



AI/ML & Data Analytics

Services Introduction

Sebastian Gebski

Principal SA, Startups CEE

Machine learning on AWS: Introduction

AI/ML on AWS

Innovation, choice, and flexibility

100,000+

Customers have used
machine learning on AWS

250+

New capabilities for
machine learning and
artificial intelligence in
just the last 12 months

92% of deep learning in the cloud runs on AWS

91% of cloud based PyTorch runs on AWS

AWS MACHINE LEARNING SOLUTIONS

Reduce training time by 50%

Provide 90% scaling efficiency

Deliver 3x faster network throughput

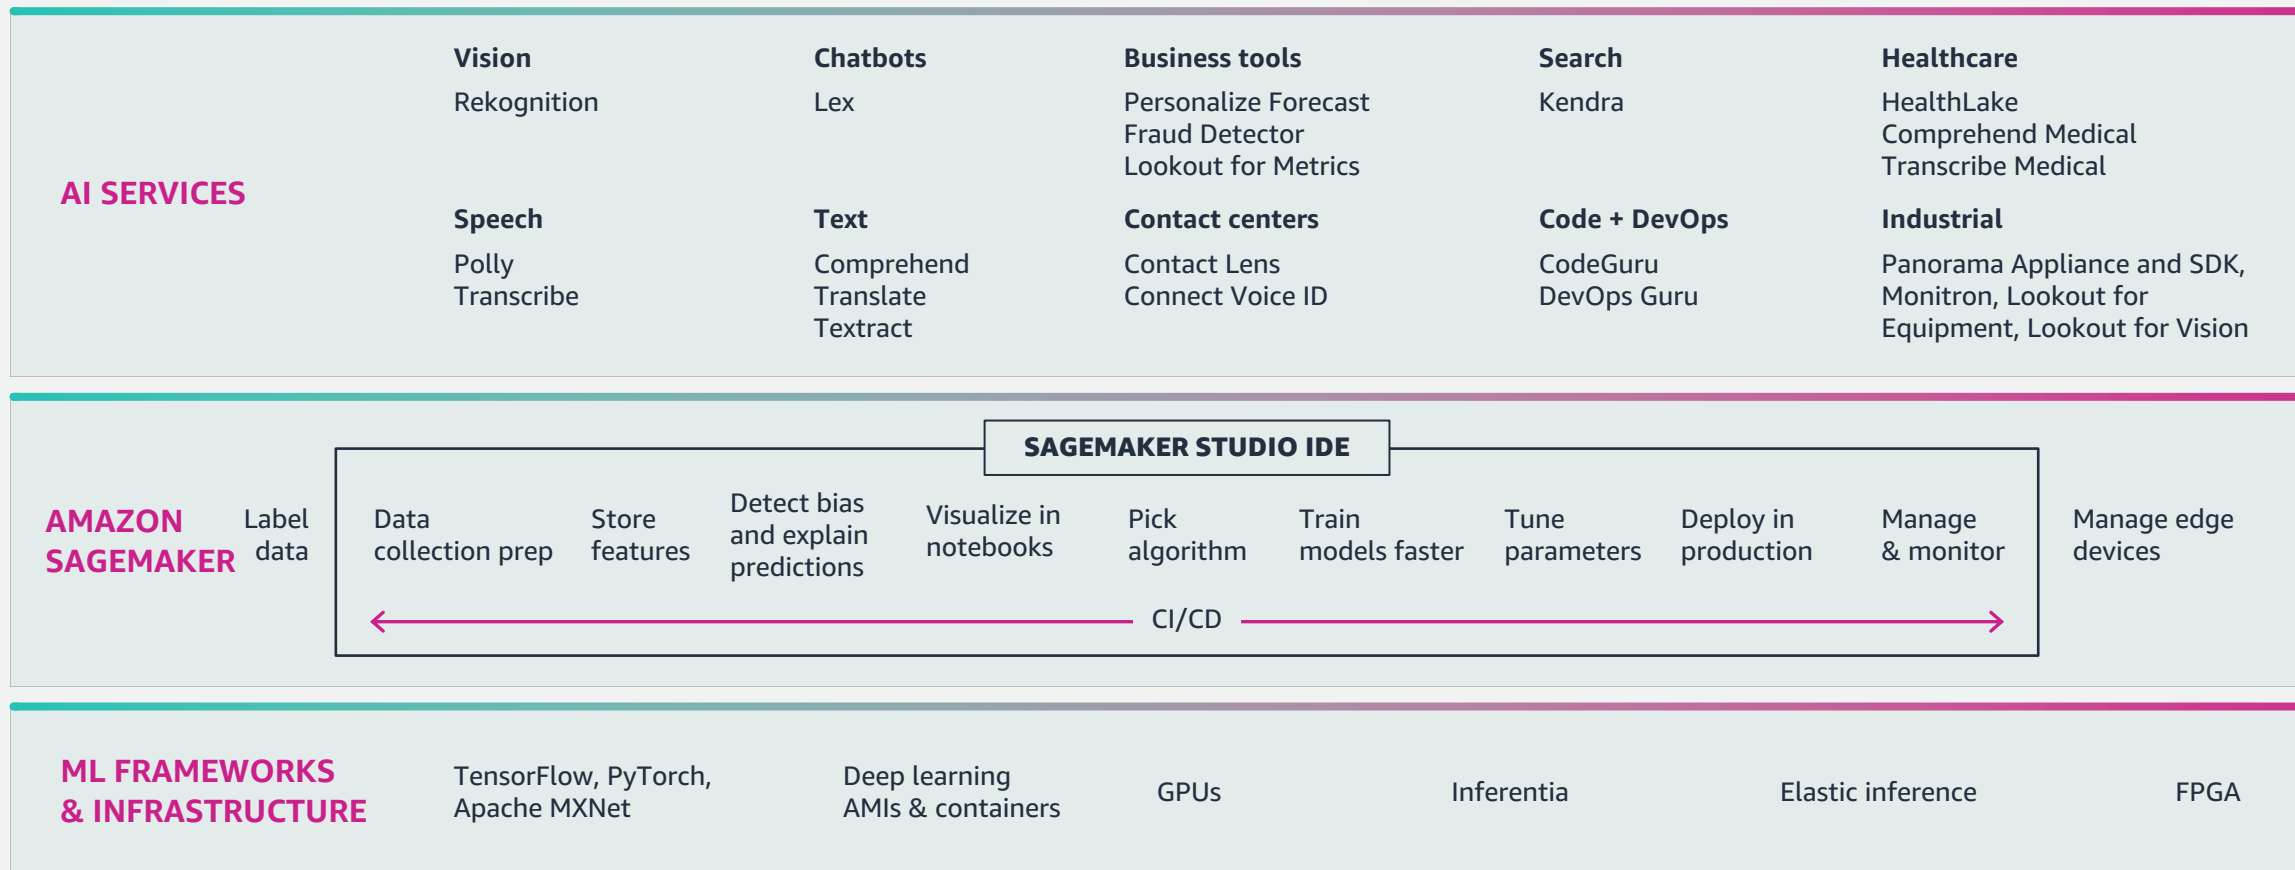
Improve price and performance by 25%

More than one hundred thousand customers use AWS for machine learning



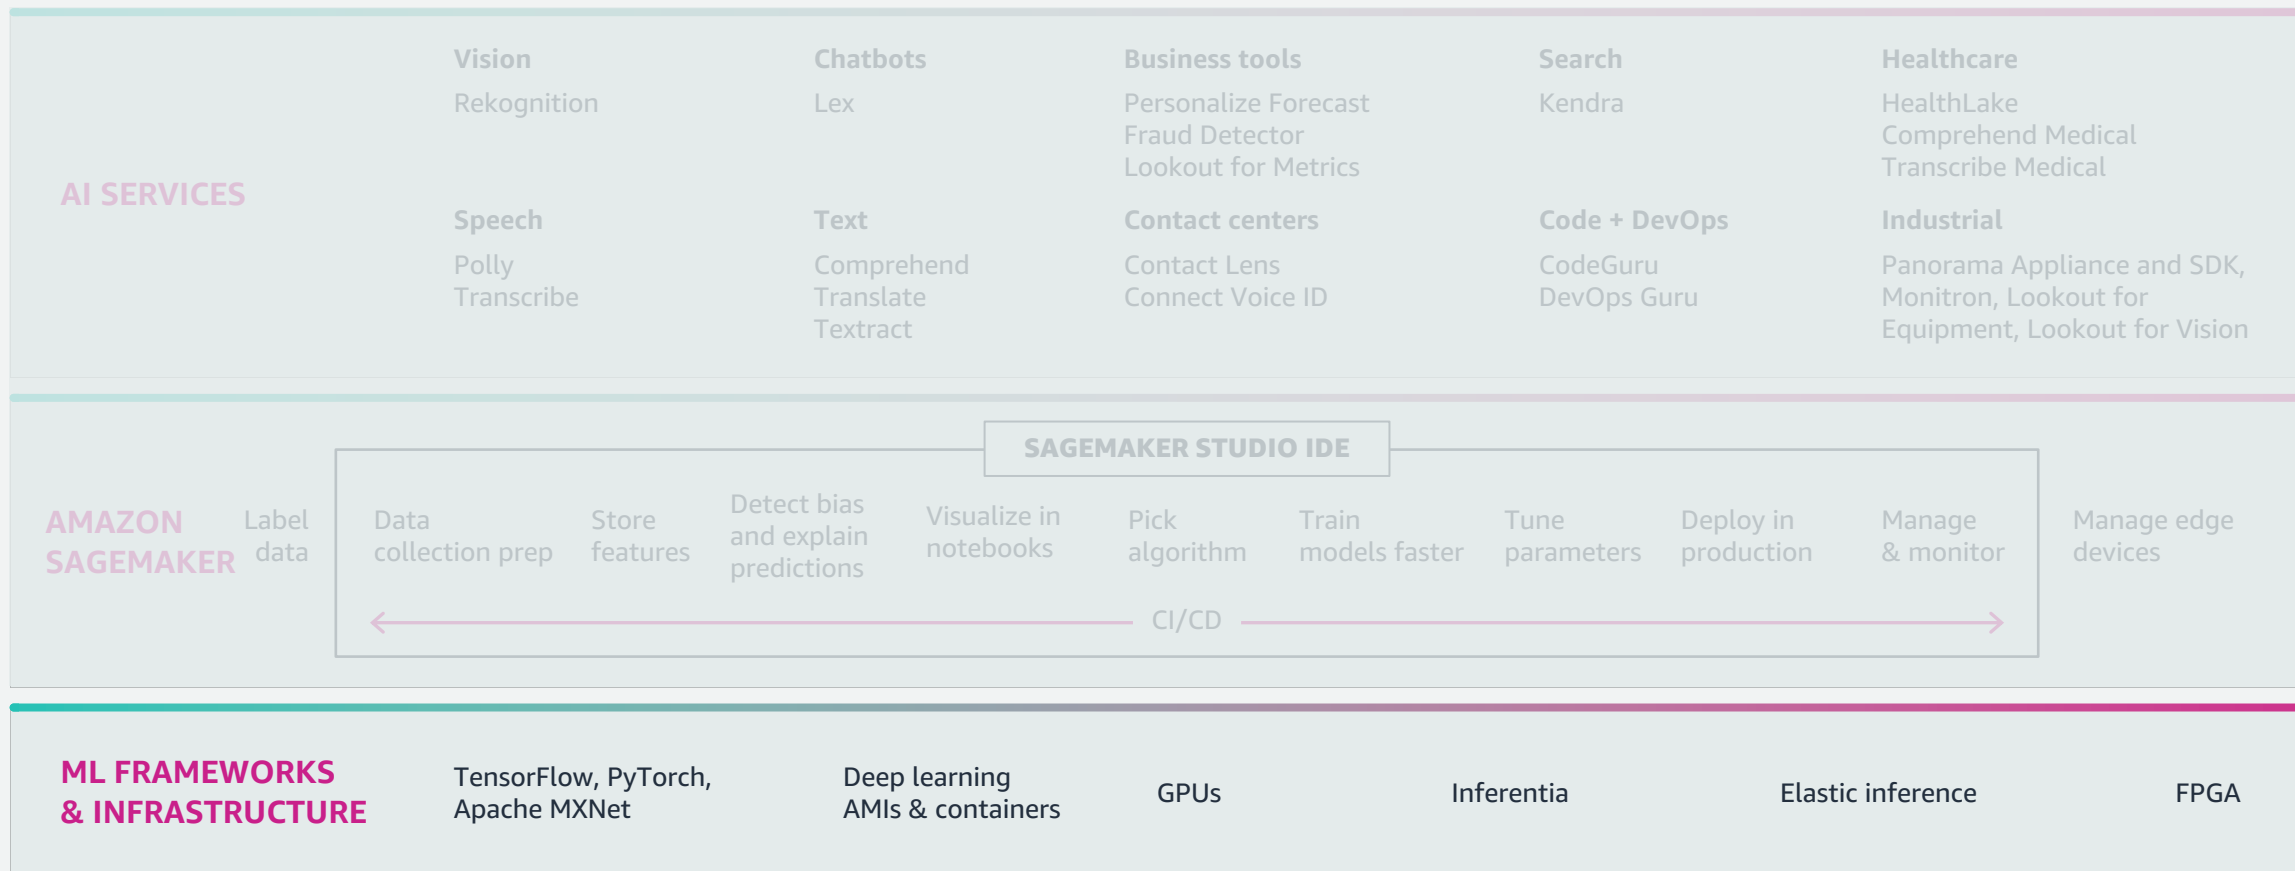
The AWS ML stack

Broadest and most complete set of machine learning capabilities



The AWS ML stack

Broadest and most complete set of machine learning capabilities



Broadest and deepest compute infrastructure for AI/ML

Choice of CPUs, GPUs, and accelerators for your performance and budget needs

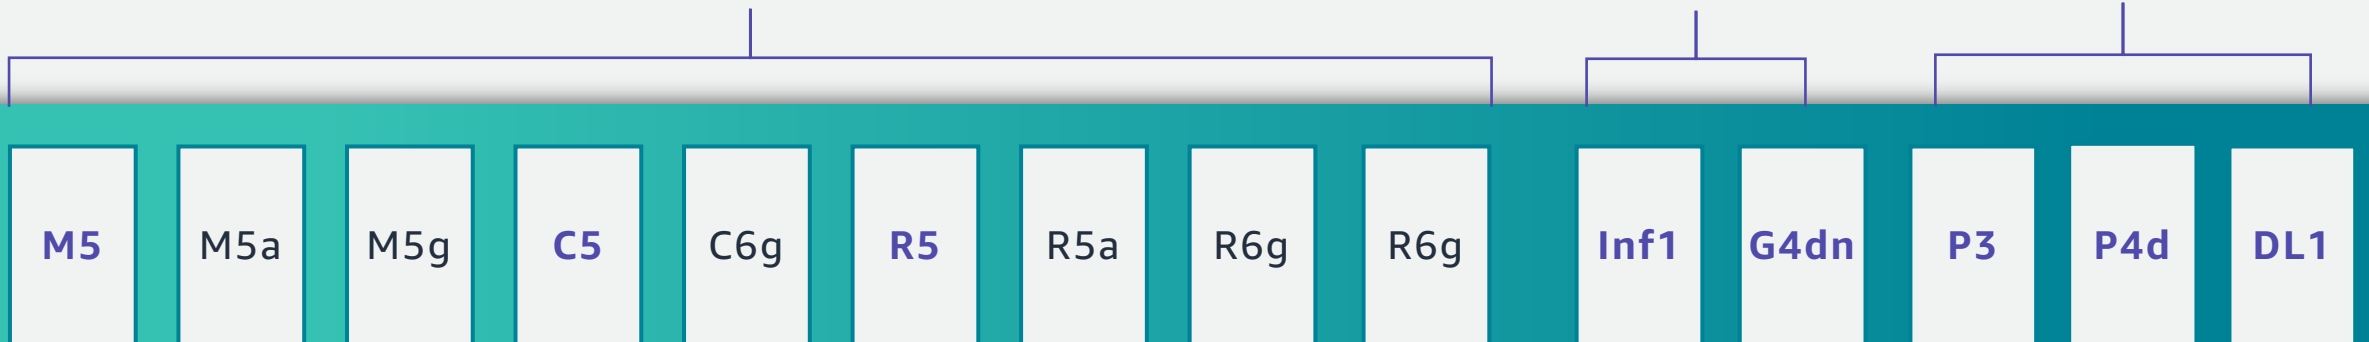
Traditional machine learning

Training + inference

Deep learning

Inference

Training



Cascade Lake CPU
Skylake CPU

Habana® Gaudi® Accelerators



EPYC CPU




Graviton CPU
Inferentia Chip



A100, V100, T4 GPUs

Introducing Amazon EC2 P4d Instances

P4d instances



