# AWS State, Local, and Education Learning Days

Arlington, VA



aws Learning Days
State, Local, and Education

# The Evolving Research Landscape & the AWS Cloud

**Greg Ritter** (he/him)

Strategy Lead – Research Data, Scholarly Communication, & Libraries

AWS Higher Education

gritter@amazon.com



aws Learning Days
State, Local, and Education

# Agenda

- The evolving research landscape
- Why AWS for research?
  - Security & Compliance
  - Research Computing
  - Research Data Management
  - Research Administration
- Questions & discussion

quickly

# The evolving research landscape

# The elephant in the ~~room~~ lab

Generated by Amazon Nova Canvas

# New challenges

- Review of new grants frozen or delayed
- Hundreds of grants have been terminated, sometimes leaving labs and staff with no way forward
- University administrations threatened with massive loss of Federal funding
- Caps threatened for indirect cost recovery
- Reductions in overall Federal support for research

# How AWS can help

***AWS helps institutions stand up the most flexible, price performant, secure, and compliant environments for researchers, and provides support for institutions and researchers through the transition from on-premises to cloud-based research.***

## Infrastructure Costs

- Cloud costs can be included as direct costs in grant budgets
- We can collaborate on chargeback models for cloud
- We can help evaluate sustainable funding approaches

## Operational Efficiency

- Automate administrative workflows
- Centralize and strengthen security and compliance
- Simplify IT management, resource allocations and monitor cloud usage and expenses

## Research Continuity

- Provide on-demand access to scalable computing resources
- Conduct research securely, in compliance with NIST, CMMC, HIPAA, and other frameworks
- Access leading AI tools & models for research workloads

# AWS Programs for Research

## AWS Cloud Credits for Research

The AWS Cloud Credit for Research are a public-sector program to support research projects that

- build cloud-hosted services, software, or tools
- migrate a research process or open data to the cloud.

## Amazon Research Awards

Amazon Research Awards (ARA) are one-year unrestricted gifts to a PI's institution or organization.

ARA support research relevant to Amazon, such as applied machine learning, automated reasoning, computer vision, fairness in AI, security, sustainability and more.

## Support for Grant Proposals

AWS supports researchers' grant development cycle by providing consultative support to help design and architect a technical solution or Proof of Concept (PoC).

The PI can include the technical architectures and diagrams along with the cost estimates in their grant proposals.

## Global Data Egress Waiver (GDEW)

GDEW helps academic institutions use AWS cloud storage, computing, and database services by waiving data egress fees.

