EXECUTIVE TRACK

Generative AI/ML and AI governance for the public sector

Sergio Ortega

AI/ML BD Lead for State and Local Governments and GovTechs

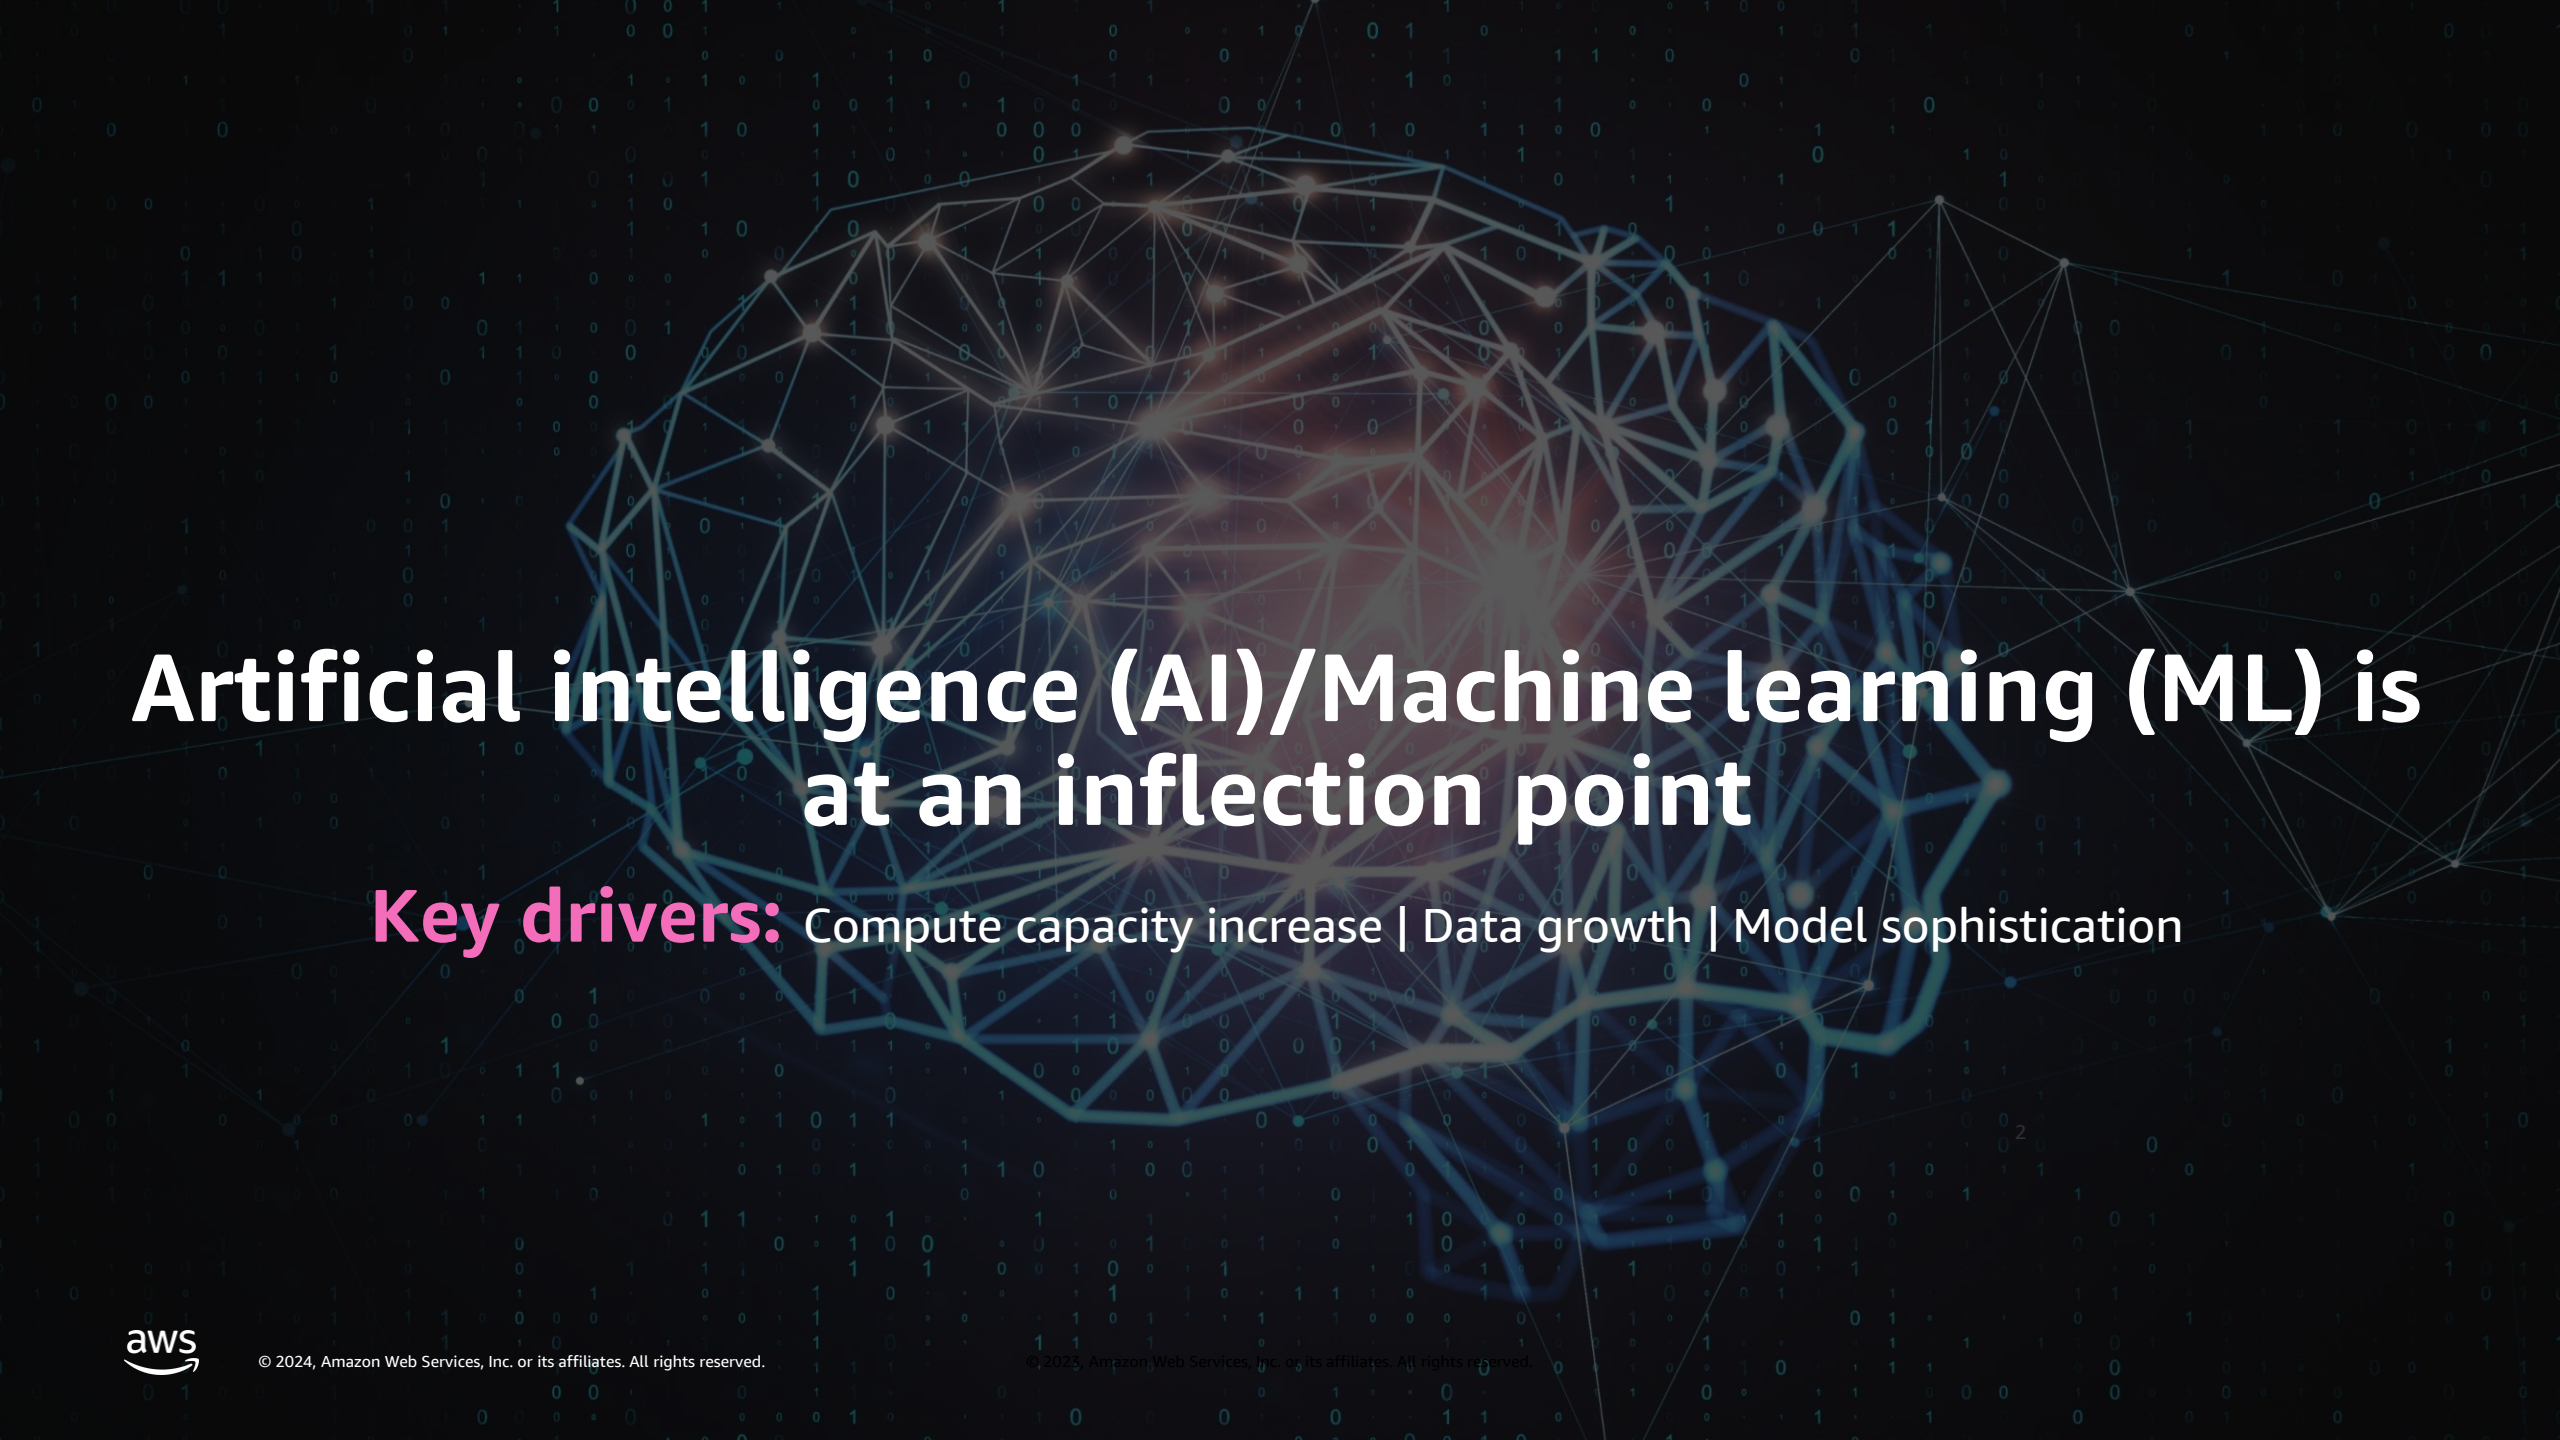

Artificial intelligence (AI)/Machine learning (ML) is at an inflection point

Key drivers: Compute capacity increase | Data growth | Model sophistication

AI, ML, Deep learning?

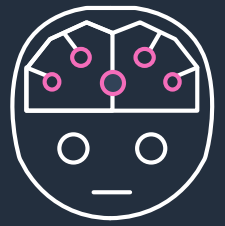

Artificial intelligence (AI)

Any technique that allows computers to mimic human intelligence using logic, if-then statements, and machine learning

Machine learning (ML)

A subset of AI that uses machines to search for patterns in data to build logic models automatically

Deep learning (DL)

A subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition

Generative AI

Powered by large models that are pretrained on vast corpuses of data and commonly referred to as foundation models (FMs)

Challenges we are hearing from state and local government customers

Demand for government services is rising while resources and capacity to deliver them **aren't keeping pace**

Citizens increasingly expect government to **provide modern digital experiences** for conducting online transactions

Aging infrastructure for data capture, storage, and management **creates friction** for leveraging data for analytics and machine learning

Complex security, privacy, and compliance requirements create barriers to change and block adoption of many SaaS solutions

Risk averse culture and institutional inertia slow innovation

Machine learning is going mainstream in public sector

Emergency management
Emergency management, emergency response

Finance and administration
Tax, revenue, regulatory/compliance, fraud, budget, purchasing/procurement

Health and benefits
Local hospitals and clinics, public health, child welfare, homelessness, seniors/youth, health/food, environmental

Assessments and permitting
Planning and zoning, assessment, land admin/recorder of deeds, mapping, watershed, housing

Economic development
Grants and benefits administration, workforce development, corporate relations, land clearance/redevelopment

Constituent engagement
Contact center, website, mobile

Enforcement
Police, jails, courts, animal control

Elections
Registration, voter management, polling, voting management, candidate management

