

The background features a dark blue gradient on the left, transitioning into a large, abstract geometric shape on the right. This shape is composed of several overlapping planes in shades of purple, magenta, and blue, creating a sense of depth and movement. A bright orange-yellow curved line runs along the bottom edge of the abstract shape.

# aws SUMMIT

TORONTO | JUNE 22-23, 2022

ANT205

# How StellarAlgo uses AWS to deliver scalable, data-driven insights

Sean Fynn (he/him)  
CTO and Co-Founder  
StellarAlgo

# Agenda

Our product platform and customer use cases

Our technical product objectives (for growth)

How we use AWS to scale

What's next for us on AWS



# stellaralgo

The sports and entertainment industry customer data platform



# Solving the **industry's** problem

Today's fans are harder to engage than ever

90% of fans never set foot inside an arena and it has become increasingly difficult to earn their attention and loyalty

Billions in lost revenue

An increasingly federated fan experience has led to a poor understanding of the fan and an inability to meet their needs

“StellarAlgo’s retention scoring played a vital role in realizing a 12% improvement in ticket package renewals for the Florida Panthers’ 2022–23 season.

Our membership services team used the model’s scoring of key package buyers and feature importance to drive a more effective retention strategy, ensuring our time was best spent on the accounts that we deemed most at-risk, making our business that much more efficient and effective during one of the busiest times of year for our staff. ”

Cole Cook

Vice-President, Business Intelligence  
Florida Panthers, NHL  
StellarAlgo customer

# SaaS (technical) product objectives

# Scaling our SaaS product

## Technical challenges to overcome

- Productize data processing (ingestion, cleansing, aggregation)
- Maintain our SLO, BCP, and DR
  - parallelism
  - repeatable installation/deployments
- Cost management
- CI/CD for ML
- Maintain InfoSec and compliance

# Scaling our SaaS product

## Our focus to achieve this

- Adopt and grow with AWS releases
  - AWS Lambda, AWS Step Functions, AWS Glue, AWS DMS, Amazon SageMaker, and more
- Commit to infrastructure as code (IaC)
- Migrate to serverless patterns and managed services
- Build robust and flexible ML DevOps

