

The background features a dark blue gradient with large, overlapping, semi-transparent shapes in shades of purple and magenta. Two thin, light blue lines cross the scene diagonally. The text is positioned on the left side of the image.

AWS re:Invent

DECEMBER 2 - 6, 2024 | LAS VEGAS, NV

ANT206

Next-generation Amazon SageMaker: The center for data, analytics & AI

Chanu Damarla

(he/him)

Principal Product Manager
Amazon Web Services

Christopher Jackson

(he/him)

Head of Digital Data & Analytics
olympics.com

Anchit Gupta

(she/her)

Sr. Product Manager
Amazon Web Services



Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



Times are changing

Technology



Changing faster
than ever

Work



Work and interact
with technology

Organizations



Adopt and adapt

Organizations understand that

Analytics & AI

tailored to their business

fuels

Their Data



ANT302 | Data foundation in the age of generative AI



Organizations are struggling

74% of businesses fail to turn data into insights



82%

of organizations have appointed a Chief Data Officer (CDO)



94%

of organizations increased data investments in 2023



20%

revenue growth seen by data-insights driven businesses

Challenges in an evolving landscape

High effort AI/ML

AI and machine learning faces cost, speed, and efficiency barriers

Evolving personas

Converging personas and jobs to be done across AI/ML and analytics

Data silos

Data silos that are not interoperable with the tools

Fragmented governance

Trust, confidence, and compliance in data management

Evolving tech environment

Evolving data environment

Evolving technology is leading to complex tool-set

Organizations are struggling with data sprawl

Organizations appreciate

Comprehensive set of purpose-built services

Optimized for **performance and cost**

Experience



Act



Govern



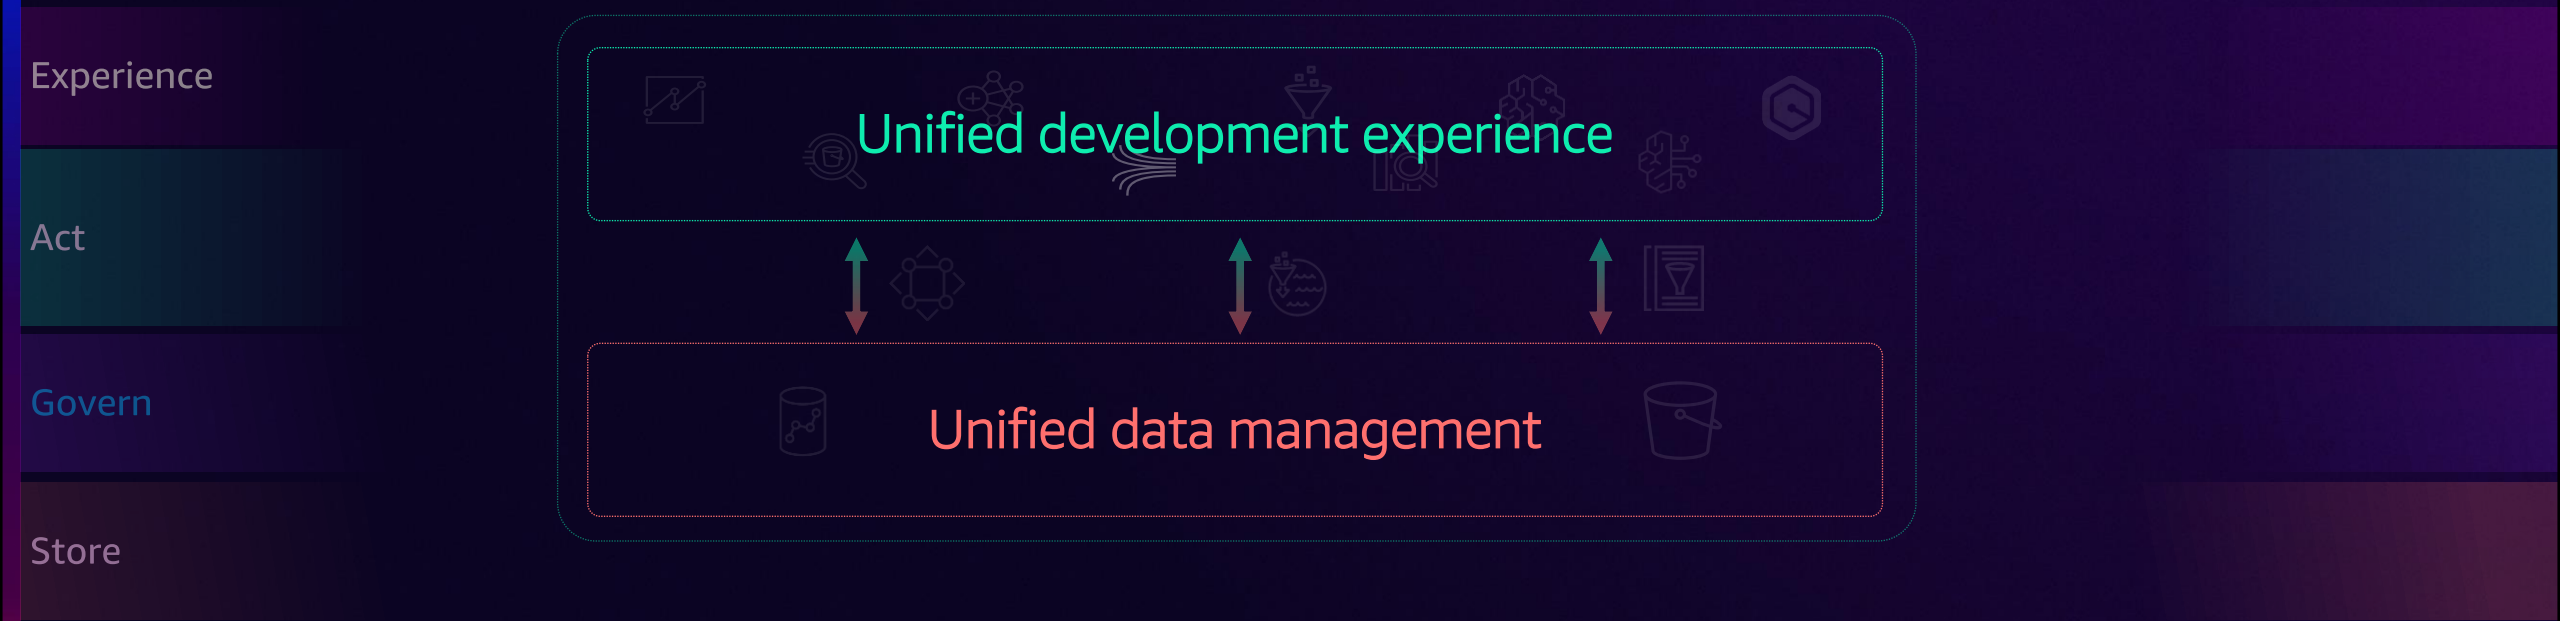
Store



Organizations want

Leverage same **rich set** of services

through **unified experiences** across them



Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com


Demo

Conclusion



The next generation of Amazon SageMaker

THE CENTER FOR ALL YOUR DATA, ANALYTICS, AND AI



SageMaker

AI

Develop and scale AI use cases
with the broadest set of tools

Persona and use case specific tools resulted in

Experience

Fragmented experience

BI Analysts Business owners

Act

Switching between tools

ML Scientists Data Engineers

Govern

Inconsistent access

Data Stewards



Amazon SageMaker Unified Studio

A SINGLE DATA AND AI DEVELOPMENT ENVIRONMENT

Preview



Amazon SageMaker Unified Studio (Preview)

Data
Processing

SQL
Analytics

Model
Development

Gen AI
Development

Business
Intelligence

Search

Streaming

coming soon

coming soon

coming soon



Balancing democratization and governance



Amazon SageMaker Catalog

DISCOVER, GOVERN, AND COLLABORATE ON DATA AND AI SECURELY

GA



Amazon SageMaker Unified Studio (Preview)