Utilities
Electric/gas/water/sewer, waste management, generation/distribution

Education and library
K–12 and early childhood, adult education, library

Transportation
Traffic, mass transit, parking, airports, ports

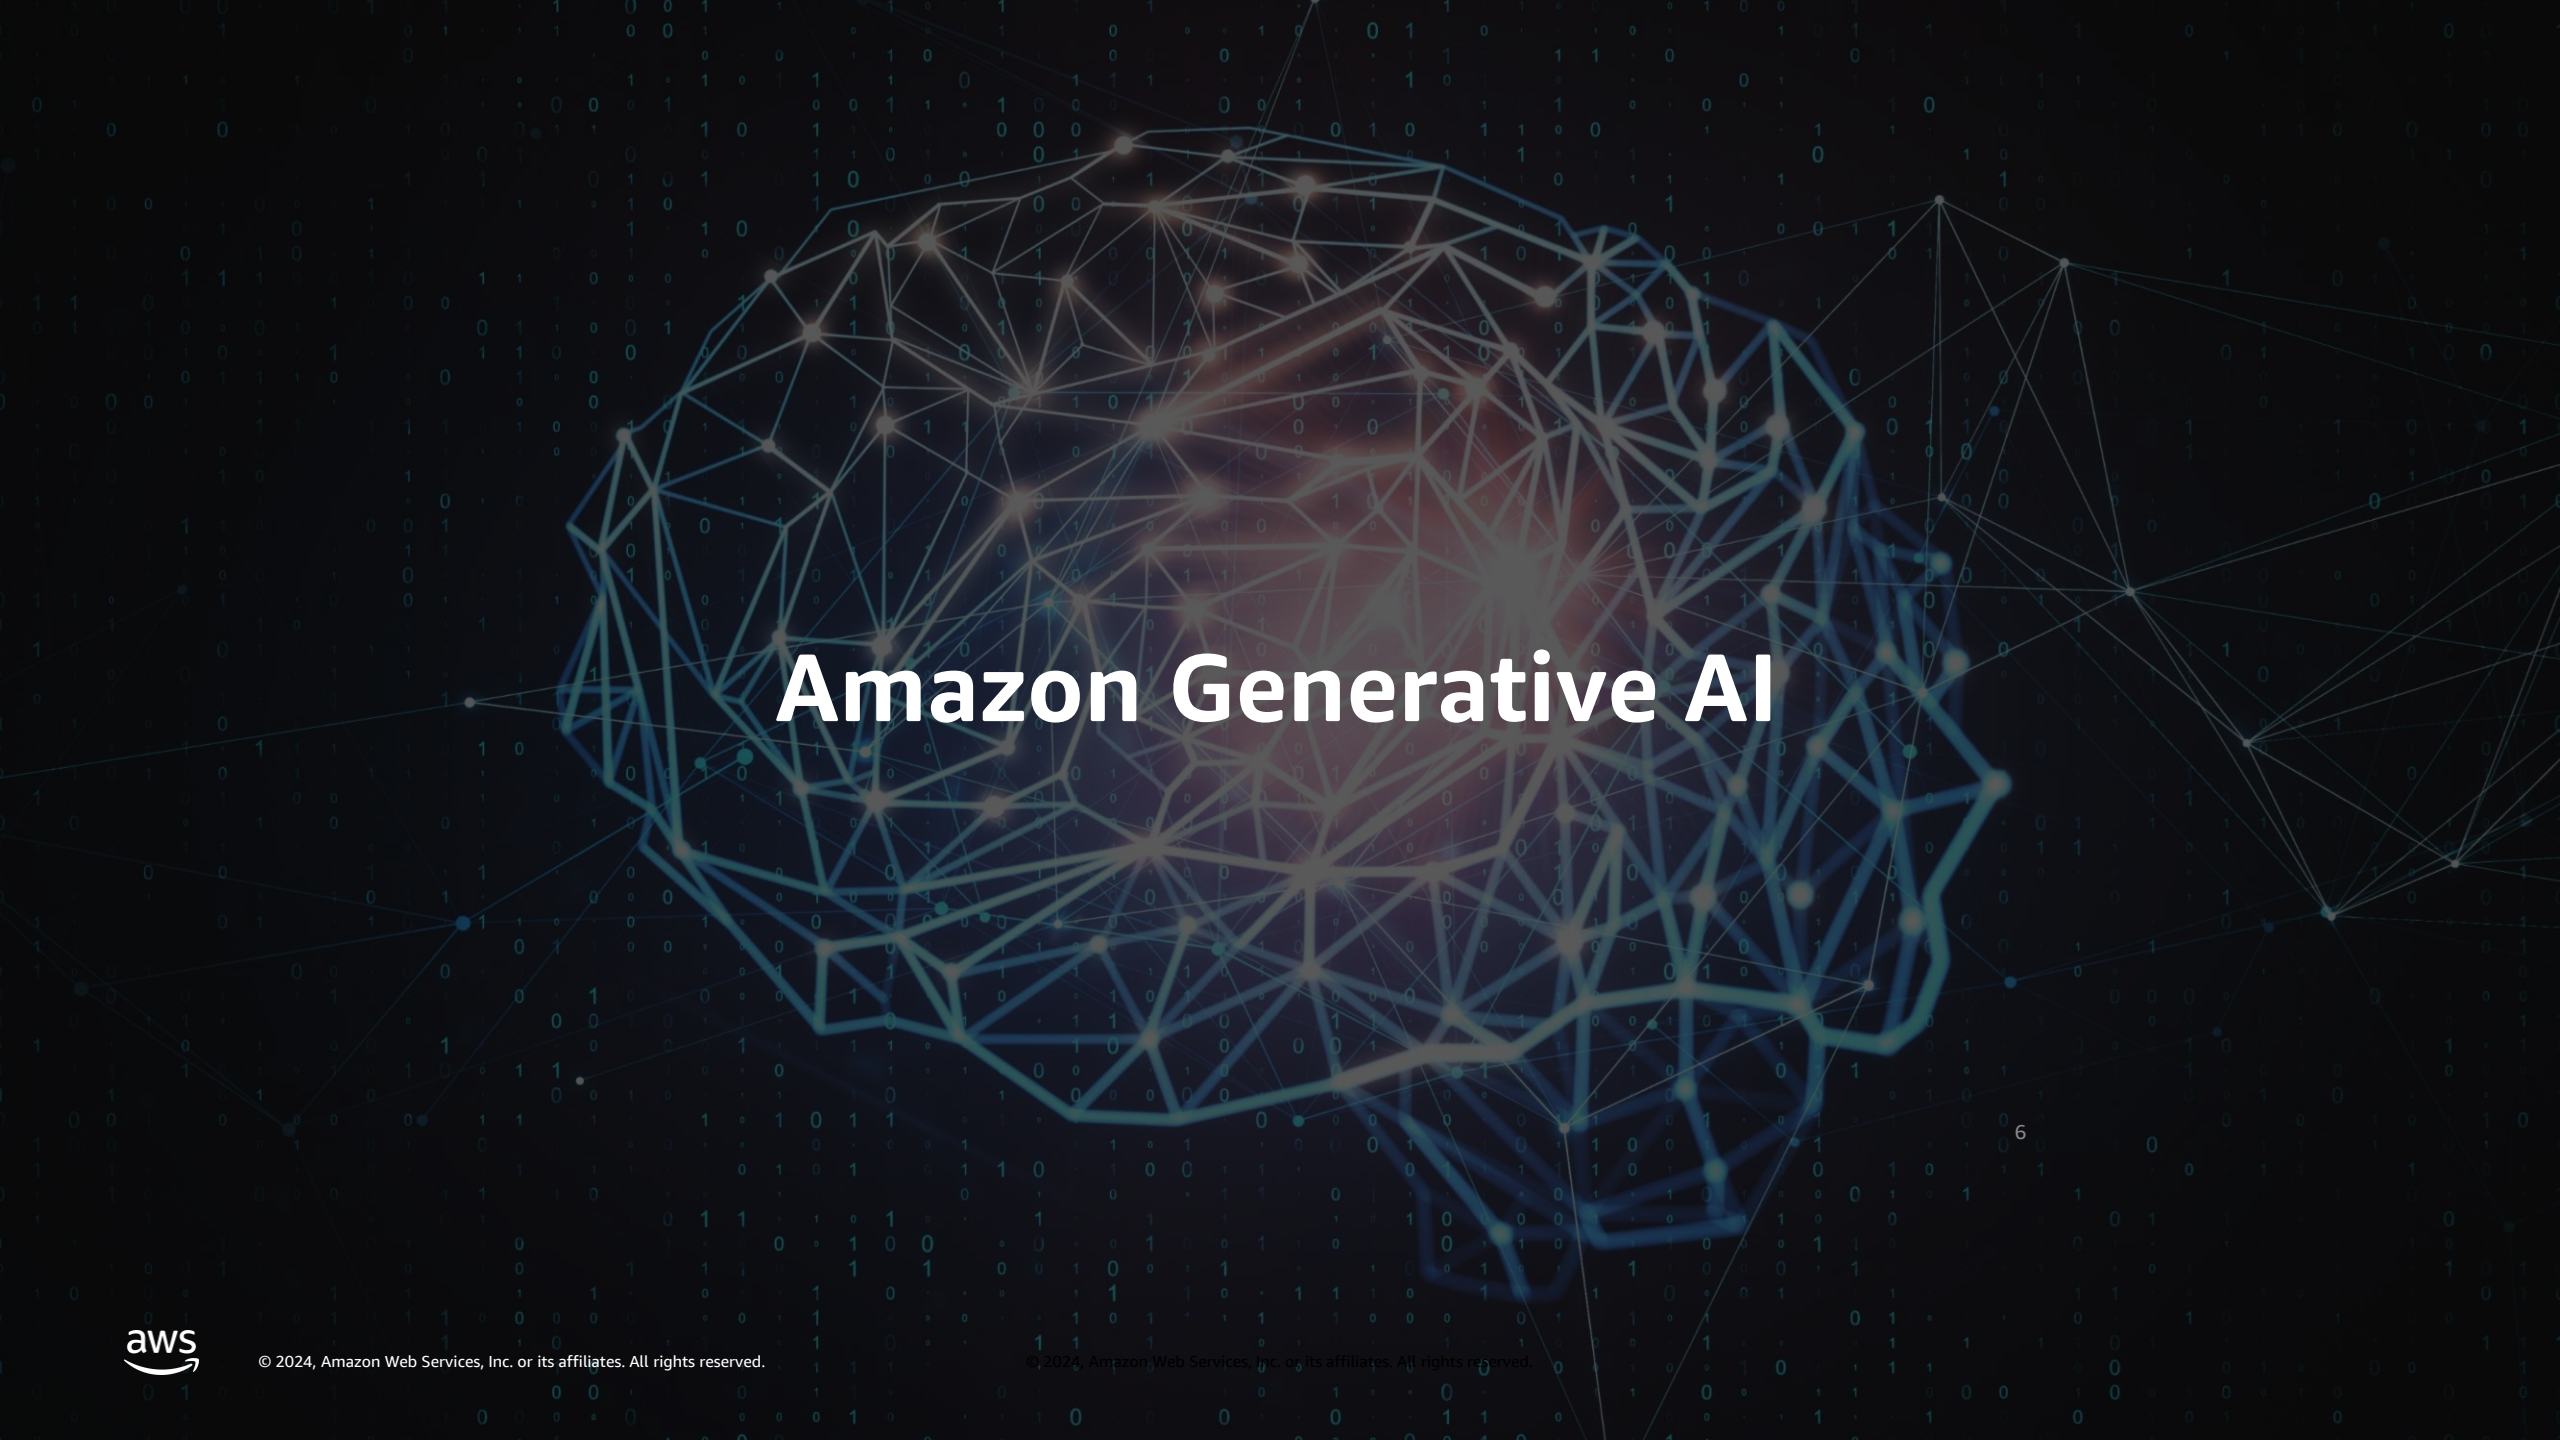

Amazon Generative AI

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

What is generative AI?

Question: What is generative artificial intelligence (AI)?

- Creates new content and ideas, including conversations, stories, images, videos, and music
- Powered by large models that are pretrained on vast corpuses of data and commonly referred to as foundation models (FMs)

AWS generative AI strategy

Flexibility

Secure customization

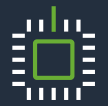

The most cost-effective infrastructure

The easiest way to build with foundation models (FMs)

Generative AI-powered solutions

Common use cases

Text generation

Q&A

Text summarization

Text extraction

Paraphrase rephrase

Search

Code generation

Image generation

Image classification

Audio generation

Video generation

Enhance Customer Experiences

CHATBOTS

VIRTUAL ASSISTANTS

CONVERSATION ANALYTICS

PERSONALIZATION

