



INNOVATE

DATA AND AI/ML EDITION

22 February 2023

Putting machine learning in the hands of every builder with AWS database and analytics services

Tom McMeekin


Enterprise Solutions Architect
Amazon Web Services



Transform customer
experience



Improve business
operations




Better and faster
decision-making



Innovation



NEXT GEN STATS

powered by 

“...AWS allows both data experts and novices to have a more seamless and intuitive experience within Next Gen Stats...”

Michael Chi, Director of NFL Next Gen Stats



Modern data strategy



Modern data strategy



MODERNIZE



UNIFY



INNOVATE

Modern data strategy



INNOVATE

Purpose-built data services

Business Intelligence & Machine Learning

Analytics

Databases

Blockchain

Data Lake

Data Movement



Purpose-built data services | the most complete family

Business Intelligence & Machine Learning

 **Amazon QuickSight**
Visualizations

 **AWS Data Exchange**
Data exchange

 **Amazon SageMaker**
ML

 **Amazon Comprehend**
NLP

 **Amazon Transcribe**
Speech-to-text

 **Amazon Textract**
Extract text

 **Amazon Personalize**
Recommendation

 **Amazon Forecast**
Forecasts

 **Amazon Translate**
Translation

 **Amazon CodeGuru**
Code reviews


 **Amazon Kendra**
Enterprise search


+Many more

Analytics


Databases


Blockchain


 **Amazon Redshift**
Data warehousing

 **Amazon EMR**
Hadoop + Spark

 **Amazon Athena**
Interactive analytics


 **Amazon Aurora**
MySQL, PostgreSQL

 **Amazon DynamoDB**
Key value, Document


 **Amazon DocumentDB**
Document

 **Amazon Neptune**
Graph


 **Amazon Managed Blockchain**

 **Amazon OpenSearch Service**
Operational Analytics


 **Amazon Kinesis Data Analytics**
Real time


 **Amazon RDS**
MySQL, PostgreSQL, MariaDB, Oracle, SQL Server,

 **Amazon ElastiCache**
Redis, Memcached

 **Amazon QLDB**
Ledger Database

 **Amazon Timestream**
Time Series

 **Amazon Keyspaces (For Apache Cassandra)**
Wide column

 **Amazon Blockchain Templates**

Data Lake

 **Amazon S3/Glacier**

 **AWS Lake Formation**
Data Lakes

 **AWS Glue**
ETL & Data Catalog

Data Movement

AWS Database Migration Service | AWS Snowball | AWS Snowmobile | Amazon Kinesis Data Firehose | Amazon Kinesis Data Streams | Amazon Managed Streaming for Kafka



