



AWS RESEARCH – 1<sup>ST</sup> OCTOBER 2024

# How to Use AWS Cloud in the Green Economy Era for Research

**AWS Education & Research  
Team**

Roberta Piscitelli, PhD, MBA - [piscitr@amazon.com](mailto:piscitr@amazon.com)

# AWS Global Infrastructure

AWS REGIONS, LOCAL ZONES, EDGE LOCATIONS, AND GLOBAL BACKBONE



## REGIONAL EXPANSION

● Available today: 32 Regions

● Coming soon: 5 Regions


# Trends We have Observed



Research is strategic for growth, recruiting and knowledge



New and different types of users need large scale computation

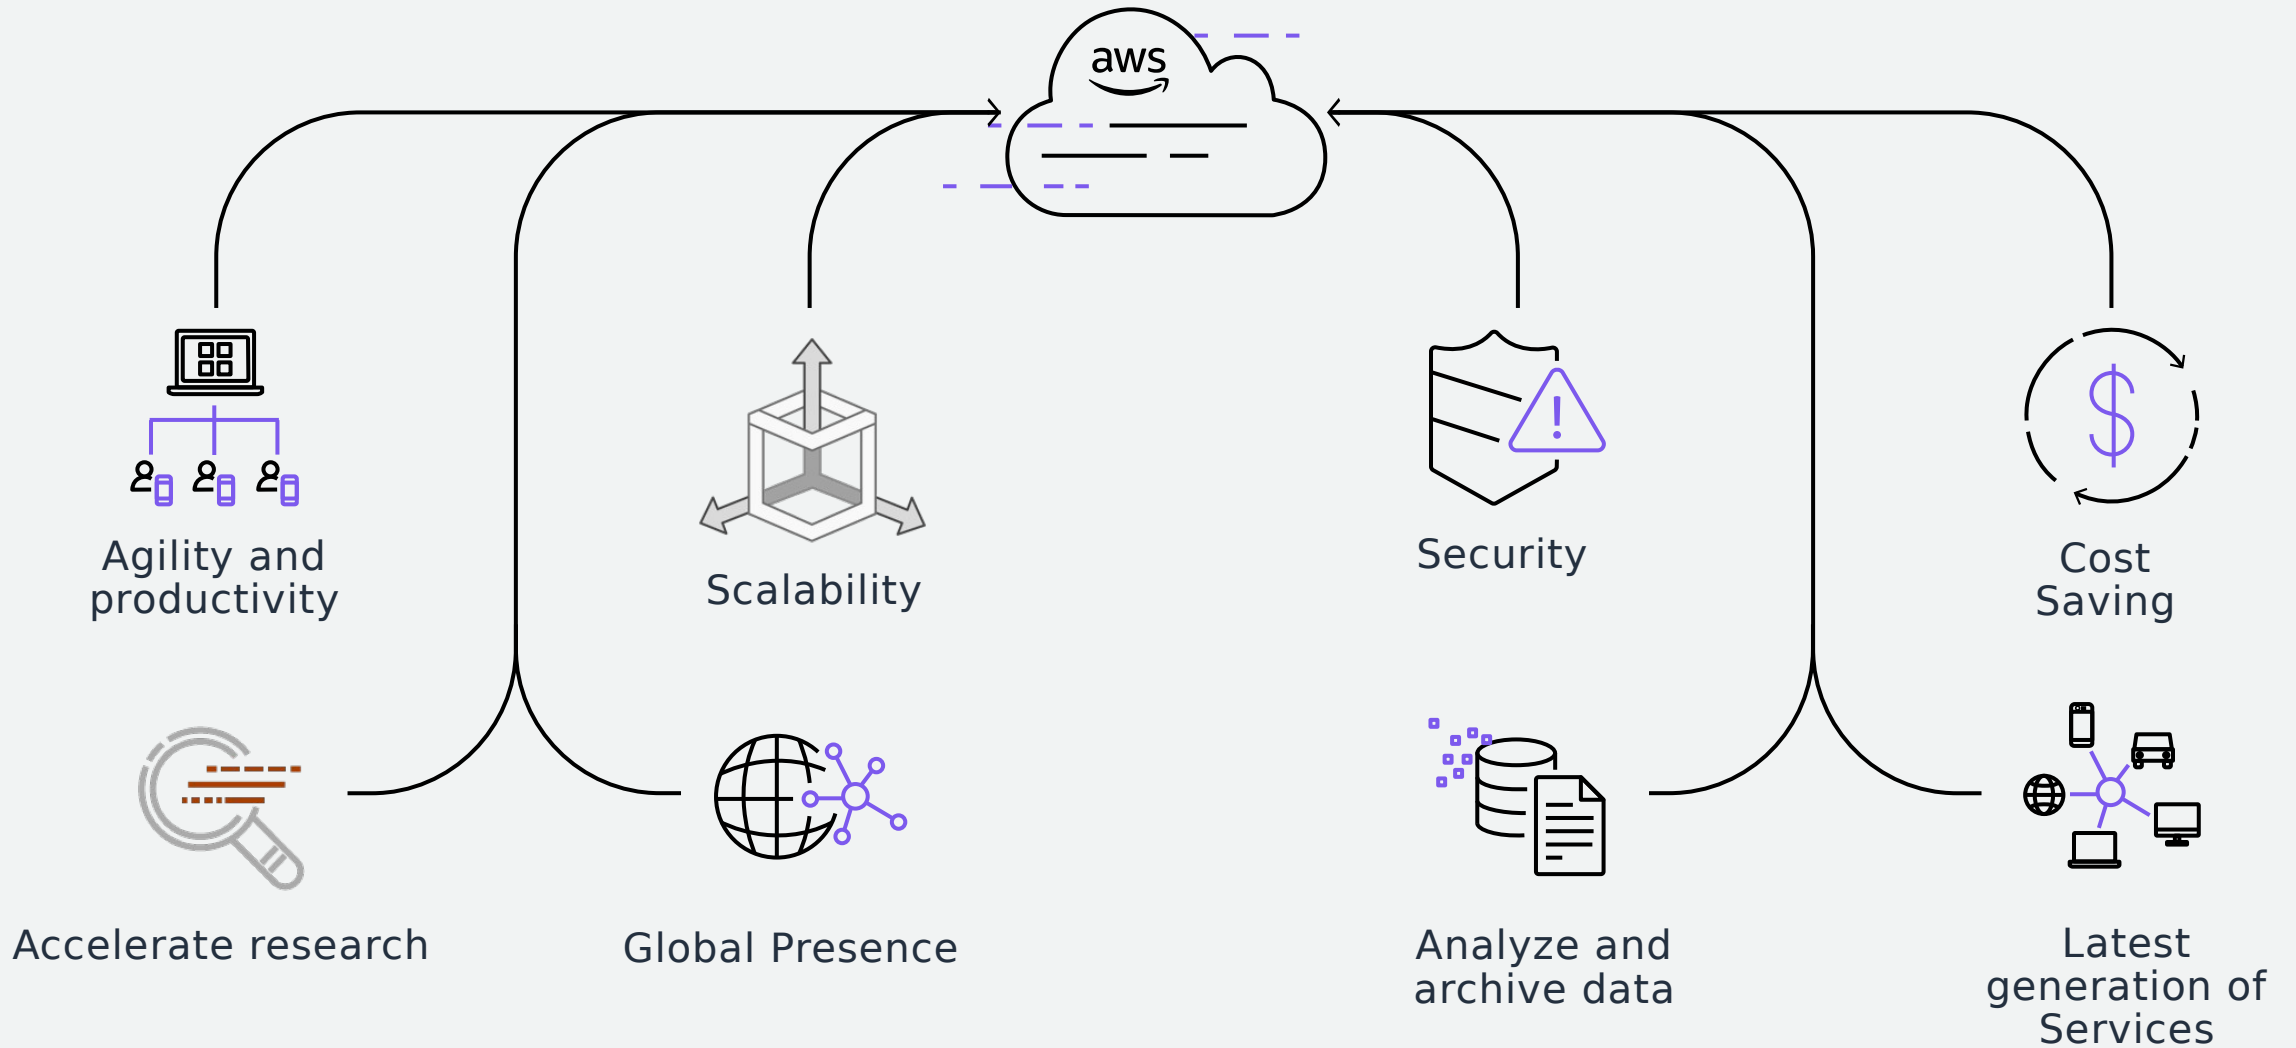


Sustaining on-premise while keeping pace with technology



Increasing grant requirements around secure research

# Why using AWS Cloud for Research





# A data-driven approach to sustainability strategy, enabled by cloud computing



Sustainability **OF** the cloud

Delivering a sustainable IT fleet – taking advantage of the cloud and AWS efficiency through migration