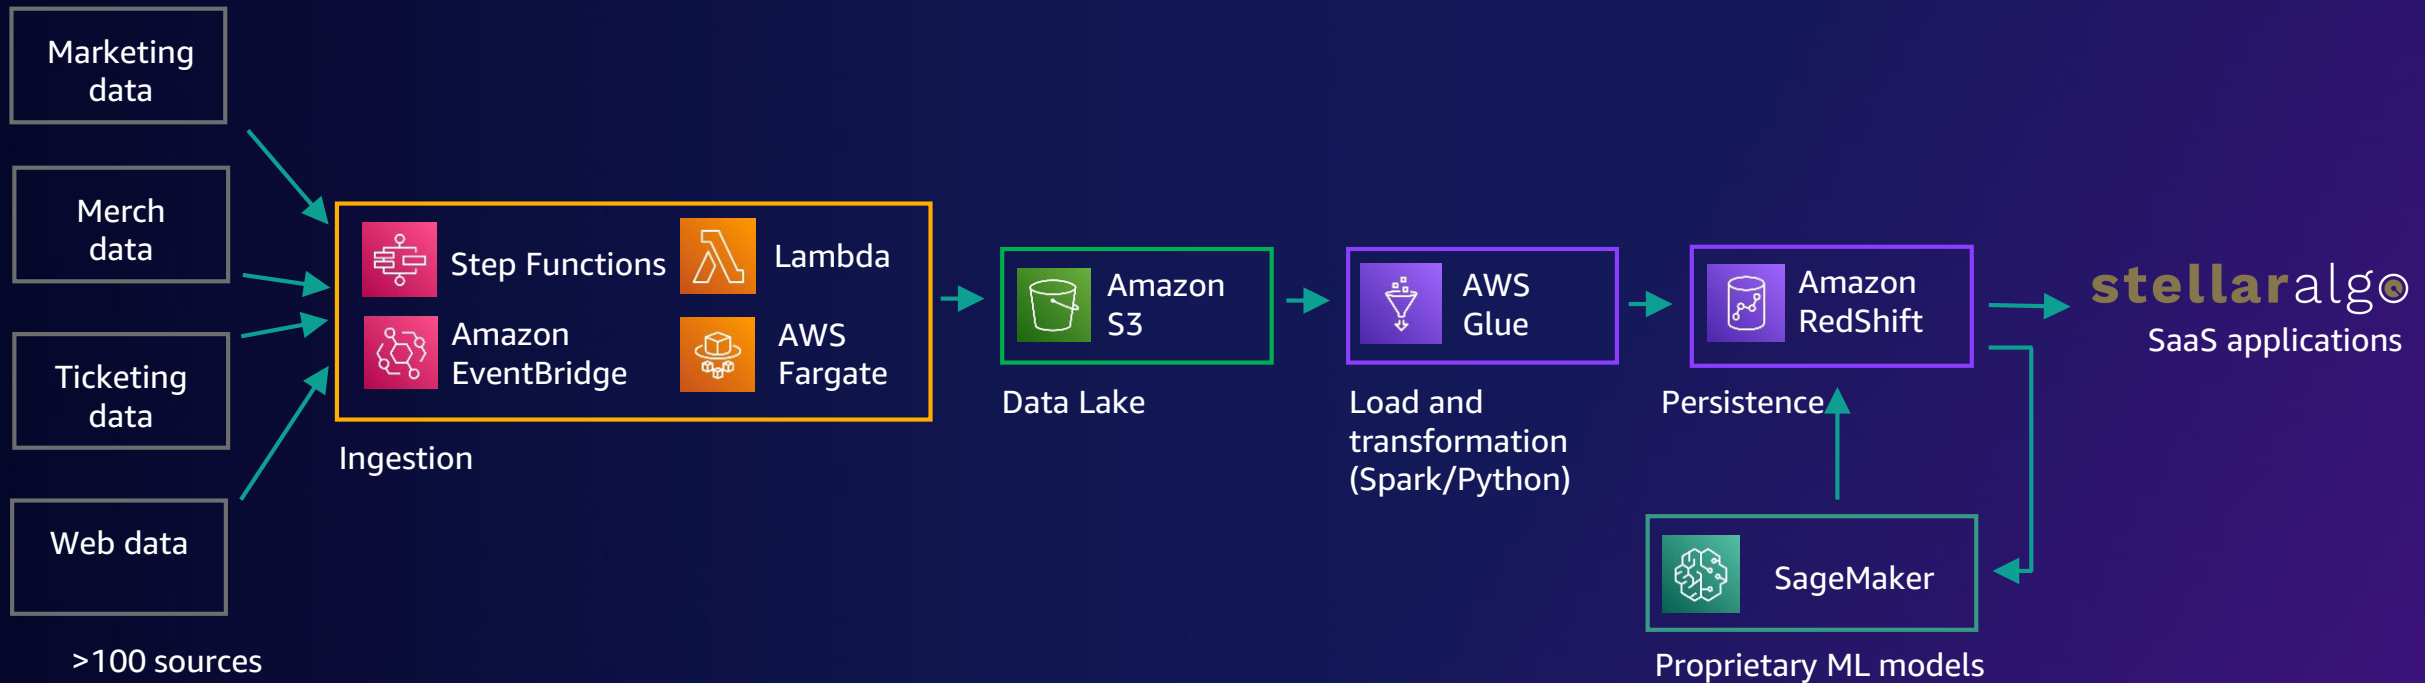
# How we use *AWS* to scale



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

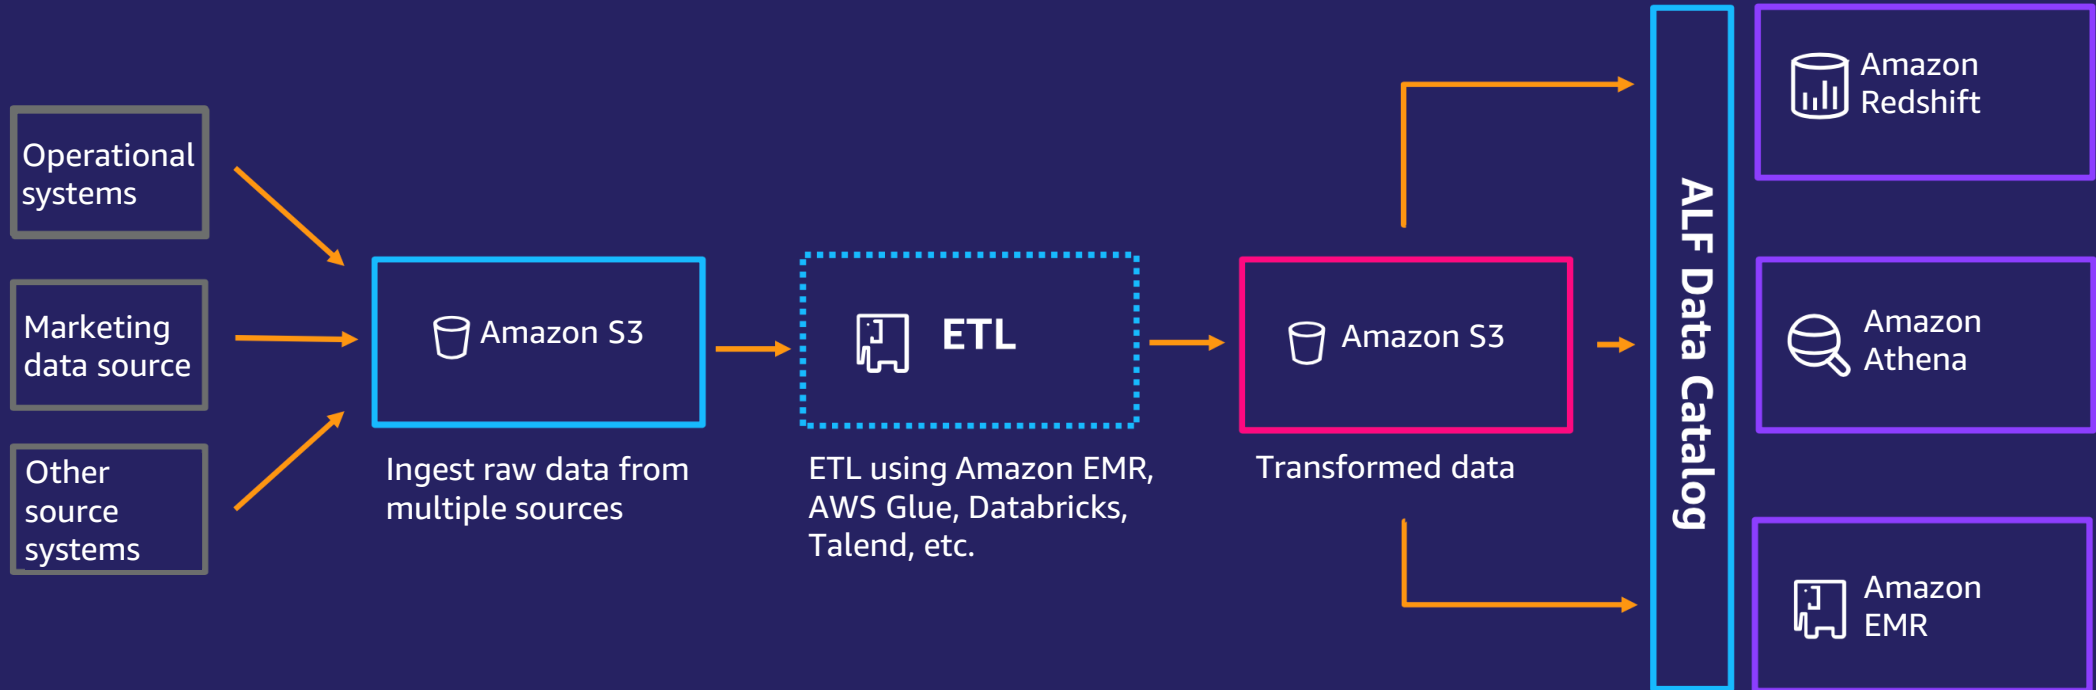


# Big picture – Data flow



# Reference architecture

PUBLISH RAW/REFINED DATA TO AMAZON S3 AND LOAD



## AWS Lake Formation Data Catalog

Note: Lake Formation leverages AWS Glue Data Catalog