Machine learning development is complex



Label data



Data collection prep



Store features



Detect bias and explain predictions



Visualize in notebooks



Pick algorithm



Train models faster



Tune parameters



Deploy in production



Manage & monitor

The AWS ML stack

ML FRAMEWORKS & INFRASTRUCTURE

PyTorch, Apache MXNet,
Hugging Face, TensorFlow

Amazon EC2

CPUs

GPUs

AWS Inferentia

AWS Trainium

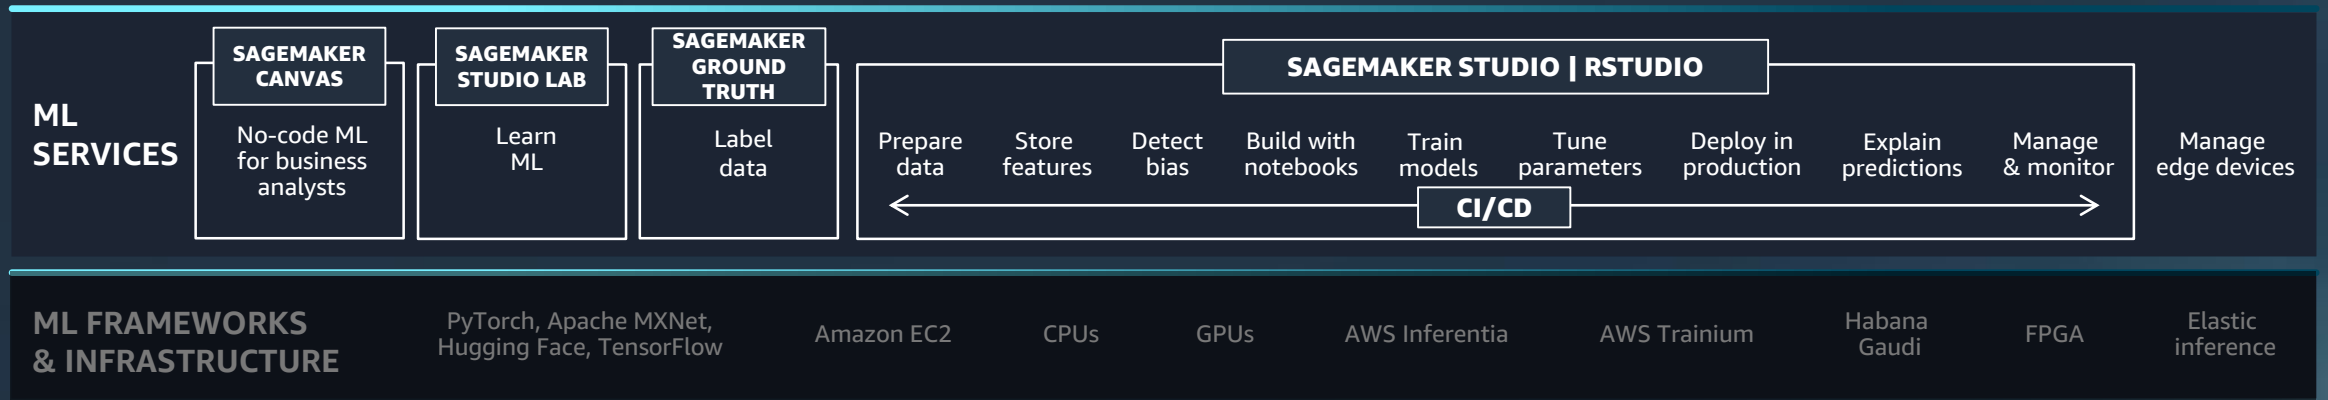
Habana
Gaudi

FPGA

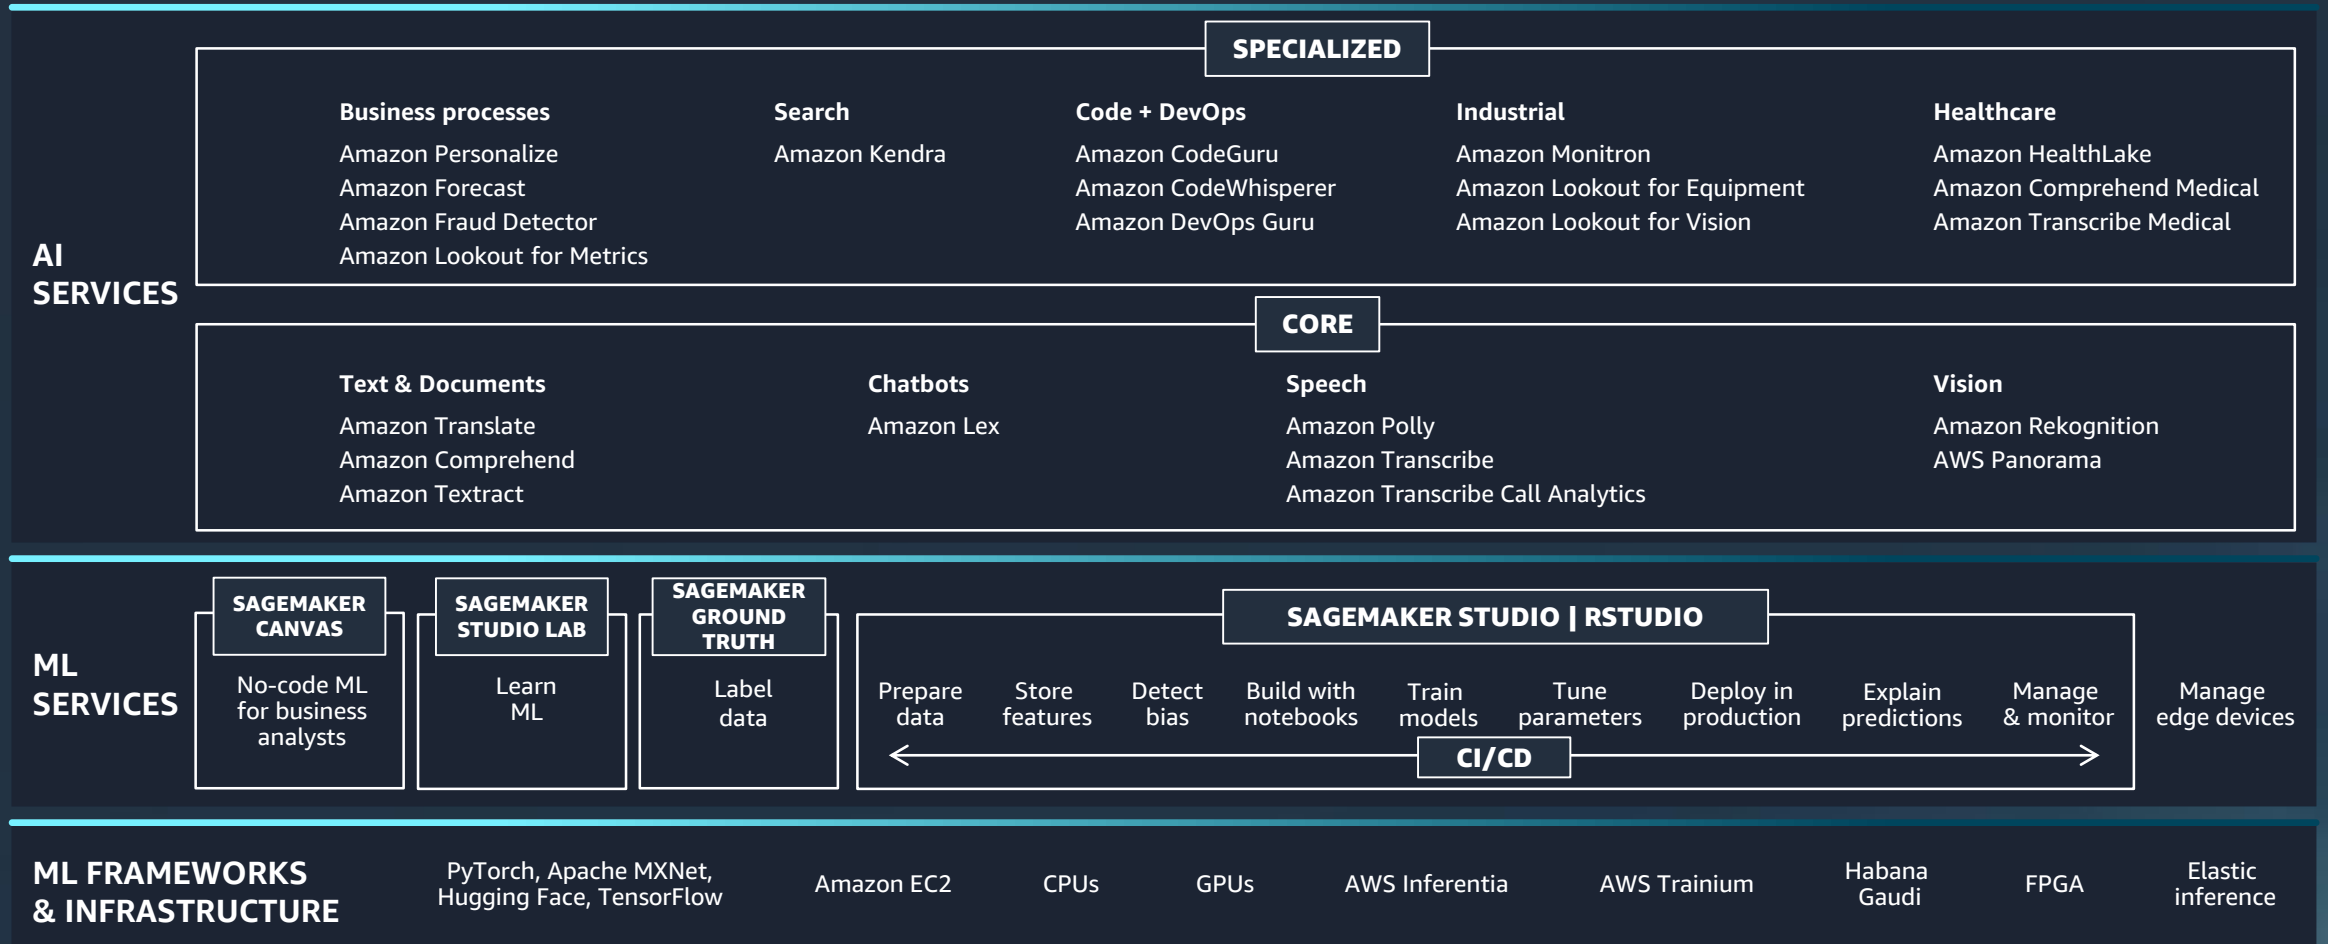
Elastic
inference



The AWS ML stack




The AWS ML stack



Machine learning is becoming accessible to all



Practitioner



Analyst



App Developer



Business User

AWS brings ML closer to data



Databases