Sustainability **IN** the cloud

Optimizing workloads on AWS with the Sustainability pillar of the Well-Architected Framework



Sustainability **THROUGH** the cloud

Deploying cloud-based solutions and advisory support to accelerate customer sustainability objectives

# Sustainability **OF** the Cloud

**Opportunity to reduce impact** through migration



# Europe: Carbon reduction opportunity

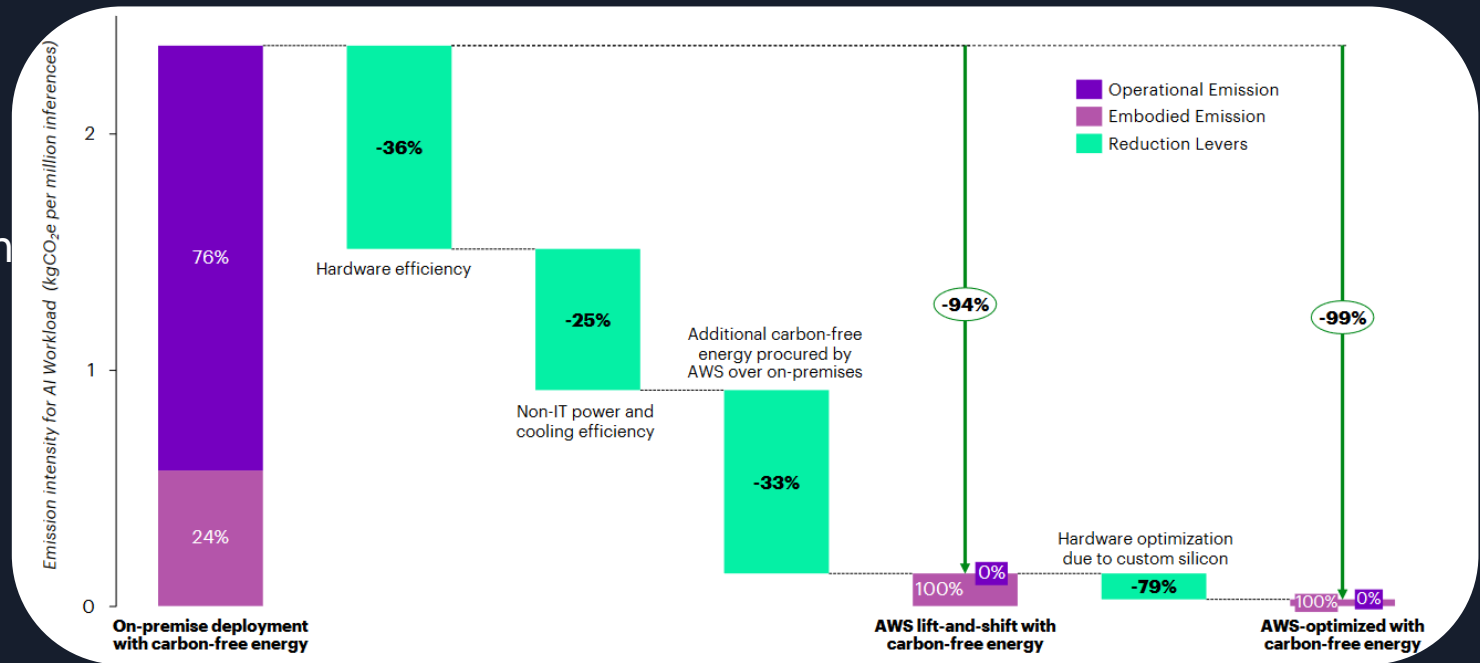
AWS enables up to **99% lower carbon emissions** when compute-heavy workloads are optimized on AWS vs. on-premises

