

AWS State, Local, and Education Learning Days

Boston, MA

Research

10:15am – 11:15am

100
level

Transforming Your Research with AWS:

AWS helps universities solve complex research challenges with scalable cloud computing and research credits.

11:30am – 12:30pm

100
level

High Performance Computing (HPC) on AWS:

Cloud bursting enables researchers to dynamically scale HPC workloads between on-premises and cloud environments

1:30pm – 3:00pm

200
level

Workshop: Research on AWS:

Hands-on introduction to AWS cloud services for scientific workflows, from basics to advanced research computing.

3:15pm – 4:15pm

200
level

Building a robust research environment:

Secure research environments help institutions meet global compliance requirements and protect sensitive data for funding agencies.





Transforming your research with AWS

Laura Boykin Okalebo, Ph.D.

Sr BDM, Research Computing Solutions
AWS, US Higher Education
lokalebo@amazon.com

Agenda

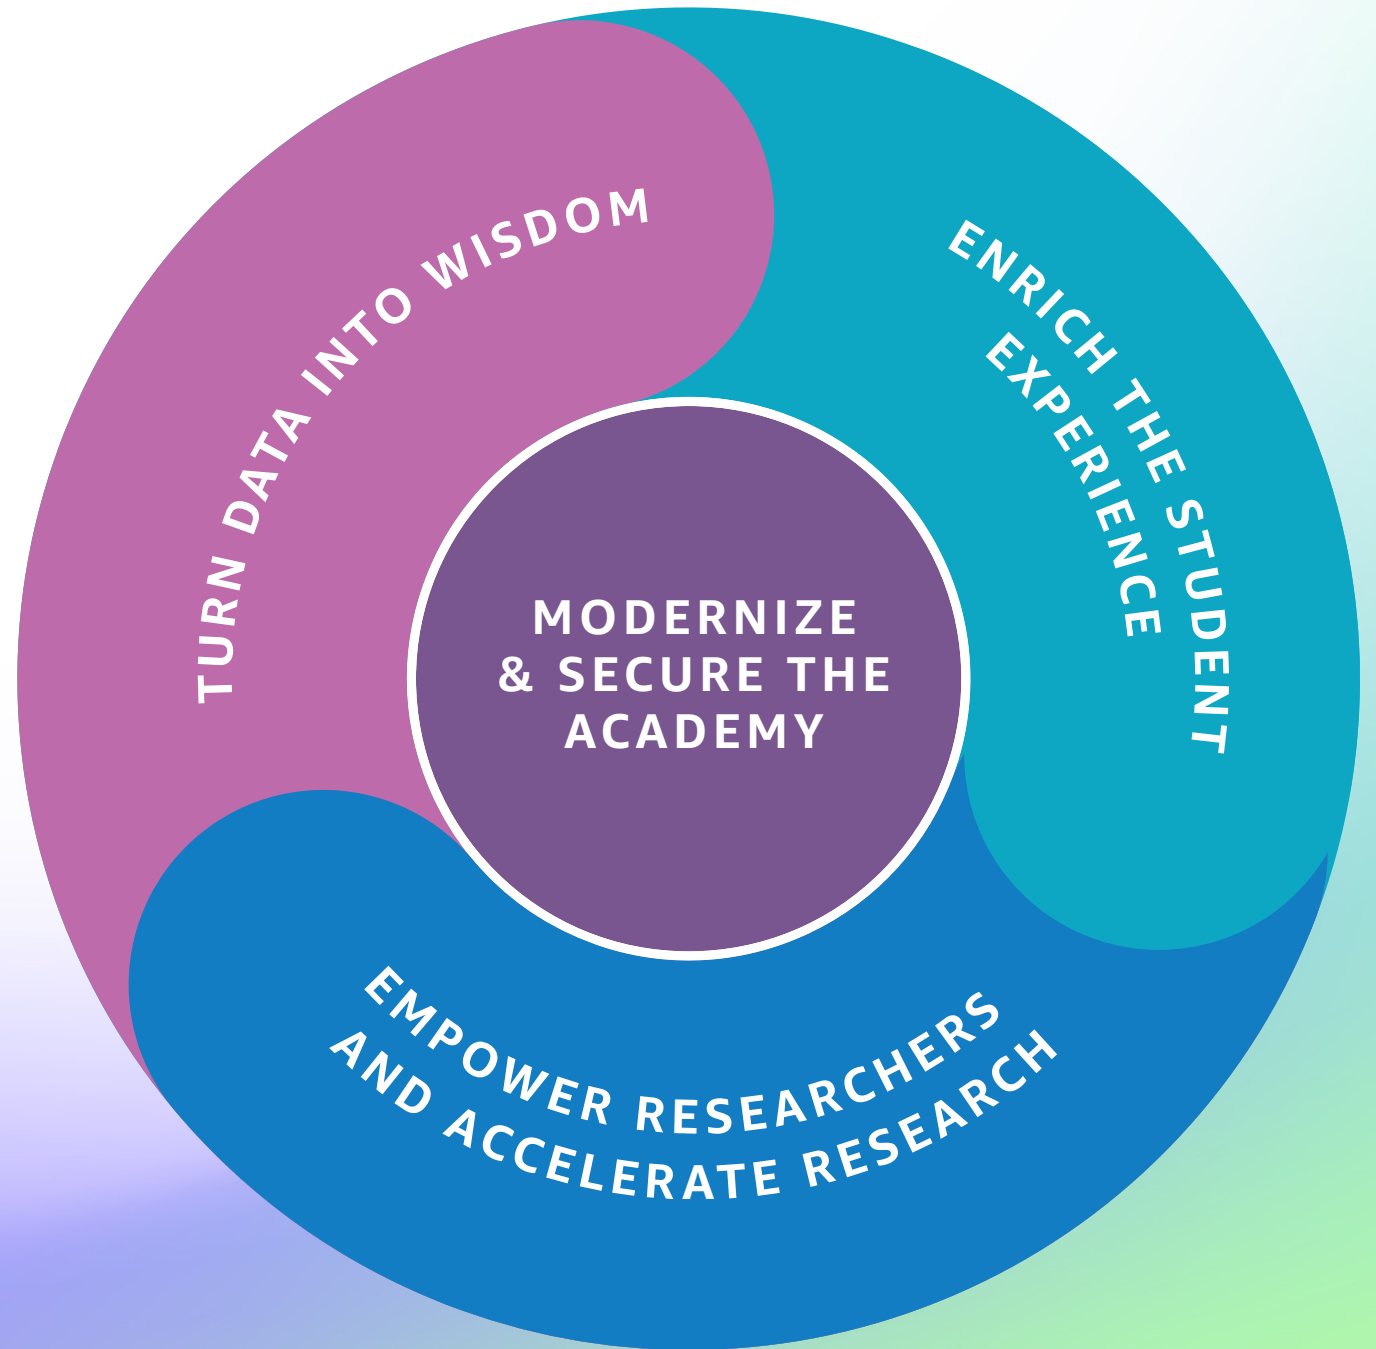
- AWS for Research
 - Customer stories
 - Activity where you tell us why AWS for your research (~25 min activity)
 - Why AWS for Research?
 - Performance and Price
 - High Performance Computing
 - Security & Compliance
- Questions & discussion