# Big picture – IaC

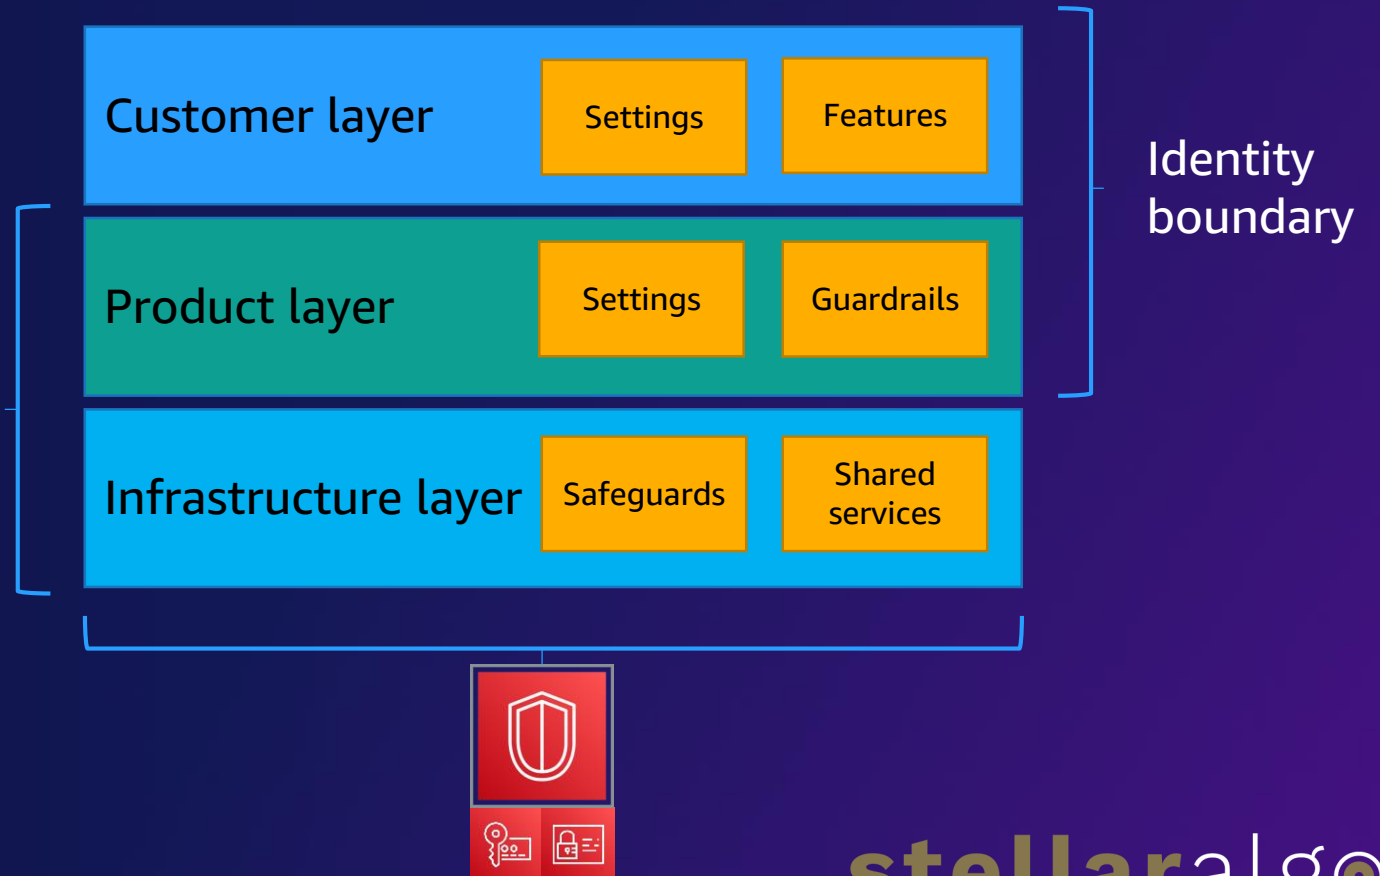
SAAS SCALABILITY FOR STELLARALGO ON AWS

## Objectives

- Secure by design and in deployment
- Reduced mean time to customer deployment
- Match performance curves for customer demand

Environment  
boundary