The most powerful GPU instances in the cloud

Up to 60% lower cost to train ML model,
2.5x more deep learning performance and
25% more GPU memory


Powered by 8 NVIDIA A100 GPUs, 400
Gbps of network bandwidth, and capable
of 2.5 PetaFLOPS of performance

Deployed in UltraClusters consisting
of thousands of tightly coupled GPUs,
ideal for ML training and HPC

Introducing Amazon EC2 DL1 Instances

Best price performance for training deep learning models in the cloud

DL1 instances



Featuring up to 8 Gaudi accelerators by Habana Labs (an Intel company)

Specifically built for training deep learning models

Up to 40% better price performance than latest GPU instances


Custom software seamlessly integrated with TensorFlow and PyTorch

Get started easily using DLC, DL AMIs, or Amazon SageMaker


Launch DL1 instances via Amazon ECS and EKS for containerized ML applications

Inf1 are built from the ground up by AWS


Purpose built for ML inference




AWS Nitro



AWS Inferentia



AWS Custom 2nd Gen Intel
Xeon Scalable Processors



100Gbps Networking



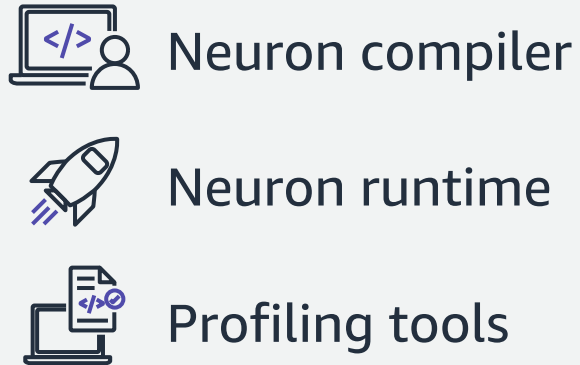
High performance



Low cost

AWS Neuron

High-performance Software Development Kit (SDK)




Easy to get started
Integrated with major frameworks



AWS Neuron

Flexibility of choice

Deploy existing models with minimal code changes.
Maintain hardware portability without dependency
on AWS software