Data
Processing

SQL
Analytics

Model
Development

Gen AI
Development

Business
Intelligence

Search

Streaming

coming soon

coming soon

coming soon

Data & AI Governance



Workload-specific storage created data silos

Store

DATA
WAREHOUSE

DATA LAKE



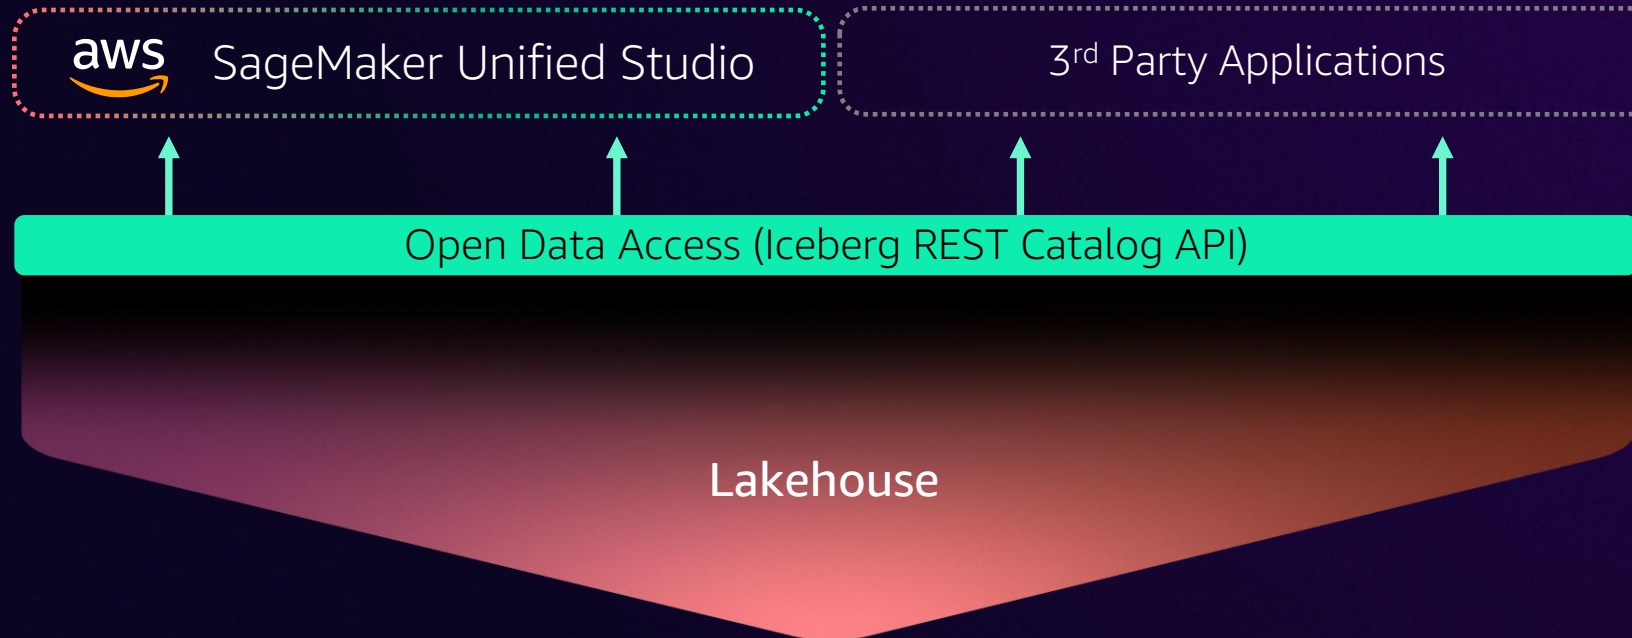
Data Owners



Amazon SageMaker Lakehouse

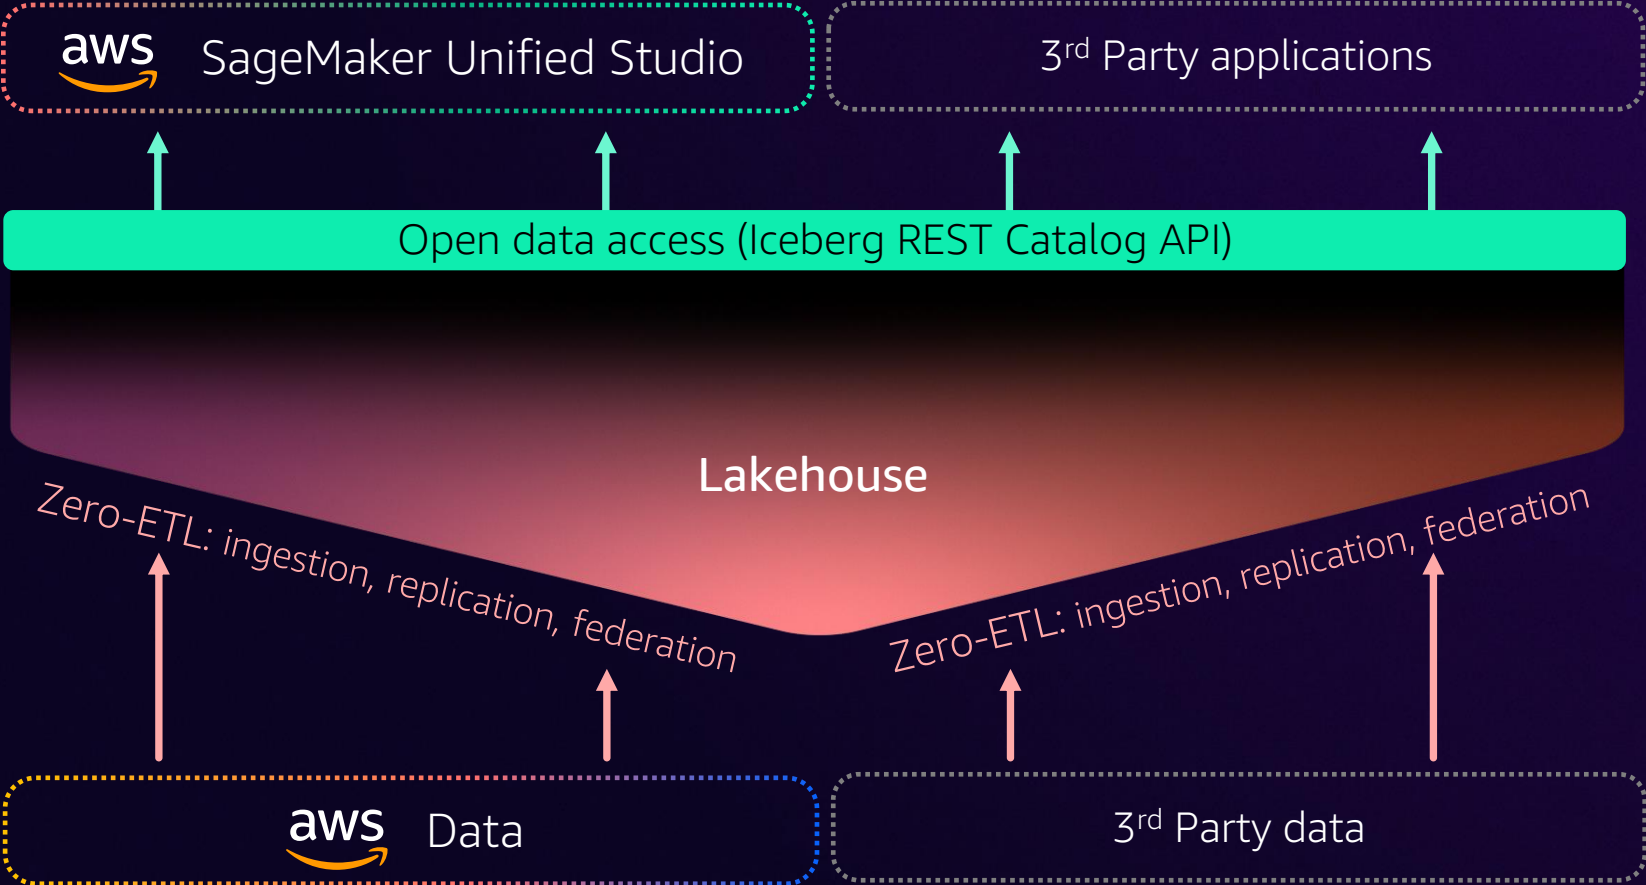
GA

SIMPLIFY ANALYTICS AND AI WITH AN UNIFIED, OPEN, AND SECURE DATA LAKEHOUSE



Zero-ETL gives you access to all your data

SIMPLIFY ANALYTICS AND AI WITH AN UNIFIED, OPEN, AND SECURE DATA LAKEHOUSE



The next generation of Amazon SageMaker

THE CENTER FOR ALL YOUR DATA, ANALYTICS, AND AI



SageMaker

AI

Develop and scale AI use cases
with the broadest set of tools

Unified Studio

Collaborate and build faster
with a single data and AI dev.
environment

Lakehouse

Reduce data silos with an open
lakehouse to unify all your data

Catalog

**Meet your enterprise security
needs** with built-in data and AI
governance



Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



Amazon SageMaker Unified Studio

A SINGLE DATA AND AI DEVELOPMENT ENVIRONMENT TO COLLABORATE AND BUILD FASTER

SageMaker Unified Studio

All of your data and tools for analytics and AI
in a single development environment