## Terraform architecture



# AWS Glue and Amazon Redshift



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



# Some statistics . . .

**240M** fan records maintained

**2.3M** new fan records monthly

**118M** fans tracked last year (Majors)

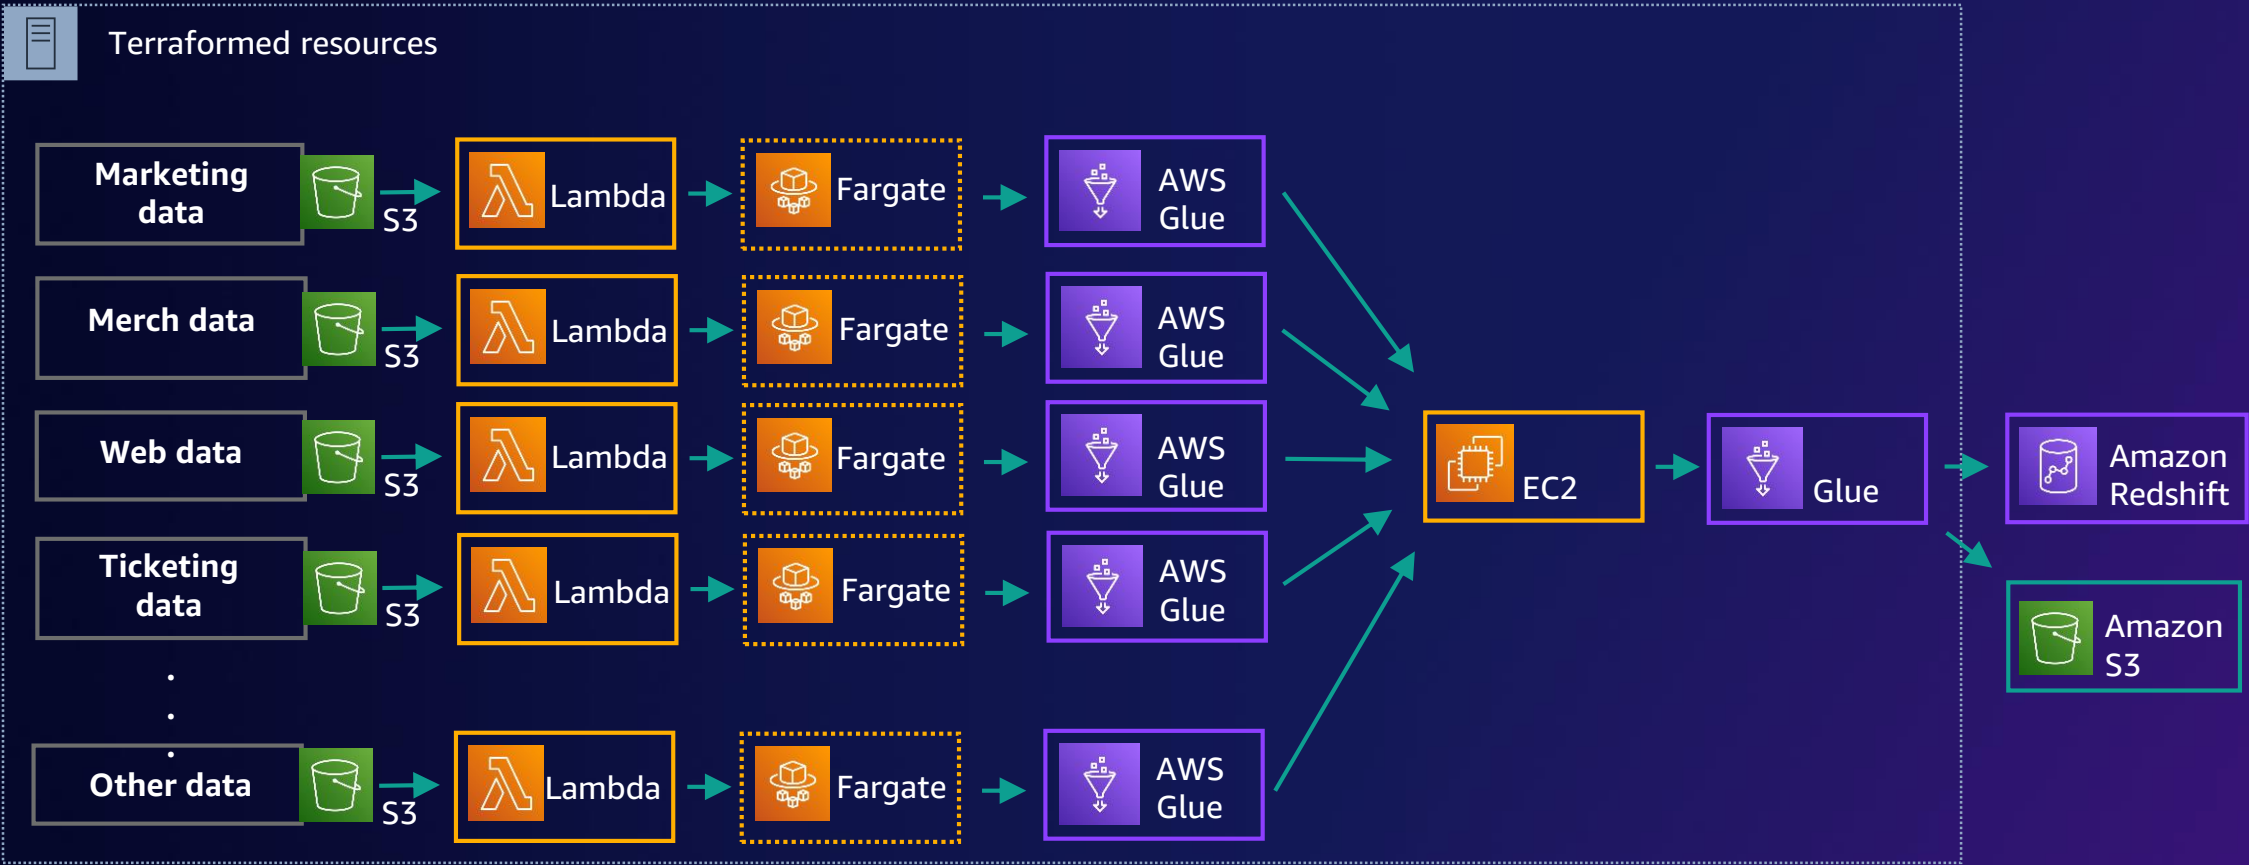
**Billions** of marketing opens, sends, and clicks tracked to the fan level last year (Majors)