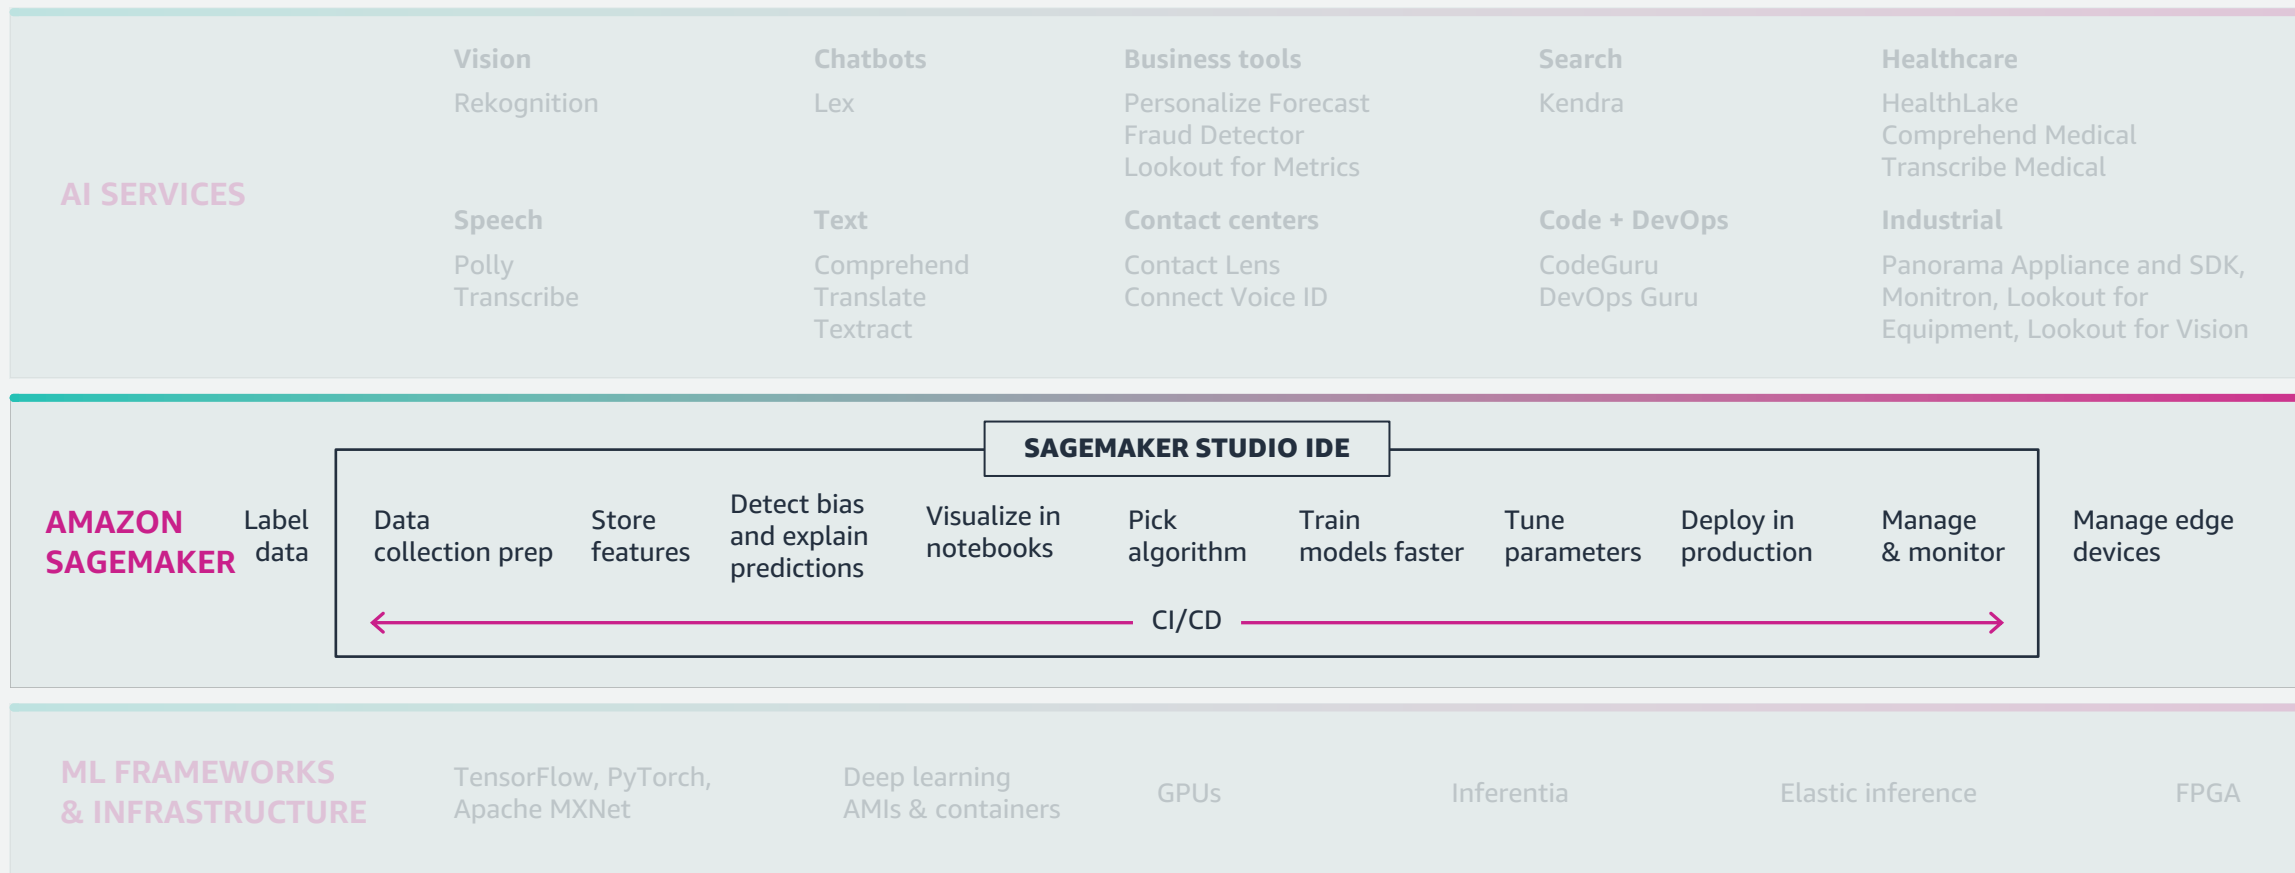


Documentation,
Examples & Support

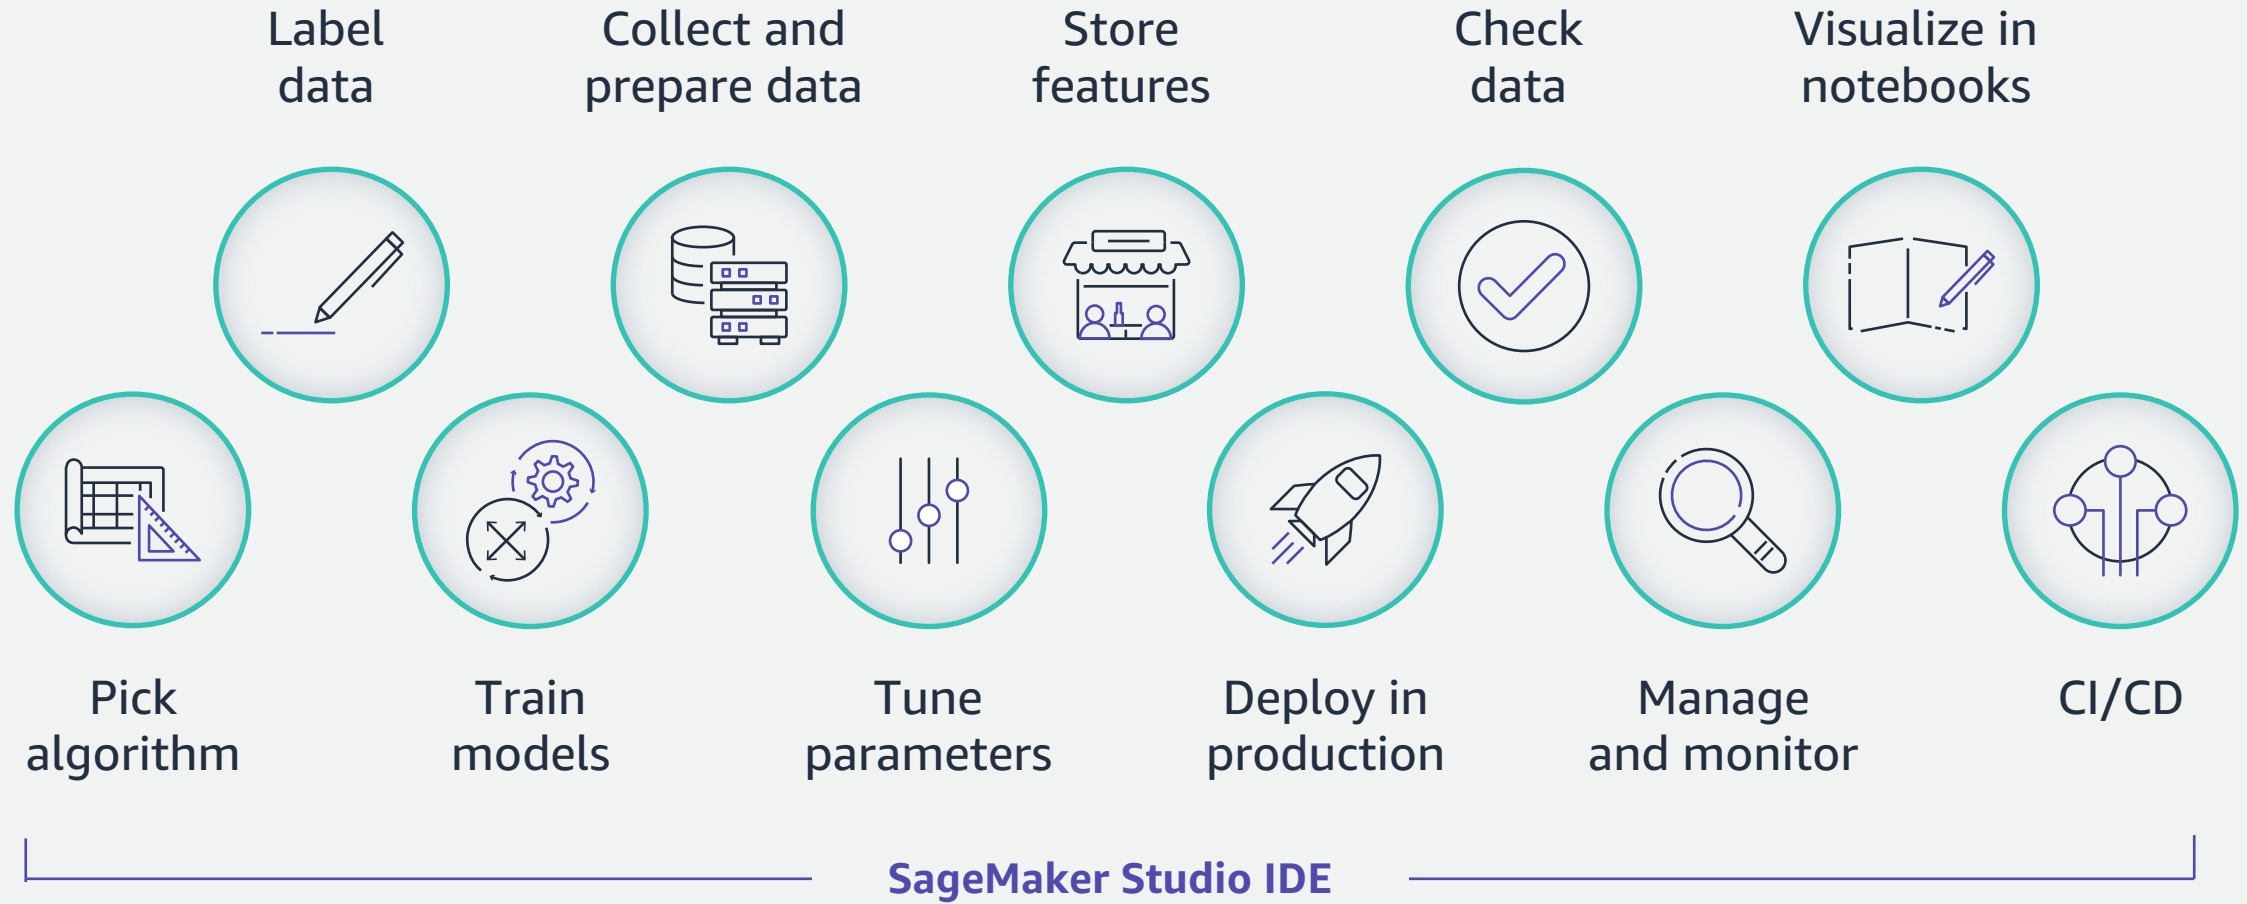
github.com/aws/aws-neuron-sdk

The AWS ML stack

Broadest and most complete set of machine learning capabilities



Amazon SageMaker: built to make ML more accessible



INTEGRATED WORKBENCH

IDE designed specifically for ML, data preparation, experiment management, and pipelines

MANAGED INFRASTRUCTURE

Designed for ultra low latency and high throughput; automatic scaling, and distributed training

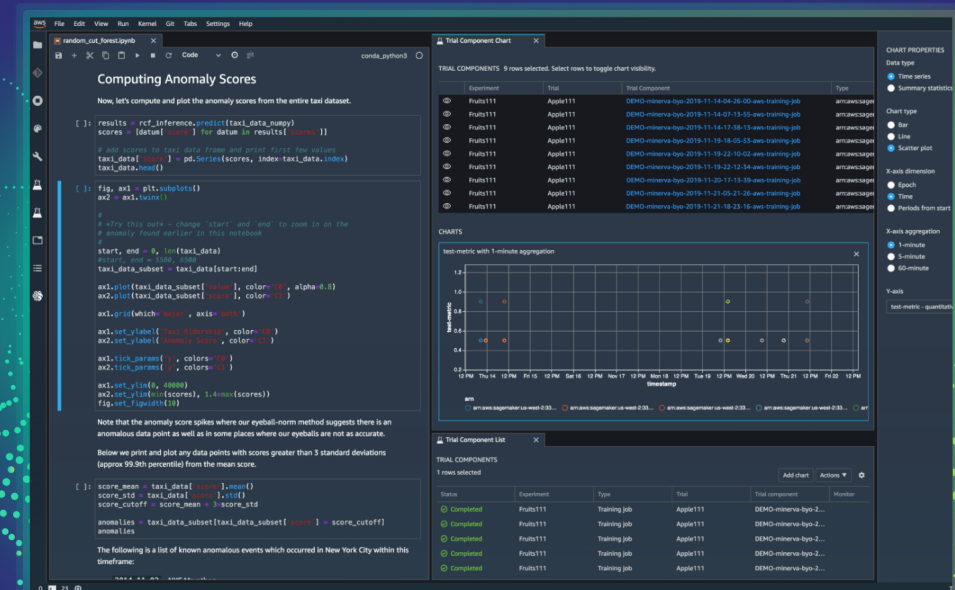
MANAGED TOOLING

Purpose-built from the ground up to work together incl. Autopilot, collaboration, notebooks, experiments, debugger, and model monitor