Our Vision:

We see a world where education is always available, personal, and lifelong for every learner.

Four Pillars



AWS for Research



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



Customer Stories

Emory University – AI.Humanity Initiative

In 2022, Emory started the AI.Humanity initiative to study AI's societal effects and shape its ethical development. The program aims to:

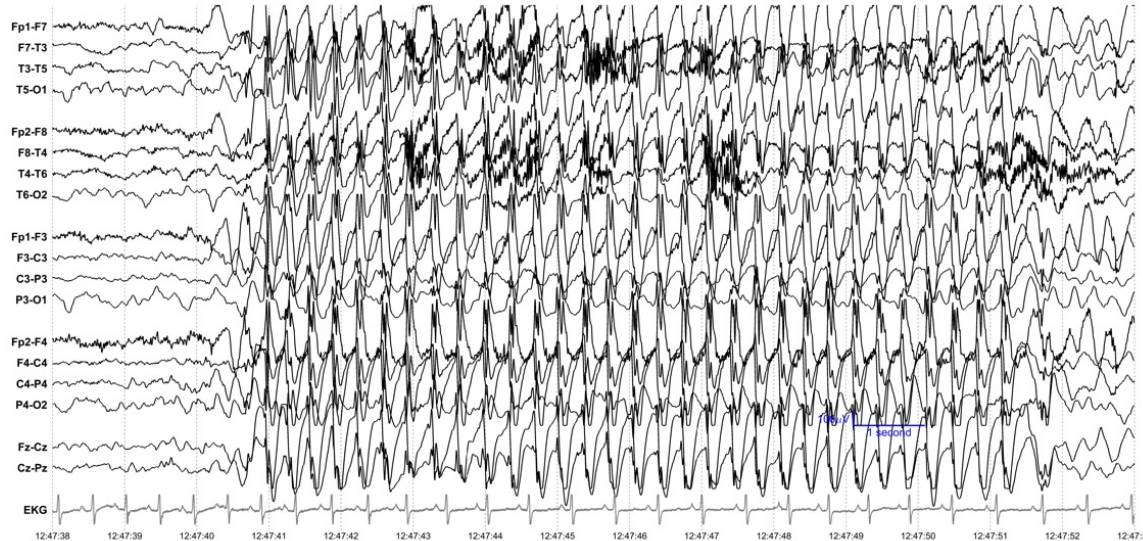
- Hire up to 60 new faculty
- Build a scholarly community
- Create AI-focused education programs
- Promote ethical AI use

Emory's goal is to become a leader in ethical AI and attract students and faculty interested in AI's transformative potential.