**4.3M** unique game attendees tracked last year (Majors)

**~100** unique productized integrations and **hundreds** of subject areas

# ELT – AWS Glue (transition state)

CUSTOMER 1



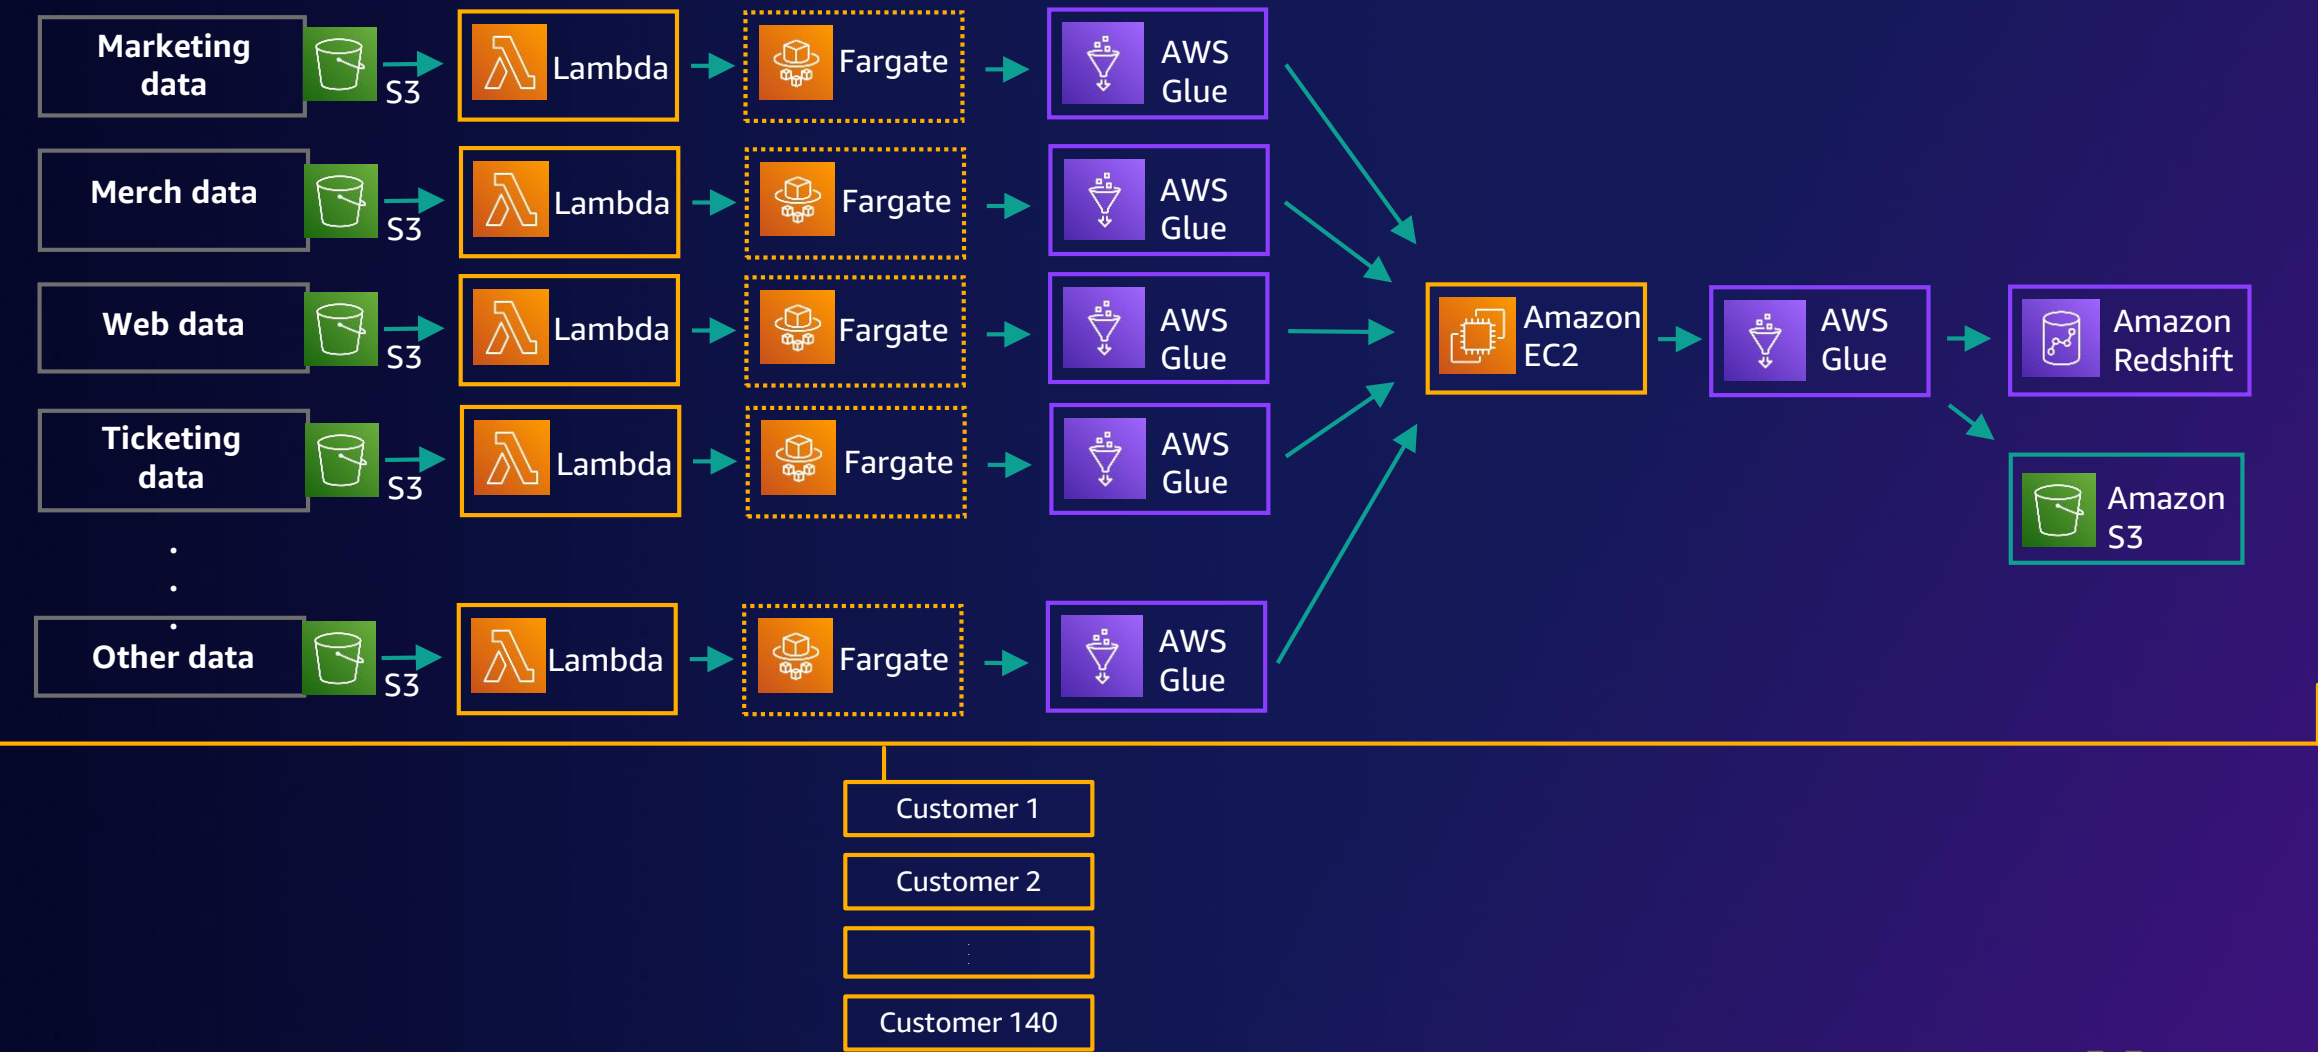
Productized integrations



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

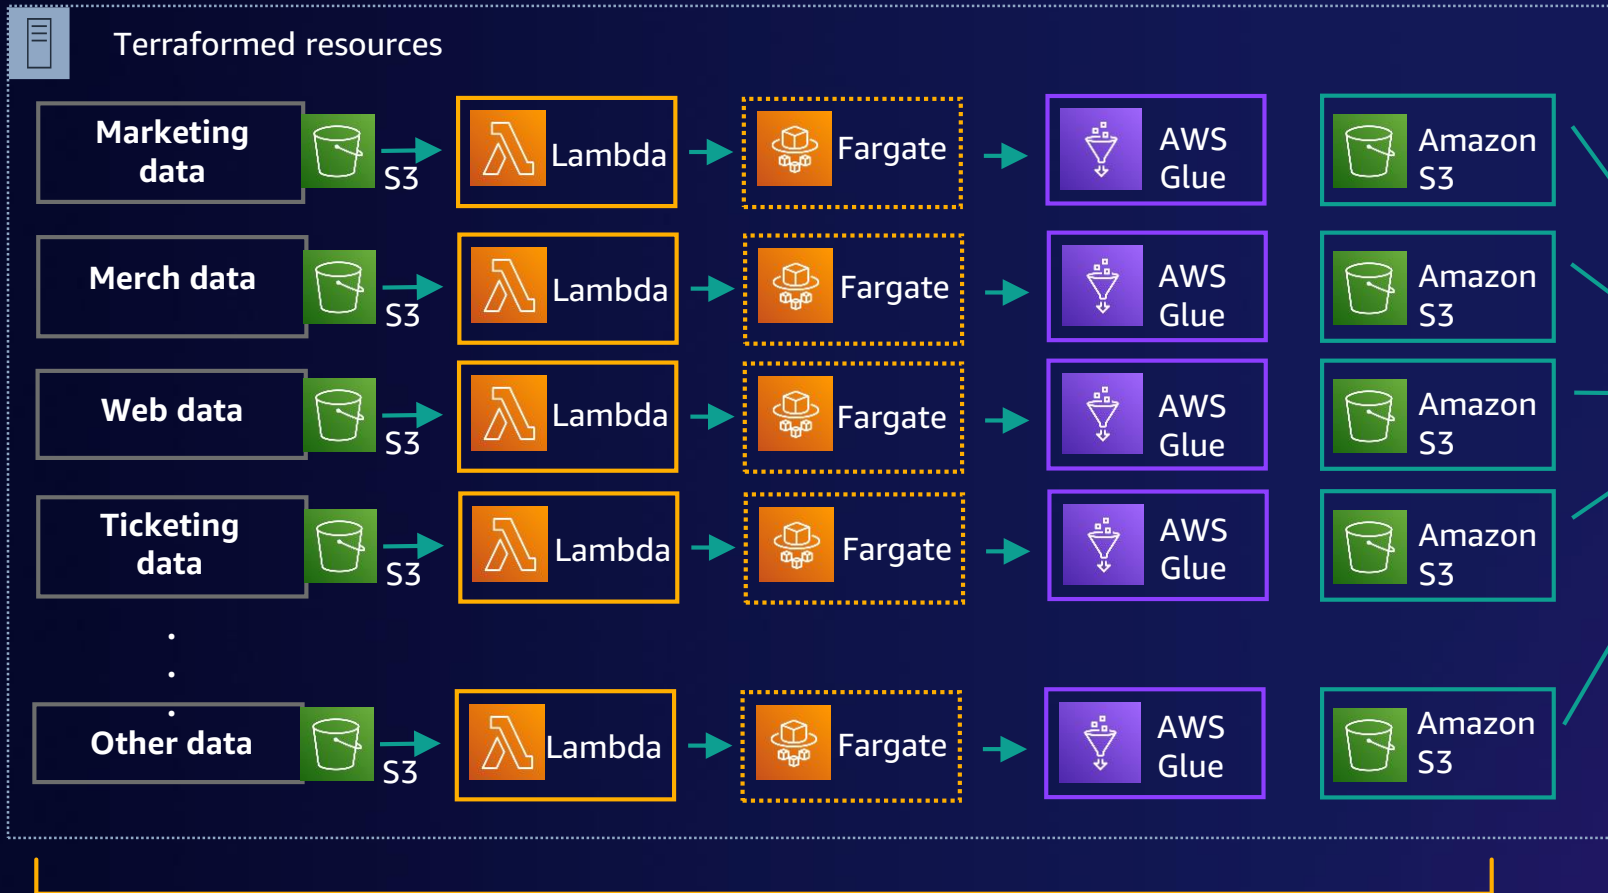
stellaralgo

# ELT – AWS Glue (transition state)



# ELT – AWS Glue (target state)

## CUSTOMER 1



*Productized integrations*

### Scalable

Full use of serverless features to allow for scalability as the client base grows; ability to scale up compute and storage based on changing requirements

### Timely

Ability to stand up and scale infrastructure on demand utilizing serverless pipelines.