The maximum discount is 15 percent of total monthly spending on AWS services

# Why AWS for research?

# *96% of US research institutions (R1) the US use AWS*

# AWS enables research

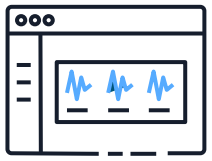

## Accelerate research

On-demand access to resources

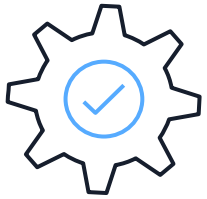

## Meet capacity demands

Scale to meet research requirements

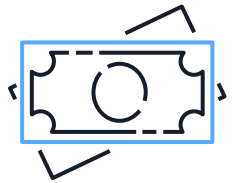

## Provide flexible resource access

Balance cost and performance

## Services beyond HPC

AI/ML, accelerators, quantum computing

## Match the pace of innovation

Scale to meet research requirements

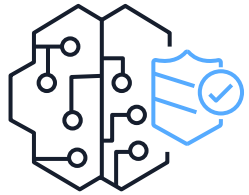

## Ensure security & compliance

Demonstrate accountability

# The research lifecycle

Archive and preserve research data, support grant close-out and reporting

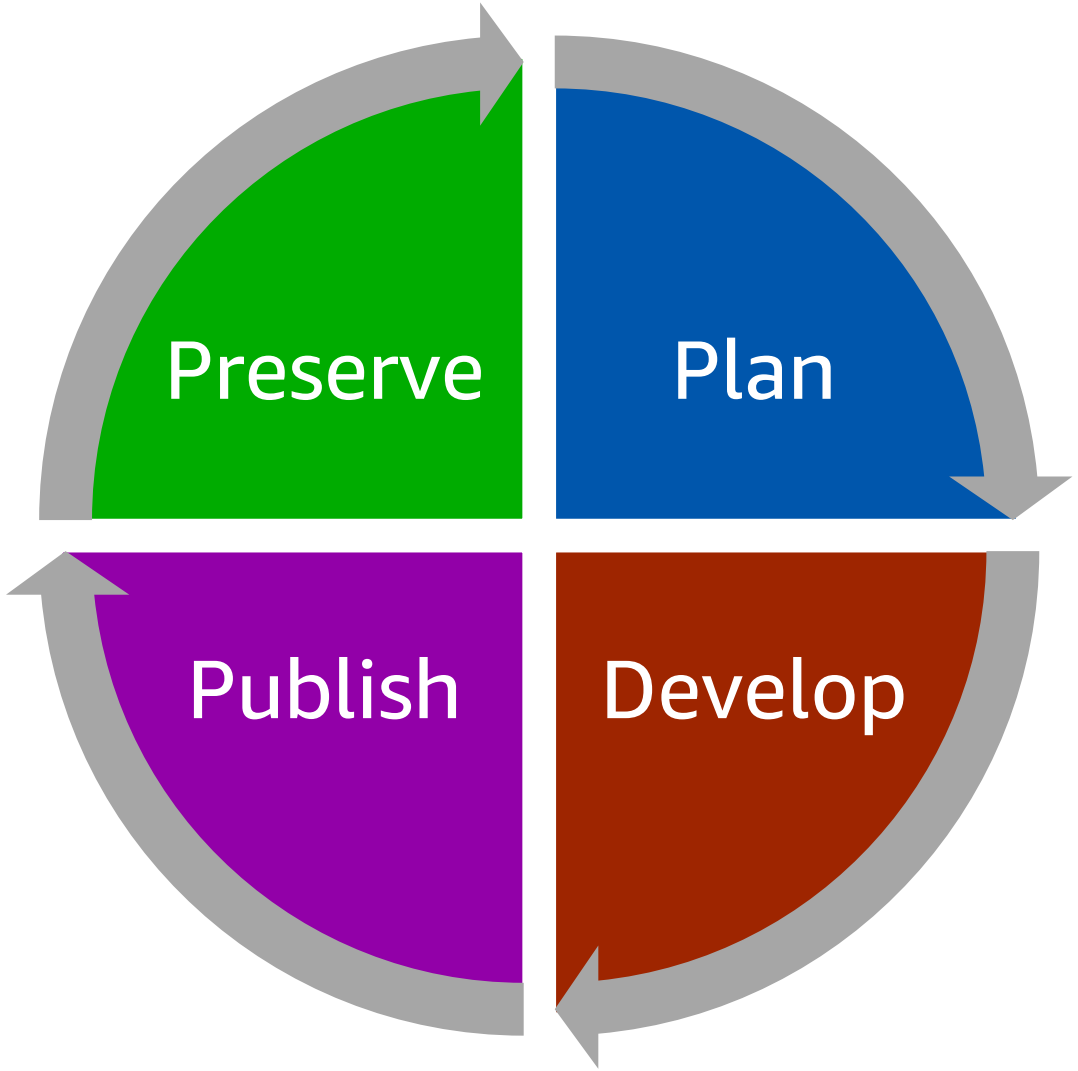


Identify research opportunities, review existing data and research, prove out key concepts, build institutional capacity for research

Promote research findings via open access and conferences, facilitate sharing & collaboration, support tech transfer & commercialization

Writing cloud-friendly proposals, building data management plans, securing research workloads, collecting and analyzing data,

# How we support researchers & institutions

We support institutions and researchers for needs common to all research and those specific to disciplines across the research lifecycle.

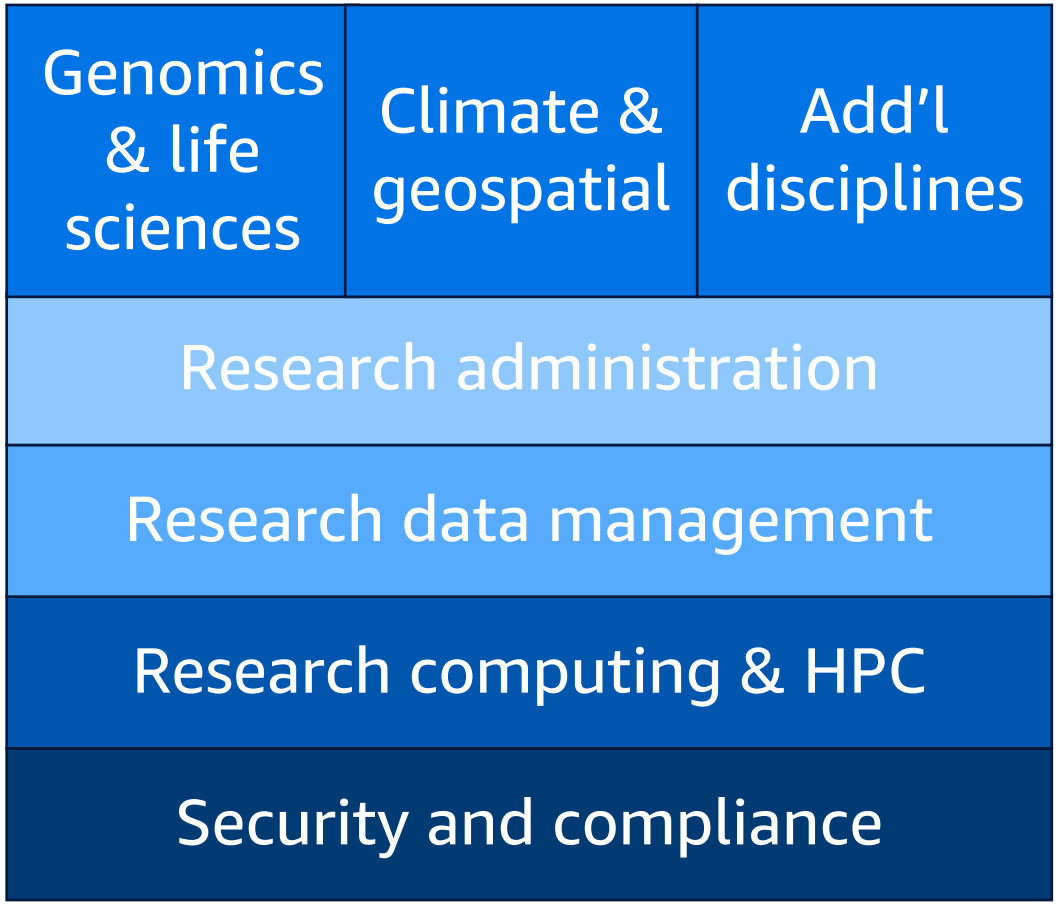

# Security & Compliance

# The “must haves” for research

- Cybersecurity Maturity Model Certification (CMMC)
  - Appearing in DoD contracts/grants: 2021-2025
  - Third-party assessed and certified solution to CMMC level 1-5
  - Certification good for three years
- NIST 800-171 Self Assessment Requirement
  - Basic Self Assessments must be reported to the Supplier Performance Risk System (SPRS) every three years
  - Random audits by Defense Industrial Base Cybersecurity Assessment Center (DIBCAC)



# Why now?

Responsible stewardship of research data is expected.
Compliance defines specific responsibilities for research data.