Boost employee productivity & creativity

CONVERSATIONAL SEARCH

SUMMARIZATION

CONTENT CREATION

CODE GENERATION

DATA TO INSIGHTS

Optimize business processes

DOCUMENT PROCESSING

DATA AUGMENTATION

CYBERSECURITY

PROCESS OPTIMIZATION

Generative AI public sector application examples

Constituent Communications

Citizen engagement and feedback, transparency

Public Health

Personalized care, population health assessments

Public Safety

Public safety and crime prevention, Emergency response and disaster management

Transportation

Trafic optimization, autonomous vehicle control, personalized transportation experiences

Finance

Budget Optimization, fraud detection, risk assessment and mitigation

Constituent Services

Urban planning, Personalized urban services

Energy and utilities

Energy management, Waste management, Smart grid optimization

Research and Engagement

Environmental monitoring

Instead of sending your data to the model, bring the model to your data.

Generative AI Stack

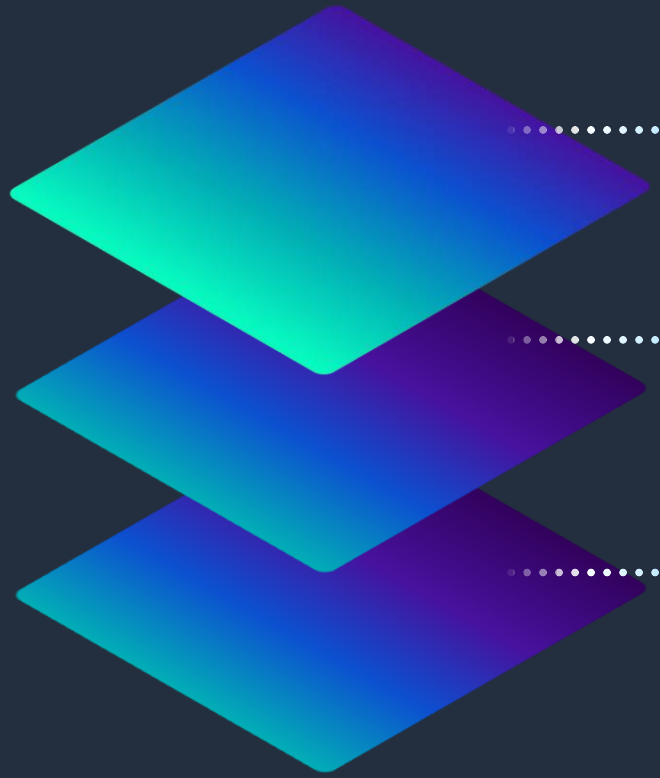

APPLICATIONS THAT LEVERAGE FMs

Amazon Q | Amazon Q in Amazon QuickSight | Amazon Q in Amazon Connect | Amazon CodeWhisperer