<https://aws.amazon.com/sagemaker>

Amazon SageMaker

Most complete,
end-to-end ML service



Amazon SageMaker overview

Amazon SageMaker

Prepare

SageMaker Ground Truth

Label training data for machine learning

SageMaker Data Wrangler

Aggregate and prepare data for machine learning

SageMaker Processing

Built-in Python, BYO R/Spark

SageMaker Feature Store

Store, update, retrieve, and share fetures

Build

SageMaker Studio Notebooks

Jupyter notebooks with elastic compute and sharing

Built-in and Bring-Your-Own Algorithms

Dozens of optimized algorithms or bring your own

Local Mode

Test and prototype on your local machine

SageMaker Autopilot

Automatically create machine learning models with full visibility

Train & tune

One-click training

Distributed infrastructure management

SageMaker Experiments

Capture, organize, and compare every step

Automatic Model Tuning

Hyperparameter optimization

SageMaker Debugger

Debug training runs

Managed Spot Training

Reduce training cost by 90%

Deploy & manage

One-click Deployment

Fully managed, ultra low latency, high throughput

Kubernetes & Kuberflow integration

Simplify Kubernetes-based machine learning

Multi-model endpoints

Reduce cost by hosting multiple models per instance

Model monitor

Maintain accuracy of deployed models

SageMaker Pipelines

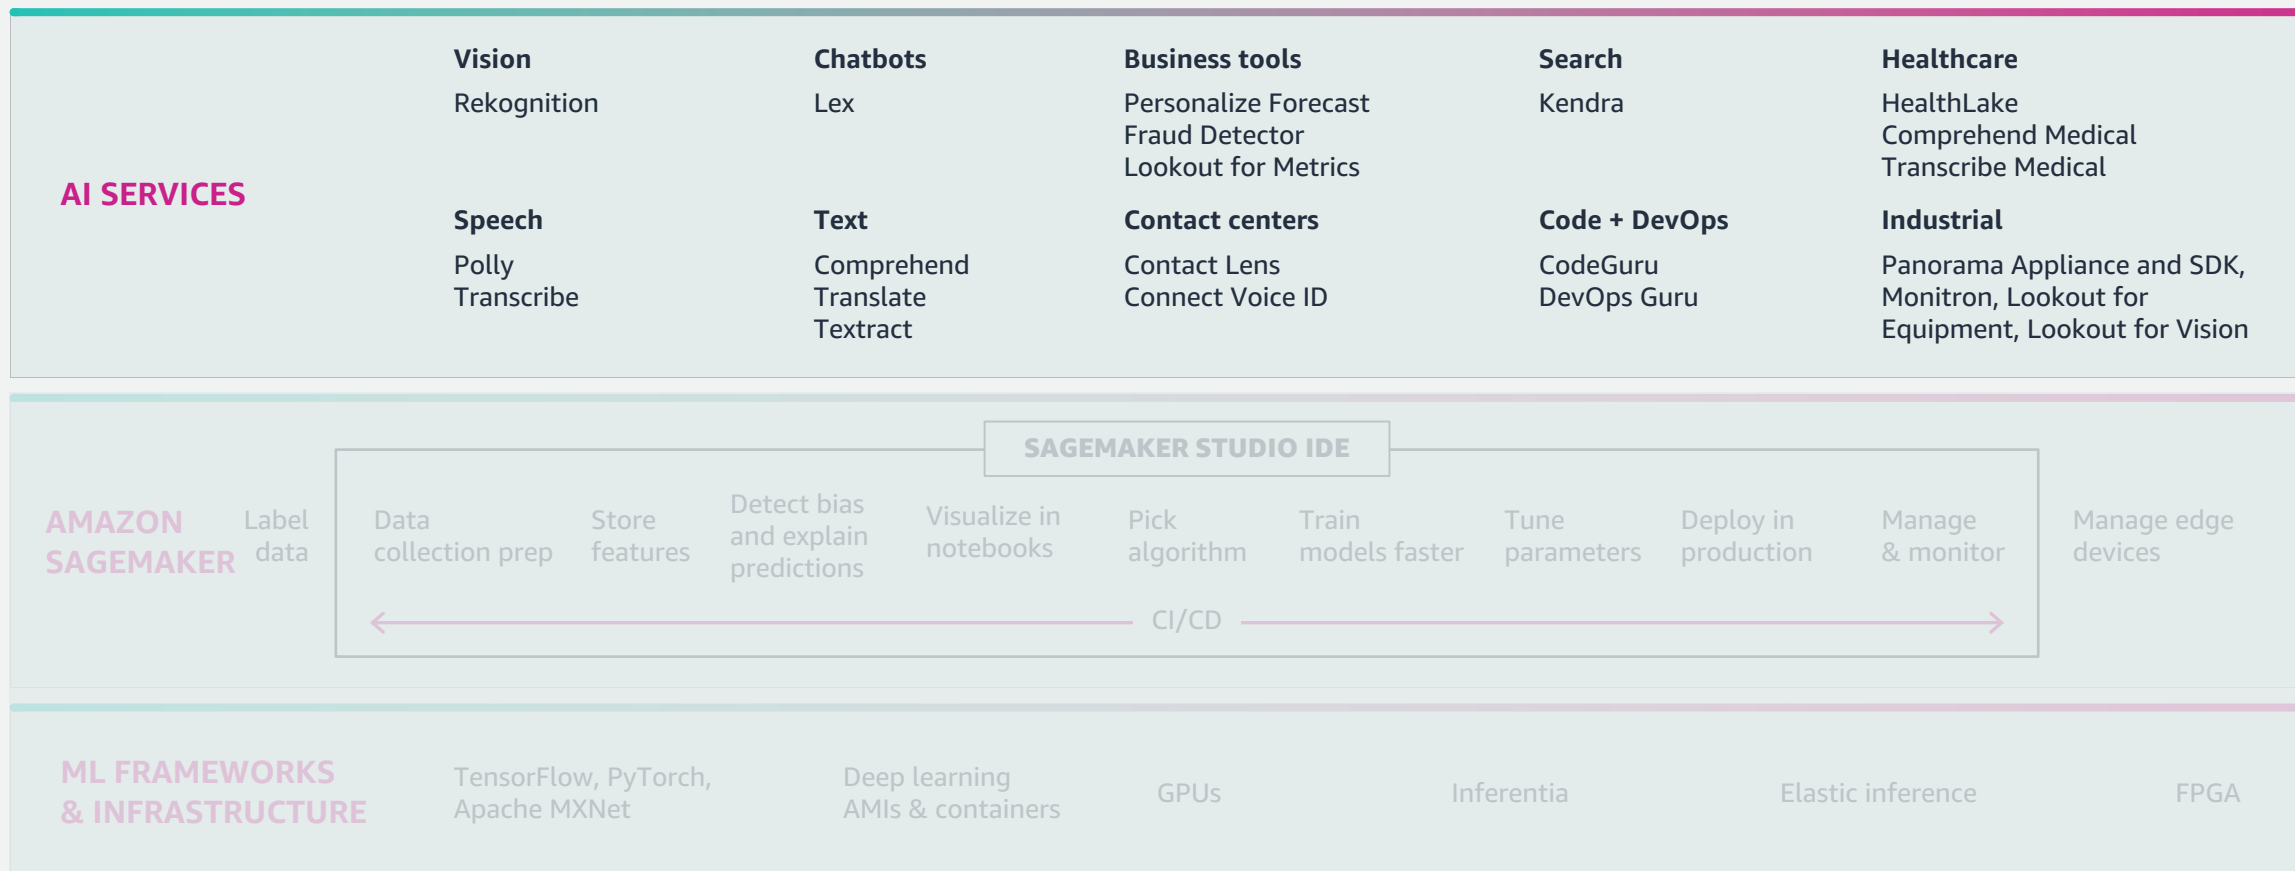
Workflow orchestration and automation

SageMaker Studio

Integrated development environment (IDE) for ML

The AWS ML stack

Broadest and most complete set of machine learning capabilities




Quickly and easily finding accurate information

Amazon Kendra

EASY TO FIND WHAT YOU ARE LOOKING FOR


QUICK SETUP

SECURE




Natural language queries

Contextual search
in unstructured
content




NLU And machine learning core

Reading comprehension
FAQ matching
Document ranking



Broad domain expertise

Pre-trained for
14 domains
and industries



Continuous improvement

Incrementally
learns from user
feedback



Native & partner connectors

Automate
ingestion
Find the
best answer
wherever it is



Secure search

Encrypted in
transit & rest
Token-based
access control


Better insights lead to better customer service with Contact Lens for Amazon Connect



Transcribes
live and
archived calls



Enhanced
search
on call and chat
Transcripts,
sentiment
scores, silence
duration, etc.



Custom
categorization
to identify
common
call types



Prioritized
list of recurring
issues based on
customer
feedback



Real time