### Maintainable

Ability to manage and monitor pipelines using standard AWS functionality, allowing engineers to focus on client deliveries

### Configurable

Meet client needs with configuration checkpoints; meet demand of specific logic while maintaining scalable, reliable pipelines

“The Rays have a small analytics team with limited in-house data engineering resources. We are able to leverage StellarAlgo’s Data Warehouse to answer many pressing questions in a manner of minutes by leveraging multiple disparate data sources that would take days to answer without the tool.

By leveraging StellarAlgo’s Data Warehouse, we have been able to save on average about 60% of the time spent collecting, compiling, and massaging the data for most of our data science projects. ”

**Juan Dominguez**

Senior Analyst, Strategy and Analytics  
Tampa Bay Rays, MLB  
StellarAlgo customer



# Amazon SageMaker

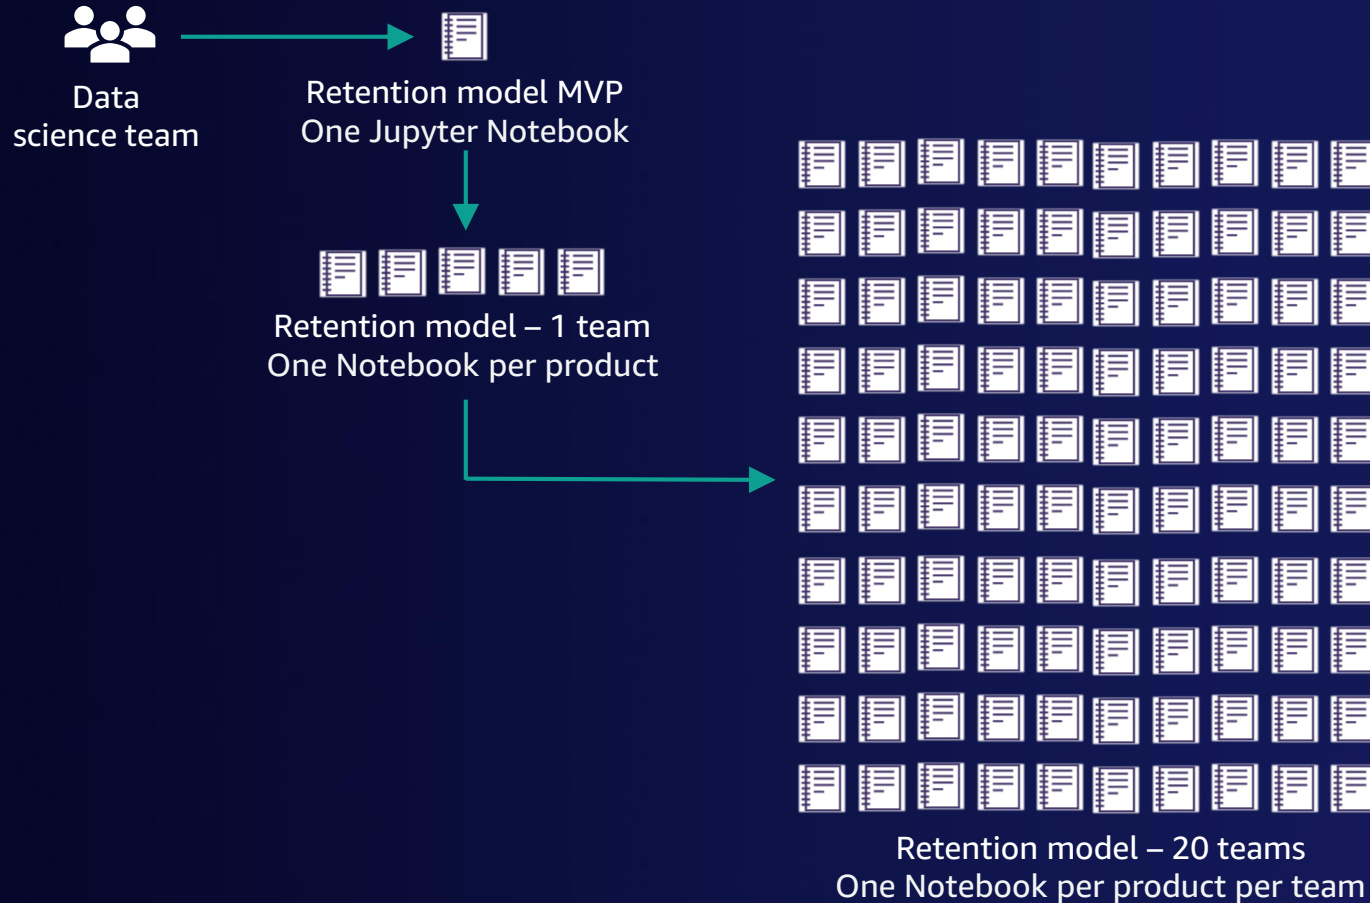


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



# SageMaker – Dev/MLOps

FROM DESKTOP NOTEBOOKS TO CLOUD PIPELINES



When can we put this in our product?

Awesome model . . .



Product management

We just sold this to 20 teams!

We'll also need this in a UAT environment . . .

When will this be ready for the teams?

One of our teams needs new features

We just sold this to 50 teams!

Make that 100 teams!

The market is asking for new type of models!