TOOLS TO BUILD WITH FMS AND LLMS

Amazon Bedrock

Guardrails | Agents | Customization capabilities

INFRASTRUCTURE FOR FM TRAINING & INFERENCE

GPUs | Trainium | Inferentia | SageMaker

UltraClusters | EFA | EC2 Capacity Blocks | Nitro | Neuron

Amazon Bedrock

The easiest way to build and scale generative AI applications with LLMs and other FMs

Choice of industry-leading FMs from AI21 Labs, Amazon, Anthropic, Cohere, Meta, and Stability AI

Customize FMs using your organization's data

Enterprise-grade security and privacy

NEW

Agents for Amazon Bedrock

Execute multi-step tasks across company systems and data sources

GENERALLY AVAILABLE

Enables generative AI applications to take action in just a few clicks

Breaks down and orchestrates tasks and executes API calls on your behalf

Securely accesses and retrieves company data

NEW

Amazon Q

A generative AI-powered assistant for work that is tailored to your business

AVAILABLE IN PREVIEW

Provides interactive answers, solves problems, generates content, and takes action

Understands your company information, code, and systems

Personalizes interactions based on your role and permissions

Built to be secure and private

Amazon Q is

AMAZON Q
YOUR AWS EXPERT

AMAZON Q IN AMAZON QUICKSIGHT
YOUR BI EXPERT

AMAZON Q
YOUR BUSINESS EXPERT

AMAZON Q IN AMAZON CONNECT
YOUR CONTACT CENTER EXPERT

Visual authoring in **QuickSight**

Use everyday language to generate and fine-tune visuals in seconds

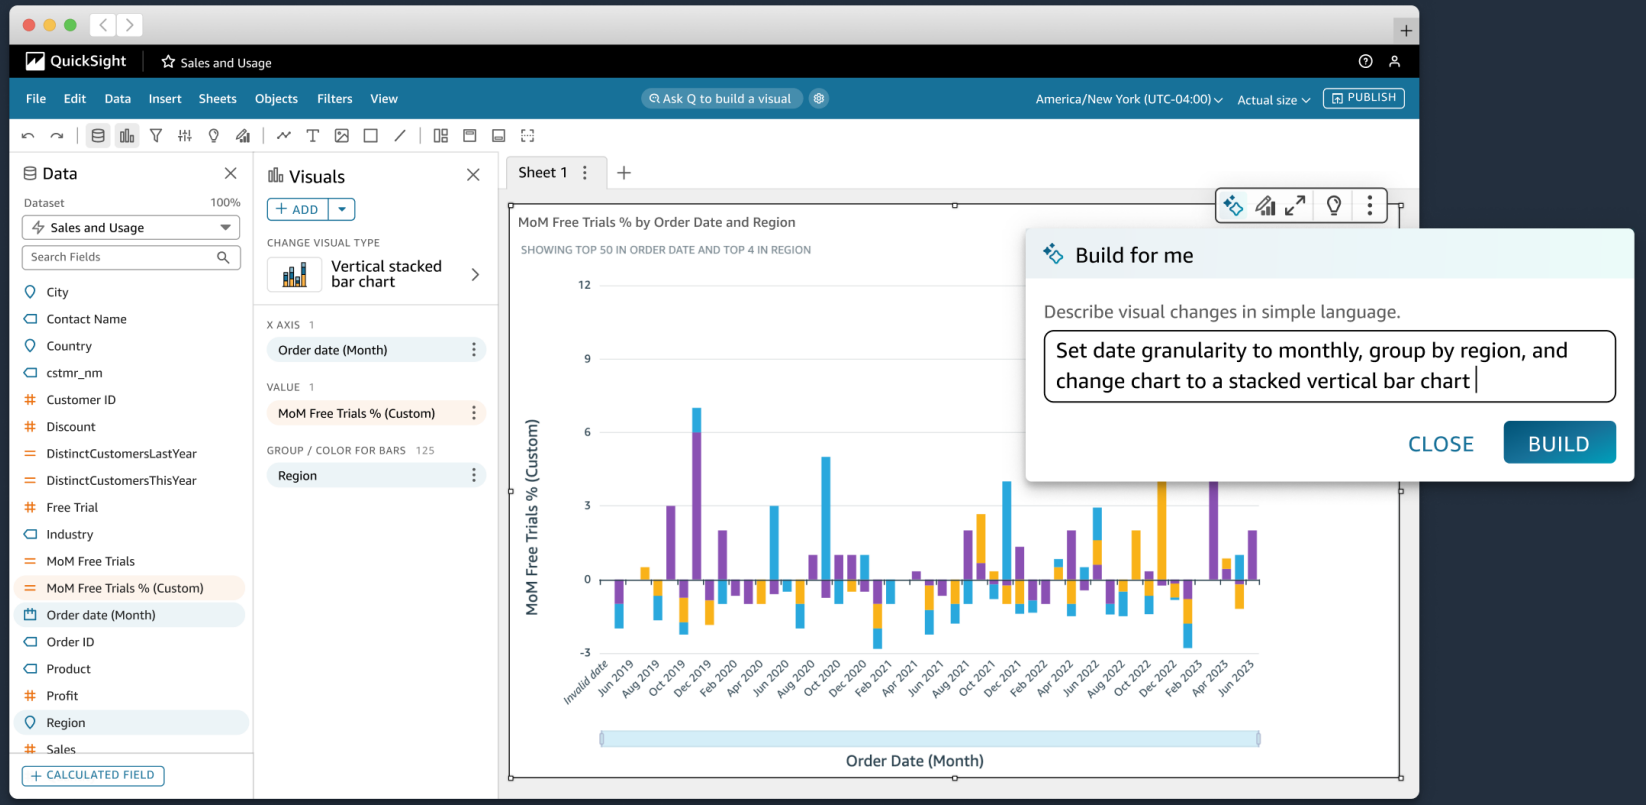

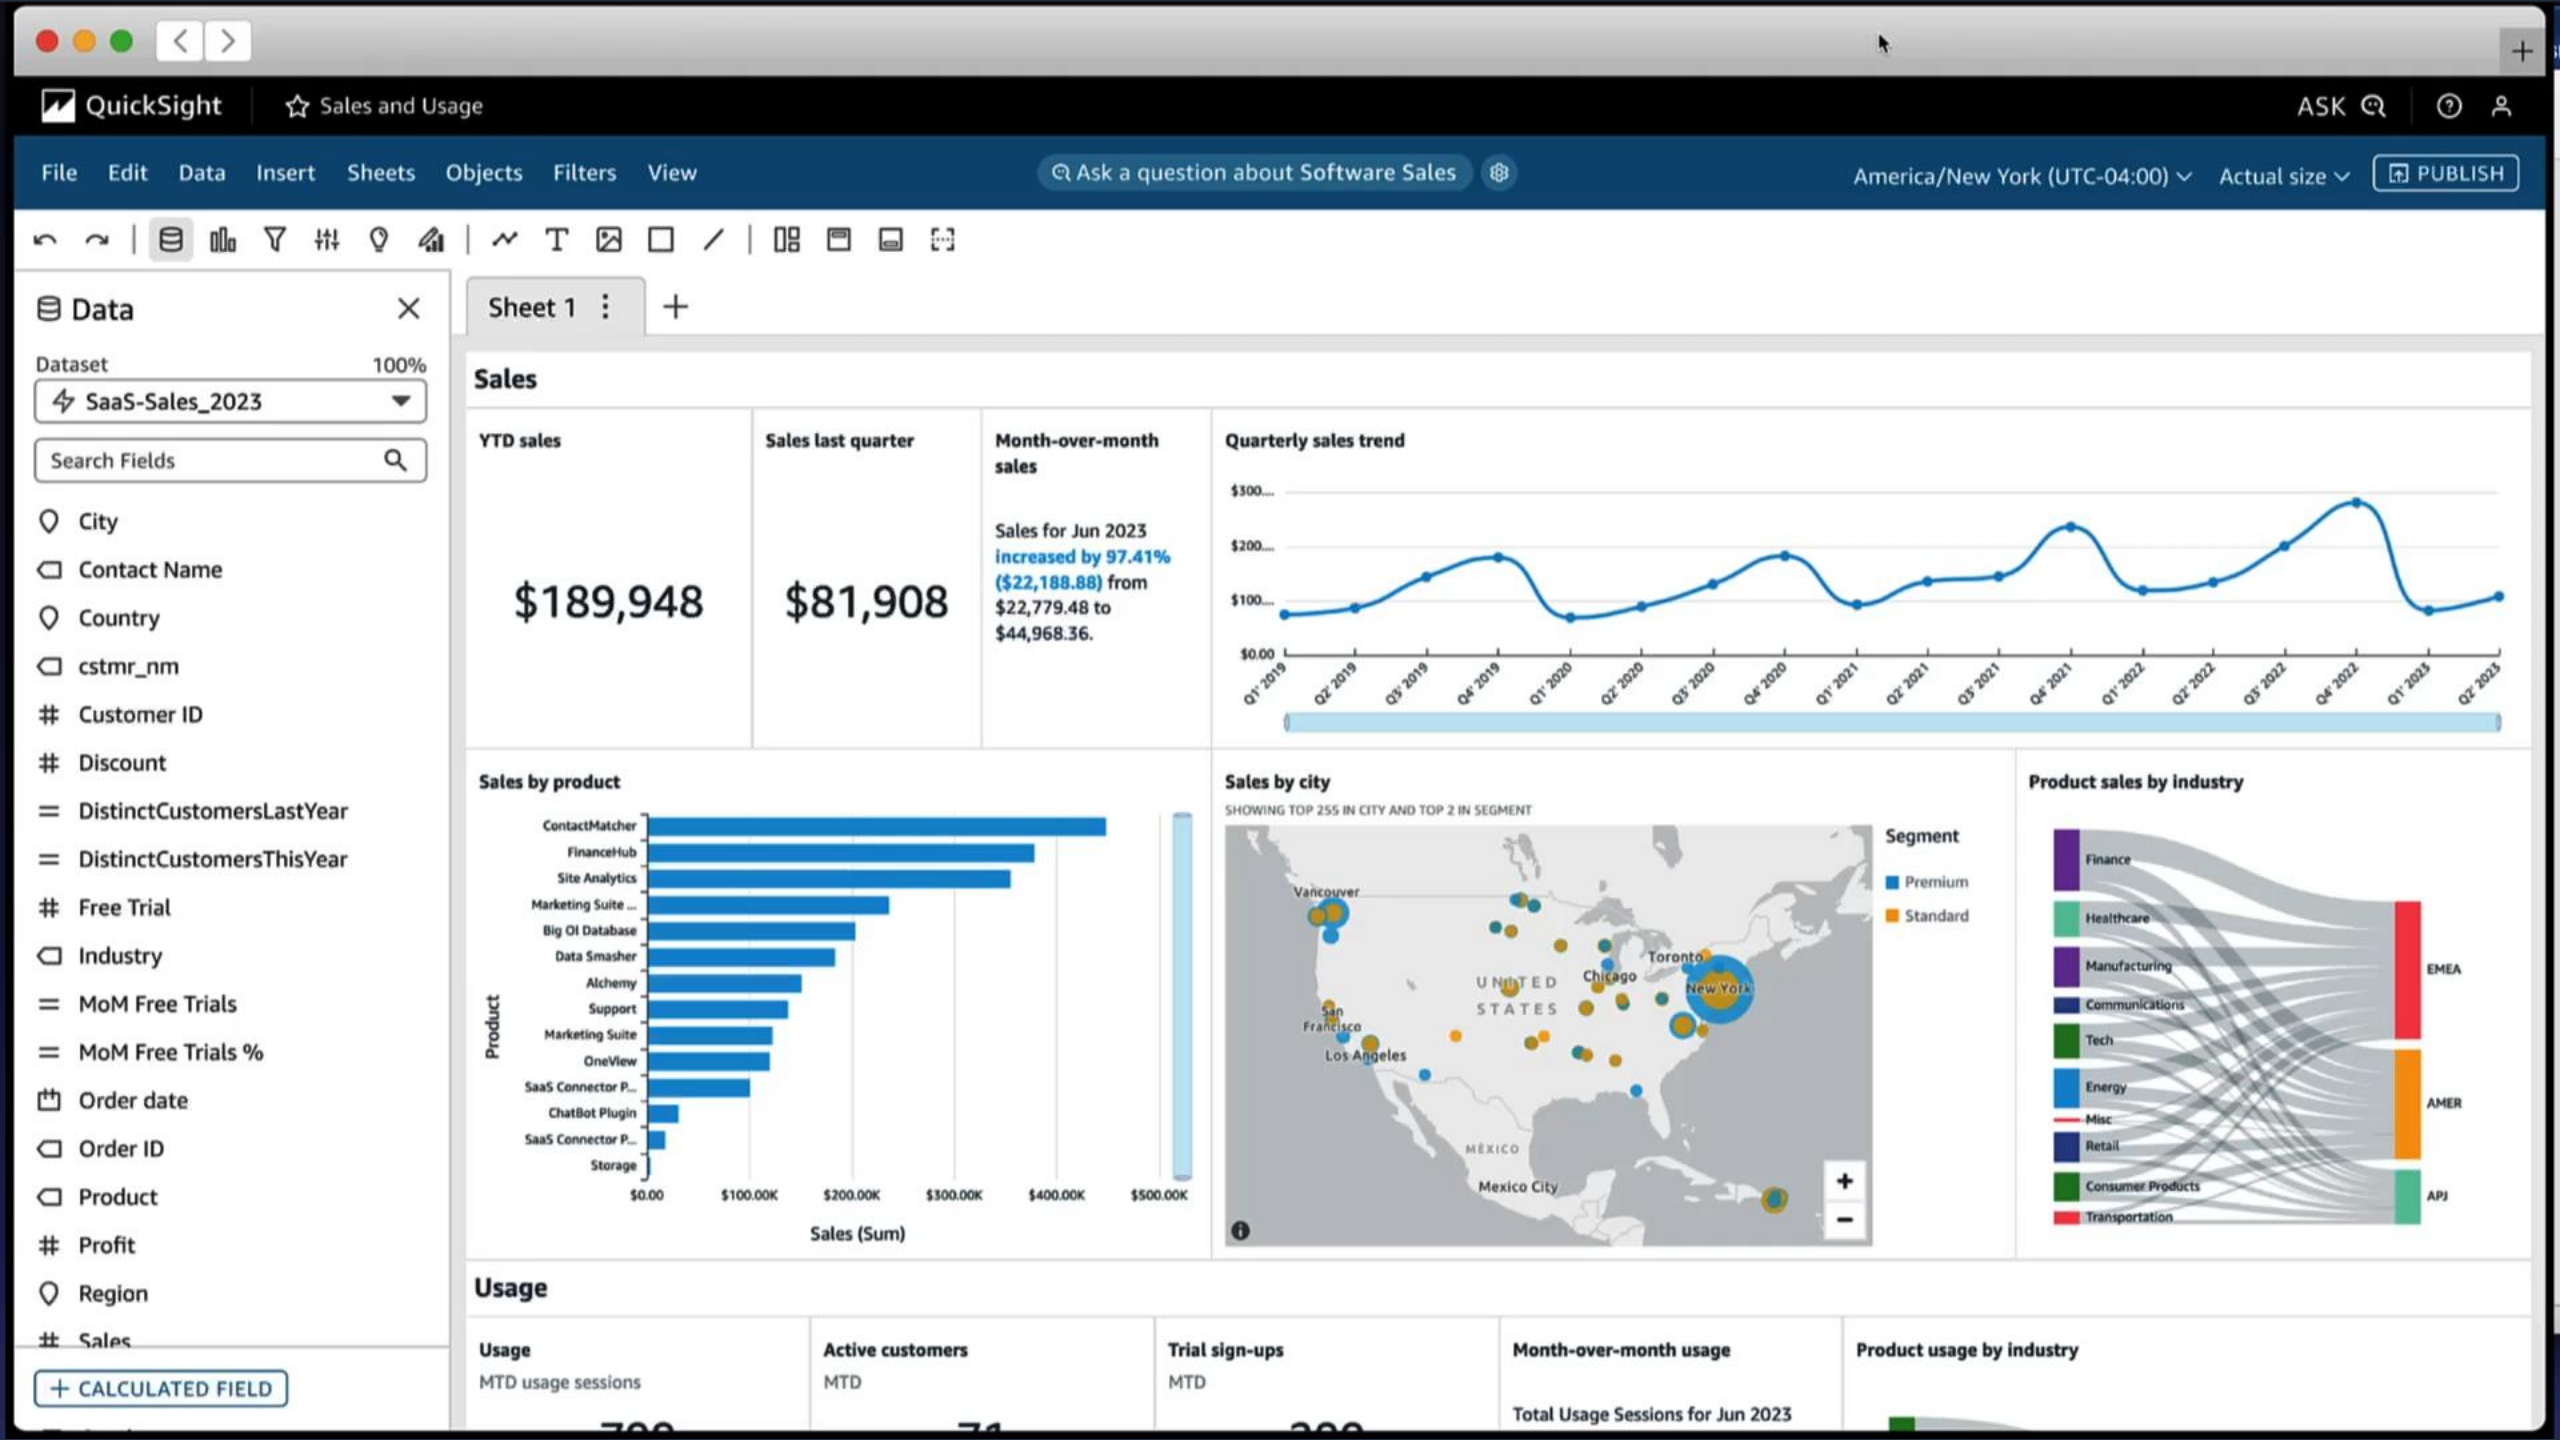
QuickSight
