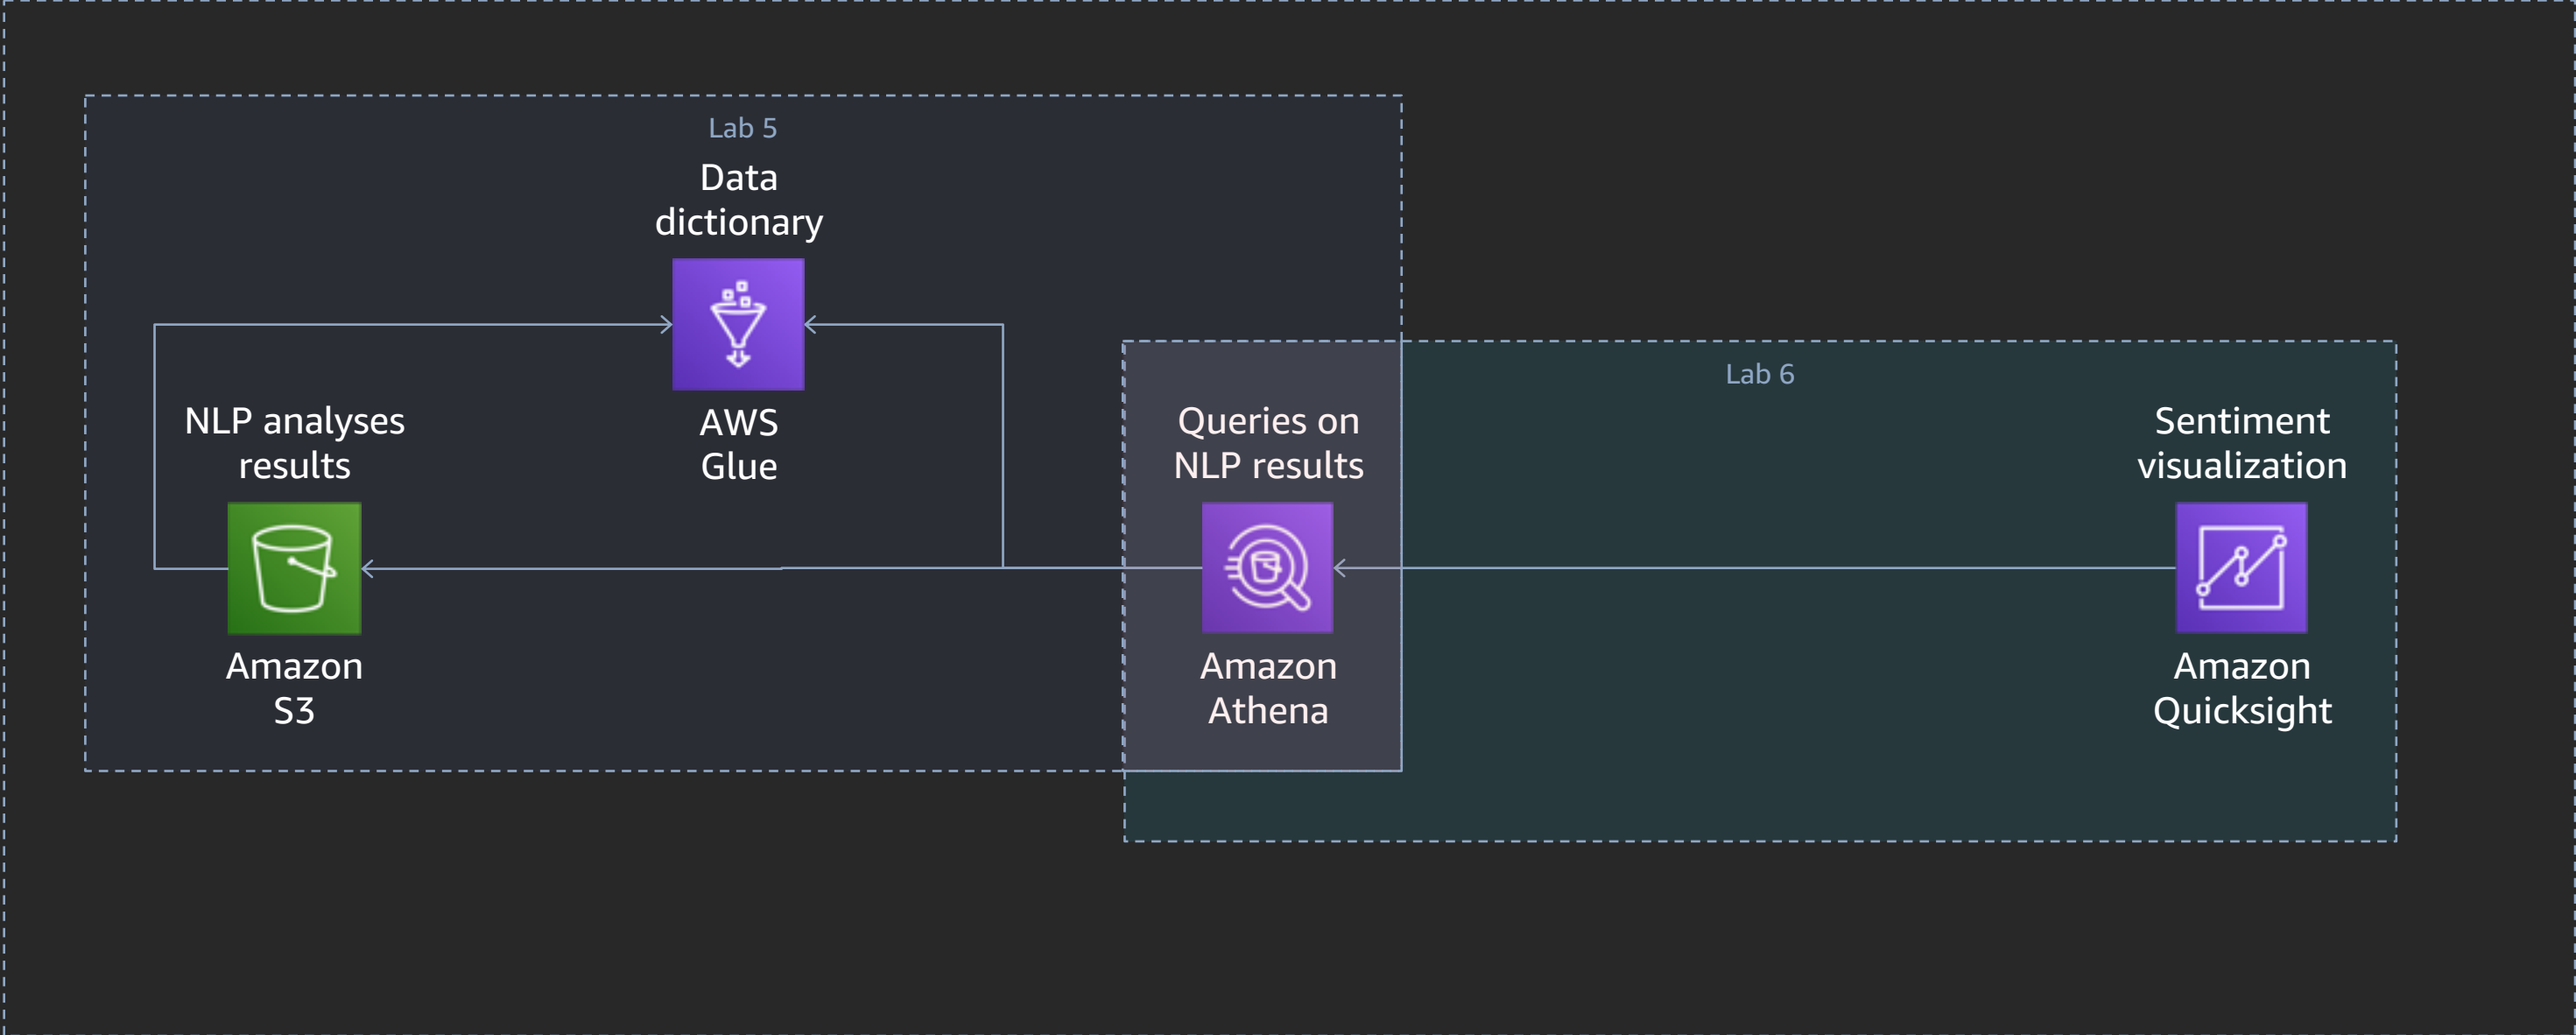
Deep dive: Labs 5–6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Deep dive: Labs 5-6

Launch a processing pipeline to automatically run ASR, NLP, and further analytics



Lab guide:

<http://bit.ly/2019-aim303>

**Go
Build!**



Wrap-up

Resources / call to action

AWS solution: AI-powered speech analytics for Amazon Connect

<http://bit.ly/2019-aim303-r1>

Blog posts that inspired this workshop

Part 1: <http://bit.ly/2019-aim303-r2>

Part 2: <http://bit.ly/2019-aim303-r3>

Where do you go from here? Think about use-case opportunities

Get continual insights about callers' sentiments

Search for and reason on top-of-call metadata

Report on agent performance / callers' happiness or frustration (agent ID, caller ID)


Find the content of this slide also on the lab guide website

Learn ML with AWS Training and Certification

The same training that our own developers use, now available on demand



Role-based ML learning paths for developers, data scientists, data platform engineers, and business decision makers



70+ free digital ML courses from AWS experts let you learn from real-world challenges tackled at AWS



Validate expertise with the
AWS Certified Machine Learning - Specialty exam

Visit <https://aws.training/machinelearning>

Thank you!

Dirk Fröhner

froehner@amazon.de | @dirk_f5r

Boaz Ziniman

boazz@amazon.com | @ziniman



Please complete the session survey in the mobile app.