



aws INNOVATE

AI/ML EDITION

24 February 2022

Get started with Amazon SageMaker in minutes

Aashmeet Kalra

Principal Solutions Architect

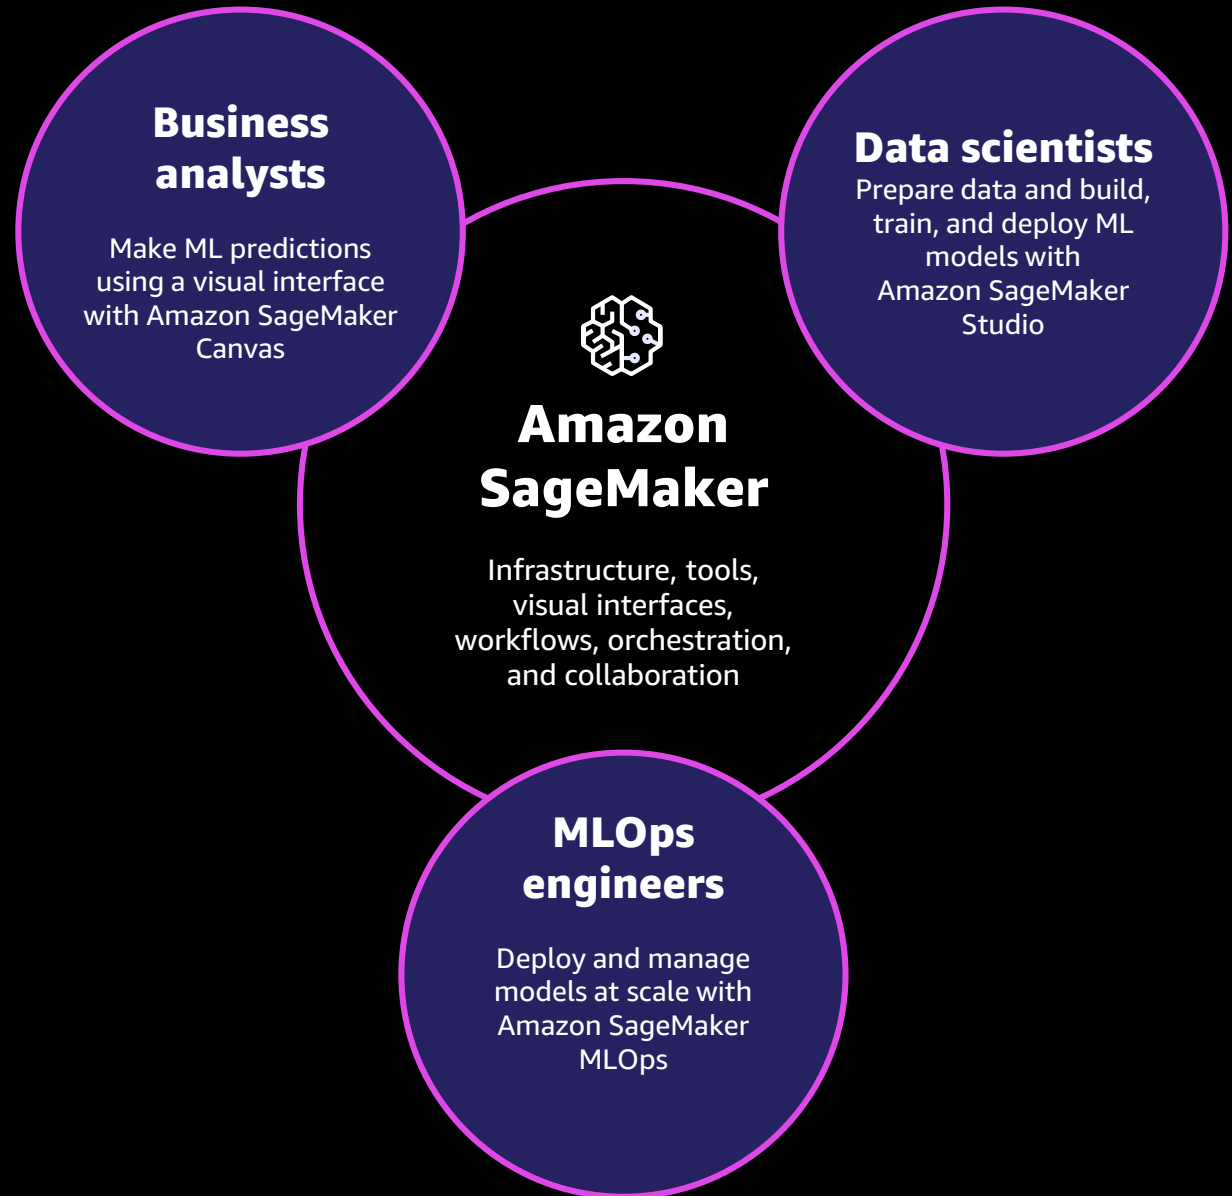
AWS



Agenda

1. Introduction to Amazon SageMaker
2. No-code ML tools
3. Purpose-built data preparation tools
4. Integrated ML tools in a single interface
5. Built-in MLOps capabilities

Amazon SageMaker helps organizations harness ML



Overcoming the barriers to ML



Not enough ML builders



No-code ML tools

Make ML predictions regardless of ML experience



Processing and labeling massive volumes of data for ML



Purpose-built data preparation tools

Label and prepare data for ML



Disparate data science tools



Integrated ML tools in a single interface

Build, train, and deploy models using IDEs



Tedious, manual ML operations



Built-in MLOps capabilities

Implement MLOps practices to streamline the ML lifecycle

No-code ML tools

Amazon SageMaker Canvas

Amazon SageMaker Canvas

Build ML models and generate accurate predictions—no code required



Quickly access and prepare data for Machine Learning



Built-in AutoML to build models and generate accurate predictions



Share ML models and collaborate with data science teams



Usage-based pricing to avoid licensing fees and reduce TCO

Purpose-built data preparation tools

Data preparation

Amazon SageMaker Data Wrangler


No-code data preparation



Single visual interface for common data prep techniques



Select data from multiple sources




300+ built-in transformations to prepare data without writing code

Amazon SageMaker Feature Store

Securely store, discover,
and share features for ML



Online and off-line



Millisecond latency



Consistent features



Visual search



Sharing and collaboration

Amazon SageMaker Ground Truth Plus

Deliver high-quality training datasets fast and reduce data labeling costs



Increase data quality through ML-powered data labeling



Access expert data labelers



Reduce data labeling costs with assistive labeling features




Improve operational efficiency by reviewing project metrics



Make data labeling accessible to data operations and program managers