Sales and Usage
ASK
File Edit Data Insert Sheets Objects Filters View
Ask a question about Software Sales
America/New York (UTC-04:00)
Actual size
PUBLISH
Data
Dataset 100%
SaaS-Sales_2023
Search Fields
City
Contact Name
Country
cstmr_nm
Customer ID
Discount
DistinctCustomersLastYear
DistinctCustomersThisYear
Free Trial
Industry
MoM Free Trials
MoM Free Trials %
Order date
Order ID
Product
Profit
Region
Sales
+ CALCULATED FIELD
Sheet 1
Sales
YTD sales
$189,948
Sales last quarter
$81,908
Month-over-month sales
Sales for Jun 2023 increased by 97.41% ($22,188.88) from $22,779.48 to $44,968.36.
Quarterly sales trend
$300...
$200...
$100...
$0.00
Q1' 2019 Q2' 2019 Q3' 2019 Q4' 2019 Q1' 2020 Q2' 2020 Q3' 2020 Q4' 2020 Q1' 2021 Q2' 2021 Q3' 2021 Q4' 2021 Q1' 2022 Q2' 2022 Q3' 2022 Q4' 2022 Q1' 2023 Q2' 2023
Sales by product
ContactMatcher
FinanceHub
Site Analytics
Marketing Suite ...
Big OI Database
Data Smasher
Alchemy
Support
Marketing Suite
OneView
SaaS Connector P...
ChatBot Plugin
SaaS Connector P...
Storage
Product
$0.00 $100.00K $200.00K $300.00K $400.00K $500.00K
Sales (Sum)
Sales by city
SHOWING TOP 255 IN CITY AND TOP 2 IN SEGMENT
Segment
Premium
Standard
Vancouver
Toronto
Chicago
New York
UNITED STATES
San Francisco
Los Angeles
MÉXICO
Mexico City
Product sales by industry
Finance
Healthcare
Manufacturing
Communications
Tech
Energy
Misc
Retail
Consumer Products
Transportation
EMEA
AMER
APJ
Usage
Usage
MTD usage sessions
Active customers
MTD
Trial sign-ups
MTD
Month-over-month usage
Total Usage Sessions for Jun 2023
Product usage by industry