**94%** lower carbon emissions when compute-heavy workloads are run as-is on AWS vs. on-premises, due to:

**36%** from efficient hardware with improved utilization

**25%** from power and cooling efficiencies

**33%** from additional carbon-free energy procurement



Find the report at [https://sustainability.aboutamazon.com/carbon\\_reduction\\_aws.pdf](https://sustainability.aboutamazon.com/carbon_reduction_aws.pdf)



# Reducing carbon emissions across our global operations

**67** data centers constructed using either or both lower-carbon concrete and steel

---


We use hydrotreated vegetable oil, a renewable fuel, to power our backup generators. HVO can reduce emissions by up to **90%** over the fuel's lifecycle compared with fossil diesel.

---

Lower lifecycle impacts of silicon-based devices from materials to manufacturing

---

In 2023, electricity consumed in **22** AWS Regions was attributed to **100% renewable energy**





# Embracing a circular economy

**Design.** We eliminate excess materials, increase recycled and bio-based content, and plan for reuse from the start.

---

**Operate.** We extended the life of our servers and networking equipment from **5 years to 6**.

---

**Recover.** In 2023, **14.6M hardware components** were diverted from landfills by being recycled or sold into the secondary market for reuse.





# AWS water positive commitment

By 2030, AWS will return more water to communities than we use in our direct operations.

As of 2023, AWS is **41%** of the way toward water positive.

A collection of stylized water droplets of various sizes, some light blue and some white with a blue outline. Each droplet contains a small orange AWS logo. The background is dark blue.

# Water stewardship

**Water efficiency.** 0.18 liters of water per kilowatt-hour water use efficiency for AWS data centers

---

**Sustainable water sources.** 24 data centers globally use recycled water for cooling

---

**Water reuse.** Reuse discharged water from our data centers in communities

---

**Water replenishment.** Amazon returns 3.9 billion liters of water to communities each year from replenishment projects completed or underway