Amazon SageMaker brings tools for every step of the ML lifecycle under one unified visual user interface



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 features

SageMaker Clarify

Detect bias and understand model predictions

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

SageMaker JumpStart

Pre-built solutions for common use cases

TRAIN & TUNE

Managed training

Distributed infrastructure management

SageMaker Experiments

Capture, organize, and compare every step

Automatic model tuning

Hyperparameter optimization

Distributed training libraries

Training for large datasets and models

SageMaker Debugger

Debug and profile training runs

Managed Spot training

Reduce training cost by 90%

DEPLOY & MANAGE

Managed deployment

Fully managed, ultra low latency, high throughput

Kubernetes and Kubeflow Integration

Simplify Kubernetes-based machine learning

Multi-model endpoints

Reduce cost by hosting multiple models per instance

SageMaker Model Monitor

Maintain accuracy of deployed models

SageMaker Edge Manager

Manage and monitor models on edge devices

SageMaker Pipelines

Workflow orchestration and automation

SageMaker Studio
Integrated development environment (IDE) for ML

Build ML models

Fully managed
shareable notebooks on
Amazon EC2




Fully managed, sharable Jupyter notebooks

Run notebooks on elastic compute resources




Built-in algorithms

15 built-in algorithms available in prebuilt container images




Prebuilt solutions and open-source models

Over 150 popular open-source models



AutoML – SageMaker AutoPilot

Automatically create ML models with full visibility




Support for major frameworks and toolkits


Optimized for popular deep learning (DL) frameworks such as TensorFlow, PyTorch, Apache MXNet, and Hugging Face

Amazon SageMaker Studio Notebook


Perform data engineering,
analytics, and ML workflows
in one notebook



**Connect with Amazon EMR,
Amazon S3, and more**




**Interactively access, transform, and
analyze a wide range of data**



**Build, train, and deploy models using
your preferred framework**

Train ML models

**Fast and cost-effective
ML model training**




Experiment management and model tuning
Save weeks of effort by automatically tracking training runs and tuning hyperparameters