Amazon Generative AI Portfolio

Choice of many Foundation Models

AI SERVICES

Bedrock

In Preview only

Amazon

Titan Text

Titan Embeddings

3rd Party

Jurassic-2

Claude

Stable Diffusion

Code Whisperer

Amazon

AI coding companion

ML SERVICES

Sagemaker

Jumpstart

Publicly available

stability.ai

Text2Image

Upscaling

alexa

AlexaTM 20B

Flan T-5 models

DistilGPT2, GPT2

Bloom models

Proprietary models

co:here

Cohere generate-med

Lighton

Lyra-Fr10B

AI21labs

Jurassic-1 Grande 17B

In Preview only

ML FRAMEWORKS & INFRASTRUCTURE

Self Managed ML

Meta PyTorch

3-way collaboration to move models to production on EC2 and Sagemaker

Hugging Face

Amazon Bedrock keeps data secure & private

None of the customer's data is used to train the underlying model

All data is encrypted in transit and at rest

Data used to customize models remains within your VPC

Support for standards, including GDPR & HIPAA

Emerging risks and challenges with generative AI

Veracity (e.g., hallucinations)

Toxicity and safety

Intellectual property

Data privacy

Innovate responsibly with generative AI

Pillars of responsible AI

Value alignment
Systems should be designed and used in ways that align with the organization mission, social norms, and legal regulations

Inclusion
Inclusion of diverse and unique skills, experiences, perspectives, and cultural backgrounds