HIPAA

GDPR

CCPA

Data privacy

NIST / CUI

DFARS

CMMC

Responsible use

# Why now?

Responsible stewardship of research data is expected.
Compliance defines specific responsibilities for research data.

**Penalties and fines**
data breach

HIPAA

GDPR

CCPA

Data privacy

**Reputational Damage**

**Funding withheld**
non-compliance

NIST / CUI

DFARS

CMMC

Responsible use

# Why AWS?

*AWS elevates your institution's security and compliance posture along with its research capabilities*

Inherit global security and compliance controls

Scale with superior visibility and control

Highest standards for privacy and data security

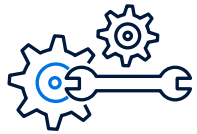

Automate and reduce risk with deeply integrated services

Largest community of security partners and solutions

# How AWS delivers

## The AWS Landing Zone Guardrails Accelerator

Leverage AWS best practices to establish an **AWS Landing Zone** with security guardrails tailored to your institutional policies and governance.

With an AWS Landing Zone you can establish consistent technical controls which support compliant, secure, and auditable provisioning and operation of research workloads on AWS.

# Research Computing & HPC

# AWS flexible compute price/performance selection

## Flexible compute to maximize performance

Memory & compute optimized
2.5-3.5Ghz, 2-16GB/core, 100Gbps, EFA

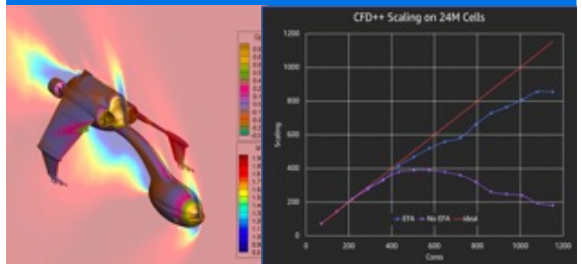


Graphics and rendering
1-8 GPUs, up to 384GB RAM, SSD

High clock speed
4.5Ghz, 192GB RAM, 100Gbps, EFA

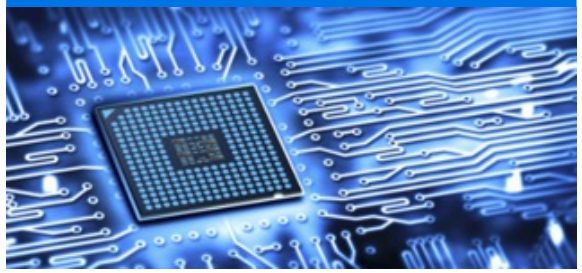

Accelerated computing
8 A100 GPUs, 1.1TB RAM, SSD, 400Gbps

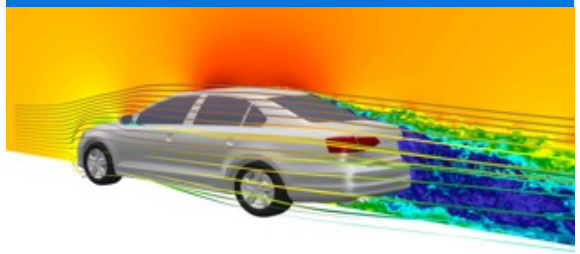

## Flexible pricing models to optimize cost

**On-Demand**

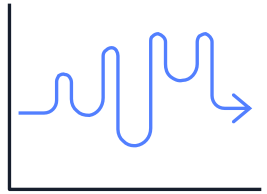

Pay for compute capacity by the second with no long-term commitments.

**Savings Plans**

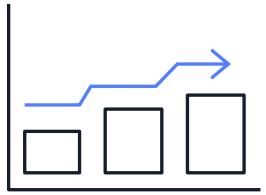

Make a commitment and to save up to 72% off compute or storage.

**Spot Instances**

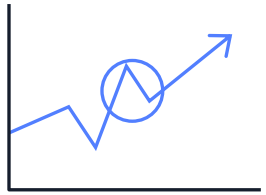

Spare EC2 capacity at savings of up to 90% off On-Demand prices.

# AWS compute instance selection

## CATEGORIES

- General purpose
- Burstable
- Compute intensive
- Memory intensive
- Storage (High I/O)
- Dense storage
- GPU compute
- Graphics intensive

+

## CAPABILITIES

- Choice of processor (AWS, Intel, AMD)
- Fast processors (up to 4.0 GHz)
- High memory footprint (up to 12 TiB)
- Instance storage (HDD and SSD)
- Accelerated computing (GPUs and FPGA)
- Networking (up to 400 Gbps)
- Bare Metal
- Size (Nano to 32xlarge)

+

## OPTIONS

- Amazon EBS
- Amazon Elastic Inference

= **800+**

# Clemson University - Natural Language Processing

Researchers conducted nearly half a million topic modeling experiments to study how human language is processed by computers.

The 1.1 Million vCPU count usage is comparable to the core count on the largest supercomputers in the world.