Sales team

# SageMaker – Dev/MLOps

FROM DESKTOP NOTEBOOKS TO CLOUD PIPELINES

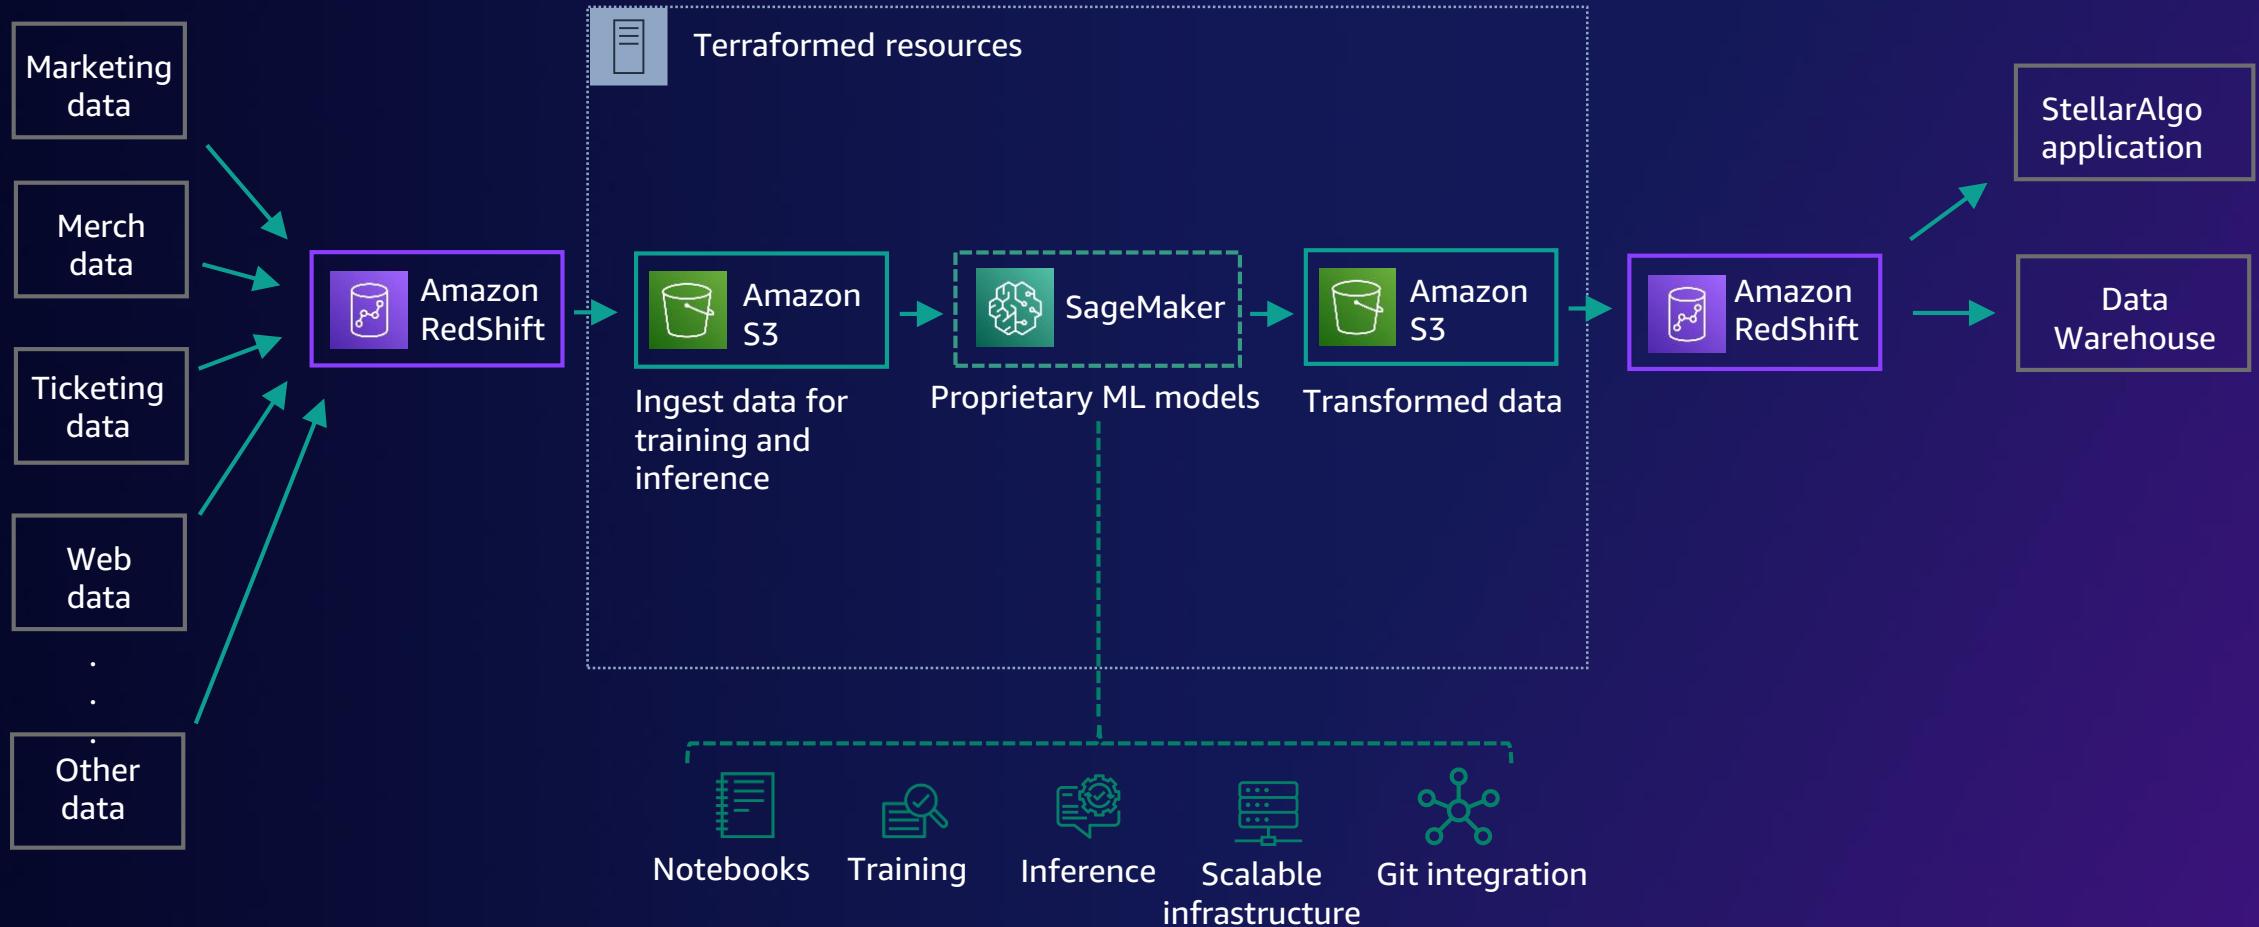
## Why SageMaker?

A framework to ingest large volumes of customer data, provide robust monitoring, and offer the ability to bring your own models to deliver predictions to hundreds of customers and tens of millions of fans – every single day

- ✓ **Spend less time running local (Amazon EC2) notebooks**  
Focus on feature engineering, model accuracy, and performance
- ✓ **Flexible and relevant to our customers strategic needs**  
Provide timely and consistent results to drive decision making
- ✓ **Faster implementation, shorter time to value, and faster iteration**  
Serverless scalability, IaC
- ✓ **Security and ability to maintain compliance from start to finish**  
Dedicated streams