Use **best-in-class** tools, no matter the job

Train, customize, and deploy AI models at **scale**

Rapidly build custom generative AI applications

Accelerate your data journey with Amazon Q Developer

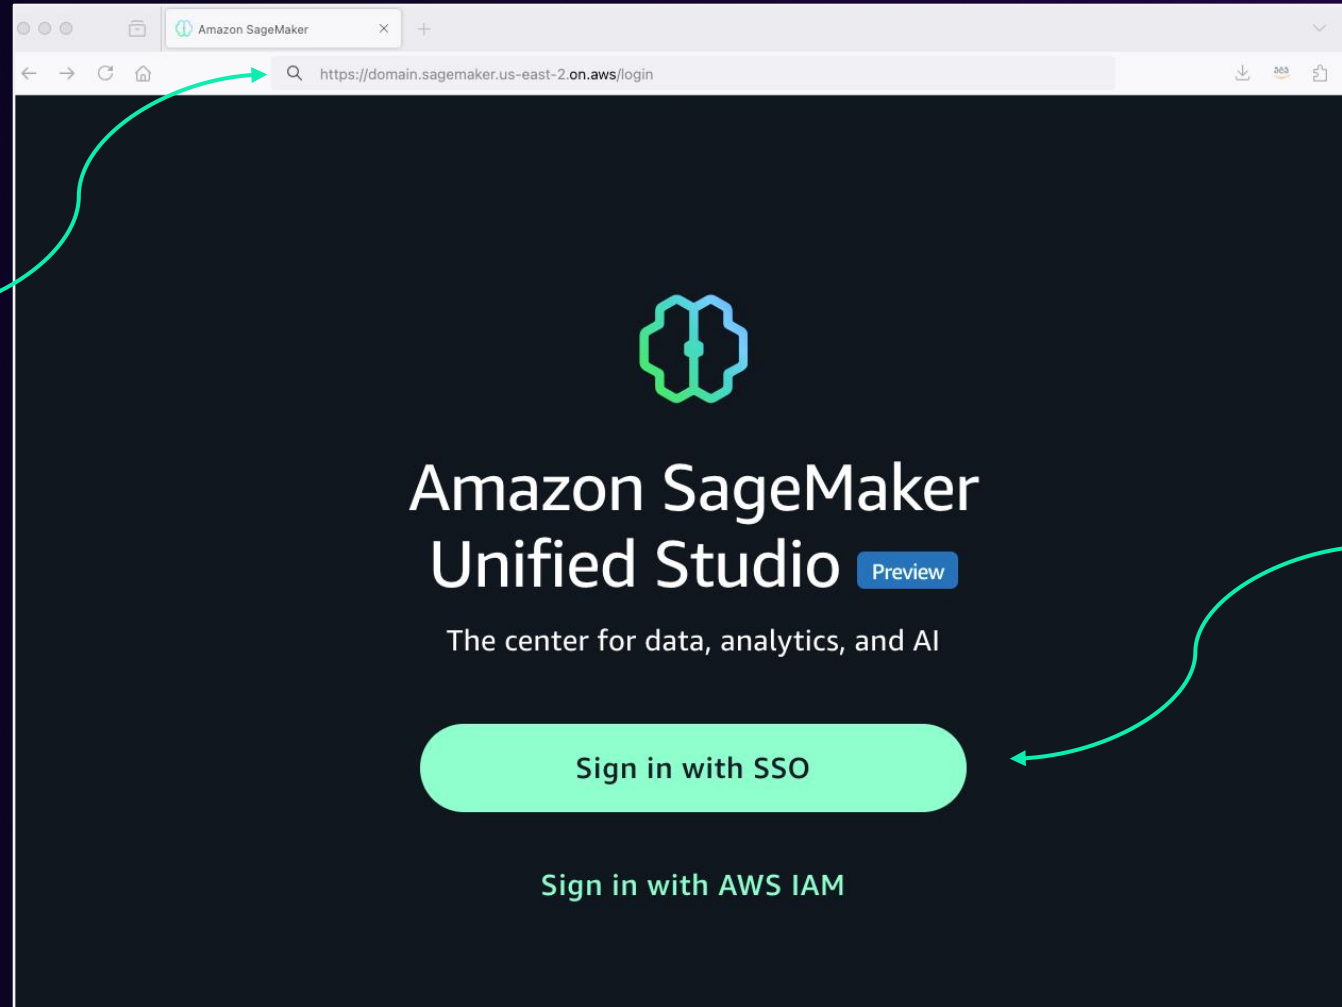


Embedded data and AI governance

An integrated experience for all your data and AI

Preview

ACROSS REGIONS, ACCOUNTS, TOOLS, AND USE CASES



Global URL; No need to remember workload specific URLs for your projects

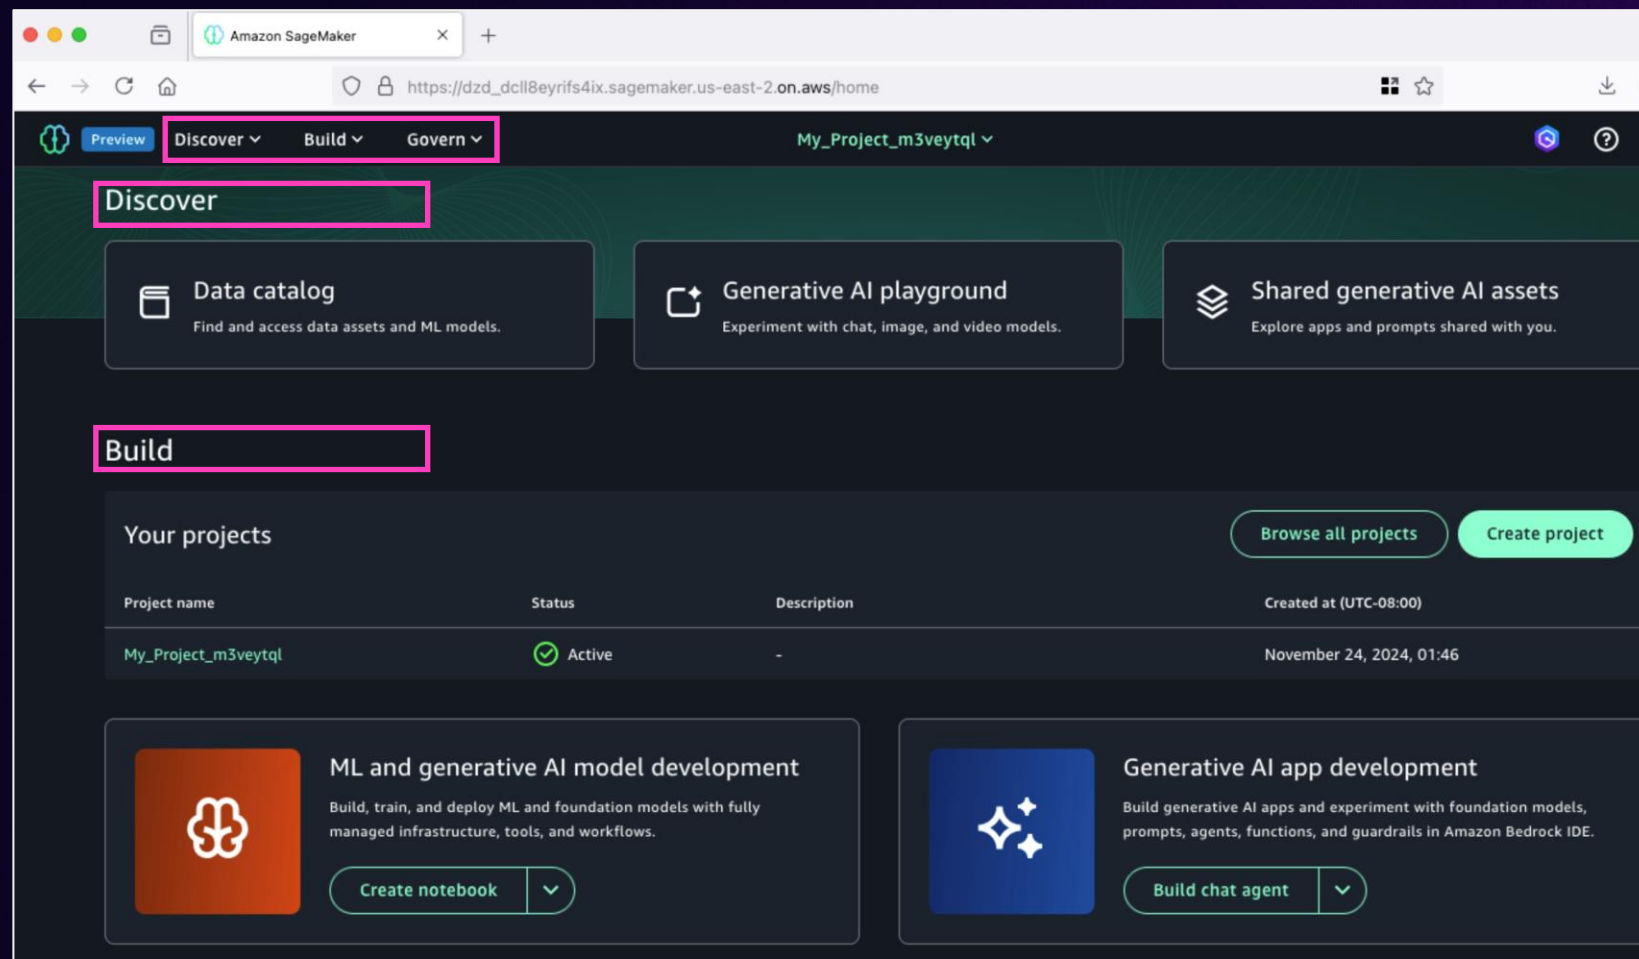
Login with IDC (preferred) support for Direct SAML & IAM