CLEMSON UNIVERSITY

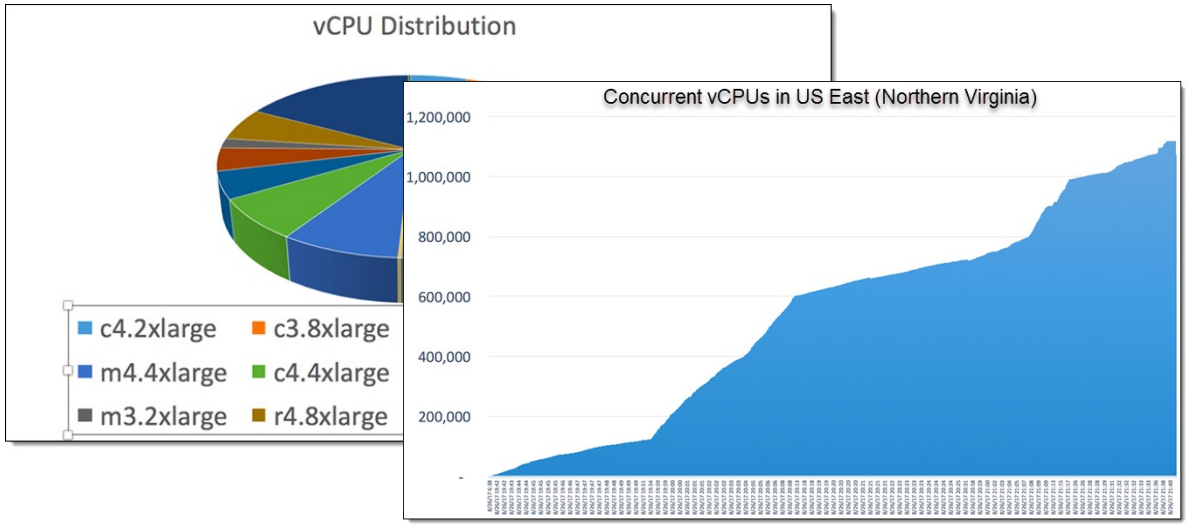


*"I am absolutely thrilled with the outcome of this experiment. The graduate students on the project […] used resources from AWS and Omnibond and developed a new software infrastructure to perform research at a scale and time-to-completion not possible with only campus resources."*

– Prof. Amy Apon, Co-Director of the Complex Systems, Analytics and Visualization Institute

# AWS bursting from on-premises


Using existing on-premises scheduler and software stack to burst jobs into AWS providing availability, scale, and capacity for and institution’s research.

# Research Data Management

# Research data management on AWS

**Find data across silos** | Discover and manage data across your entire storage estate, globally

**Easily curate and share** | Easily prepare data for collaborative analysis, preservation, and sharing

**Capability at cloud scale** | Get access to capabilities at scale for fast, secure data processing & sharing, and AI/ML capabilities

# Amazon SageMaker

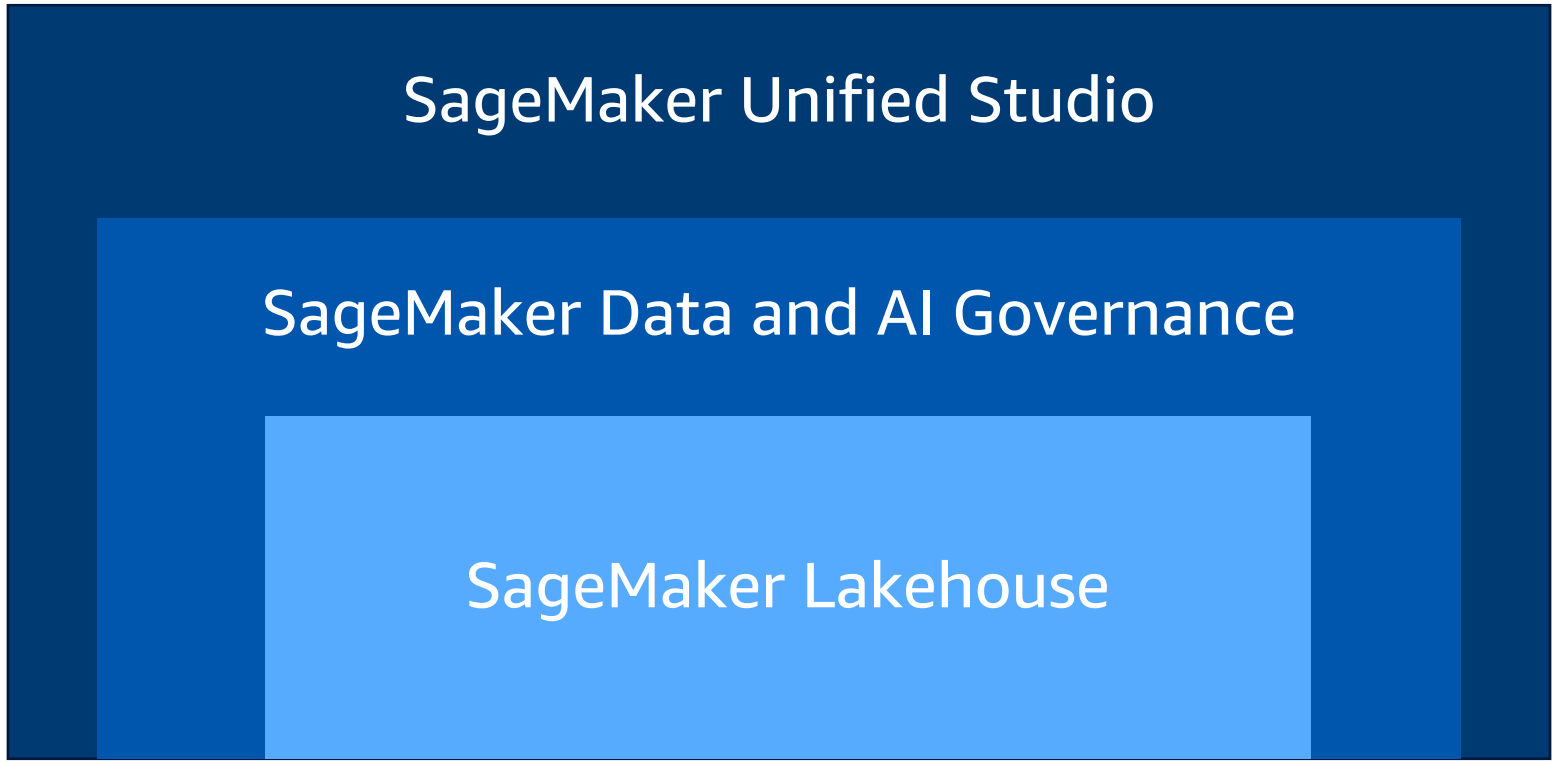

# Amazon SageMaker

## SageMaker Lakehouse

Unifying data across data warehouses, data lakes, operational databases, and applications with a single copy of data