Data warehouses
+ data lakes



Business
intelligence tools

AMAZON
AURORA ML



AMAZON
NEPTUNE ML



AMAZON
REDSHIFT ML



AMAZON
ATHENA ML



AMAZON
QUICKSIGHT ML





Analyst

Amazon Redshift ML

CREATE, TRAIN, AND DEPLOY ML
MODELS USING FAMILIAR SQL
COMMANDS

TRAIN



```
CREATE MODEL customer_churn
FROM (SELECT c.age, c.zip, c.monthly_spend,
c.monthly_cases, c.active AS label
FROM customer_info_table c)
TARGET label
FUNCTION predict_customer_churn
```

PREDICT



```
SELECT n.id, n.firstName, n.lastName,
predict_customer_churn(n.age,c.zip,..)
AS activity_prediction
FROM new_customers n
WHERE n.marital_status = 'single'
```




“...We were immediately struck by the fact that Amazon Redshift ML makes fitting an XGBoost model or feed-forward neural network as easy as writing a SQL query...”





Analyst

Amazon Athena ML

INVOKE ML MODELS FOR
INFERENCE DIRECTLY FROM A
SQL QUERY

P R E D I C T



```
USING EXTERNAL FUNCTION
predict_customer_churn (age INTEGER)
  RETURNS DOUBLE
  SAGEMAKER 'xgboost-2019-09-20-04-49-29-
303'

SELECT n.id, n.firstName, n.lastName,
predict_customer_churn(n.age)
AS activity_prediction
FROM new_customers n
WHERE n.marital_status = 'single'
```



App Developer

Amazon Aurora ML

SIMPLE, OPTIMIZED, AND
SECURE INTEGRATION BETWEEN
AURORA AND AWS ML SERVICES
USING SQL



```
1 select customer_id, review_id, review_body
2 from review_simple;
3
```

Successfully run. Total query runtime: 1 secs 609 msec.
49 rows affected.