Discover, build, and govern data, analytics, and AI

Preview

DATA, AI/ML ARTIFACTS, TOOLS, CODE, AND USERS IN ONE PLACE

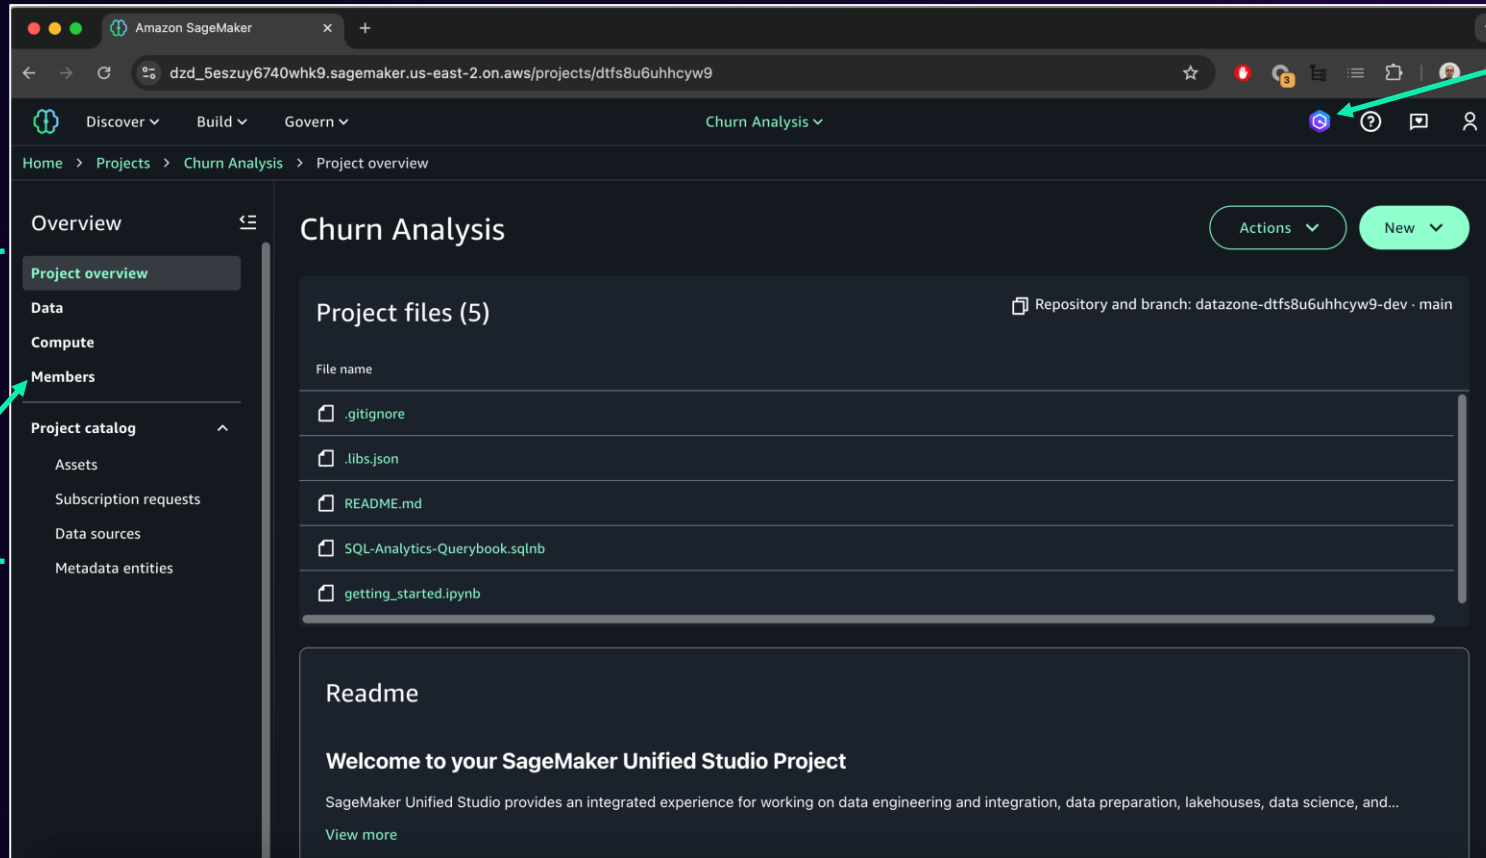


Builder experience is centered around projects

DATA, AI/ML ARTIFACTS, TOOLS, CODE, AND USERS IN ONE PLACE

See the code, data, ML assets, and compute for your projects in one view

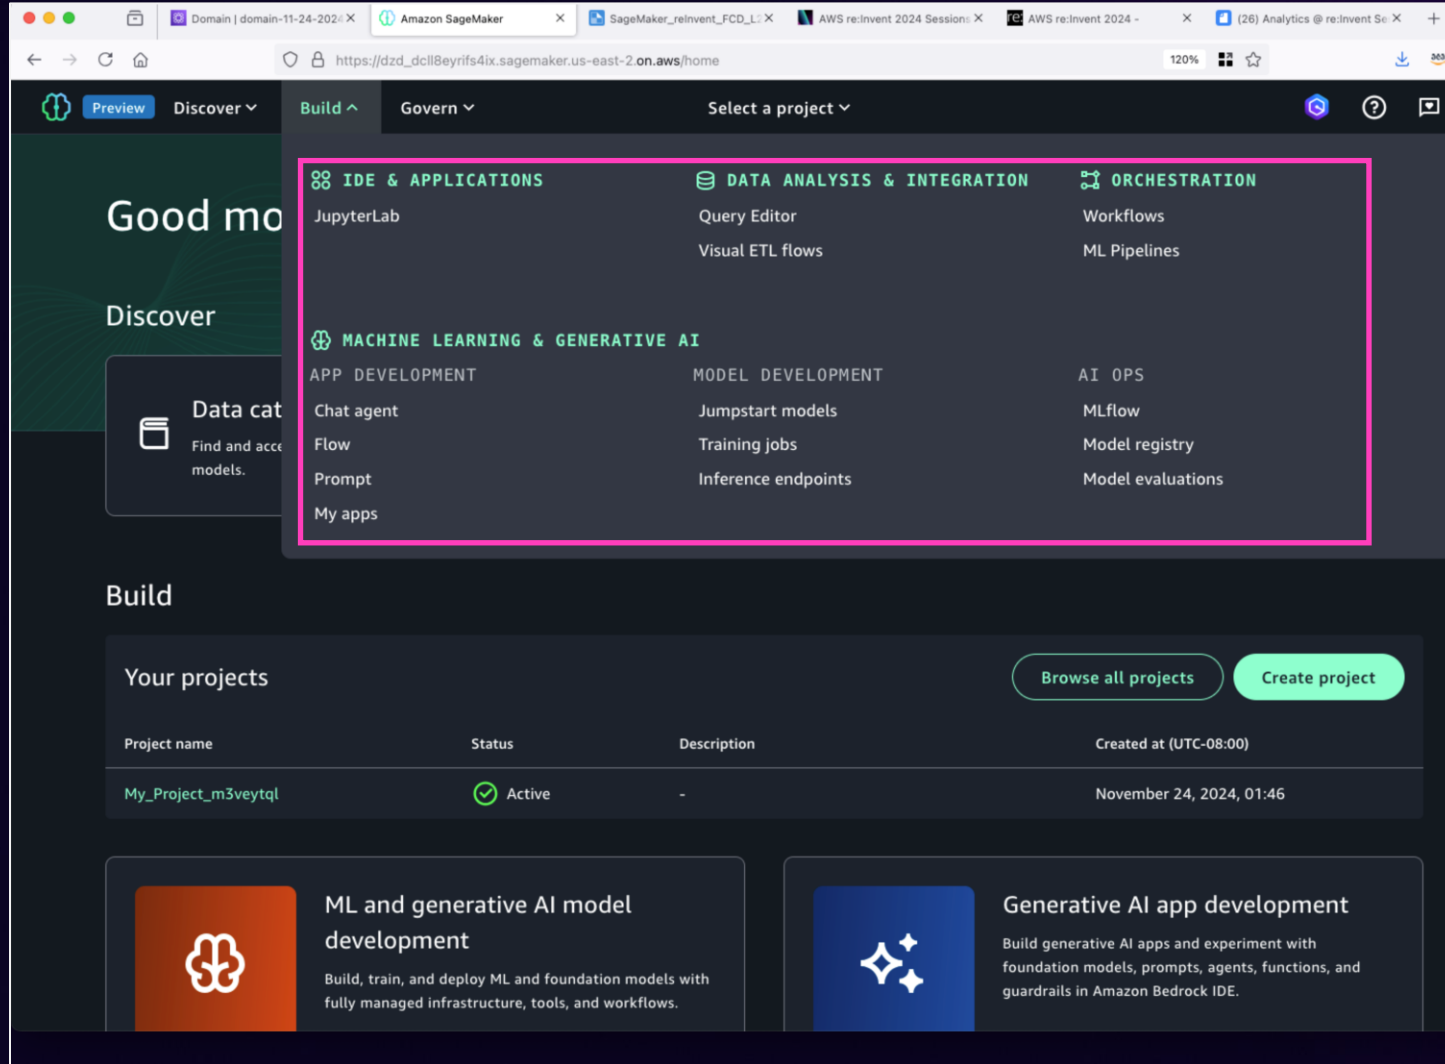
Efficiently onboard variety of users so they can easily collaborate



Gen AI-assisted experience powered by Amazon Q

Work with notebooks, queries, models, and more from familiar AWS tools

Specialized tools to support all your use cases



All the tools you need for **data processing**, SQL analytics, **machine learning**, and **gen AI**

Tools are **integrated with git** - you don't have to set up integrations

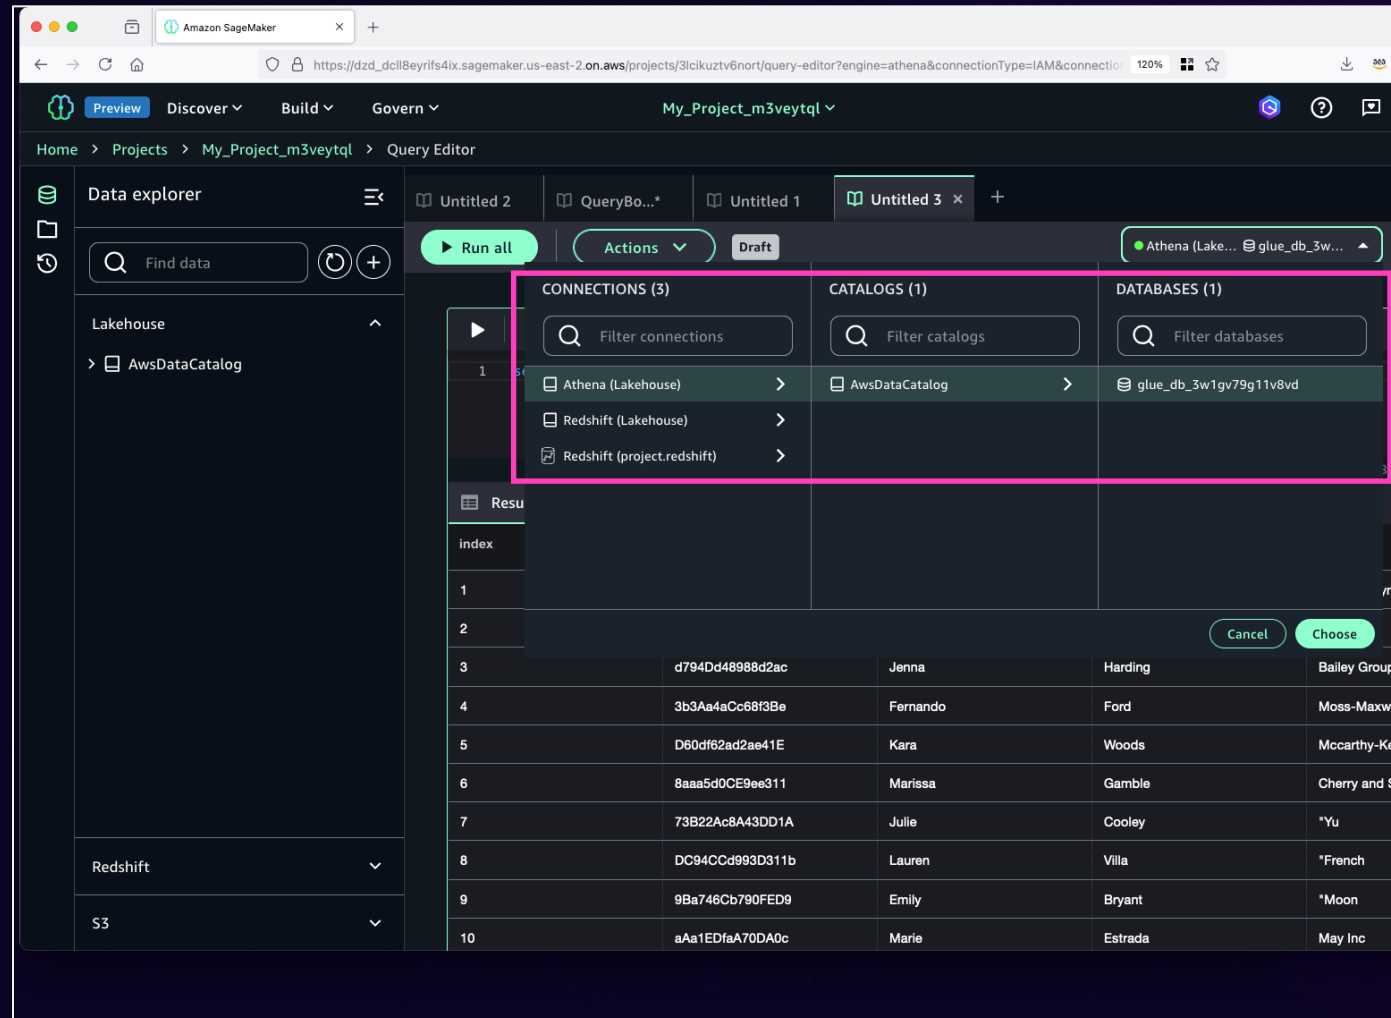
Tools share access to data; **no need to configure tool specific permissions**

Tools can **connect to work** with variety of **AWS** and **3rd party** services

Unified Studio Query Editor

FOR YOUR DATA EXPLORATION AND SQL ENGINEERING WORKLOADS

Preview



Query data lakes, data warehouses, S3 buckets, and federated sources

Data explorer lets you explore all the datasets you have access to

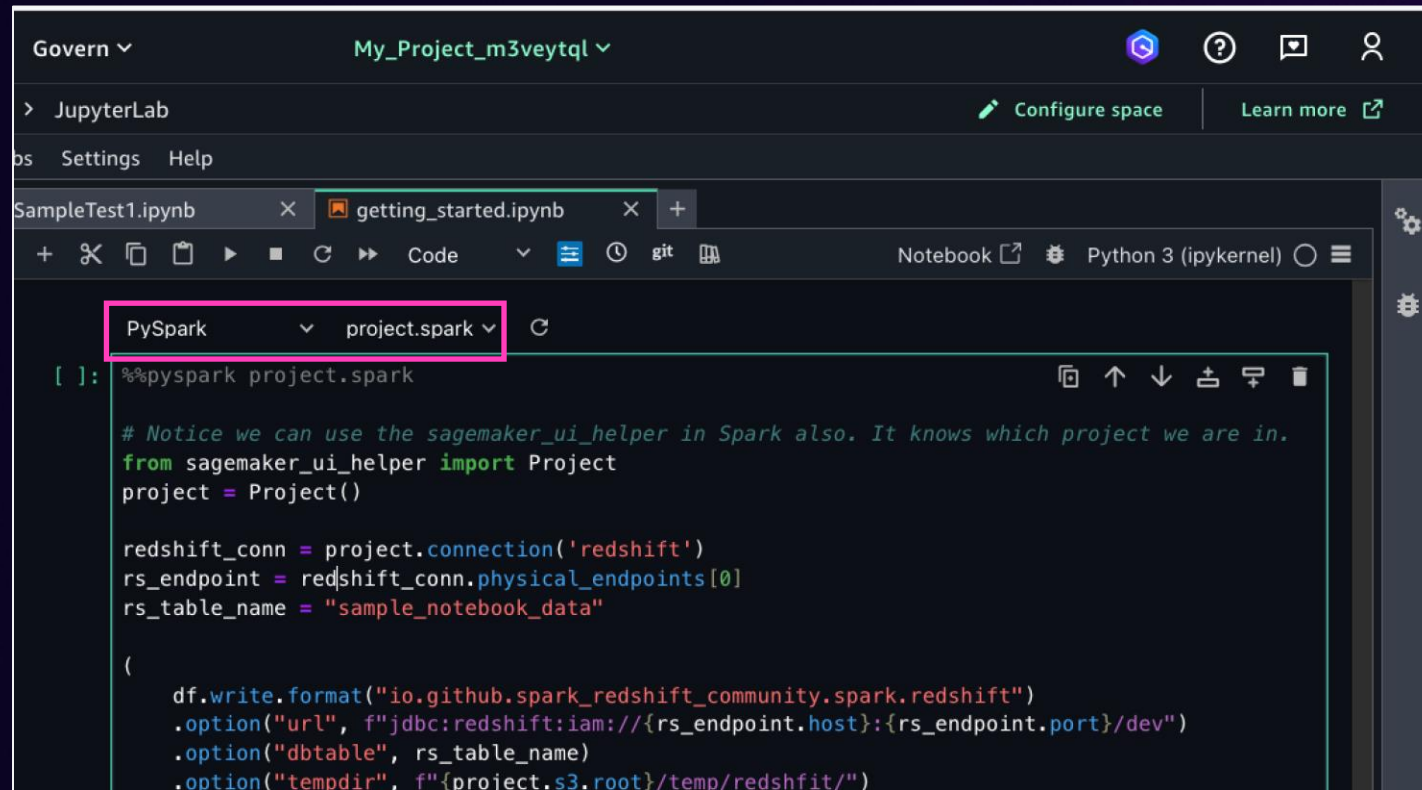
Choice of Amazon Redshift or Amazon Athena engines