**Zero-ETL, federation, sharing**


AWS DATA SOURCES | THIRD PARTY DATA SOURCES

# Amazon SageMaker

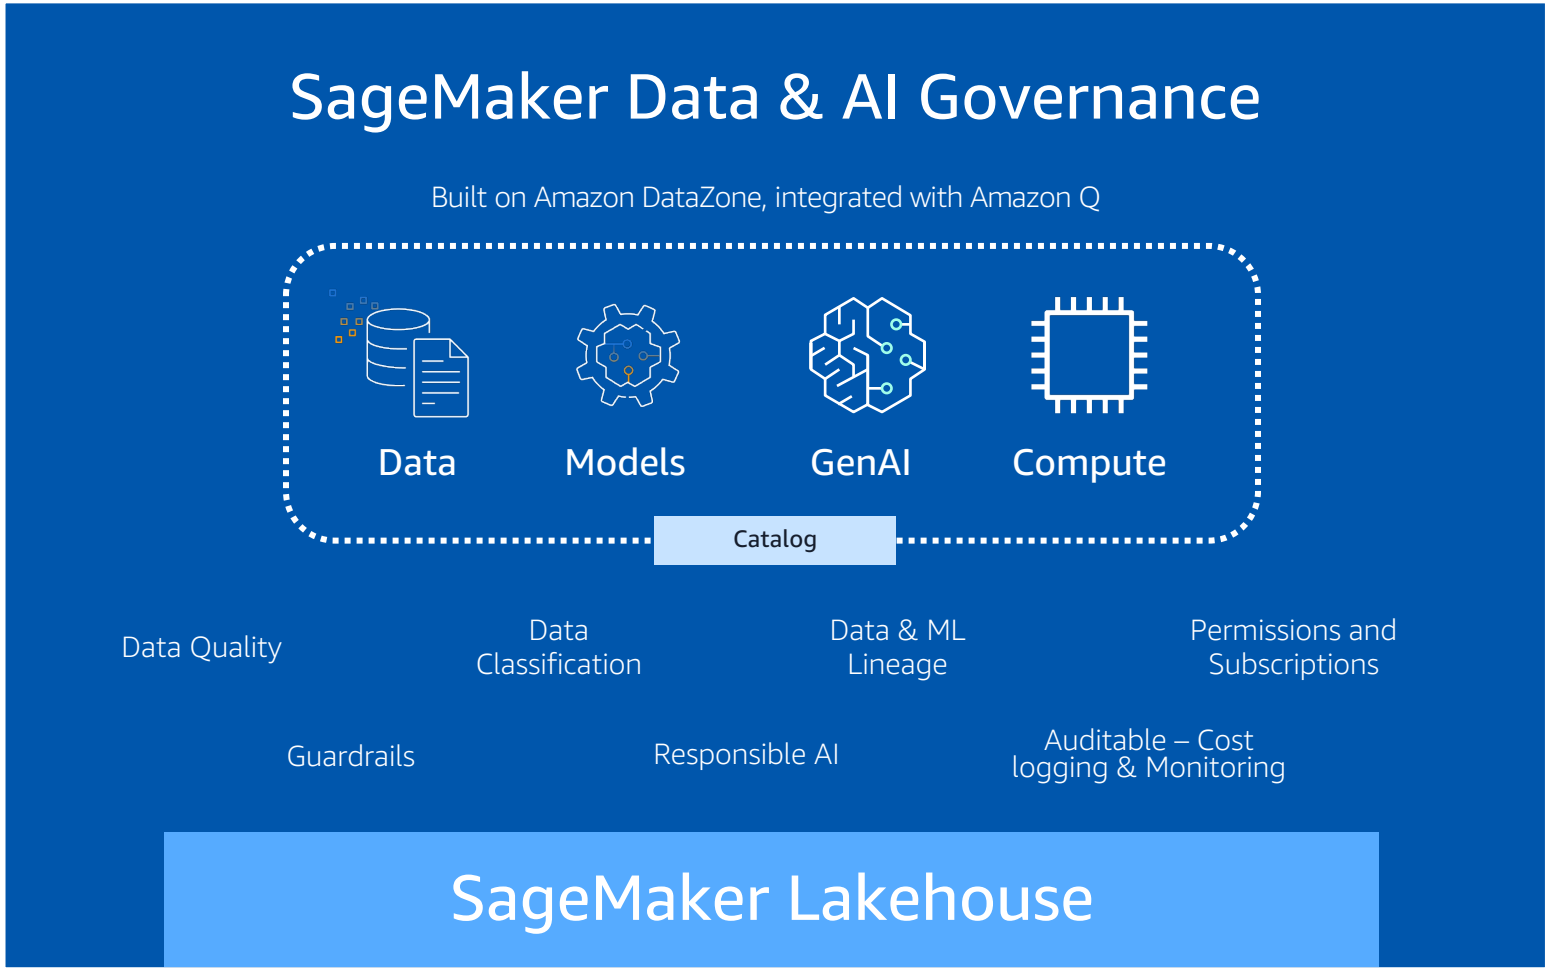

# Amazon SageMaker

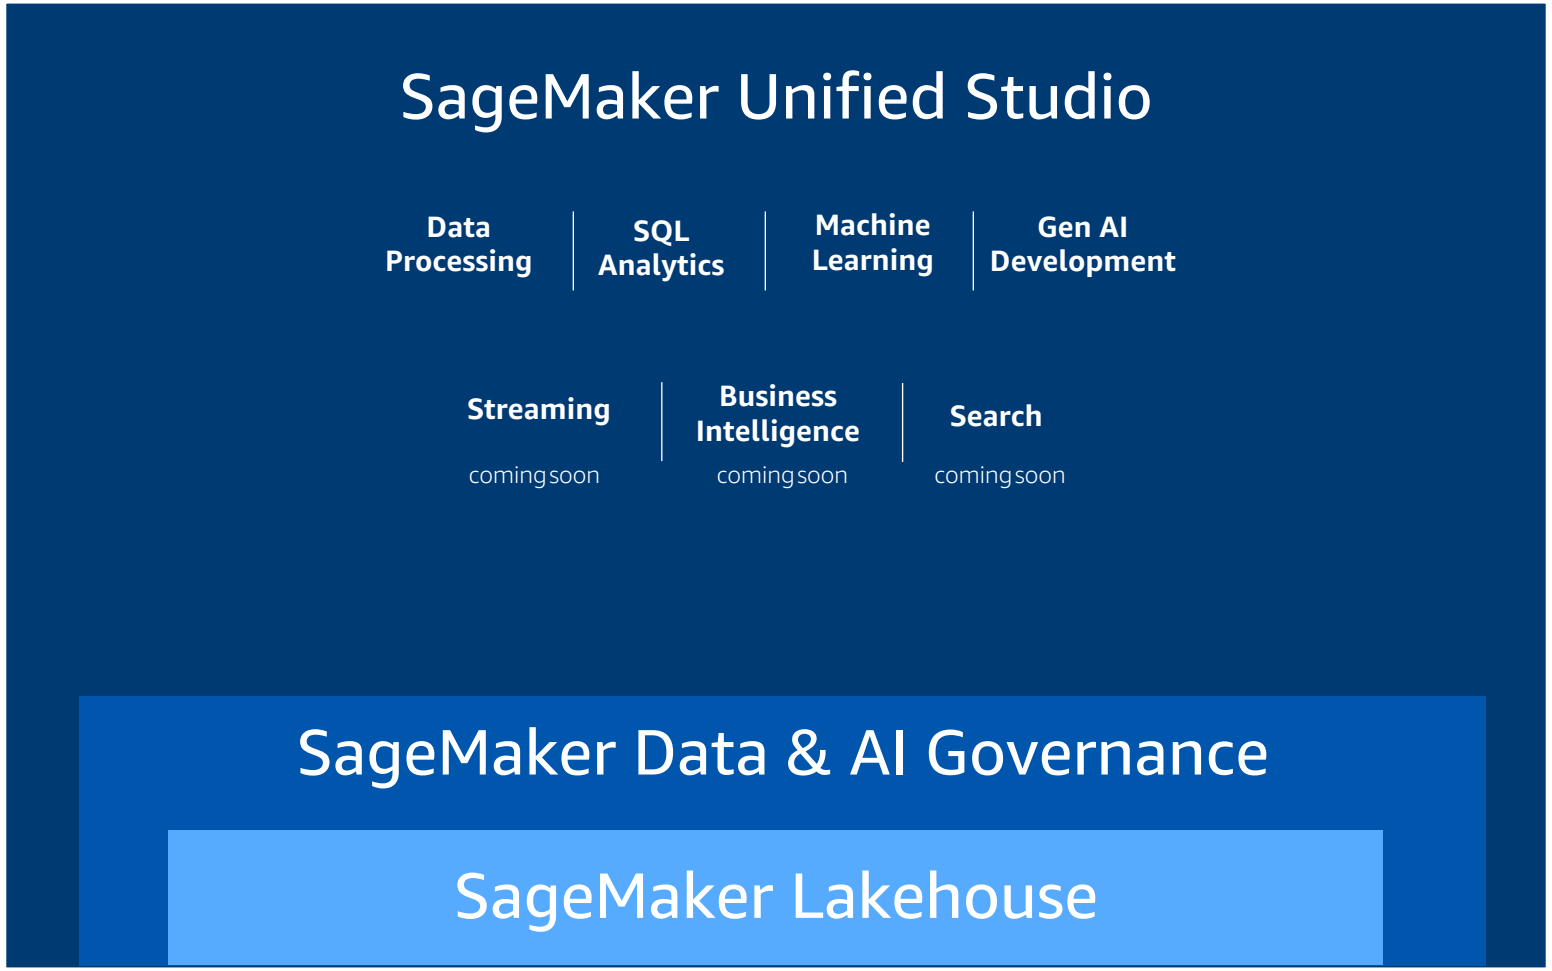

# Amazon SageMaker

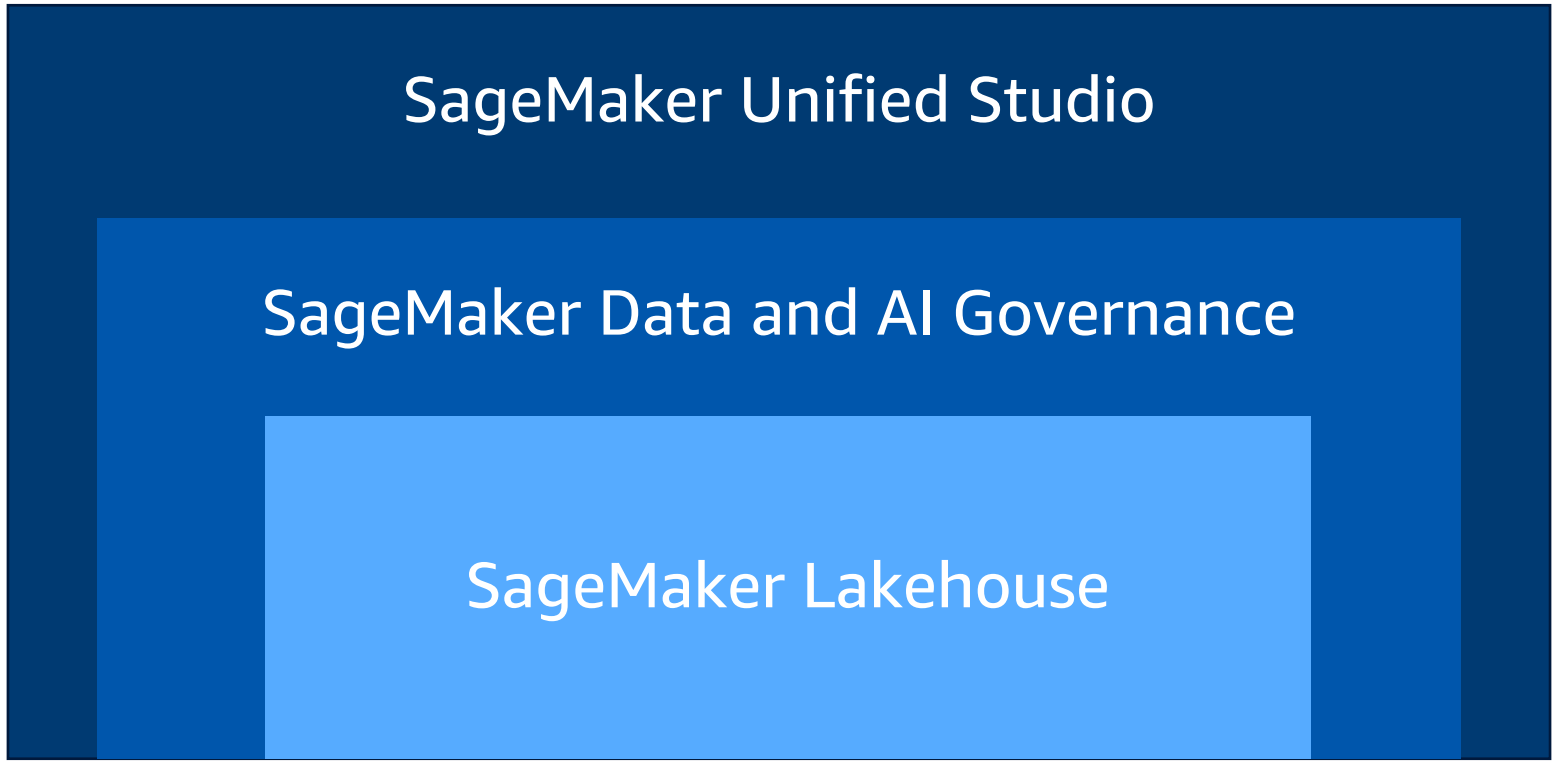

# Research Administration

# What is research administration?

*The coordination, management, & oversight of research activity across the sponsored research projects lifecycle -- from opportunity identification to project closeout.*

## People

Research and IT Leaders
Department Leaders & Staff
Compliance Officers
Pre- and Post-Award Admins
IT Staff

## Processes

Research Development
Pre-Award Management
Regulatory Compliance
Post-Award/Financial Management
Tech transfer & data sharing

## Systems

Grant management platforms
Financial systems & tools