dashboard
and alerting
for supervisors



Provide agents
with answers
to questions
as they are
being asked

Identify fraud faster with Amazon Fraud Detector



Enhance
fraud
detection
with ML



Any level of ML
expertise can
build ML fraud
models



ML boost
from Amazon
experience,
enrichments



Fewer false
positives,
manual reviews



Fraud staff
self-service
to address
threats faster



Lower TCO,
faster TTV

Amazon CodeGuru to build and run high-performing software



WRITE + REVIEW

Built-in code reviews
with intelligent
recommendations

BUILD + TEST

Detect and optimize
the expensive lines of
code pre-prod

DEPLOY


MEASURE

Easily identify application
inefficiencies in production
environment

IMPROVE

Amazon DevOps Guru

ML-powered cloud operations service to improve application availability



DevOps Guru is an ML-powered service that makes it easy for developers and operators to automatically detect issues to improve application availability and reduce expensive downtime—no machine learning experience required.


Delight your customers and improve customer experience with Amazon Personalize



Deliver
high quality
recommendations



Adapt to changes
in customer intent
in real time



Train a
recommendation
model with a
few clicks



Generate
recommendations
for almost any
product or content


Industrial AI

AI and ML services for asset-intensive industry use cases



AWS Panorama Appliance

Hardware appliance to add computer vision to existing onsite cameras



AWS Panorama Device SDK

Build new cameras and devices that run computer vision applications at the edge



Amazon Monitron

End-to-end system for equipment monitoring to detect abnormal machine behavior and enable predictive maintenance



Amazon Lookout for Equipment

Detect abnormal machine behavior using existing industrial sensor data



Amazon Lookout for Vision

Spot defects and anomalies in manufacturing using computer vision

Purpose-built and HIPAA-eligible services

AI and ML services for healthcare & life sciences use cases



Amazon HealthLake

Store, transform, query, and analyze health data in minutes




Amazon Comprehend Medical

Understand medical context with advanced text analytics



Amazon Transcribe Medical

Automatically convert medical speech to text



Amazon Textract

Easily extract text and data from virtually any medical document




Amazon Rekognition

Automate medical image and video analysis

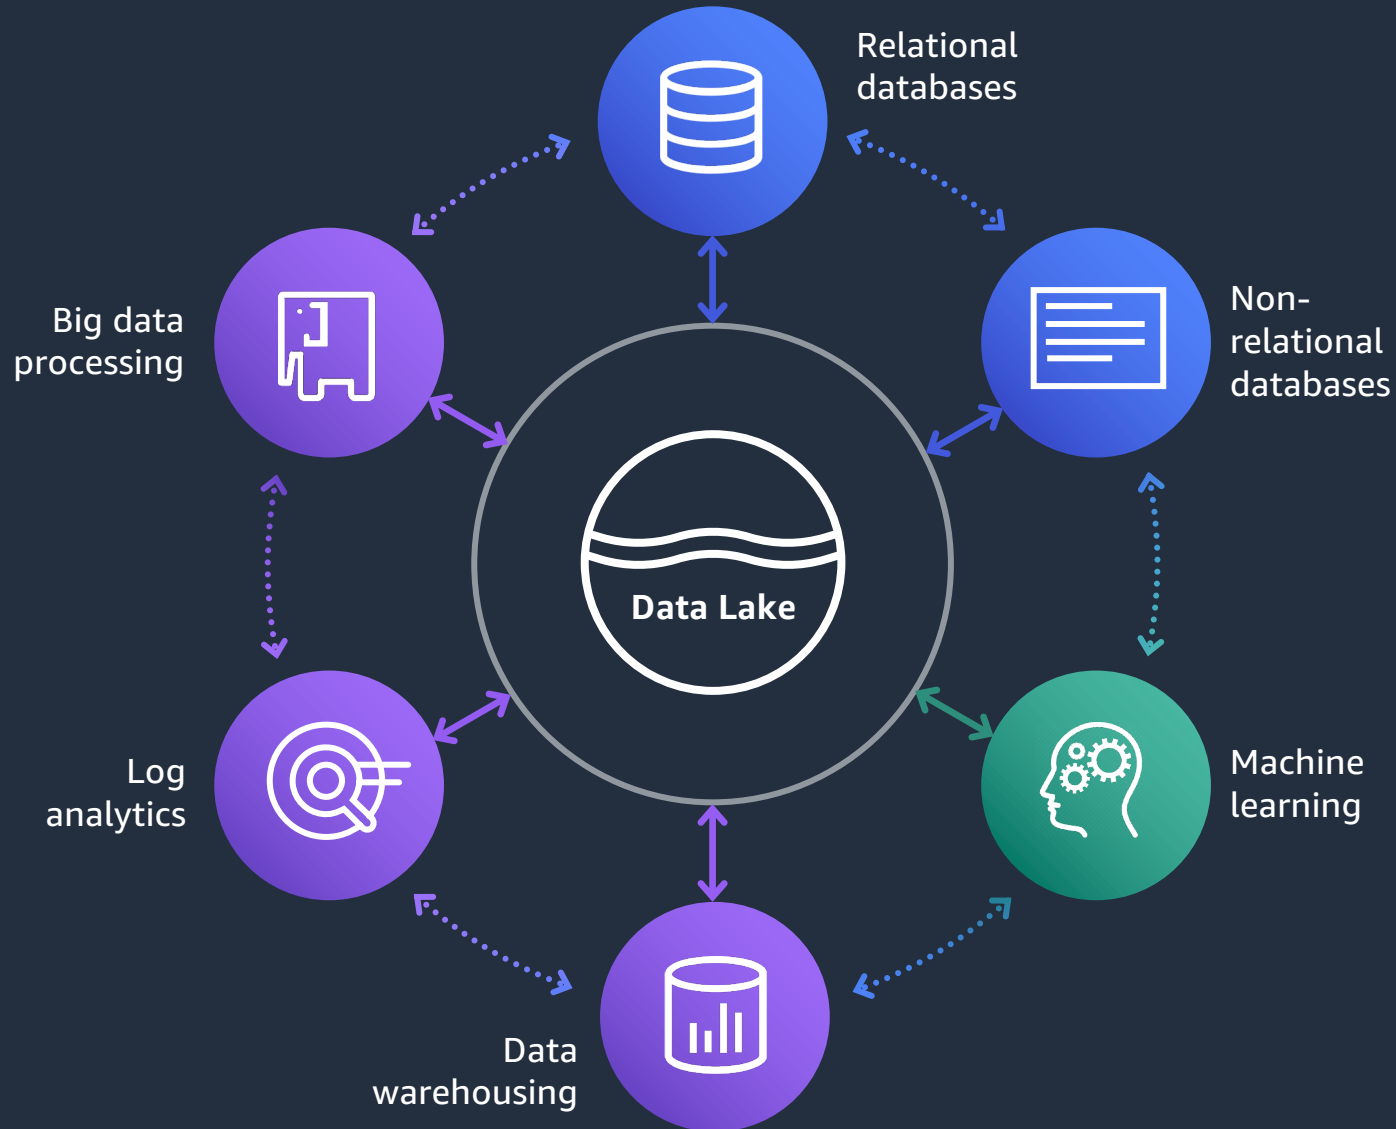


Data Analytics on AWS:

Introduction



Modern Data architecture (expectations)



Scalable data lakes

**Purpose-built
data services**

**Seamless
data movement**

Unified governance

**Performant and
cost-effective**

Modern Data architecture on AWS



Scalable data lakes

**Purpose-built
data services**

**Seamless
data movement**

Unified governance

**Performant and
cost-effective**

Scalable data lakes



Amazon S3 is the most popular choice for data lakes



More data lakes run on AWS than anywhere else

Tens of thousands of data lakes run on AWS across all industries

AMGEN

BMW
GROUP

CHANGE
HEALTHCARE

COMCAST

Continental

coursera

DATADOG

Discovery
CHANNEL

duolingo

EA

ENGIE

EPIC
GAMES

expedia
group

experian

Fannie Mae

FICO

airbnb

FOURSQUARE

INCHCAPE
SHIPPING SERVICES

INNOVO

intuit

INVISTA

ironSource

Klarna

Liberty Mutual
INSURANCE

mercado
libre

Movable Ink

GE

GRUBHUB

Firefox

docomo

New Relic

NuData Security
mastercard

OLX

PADDYPOWER
betfair

PennyMac
Correspondent

Pinterest

REDFIN

robinhood

ROBLOX

salesforce

SIEMENS

slack

Snap Inc.

Sysco

Goldman
Sachs

ANA

theTradeDesk

TOYOTA

Vanguard

vyaire
MEDICAL

yelp

Zillow

ZipRecruiter

Purpose-built data services

Optimize performance, cost, and scale for your use cases



Amazon
Athena

Interactive query



Amazon
EMR

Big data processing



Amazon OpenSearch
Service (successor to
Amazon Elasticsearch
Service)

Log and
search analytics



Amazon
Kinesis and
Amazon MSK

Real-time analytics



Amazon
Redshift

Data warehousing

Amazon Athena

Query data in S3 using SQL



Serverless

Quickly query S3 data without managing infrastructure, and pay only for the queries you run



Open and standard

Use ANSI SQL for querying with support for Parquet, CSV, JSON, Avro and other standard data formats



Fast interactive performance

Parallel execution to deliver most results within seconds, with no cluster management required



Cost effective

Pay only for queries run and save 30–90% by compressing, partitioning, and converting your data into columnar formats

Amazon EMR

Easily run Spark, Hadoop, Hive, Presto, HBase, and other big data frameworks



Automate provisioning, configuring, and tuning
Easy setup, management, and monitoring, with latest open-source framework updates within 30 days



Run workloads faster and more cost-effectively
1.7x faster than standard Apache Spark 3.0 at 40% of the cost, and 2.6x faster than open-source Presto 0.238 at 80% of the cost



Automatically scale up and down
Manage cluster size based on utilization to reduce costs



Simple and predictable pricing
Per-second pricing, and save 50%–80% with Amazon EC2 Spot and Reserved Instances

Amazon EMR differentiated performance

5



1.7x faster performance than standard
Apache Spark 3.0 at 40% of the cost

Up to 2.6x faster performance than open-source
Presto 0.238 at 80% of the cost

11.5% average performance improvement
with Graviton2

25.7% average cost reduction with Graviton2

Amazon OpenSearch Service (successor to Amazon Elasticsearch Service)

Search, visualize, and analyze up to petabytes of text and unstructured data



Fully managed

Operate OpenSearch with the leading contributor of the community-driven, open source software.



Easily accessible

Quickly search and analyze your unstructured and semi-structured data to easily find what you need.



Cost-effective

Eliminate operational overhead and reduce cost with automated provisioning, software installation, patching, storage tiering, and more.


The OpenSearch Project

An Apache 2.0-licensed search and analytics suite



100% open source

Providing you the freedoms, so you can freely view, use, change, and distribute the code



Enterprise-grade

Delivering security and advanced capabilities such as alerting, SQL, and cluster diagnostics



Community-driven

Providing individuals and organizations the freedom to easily contribute changes



Amazon MSK

Fully managed, highly available, and secure Apache Kafka service



Fully compatible

Run your existing Apache Kafka applications on AWS without changes to source code



Fully managed

Focus on creating applications not managing your Apache Kafka environment



Elastic stream processing

Run Apache Flink applications written in SQL, Java, or Scala that elastically scale to process data streams



Highly available

Take advantage of multi-AZ replication within an AWS region



Highly secure

Protect your data with multiple levels of security, including VPC network isolation, encryption at-rest and in-transit, and more



Amazon Kinesis

Easily collect, process, and analyze data and video streams in real time



Kinesis Data Analytics

Analyze data streams with serverless Apache Flink or SQL



Kinesis Data Streams

Capture, process, and store data streams



Kinesis Data Firehose

Load data streams into AWS data stores

Amazon Redshift

Analyze all your data with the fastest and most widely used cloud data warehouse



Analyze all your data

Deepest integration with your data lake



Performance at any scale

Up to 3x better price performance than other cloud DW



Lower your costs

At least 50% less expensive than other cloud DW

Amazon Redshift innovates to meet your needs



Analyze all your data

Modern Data with
AWS integration



Performance & scale

Fast and self-tuning



Low cost & best value

Predictable costs



Amazon Redshift
Spectrum + Lake
Formation



Data lake
export



Federated
query



Data sharing



Amazon
Redshift ML



Concurrency
scaling



RA3 nodes &
managed
storage



AQUA



Materialized
views



Automated
perf. tuning



On-demand
and RIs



Cross-AZ cluster
recovery



Pause and
resume



Cost controls



Built-in security
features



Automatic
workload
manager

Seamless data movement



Seamless data movement

Move your data, at scale, to where you need it the most



Extract,
transform, load



Visual data
preparation



Data
replication



Data warehouse
to/from data lake



Federated
query

AWS Glue

Simple, scalable, and serverless data integration



Connect to more sources

Easily ingest data from hundreds of popular data sources



Simplify workflow orchestration

Easily run and manage thousands of data integration jobs



No servers to manage

Pay only for the resources your jobs consume

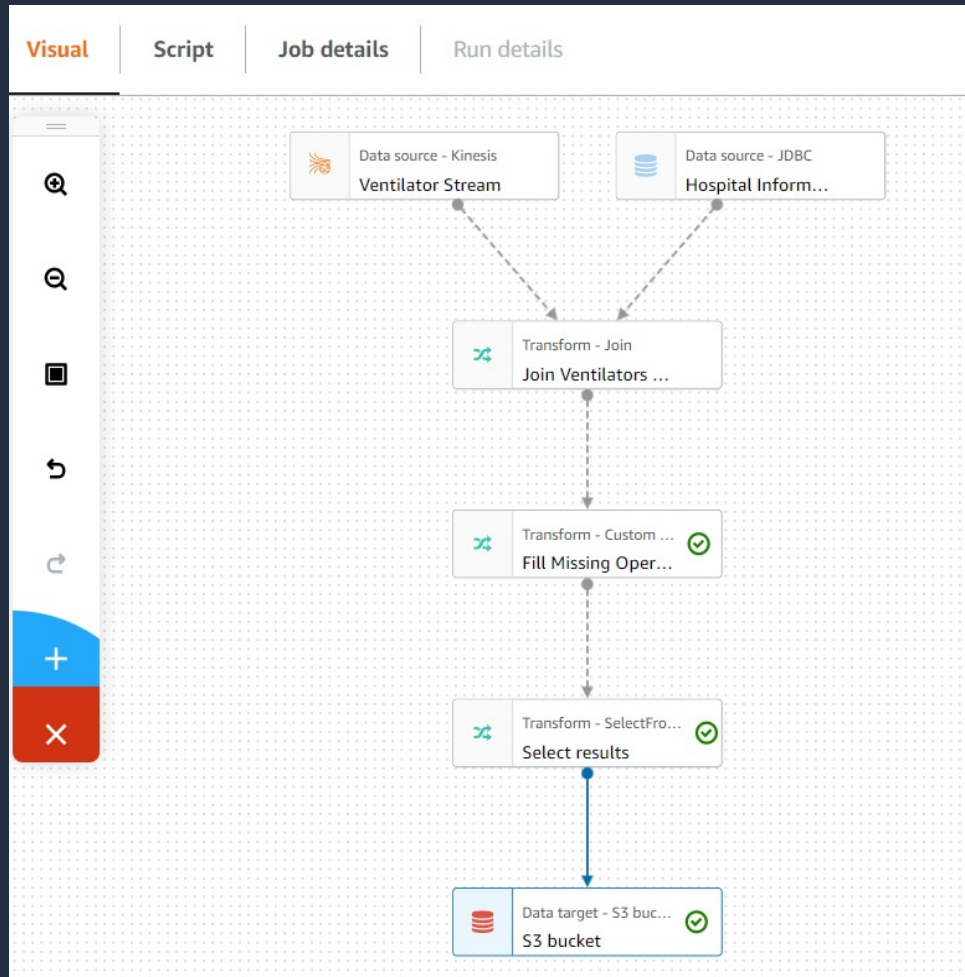


Simplify development

Visually develop and manage data integration jobs

AWS Glue Studio

Easily author, run, and monitor AWS Glue ETL jobs



Author AWS Glue jobs **visually** without coding


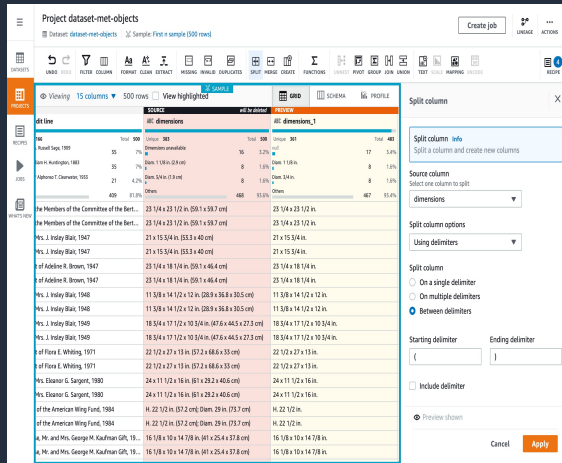
Monitor **1,000s** of jobs through a **single pane of glass**