Unified Studio notebooks

FOR YOUR DATA ENGINEERING AND SPARK WORKLOADS

Preview



The screenshot shows the Unified Studio notebook interface. At the top, there's a navigation bar with 'Govern' and 'My_Project_m3veytql'. Below that, there's a 'JupyterLab' section with 'Configure space' and 'Learn more' links. The main area shows a notebook titled 'getting_started.ipynb' with a 'PySpark' kernel selected. The code cell contains the following Python code:

```
[ ]: %%pyspark project.spark

# Notice we can use the sagemaker_ui_helper in Spark also. It knows which project we are in.
from sagemaker_ui_helper import Project
project = Project()

redshift_conn = project.connection('redshift')
rs_endpoint = redshift_conn.physical_endpoints[0]
rs_table_name = "sample_notebook_data"

(
    df.write.format("io.github.spark_redshift_community.spark.redshift")
    .option("url", f"jdbc:redshift:iam://{rs_endpoint.host}:{rs_endpoint.port}/dev")
    .option("dbtable", rs_table_name)
    .option("tempdir", f"{project.s3.root}/temp/redshfit/")
)
```

Use **Spark**, **SQL**, **Python**, and **PySpark** to process data

Choice of Amazon **EMR** (serverless and EC2), **AWS Glue**, and **local engines**

Configure instance and memory sizes based on your needs

Advanced **debugging capabilities** like **SparkUI** included



Unified Studio machine learning

Preview

FOR YOUR MODEL DEVELOPMENT AND MACHINE LEARNING WORKLOADS

Amazon SageMaker

Discover ▾ Build ▲ Govern ▾ Churn Analysis ▾

- IDE & APPLICATIONS
 - JupyterLab
 - Partner AI Apps
- DATA ANALYSIS & INTEGRATION
 - Query Editor
 - Visual ETL flows
- ORCHESTRATION
 - Workflows
 - ML Pipelines
- MACHINE LEARNING & GENERATIVE AI
 - APP DEVELOPMENT
 - Chat agent
 - Flow
 - Prompt
 - My apps
 - MODEL DEVELOPMENT
 - Jumpstart models
 - Training jobs
 - Inference endpoints
 - HyperPod
 - AI OPS
 - MLflow
 - Model registry
 - Model evaluations

Your projects

Project name	Status	Description
Gaming Analysis	Active	-
Churn Analysis	Active	-

Existing SageMaker AI tools in a unified data, analytics, and ML env.

Felxible and powerful: customize 400+ foundation models or deploy models on 70+ instance types

Hyperpod for largescale distributed training and model tuning

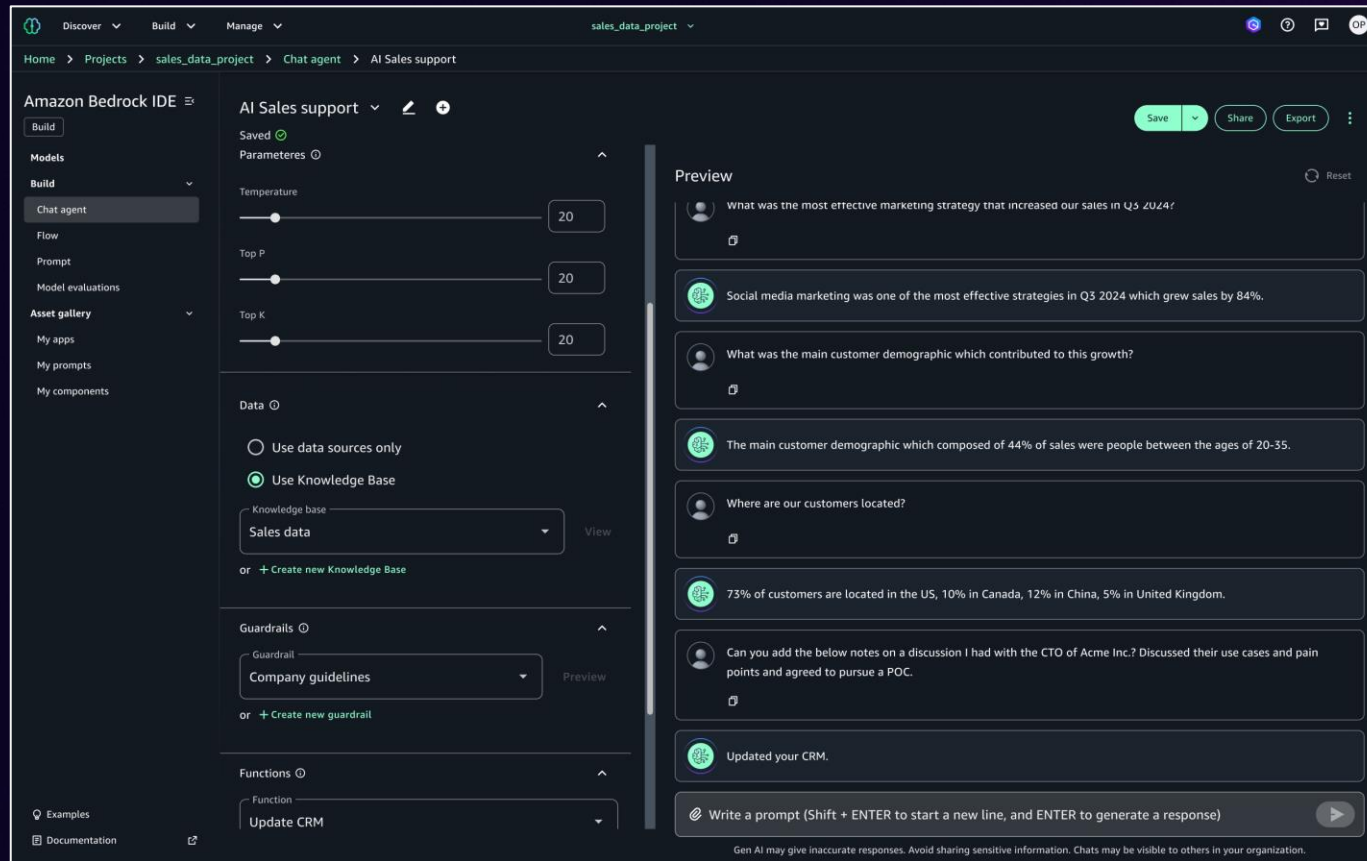
Integrated ML governance tools



Unified Studio includes Amazon Bedrock IDE

Preview

ACCELERATE GENERATIVE AI APPLICATION DEVELOPMENT



Collaboratively build gen AI apps with stakeholders of all skill levels

Quickly compare text, image, or video foundation models

Visually create apps with a drag-and-drop flow designer

Customize with your data for your use cases with knowledge bases

Ensure responsible responses with guardrails and filters



Amazon Q integrations to boost productivity

USE NATURAL LANGUAGE TO DISCOVER, BUILD, AND GOVERN



Q & A

Ask questions about SageMaker, understand best practices, quickly get info about your AWS resources, and more!



Data discovery

Search and discover data from your BDC using natural language



Code authoring

Generate code snippets or entire functions, identify and fix errors, upgrade apps to newer versions



SQL querying

Simplify query authoring and increase productivity. Express queries in natural language and receive SQL code recommendations

AMAZON Q ASSISTANT FOR AWS PRODUCT DOCUMENTATION, HELP, AND RESOURCES



Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



Domains: Model your organization

Preview

The screenshot displays the Amazon SageMaker Domains console for a specific domain. The domain name is 'domain-11-23-2024-224905'. The console shows the following sections:

- Domain details:** Includes the Amazon SageMaker Unified Studio URL (https://dzd_5wxc56oagcl46h.sagemaker.us-east-2.on.aws/), Domain ARN (arn:aws:datazone:us-east-2:785789292865:domain/dzd_5wxc56oagcl46h), Created on (November 23, 2024, 22:49 (UTC-08:00)), and Domain version (Amazon SageMaker platform).
- Next steps for your domain:** A section with a right-pointing arrow.
- User management:** A navigation bar with tabs for Account associations, Blueprints, Project profiles, Amazon Bedrock models, Connections, Amazon Q, and Tags.
- SSO user access:** Shows the Login method (IAM Identity Center) and User and group assignment method (Require assignments).
- Users (2):** A table listing users with columns for Name and Status. The users listed are 'avijain' (Assigned) and 'cd2k' (Activated).

Unified Studio with a unique URL per domain

Provide consistent management of studio, accounts, users, network...