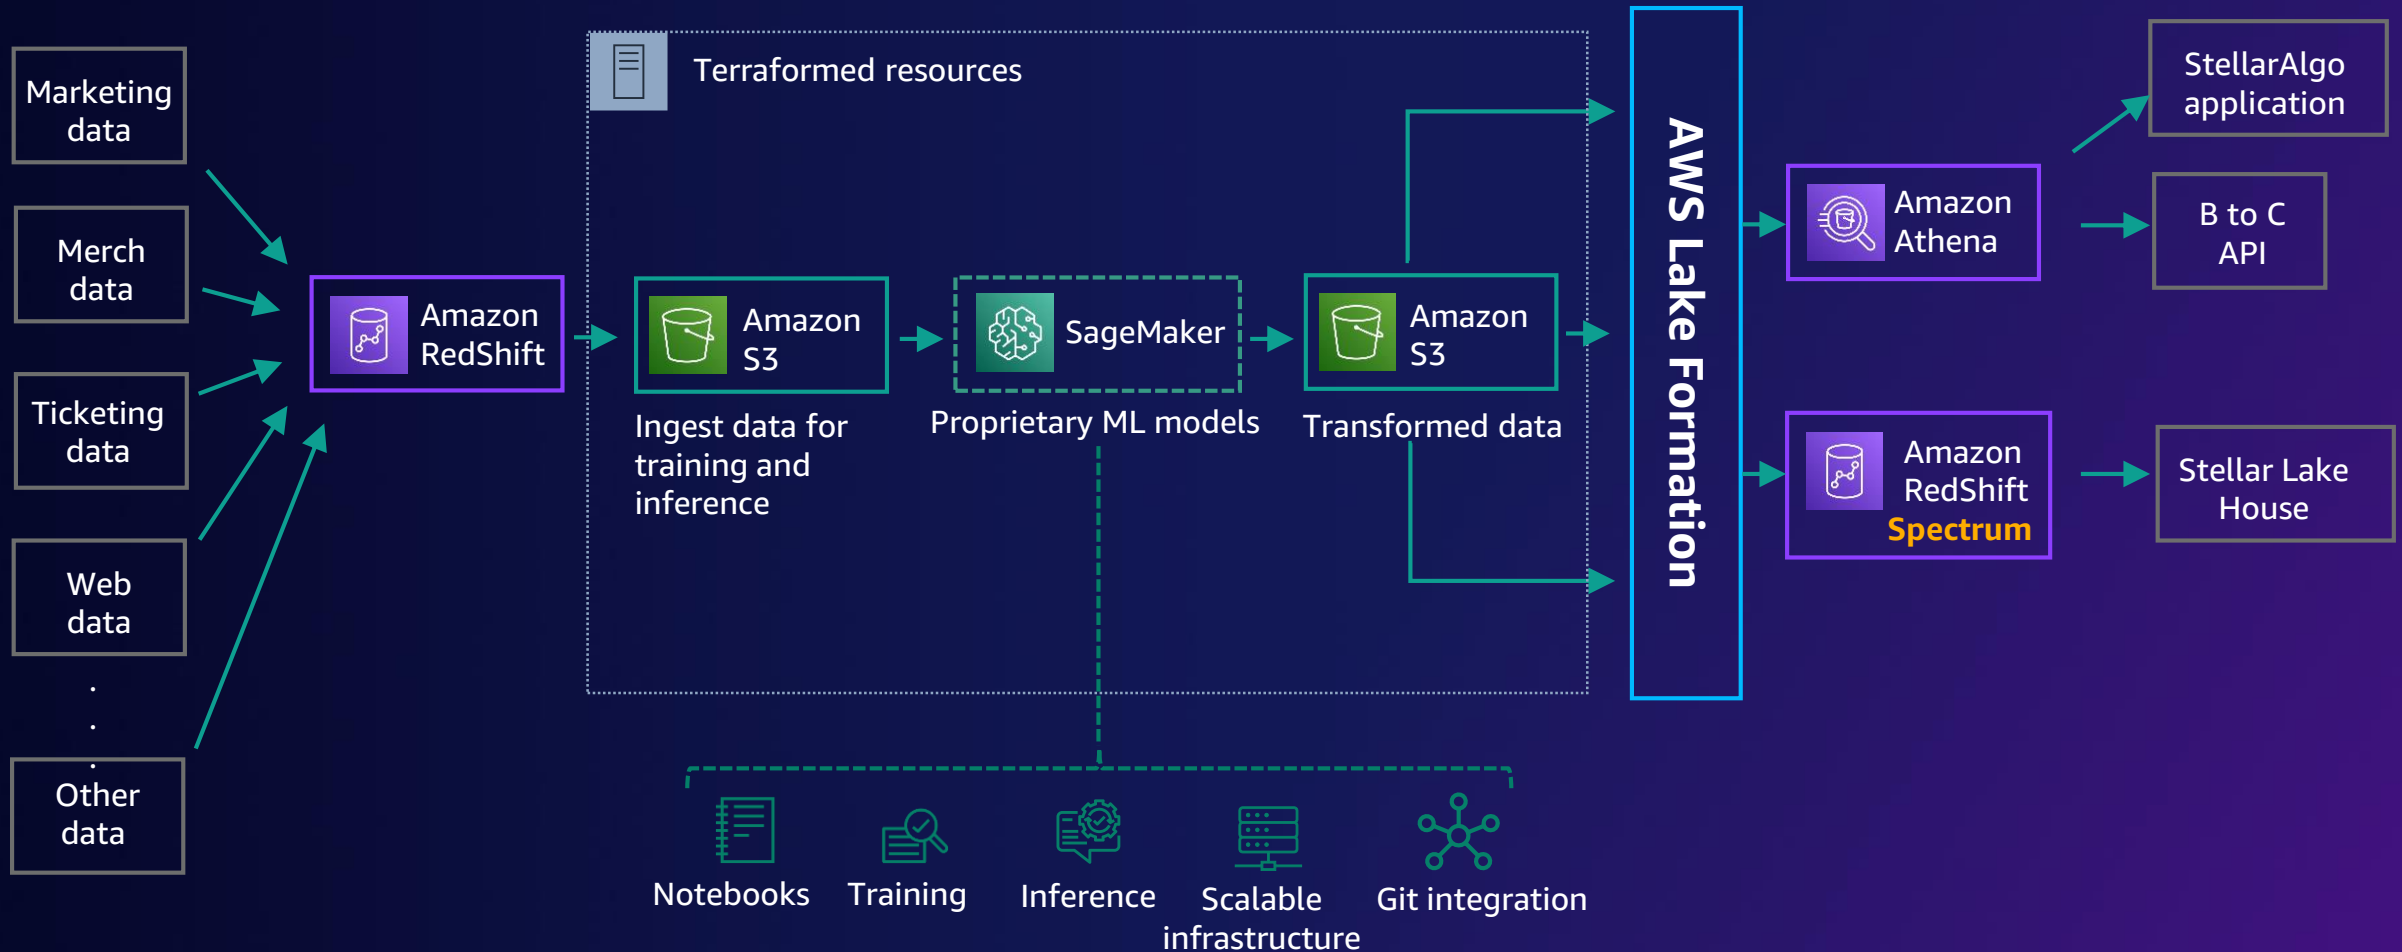
# SageMaker – Current state

ML ARCHITECTURE – TRAINING, INFERENCE AND SERVING PREDICTIONS



# SageMaker – Future state

ML ARCHITECTURE – TRAINING, INFERENCE AND SERVING PREDICTIONS



# What's next for StellarAlgo on AWS?



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



# Doubling down on AWS

**Our data processing volumes are expected to at least double again this year!**

- More serverless and microservices refactoring
- More use of AWS (Amazon Kinesis Data Firehose, Athena, Lake Formation . . . )
- More data transformations and aggregation within AWS Glue (Spark/Python)
- Expansion of Amazon Redshift Spectrum and utilization of new features coming for R3!
- Expand SageMaker use for more models and new access patterns

# Thank you!

Sean Fynn

[sfynn@stellaralgo.com](mailto:sfynn@stellaralgo.com)

 [linkedin.com/in/seanfynn/](https://www.linkedin.com/in/seanfynn/)