Compliance systems
BI and reporting tools

# Modernizing RA with AWS

Strategy & Stakeholder Alignment | Data Infrastructure Modernization | Data Security & Governance | System Integrations | Identity & Access Management

## Grant proposal development

Recommender engines for collaborator and opportunity identification

GenAI for scientific literature reviews, proposal development, & evaluation

Intelligent document processing

## Contract & agreement reviews

Automated compliance analysis and risk assessment

GenAI for contract summarization

GenAI for contextual redlining and remediation recommendations

## Sponsored research analytics

Dashboards & visualizations with natural language query

Process metrics, funding trends, and impact tracking

Predictive decision support using historical insights, market trends, & scenario building

## Sponsored research Q&A chatbots

Federal, state, sponsor, and university policy evaluation

Process, procedure, form, and points of contact discovery

System & tool identification for proposal development and grant administration

## Grant financial management

Grant budget tracking and forecasting

Automated invoice generation, reconciliation, and collections

Intelligent document processing

Financial report generation

# The Art of the Possible: GROW

**GROW: Grant Research Opportunity Wizard**

Streamline the process of searching for grants & getting started with initial proposal

Leverages Amazon Bedrock with Anthropic Claude and other AWS services.

*"Identifying federal agency opportunities through targeted search is easier… and the tool allows for further ideation on existing proposals"*