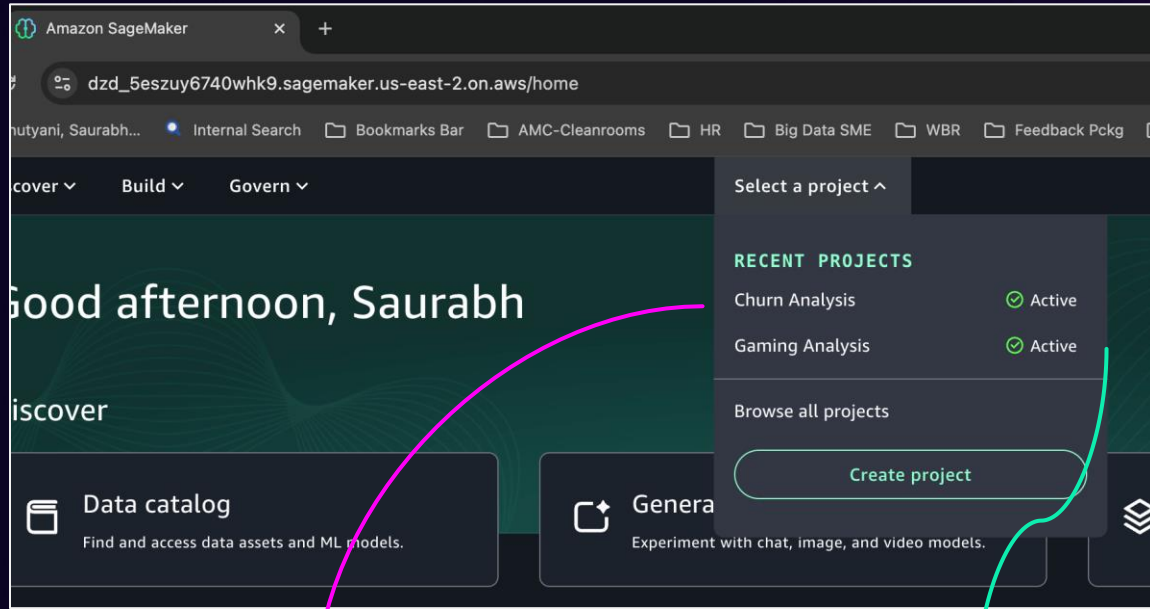
Start simply with single account and associate accounts to expand

Flexible to support your account and data sharing architectures



Projects: Simplify development and collaboration

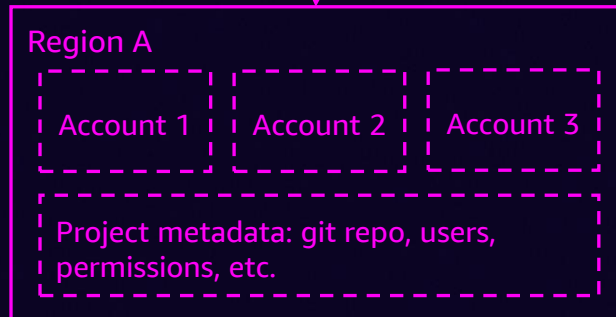
ACROSS REGIONS, ACCOUNTS, TOOLS, AND USE CASES



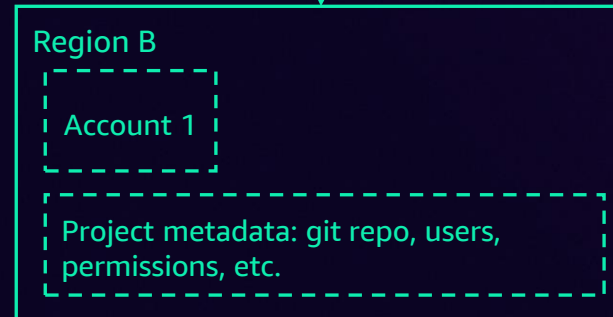
Access projects across regions and accounts from the Unified Studio

Mental model: you are federating into different accounts in different regions and assuming a role

Churn Analysis



Gaming Analysis

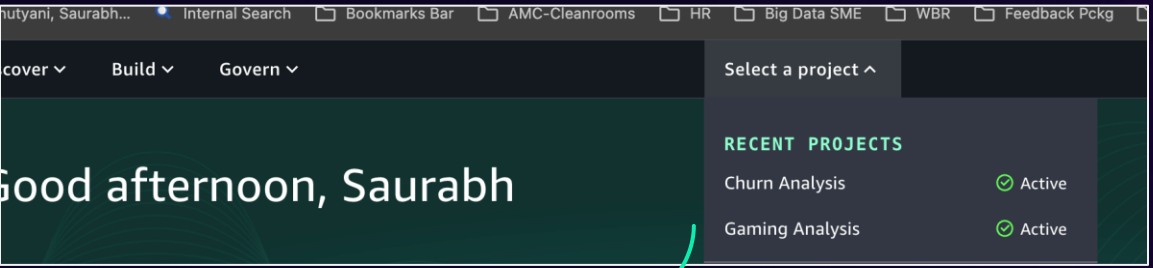


Project metadata includes regions, accounts, users, and git repo



Project Roles: Provide access to resources

ACROSS REGIONS, ACCOUNTS, TOOLS, AND USE CASES



Gaming Analysis



Projects contain one or more project roles

There is at least one project role per account the project is in


This role has access to tools, compute, data, and ML assets

To deploy environments use blueprints

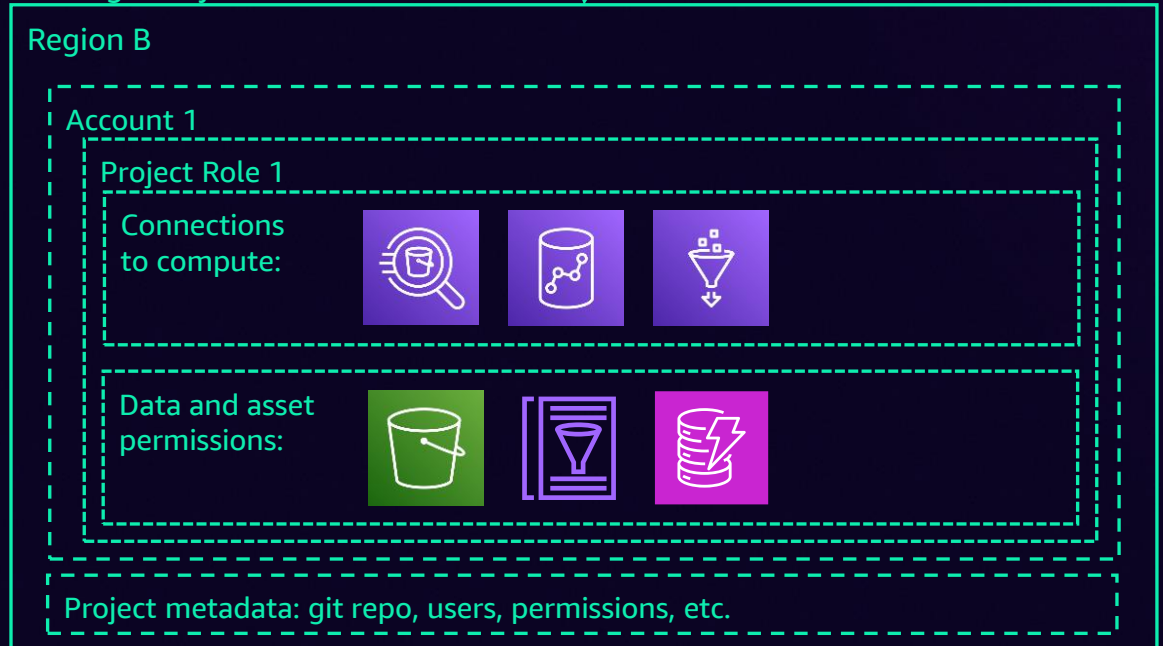


Project Roles: Provide access to resources

ACROSS REGIONS, ACCOUNTS, TOOLS, AND USE CASES



Gaming Analysis



Projects contain one or more project roles

There is at least one project role per account the project is in

This role has access to tools, compute, data, and ML assets

To deploy environments use blueprints



Blueprints: Enable responsible self-service

Gaming Analysis



Add compute ← ×

This resource will be shared between all project members.

1 Select an option to add compute — 2 Choose type — 3 Enter details

Select an option to add compute

Create new compute resources Connect to existing compute resources

Add compute ← ×

This resource will be shared between all project members.

✓ Select an option to add compute — 2 Choose type — 3 Enter details

All compute types (2)

Search

EMR on EC2 cluster EMR Serverless

Bring **existing** resources: Redshift, EMR, SageMaker AI, and Amazon Bedrock

Create ad-hoc resources as needed within projects

Redshift & EMR endpoints, Amazon Bedrock components, federated connectors

With appropriate governance guardrails and lifecycle management



Blueprints: Enable responsible self-service

Gaming Analysis



Add compute

This resource will be shared between all project members.

1 Select an option to add compute — 2 Choose type — 3 Enter details

Select an option to add compute

Create new compute resources

Connect to existing compute resources

Add compute

This resource will be shared between all project members.

1 Select an option to add compute — 2 Choose type — 3 Enter details

All compute types (2)

Search

EMR on EC2 cluster

EMR Serverless

Bring **existing** resources: Redshift, EMR, SageMaker AI, and Amazon Bedrock

Create ad-hoc resources as needed within projects

Redshift & EMR endpoints, Amazon Bedrock components, federated connectors

With appropriate governance guardrails and lifecycle management



Catalog: Asset discovery and permissions mgmt.

Preview

Gaming Analysis



Add IAM, LF, or resource policies to provide the env's role access to assets

Managing these permissions can be cumbersome and error prone

Publish to catalog and allow data workers to subscribe to assets

SageMaker Catalog orchestrates and automates permissions mgmt.

Find data

Lakehouse

- AwsDataCatalog
 - glue_db_3w1gv79g11v8vd
 - customers-1000

customers-1000

Actions

- Query with Athena
- Query with Redshift
- Open in Jupyter Lab notebook
- Publish to catalog

Table details

Description	La
Table created	N
from.uploadin	2

1 search result for "customers" and selected filters: Data type (1)

customers

SORT BY: RELEVANCE

Show per page: 10

ASSET

Customers 1000

Technical name: customers-1000 Asset type: Glue Table Metadata last updated: Nov 25, 2024, 01:58:23 AM

No description

SUBSCRIBE



Catalog: Asset discovery and permissions mgmt.