Emory evaluated various HPC solutions to meet faculty needs efficiently and cost-effectively. They chose a cloud-based HPC on AWS as their initial step, offering scalable compute, modern GPUs, robust security, and cost savings.

Harvard Medical School – Brain Science




“We intend to offer the machine learning EEG interpretation models as an online service, where sites can upload their EEG data and get their results back within seconds...This is where the cloud is going to be crucial.”

– Dr. Brandon Westover, Associate Professor, Neurology, Massachusetts General Hospital and Harvard Medical School

EEGs are vital diagnostic tools for brain disorders like seizures, stroke, and sleep issues. However, they're underutilized due to a shortage of EEG experts and limited automatic interpretation methods, despite their potential to provide crucial information.

U Penn – Enabling Alzheimer's Research

Penn researchers initially tested AWS for genomic computations with 36 exomes. This project evolved into NIAGADS DSS, one of the largest Alzheimer's genomic databases, now powered by AWS.



NIAGADS utilizes various AWS services for genomics data management, including S3, EC2, EFS, and EBS for storage and processing. It also employs AWS security services like CloudTrail, GuardDuty, and Security Hub for compliance and monitoring.

St. Louis University – Applied Economic Research



“We were trying to understand how we could put this data to use efficiently and quickly, because we didn’t want researchers wasting their time in pre-processing, They could build out use cases and continue their research if we could solve this problem for them. AWS gave us a dedicated solutions architect who would do whatever they could to help us succeed. This was key.”

– Shruthi Sreenivasa Murthy, Assistant Director of Research Computing

Social scientists need large datasets, like anonymized cell phone data, to study major social issues. While obtaining data was straightforward, the team faced challenges in managing 450 terabytes of data for storage, cleaning, and processing.

Children's Hospital of Philadelphia – FPGA

1,000 diverse pediatric genomes were processed into useable data files in two hours and twenty-five minutes



World Record for processing Genomes using FPGAs

Deployed on 1,000 Amazon EC2 F1 instances

One of the largest cohorts for this demographic that has been sequenced to date

Utilized Edico Genome's DRAGEN™ Genome Pipeline

Baylor U/Pontifica Universidad Católica del Perú (PUCP) – NLP for the Hearing Impaired




“We are training a machine learning model that will recognize the sign and translate to Spanish and vice versa. Using the cloud lets our team accelerate their progress which will allow the project to scale and incorporate more and more vocabulary when we launch our application..”

– Gissella Bejarano, PhD Computer Science and postdoctoral researcher at Baylor University


Cloud-powered deep learning has improved AI applications for vision and speech, potentially aiding 70 million people with disabilities. AWS is supporting a project at Pontifica Universidad Católica del Perú (PUCP) in Peru to translate Peruvian Sign Language using AI technology.

Why AWS for Research?


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

Activity: Why AWS for your research?



Individual Reflection (5 min)

- 🤔 Grab a post-it note
- ✍️ Write down your reasons for using AWS in research



Table Discussion (10 min)

- 💭 Share your thoughts
- 👥 Discuss various reasons with your group



Group Share-Out (10 min)

- 📣 Present your table's key points
- 🔄 Learn from other groups

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.

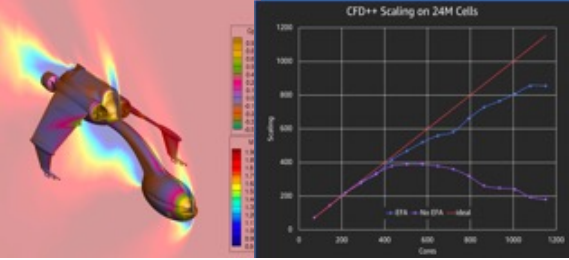
Genomics & life sciences	Climate & geospatial	Add'l disciplines
Research administration		
Research data management		
Research computing & HPC		
Security and compliance		