Practitioner

Amazon Neptune ML

MAKE PREDICTIONS ON GRAPH
DATA WITHOUT ML EXPERTISE

TRAIN



```
neptune_ml.train_model (  
    export_file_location,  
    TRAINING_JOB_NAME,  
    PROCESSING_PARAMS,  
    TRAINING_PARAMS,  
    ENDPOINT_PARAMS  
)
```


PREDICT



```
%%gremlin  
  
g.with (  
    "Neptune#ml.endpoint", "<ENDPOINT_NAME>").  
V().has('title', 'MOVIE_NAME(YEAR)').properties  
("genre").with("Neptune#ml.classification").value()
```



Detecting identity fraud using Amazon Neptune ML





Business User

Amazon QuickSight ML

ML-POWERED INSIGHTS AND
DATA VISUALIZATIONS



ML predictions



Machine learning powered insights



Natural language querying capability



Auto narrative



MoM Sales

\$2,609 ↑

MoM Customer Orders

-26 ↓

Anomalies Detected

The top anomaly detected on Mar 2022 was:
 • Sales for **Mondelez International** at \$14,052, which was **higher** than the expected \$505


Region

All

Customer Segmentation



Sales by Industry



Forecast

Total sales is forecasted to be **\$48,657** for Aug 2023

Period over period

Total sales for Jun 2022 **increased by 5.71%** (\$2,609) from \$45,651 to \$48,260.



SHOWING TOP 48 IN LONGITUDE, LATITUDE AND TOP 3 IN REGION



Quarter	Strategic	SMB	Enterprise
Q2' 2022	\$51,000	\$49,292	\$33,731
Q1' 2022	\$33,551	\$63,530	\$21,814
Q4' 2021	\$94,031	\$100,846	\$41,016
Q3' 2021	\$43,915	\$79,942	\$20,757
Q2' 2021	\$42,418	\$69,477	\$23,475
Q1' 2021	\$26,578	\$46,031	\$19,987

Top 25 Customers




Profit

\$10K

Putting machine learning in the hands of every builder



Practitioner



Analyst



App Developer



Business User

AWS brings **ML** closer **to the data**



AMAZON
AURORA ML



AMAZON
NEPTUNE ML



AMAZON
REDSHIFT ML



AMAZON
ATHENA ML



AMAZON
QUICKSIGHT ML

Visit the Data & AI/ML resource hub

Dive deeper into these resources, get inspired and learn how you can use AI and machine learning to accelerate your business outcomes.

- 6 steps to machine learning success e-book
- 7 leading machine learning use cases e-book
- Machine learning at scale e-book
- Achieving transformative business results with machine learning e-book
- Tackling our world's hardest problems with machine learning e-book
- Accelerating machine learning innovation through security e-book
- ... and more!



<https://bitly.co/FqdC>

[Visit resource hub](#)



AWS Training and Certification

Access the AI & ML learning plan courses built by AWS experts on AWS Skill Builder

- Get started with digital self-paced, on-demand training and ramp-up guides to help you grow your technical skills
- Learn how to apply machine learning, artificial intelligence, and deep learning to unlock new insights and value in your role
- Take the steps today, towards validating your expertise with an AWS Certified Machine Learning – Specialty Certification



<https://bit.ly/3FnxDH7>

Learn your way [explore.skillbuilder.aws](https://skillbuilder.aws) »



Thank you for attending AWS Innovate – Data & AI/ML Edition

We hope you found it interesting! A kind reminder to **complete the survey**.
Let us know what you thought of today's event and how we can improve the event experience for you in the future.



aws-apj-marketing@amazon.com



twitter.com/AWSCloud



facebook.com/AmazonWebServices



youtube.com/user/AmazonWebServices



slideshare.net/AmazonWebServices



twitch.tv/aws

Thank you!