Preview

Gaming Analysis

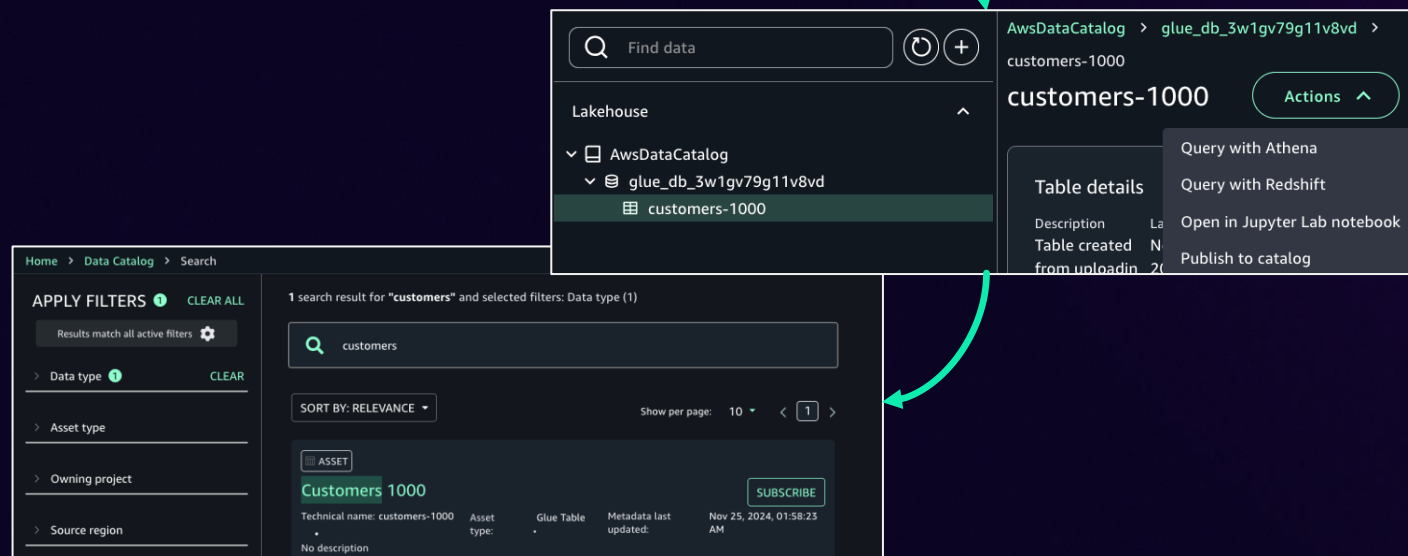


Add IAM, LF, or resource policies to provide the env's role access to assets

Managing these permissions can be cumbersome and error prone

Publish to catalog and allow data workers to subscribe to assets

SageMaker Catalog orchestrates and automates permissions mgmt.



Project profiles: create projects

Preview

Govern ▾ My_Project_m3veytql ▾

Create project

SageMaker Unified Studio projects are containers for code artifacts, data assets, workflows, and other AWS resources. Project members can collaborate in projects to build and deploy applications.

1 Step 1 Choose name and project profile 2 Step 2 Customize parameters 3 Step 3 Review selections

Project name and description

Project name

My_Project_m41wfbfh

Project name must be 1 - 64 characters. Valid characters are a-z, A-Z, 0-9, and - (dash).

Description - *optional*

Share the business context or purpose of the project with other project members.

Characters: 0/2048

Project profile

Project profiles simplify project creation with predefined infrastructure.

Find project profile

- Data analytics and AI-ML model development**
Analyze data and build ML and Gen AI models powered by Amazon EMR, AWS Glue, Amazon Athena, Amazon SageMaker and Lakehouse
Tooling Data Lake Redshift Serverless
+ 5 more
- Generative AI application development**
Build generative AI applications powered by Amazon Bedrock
Tooling Bedrock Chat Agent
Bedrock Knowledge Base + 5 more
- SQL analytics**
Analyze your data in SageMaker Lakehouse using SQL
Tooling Data Lake Redshift Serverless
+ 1 more

Cancel Continue

Quickly create projects for specific use cases with project profiles

Project profiles are templates that:
set up environments
provision and connect to compute
enable Blueprints for self-server resources

Use one of the 3 included profiles
or create your own





We want to streamline the process of data assessment, so our **data analysts, ML scientists, and data engineers can work efficiently** . . . we are excited about the launch of Amazon SageMaker Unified Studio and its ability to simplify data access and enhance collaboration.

Oscar Chang

Chief Development Officer, Trend Micro

Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



Intro — Chris Jackson

Head of Digital Data and Analytics for the Olympics

- Across olympics.com, @olympics and Olympics apps
- Leads teams responsible for Analytics, AI, and Data Management

Before the Olympics

- Strategy: Manager at McKinsey, Head of Strategy at the BBC
- Tech: Founded a data startup and, engineer on early speech to text AI
- Sport: Produced sports radio shows
- *And inadvertently the first person to tweet the term "Big Data"*