Distributed processing without the learning curve

Advanced transforms **though code snippets**

AWS Glue DataBrew

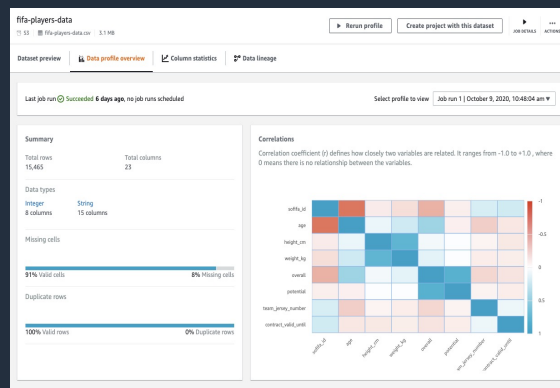
Visual data preparation for analytics and machine learning



Clean and normalize data with
a rich visual interface

Choose from 250+ built-in
transformations to automate tasks

Profile data to understand data
patterns and anomalies

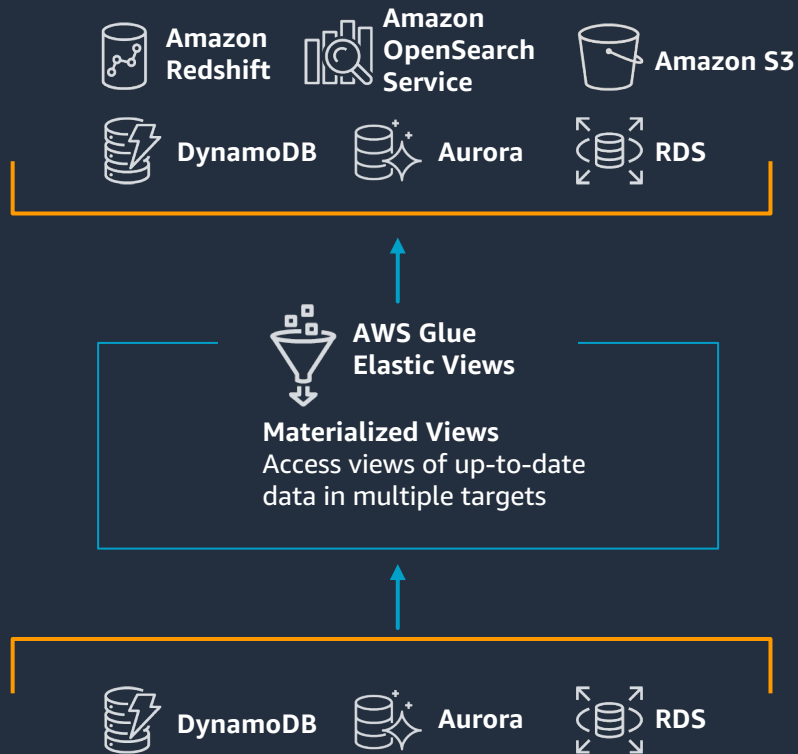


Work on large datasets at scale

AWS Glue Elastic Views

Easily combine and replicate data across multiple data stores

NEW
PREVIEW



Create materialized views across a wide variety of databases and data stores using familiar SQL

Continually monitors source databases for changes and updates targets within seconds

Serverless and automatically scales capacity up and down to accommodate your workloads

Handles the heavy lifting of copying and combining data without requiring custom code

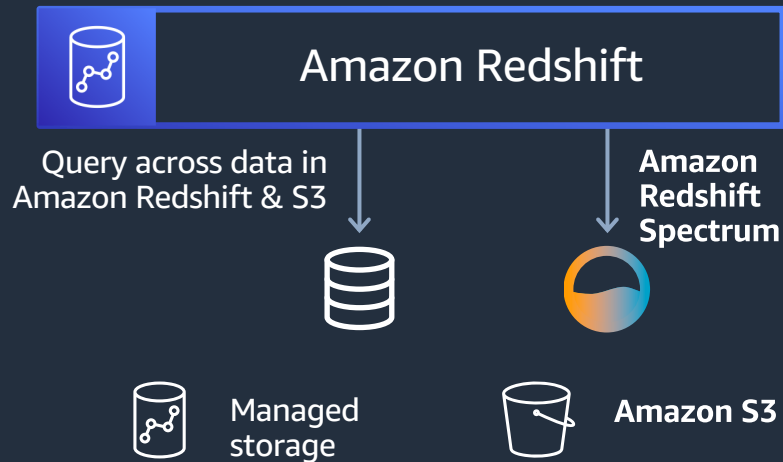
Moving data to and from the data lake

Extend the data warehouse to exabytes of data in Amazon S3 data lakes

Directly query data stored in Amazon S3

Parquet, ORC, Avro, JSON, and CSV data formats

Any scale of data; pay for what you use




Unload Amazon Redshift data as Parquet to Amazon S3 data lakes for faster sharing and analytics

Parquet is an open data format supported by Amazon EMR, Athena, and Amazon Redshift

Amazon Redshift now supports exporting data to Amazon S3 in Parquet format

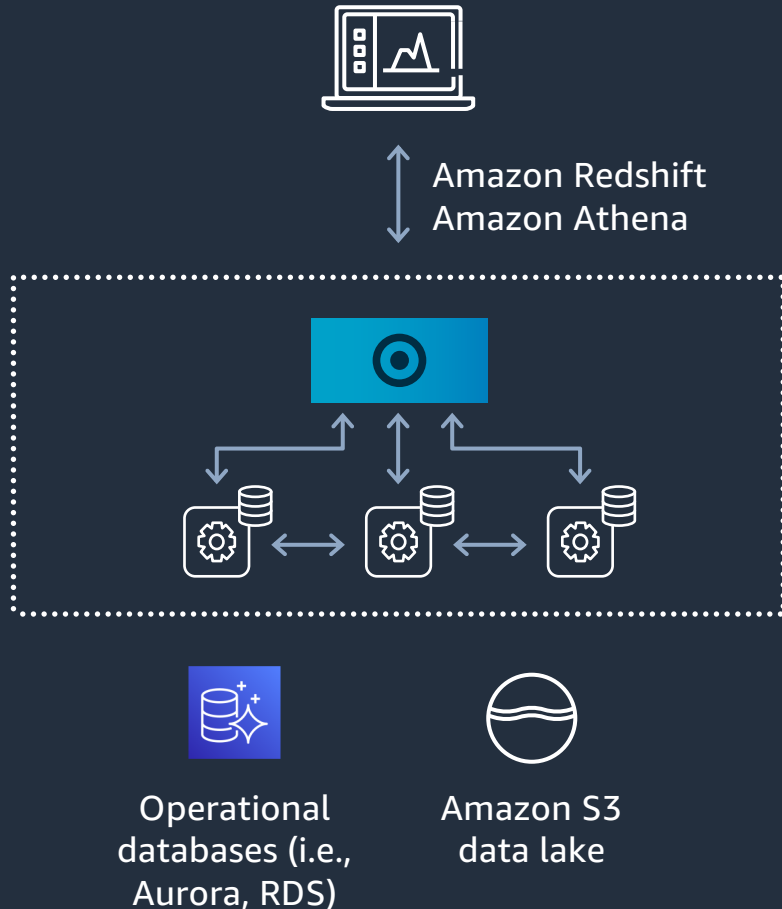
Use SQL with Amazon Redshift's Unload command to export data in Parquet format