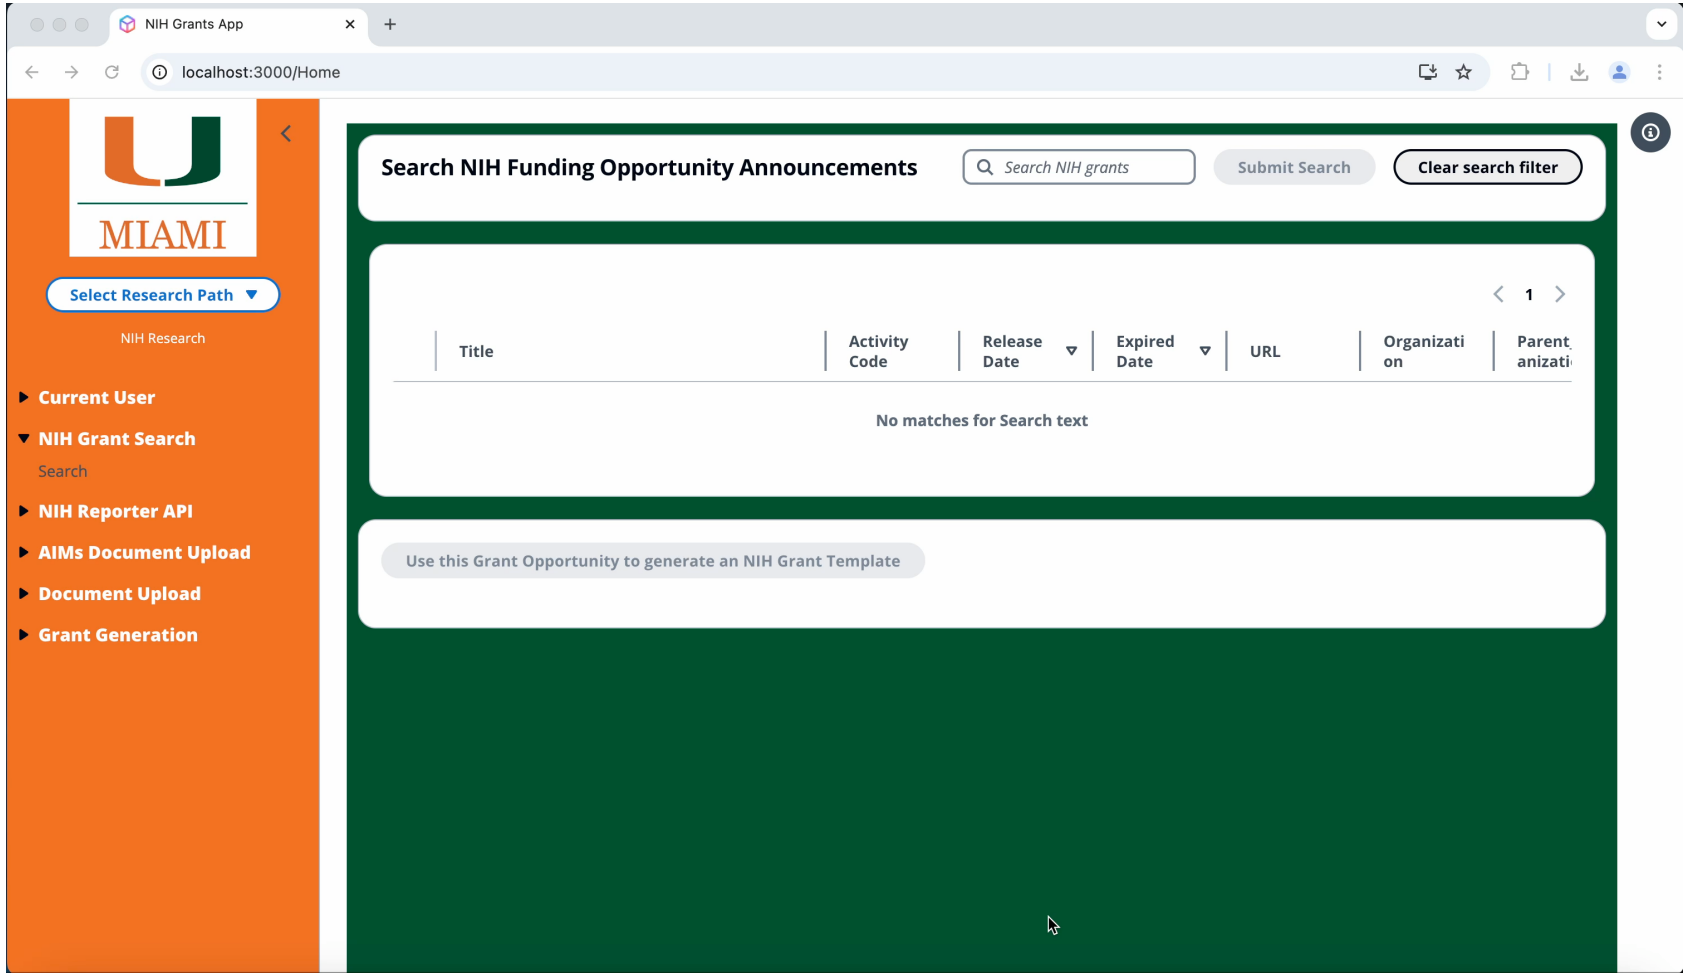


*GROW is on track to be open-sourced in Q2 2025*

# Wrapping Up

# Why AWS? Because AWS enables research

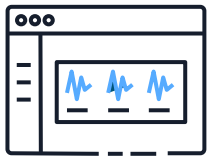

## Accelerate research
On-demand access to resources

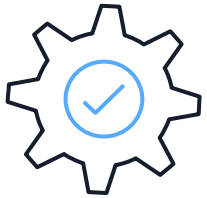

## Meet capacity demands
Scale to meet research requirements

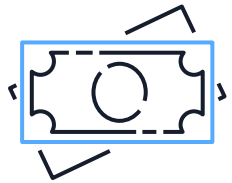

## Provide flexible resource access
Balance cost and performance

## Services beyond HPC
AI/ML, accelerators, quantum computing

## Match the pace of innovation
Scale to meet research requirements

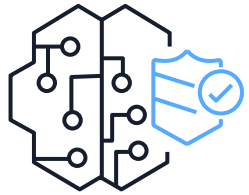

## Ensure security & compliance
Demonstrate accountability

# Thank you!

**Greg Ritter** (he/him)

Strategy Lead – Research Data, Scholarly Communication, & Libraries

AWS Higher Education

gritter@amazon.com

**Please complete the survey for this session**

## Research Track

The Evolving Research Landscape and the AWS Cloud



aws Learning Days
State, Local, and Education