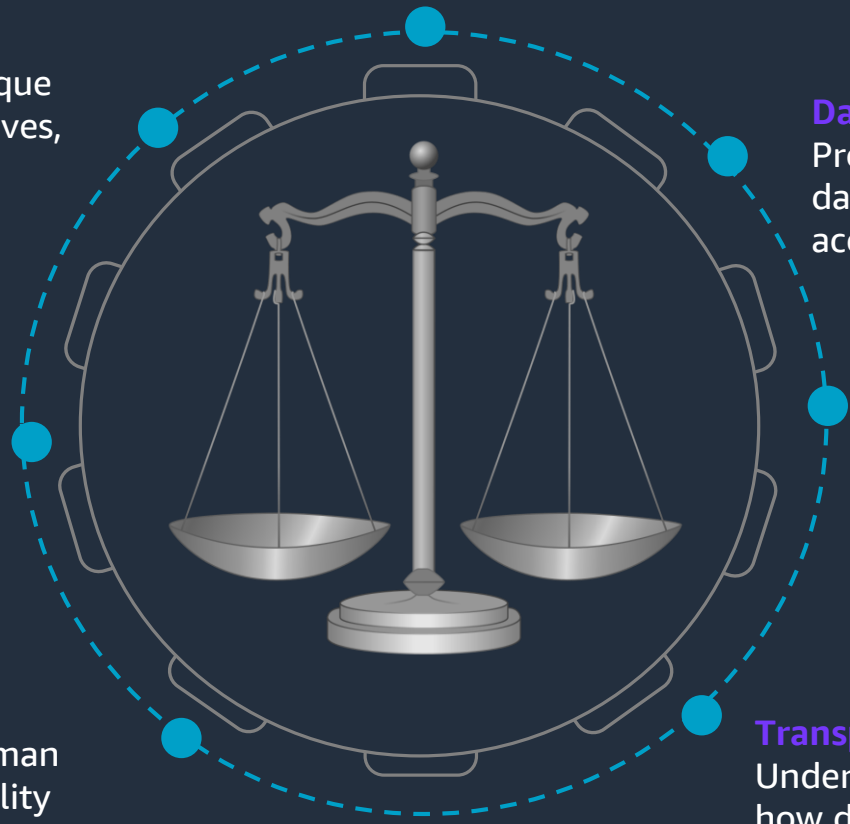

Data privacy and protection
Protects the quality and integrity of data used as well as its relevance, access, and processing

Training and education
Appropriate knowledge sharing and education to understand purpose, use, and impact

Fairness
Systems must be designed to minimize bias and promote inclusive representation

Accountability
Structured, maintaining human involvement and responsibility for design, development, decision processes, and outcomes

Transparency and explainability
Understanding how data is used and how decisions and outcomes are made understandable to a human

Our commitment…
…and how we drive adoption and improvement

Developing AI in a responsible way is integral to our approach

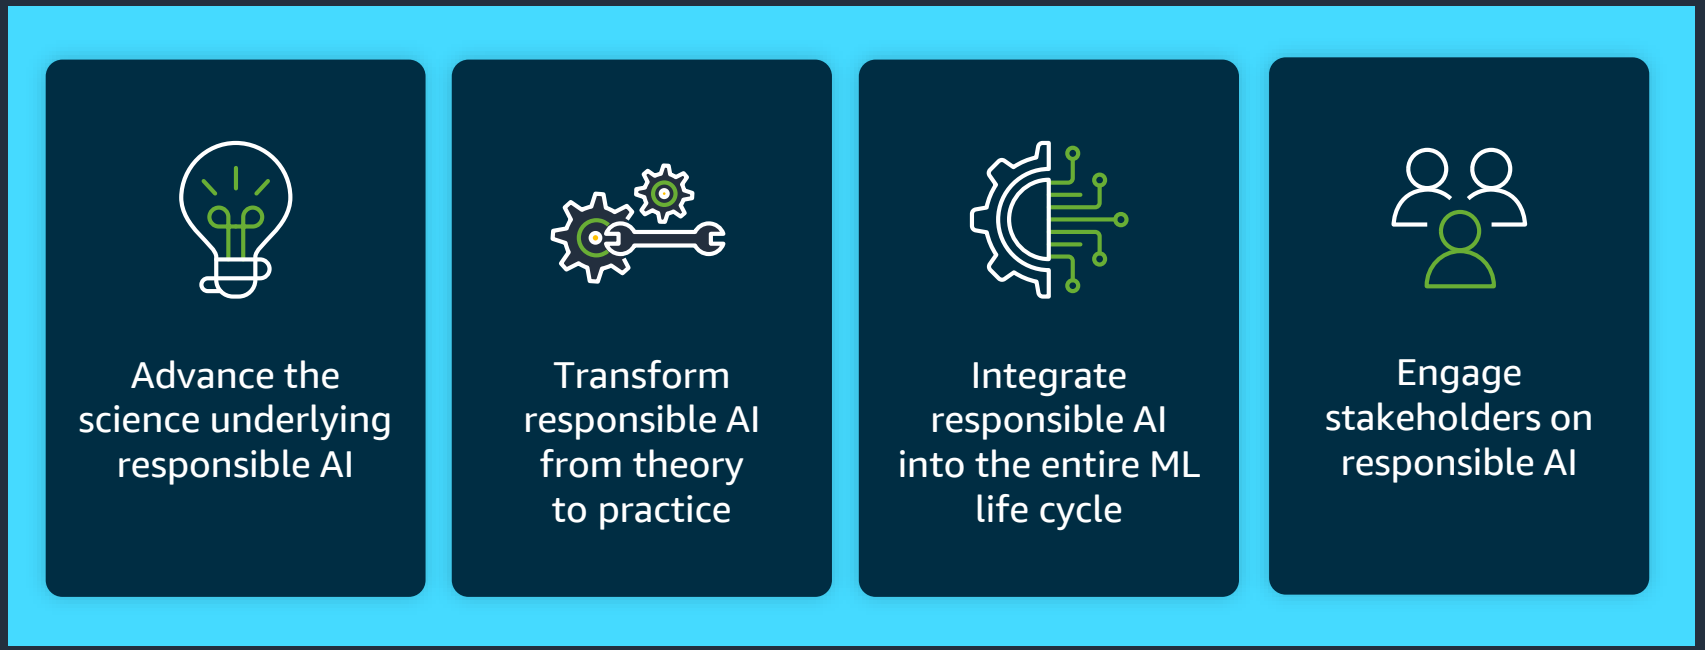

AWS supports a technical response to responsible AI

Flexible architecture: Choice of foundation models (FMs)
Domain expertise, regulatory compliance, accuracy, cost and sustainability

Guardrails for Amazon Bedrock

Amazon SageMaker Clarify

Watermarking with Amazon Titan

AWS AI Service Cards

NEW

Guardrails for Amazon Bedrock

Safeguard your generative AI applications with your responsible AI policies

AVAILABLE IN PREVIEW

Easily configure harmful content filtering based on your responsible AI policies

Apply Guardrails to any FM or agent

Redact PII information in FM responses (coming soon)

How to address generative AI in Public Sector

- How can we support you in ensuring accuracy and authority of model outputs?
- How can we use “guardrails” to minimize inappropriate content?
- How can we maximize the public investment and minimize cost?
- What are your latency requirements for specific use cases?
- What are the potential regulatory, data privacy, & security considerations that may dictate model architecture.
- How can we ensure proprietary, copyrighted and IP concerns are addressed?

Please complete the session survey by scanning the QR code

Select: Executive track

Thank you!

Sergio Ortega
sergioai@amazon.com