Paris Olympics in numbers

THE BIGGEST DIGITAL SPORTING EVENT EVER



24,000+ media personnel



1,200+ cameras



32 sports



39 venues



205 territories represented



91 territories won a medal



300M+ users on olympics.com



12B+ engagements on @olympics

Behind the scenes

ORGANIZATIONAL COMPLEXITY



Host cities



Official partners



Sports federations



Media rights holders



National teams



Other media

Behind the scenes

ORGANIZATIONAL COMPLEXITY



Host cities



Official partners



Sports federations



Media rights holders



National teams



Other media

Data complexity

TWO WEEK PEAK EVERY
FOUR YEARS

x3 ↑
Users

X1.7 ↑
Frequency

x2 ↑
Actions

Writes
x150

x40 ↑
Users

x3.75 ↑
Actions

Read
capacity
x10

Our approach

SERVERLESS FROM BIRTH

Sources



40+ Additional Data Sources



INGEST



API Gateway



QuickSight



Databrew



Step Functions



Storage Gateway



PROCESSING



Kinesis



Lambda



Cloud Formation



Glue



Flink



Storage



Redshift



S3



Athena



Lake Formation



OpenSearch



Use

AI & ML



SageMakers



AI Tools



Power BI

Personalization

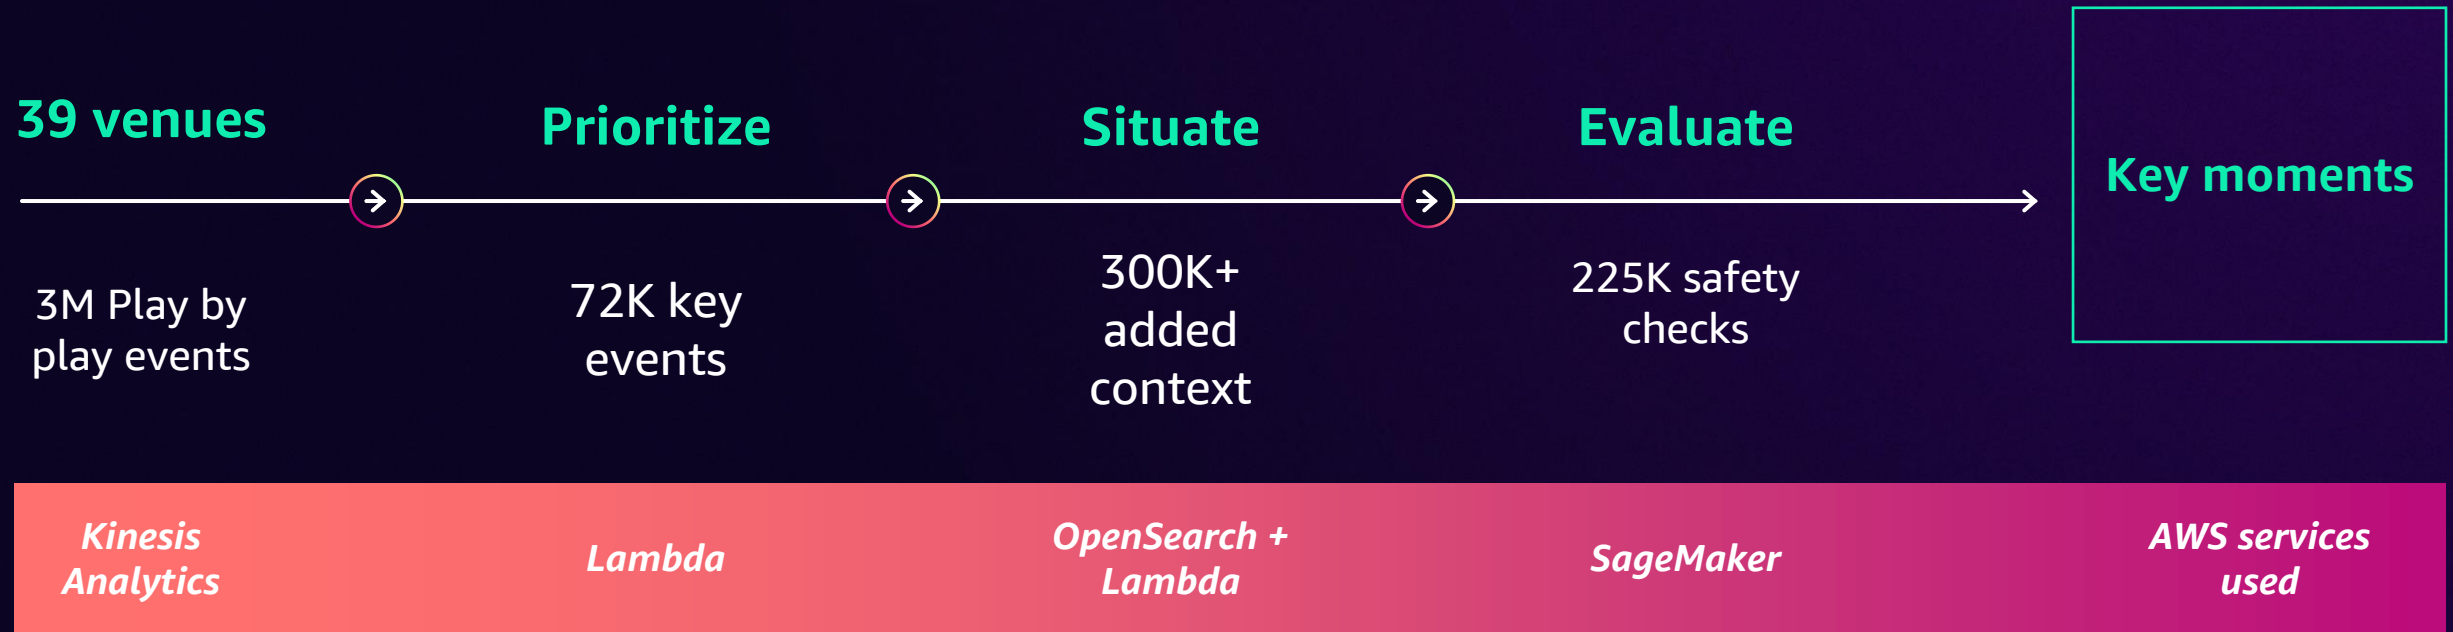


API Gateway



AI at Olympic scale

INTRODUCING OUR MOST DEMANDING USER




Towards LA 2028 — Our needs



Fast and effective onboarding
of data workers



Adding resource fast to meet
unexpected demands



Across a complex web of
organizations



... while complying with media
rights and regulations

Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

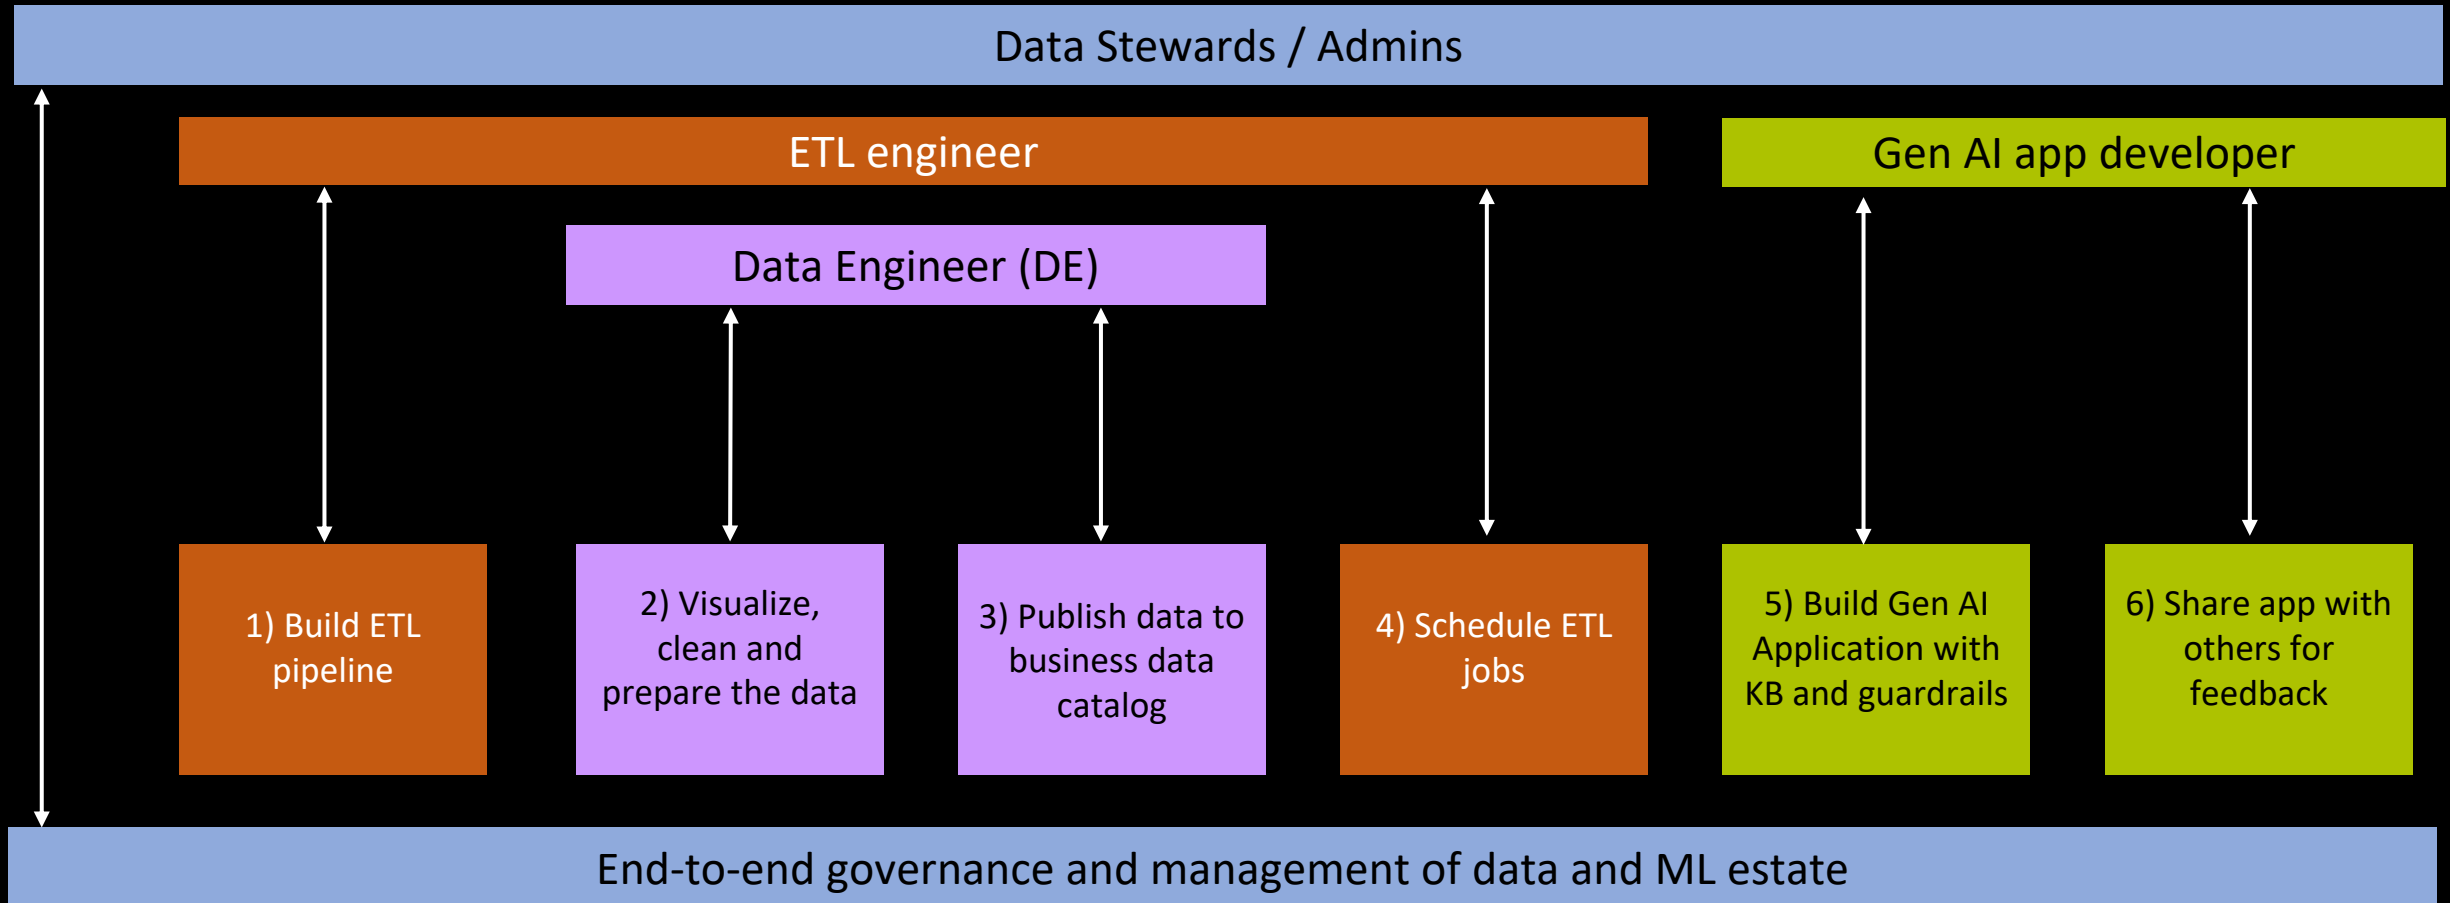
Solving for the unique challenges of olympics.com

Demo

Conclusion





SageMaker Unified Studio – SportsVision Demo




Good afternoon, Rachel

Discover


 **Data catalog**
Find and access data assets and ML models.

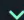
 **Generative AI playground**
Experiment with chat, image, and video models.


 **Shared generative AI assets**
Explore apps and prompts shared with you.

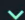
Build


[Browse all projects](#) [Create project](#)

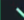
 **ML and generative AI model development**
Build, train, and deploy ML and foundation models with fully managed infrastructure, tools, and workflows.


[Create notebook](#) 

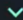
 **Generative AI app development**
Build generative AI apps and experiment with foundation models, prompts, agents, functions, and guardrails in Amazon Bedrock IDE.

[Build chat agent](#) 


 **Data processing and SQL analytics**
Analyze, prepare, and integrate data for analytics and AI using Amazon Athena, Amazon EMR, AWS Glue, and Amazon Redshift.


[Create SQL query](#) 

 **Data and AI governance**
Publish your data products to the catalog with glossaries and metadata forms. Govern access securely in the Amazon SageMaker Catalog built on Amazon DataZone.

[Publish data](#) 


Learn about Amazon SageMaker Unified Studio


[How-to videos](#) 


[Help and documentation](#) 

Good evening, Nina

Discover


 **Data catalog**
Find and access data assets and ML models.


 **Generative AI playground**
Experiment with chat, image, and video models.


 **Shared generative AI assets**
Explore apps and prompts shared with you.


Build

[Browse all projects](#) [Create project](#)

 **ML and generative AI model development**
Build, train, and deploy ML and foundation models with fully managed infrastructure, tools, and workflows.


[Create notebook](#) 


 **Generative AI app development**
Build generative AI apps and experiment with foundation models, prompts, agents, functions, and guardrails in Amazon Bedrock IDE.

[Build chat agent](#) 


 **Data processing and SQL analytics**
Analyze, prepare, and integrate data for analytics and AI using Amazon Athena, Amazon EMR, AWS Glue, and Amazon Redshift.


[Create SQL query](#) 

 **Data and AI governance**
Publish your data products to the catalog with glossaries and metadata forms. Govern access securely in the Amazon SageMaker Catalog built on Amazon DataZone.

[Publish data](#) 

Learn about Amazon SageMaker Unified Studio

[How-to videos](#) 

[Help and documentation](#) 

Good evening, George

Discover

Data catalog Find and access data assets and ML models.

Generative AI playground Experiment with chat, image, and video models.

Shared generative AI assets Explore apps and prompts shared with you.

Build

Your projects

Browse all projects Create project

Table with 4 columns: Project name, Status, Description, Created at (UTC-08:00). Row 1: SportsVision_chat_app, Active, -, November 30, 2024, 17:43

ML and generative AI model development Build, train, and deploy ML and foundation models with fully managed infrastructure, tools, and workflows. Create notebook

Generative AI app development Build generative AI apps and experiment with foundation models, prompts, agents, functions, and guardrails in Amazon Bedrock IDE. Build chat agent

Data processing and SQL analytics Analyze, prepare, and integrate data for analytics and AI using Amazon Athena, Amazon EMR, AWS Glue, and Amazon Redshift. Create SQL query

Data and AI governance Publish your data products to the catalog with glossaries and metadata forms. Govern access securely in the Amazon SageMaker Catalog built on Amazon DataZone. Publish data

Learn about Amazon SageMaker Unified Studio

How-to videos

Help and documentation

Agenda

Why the next generation of Amazon SageMaker?

The next generation of Amazon SageMaker Platform empowers organizations

Amazon SageMaker Unified Studio empowers builders

Amazon SageMaker Unified Studio key concepts

Solving for the unique challenges of olympics.com

Demo

Conclusion



The next generation of Amazon SageMaker

THE CENTER FOR ALL YOUR DATA, ANALYTICS, AND AI



SageMaker

AI

Develop and scale AI use cases with the broadest set of tools

Unified Studio

Collaborate and build faster with a single data and AI dev. environment

Lakehouse

Reduce data silos with an open lake house to unify all your data

Catalog

Meet your enterprise security needs with built-in data and AI governance



Thank you!

Chanu Damarla

Chris Jackson

Anchit Gupta



Please complete the session survey in the mobile app