Unloaded data is automatically registered in AWS Glue Data Catalog



Federated query in Amazon Redshift and Athena

Unified analytics across databases, data warehouse, and data lake



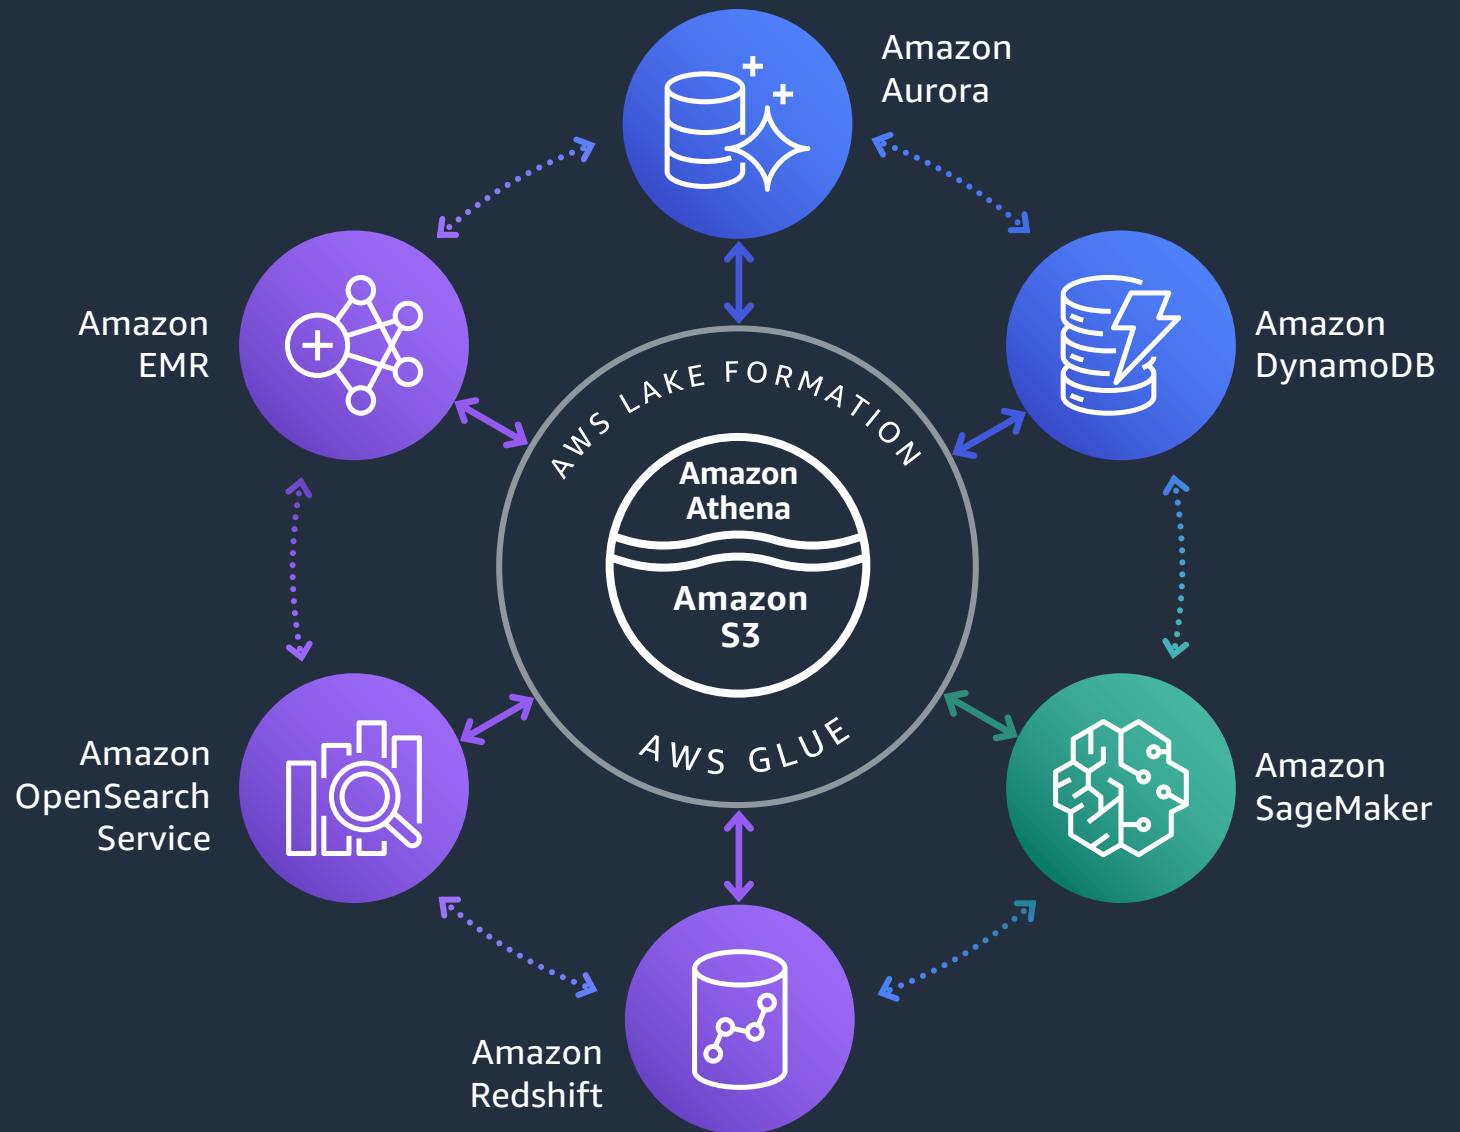
Integrate operational database with data warehouse and Amazon S3 data lake

Analytics on operational data without data movement and ETL delays

Flexible and easy way to ingest data, avoiding complex ETL pipelines

*Other sources available in Amazon Athena: Amazon ElastiCache for Redis, Amazon DocumentDB, Amazon DynamoDB, HBase in Amazon EMR

Unified governance



AWS Lake Formation

Build a secure data lake in days



Build data lakes quickly

Move, store, and catalog your data faster; simplify data management with governed storage



Simplify security management

Centrally define and enforce security, governance, and auditing policies



Provide self-service access to data

Share datasets easily and securely within your organization and with partners

Performant and cost-effective



Industry-leading choice
of 200+ instance types
to meet workload needs



100 GBPS bandwidth
network interfaces
for performance



On-demand, reserved,
and spot instances
to reduce costs



Five highly available
storage tiers and
intelligent tiering

Amazon QuickSight



A scalable, embeddable,
ML-powered BI service
built for the cloud



BI at scale

No servers to manage; pay per use billing



Embedded analytics

Quickly embed dashboards in your apps



ML-powered insights

Built-in anomaly detection and forecasting

Written narratives to interpret your data for you

Amazon QuickSight Q

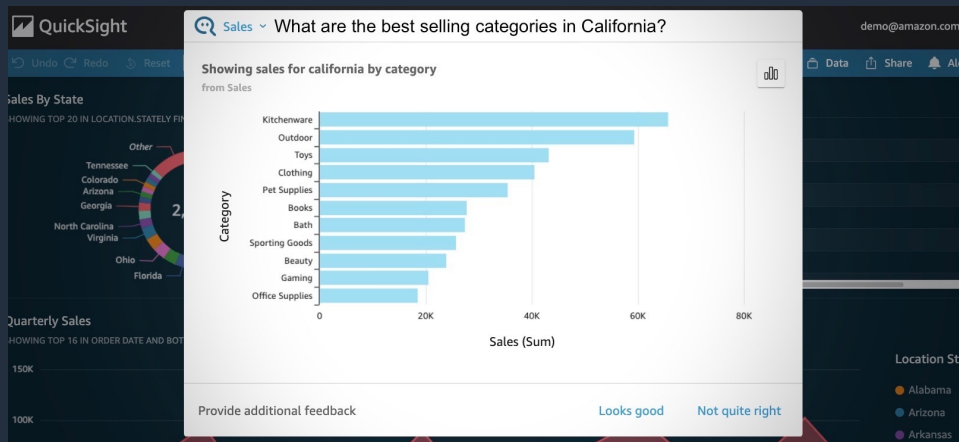
ML-powered natural language capability in Amazon QuickSight

NEW
PREVIEW

Enter business questions in search bar and get answer in seconds

ML generates data models that automatically understands meanings and relationships

Not limited to only asking a specific set of questions

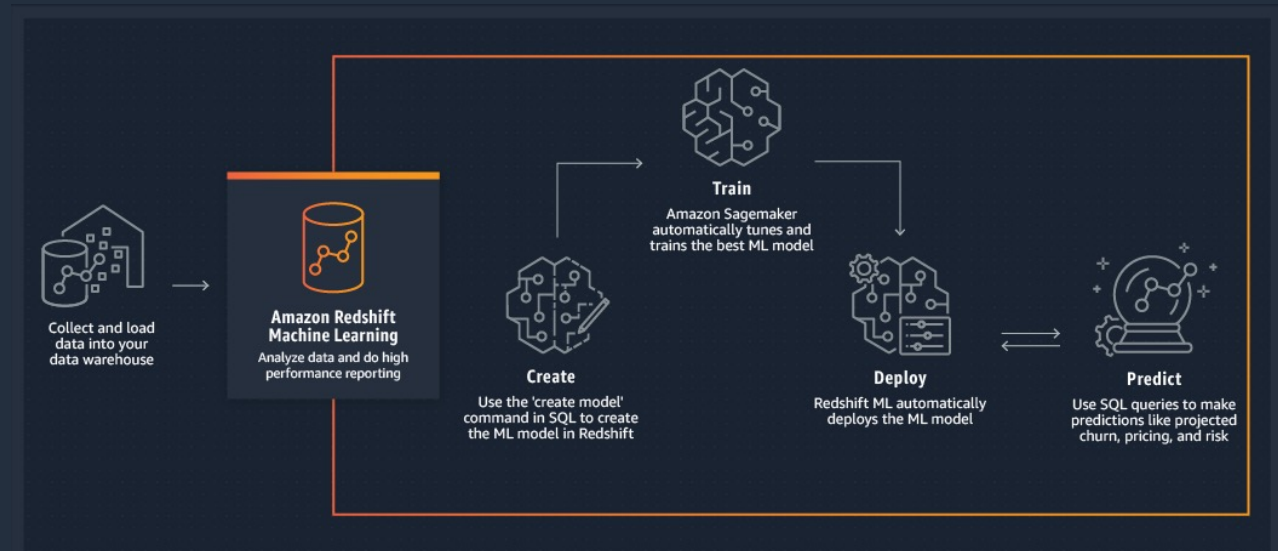


Amazon Redshift ML

PREVIEW

Create, train, and deploy machine learning (ML) models using familiar SQL commands

- ✓ Simple, optimized, and secure integration between Redshift and Amazon SageMaker
- ✓ Train and deploy an ML model using a SQL command in your data warehouse
- ✓ Embed predictions like fraud detection, risk scoring, and churn in queries and reports



Thank you