Performance and Price

AWS flexible compute price/performance selection

Flexible compute to maximize performance


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



CFD++ Scaling on 24M Cells

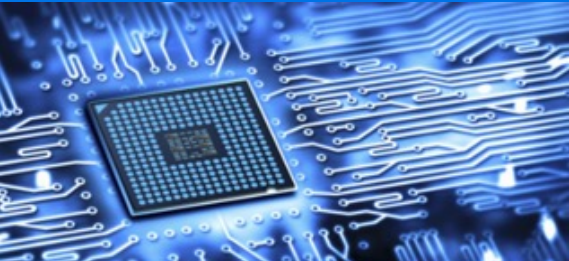
A 3D visualization of a wing with a color-coded flow field, and a line graph showing performance scaling on 24M cells. The graph plots 'Cells' on the x-axis (0 to 1200) and 'Performance' on the y-axis (0 to 1200). Three lines represent different configurations: '10k', '10k-10k', and '10k-10k-10k'. The '10k-10k-10k' line shows the highest performance, reaching approximately 1100 at 1200 cells.

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



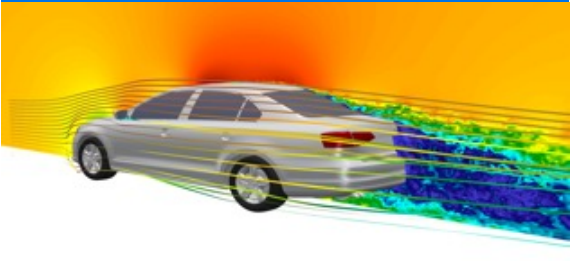
A 3D rendering of a car chassis, showing the internal structure and components. The rendering is displayed in a software interface with various toolbars and panels.

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



A close-up image of a microchip on a circuit board, with glowing blue lines representing data flow or connections.


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



A 3D rendering of a car with a color-coded airflow simulation around it, showing the aerodynamic flow patterns.


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.

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




MORE THAN

800+

INSTANCE TYPES

for virtually any
workload and
research need

High Performance Computing



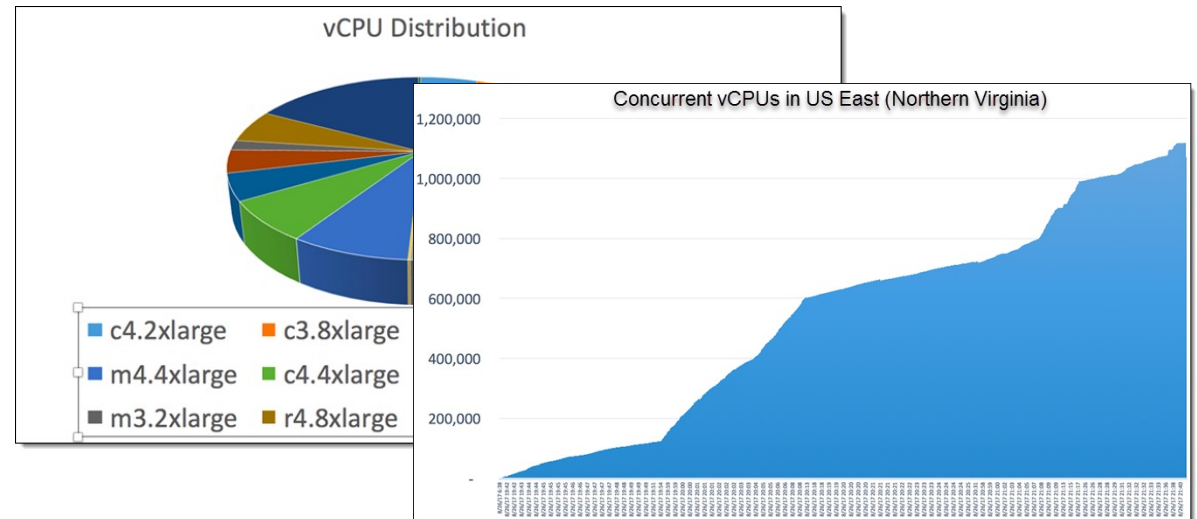
HPC on AWS

Flexible configuration and virtually unlimited scalability to grow and shrink your infrastructure as your HPC workloads dictate, not the other way around

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.



"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.

Security and 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



Funding withheld

non-compliance

NIST / CUI

DFARS

CMMC

Responsible use


Reputational Damage

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.

Summary


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!

Laura Boykin Okalebo, PhD (she/her)

Sr BDM, Research Computing
AWS, Higher Education

Email lokalebo@amazon.com

 @laura_boykin

Up Next In This Room:

11:30am – 12:30pm

100
level

**High Performance
Computing (HPC)
on AWS:**

Cloud bursting enables
researchers to
dynamically scale HPC
workloads between
on-premises and cloud
environments

**Please complete the survey
for this session**



Research Track

Transforming your research
with AWS

aws Learning Days
State, Local, and Education