Debug and profile training runs
Use real-time metrics to correct performance problems



Distributed training
Complete distributed training up to 40% faster



Training compiler
Accelerate training times by up to 50% through more efficient use of GPUs



Managed spot training
Reduce the costs of training by up to 90%

Deploy ML models

Fully managed deployment
for inference at scale



Wide selection of infrastructure

70+ instance types with varying levels of compute and memory to meet the needs of every use case



Single-digit millisecond overhead latency

For use cases requiring real-time responses



Asynchronous inference

Supports large models with long-running processing times



Cost-effective deployment

Multi-model/multi-container endpoints, serverless inference, and elastic scaling



Built-in integration for MLOps


ML workflows, CI/CD, lineage tracking, and catalog



Automatic deployment recommendations


Optimal instance type/count and container parameters, and fully managed load testing

Support the responsible use of ML throughout the model lifecycle



Build

Perform bias analysis during exploratory data analysis



Train

Conduct bias and explainability analysis after training



Deploy

Explain individual inferences from models in production



Monitor

Validate bias and relative feature importance over time

Built-in MLOps capabilities

Operational challenges with managing the ML lifecycle

**Manual
iterative
processes slow
down ML
innovation**



**Difficult to
scale the
number
of models in
production**



**CI/CD for ML
requires
writing
custom code**



Amazon SageMaker MLOps

Streamline the ML lifecycle




Automate ML workflows to scale model development




Build CI/CD pipelines for ML to accelerate model deployment



Catalog model versions, metadata, metrics, and approvals for traceability and reusability



Track lineage for troubleshooting and compliance



Maintain accuracy of predictions after models are deployed



Enhance governance and security


Amazon SageMaker key benefits

Most complete end-to-end ML service



Democratize ML innovation

Empower more groups of people, including business analysts



Prepare data at scale

Access, label, and process structured and unstructured data



Accelerate the ML lifecycle

Reduce training time from hours to minutes



Streamline ML processes

Automate and standardize MLOps processes

Amazon SageMaker

Getting started

Business analysts

**Get started with
Amazon SageMaker Canvas**

No-code ML



Data scientists

**Onboard through
Amazon SageMaker Studio**


Single IDE for full-code ML



MLOps engineers

**Implement MLOps with
Amazon SageMaker Pipelines**

Workflow automation
and CI/CD



Learn and experiment with machine learning

Amazon SageMaker Studio Lab (Preview)

**Free, no-configuration
development environment**

All you need is an email address to get started

As many CPU and GPU sessions as needed

Automatically save session data

Models ready for production

Demo



Recap & other resources

1. We covered the following topics today:

- Introduction to platform services – Amazon SageMaker
- No-code ML tools
- Purpose-built data preparation tools
- Integrated ML tools in a single interface
- Built-in MLOps capabilities

2. Additional resources:

- [Amazon SageMaker](#)
- [AWS Well-Architected Framework: Machine Learning Lens](#)

Visit the AI & Machine Learning resource hub for more resources

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

- The machine learning journey e-book
- 7 leading machine learning use cases e-book
- A strategic playbook for data, analytics, and machine learning e-book
- Accelerate machine learning innovation with the right cloud services & infrastructure e-book
- Choosing the right compute infrastructure for machine learning e-book
- Improving service and reducing costs in contact centers e-book
- Why ML is essential in your fight against online fraud e-book
- ... and more!



<https://bit.ly/3mwi59V>

Visit resource hub

AWS Machine Learning (ML) Training and Certification



AWS is how you build machine learning skills

Courses built on the curriculum leveraged by Amazon's own teams. Learn from the experts at AWS.

aws.training/machinelearning



Flexibility to learn your way

Learn online with on-demand digital courses or live with virtual instructor-led training, plus hands-on labs and opportunities for practical application.

explore.skillbuilder.aws/learn



Validate your expertise

Demonstrate expertise in building, training, tuning, and deploying machine learning models with an industry-recognized credential.

aws.amazon.com/certification

Thank you for attending AWS Innovate – 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!

Aashmeet Kalra