# Sustainability **IN** the cloud

**Optimizing for sustainability**

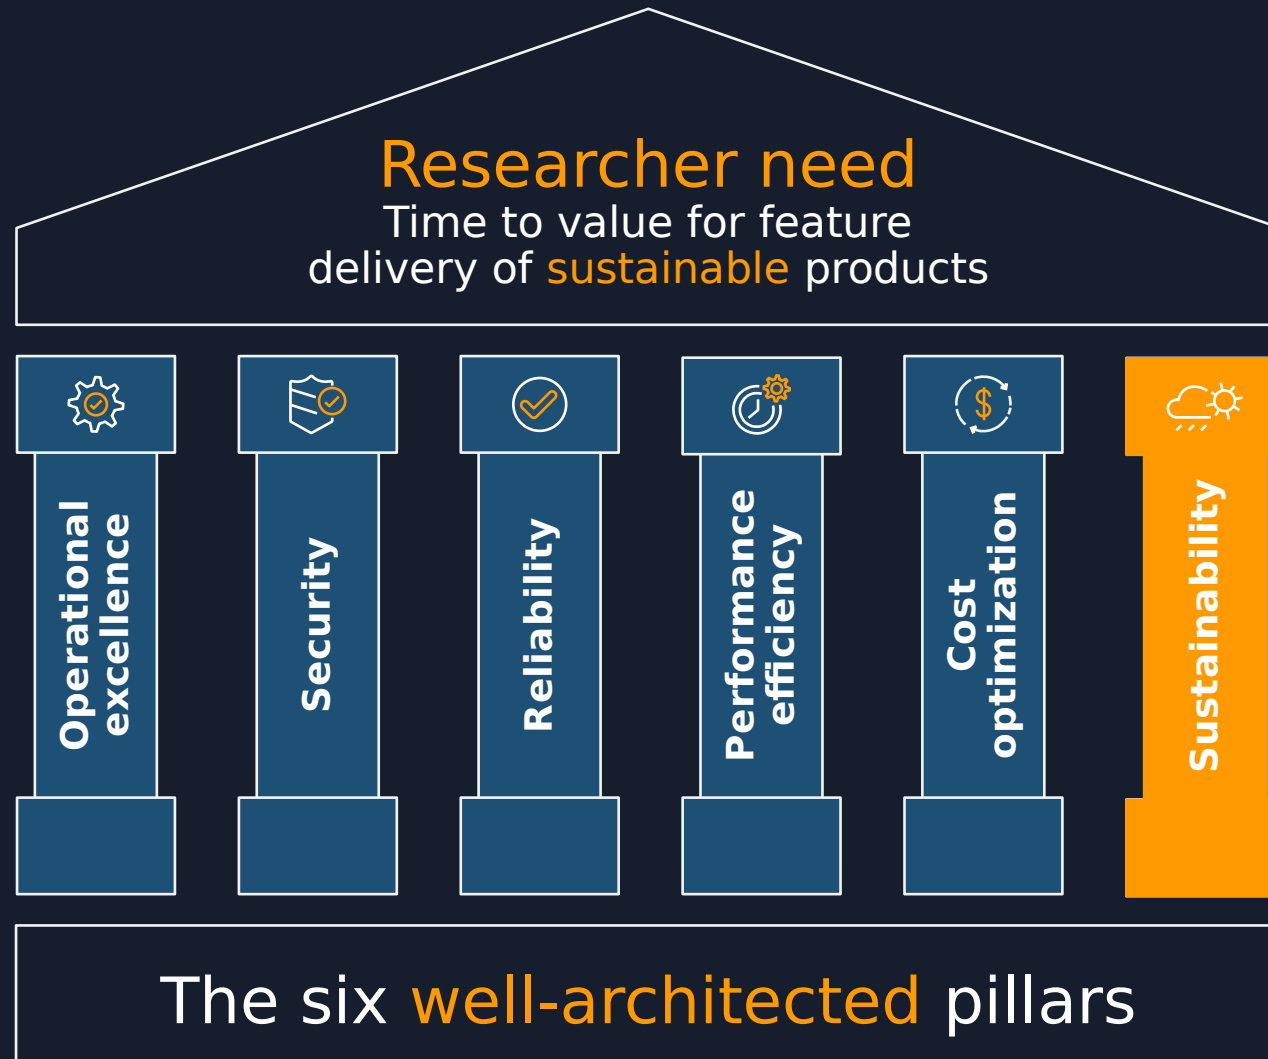
**Sustainable Pillar of the Well-Architected Framework**

**Customer carbon footprint tool**






# Sustainability Pillar of the Well-Architected Framework




# Focus domains of the AWS Well-Architected Sustainability Pillar




Region selection



Alignment to demand



Data



Hardware & services



Software & architecture



Process & culture

# Graviton, Inferentia, and Trainium


Graviton-based Elastic Compute Cloud instances use up to **60% less energy** than comparable Amazon EC2 instances

---

Models built on Trainium result in **energy-consumption reductions of up to 25%** versus comparable instances

---

Our Inferentia2 machine learning accelerator delivers up to **50% higher performance per watt**



# New LLM development and refinement UAE Technology Innovation Institute



## CHALLENGE

TII researchers needed an easy way to develop, iterate, and distribute Falcon-40B LLM. Experimenting at such massive scale was slowed by limited computational resources and complex on-prem infrastructure. Researchers sought more flexibility to efficiently enhance, evaluate, and provide access to Falcon under dynamic capacity requirements to empower their team's foundation model research.

## SOLUTION

TII utilized SageMaker's managed machine learning infrastructure and tools to train, host, and deploy versions of Falcon with greater speed and experimentation. SageMaker provided on-demand access to training clusters along with modular deployment options for easy distribution to varying client workloads—key for customizing model access.

## OUTCOME

- ✓ TII successfully implemented Falcon-180B by using SageMaker and custom innovation. This allows TII's contribution within UAE's 2031 National AI Strategy, fostering economic growth and social progress. Releasing UAE's Falcon 180B, World's Top-Ranked Open Source AI Model will further encourage AI academic research and scientific collaboration.



# Large Models built on AWS



Cibrain develops Lince Zero—the first large language model (LLM) optimized for Spanish using Amazon SageMaker



Stability AI will build AI models on compute clusters with thousands of GPU or AWS Trainium chips, reducing training time and cost by 58%



BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance, utilizing AWS infrastructure services

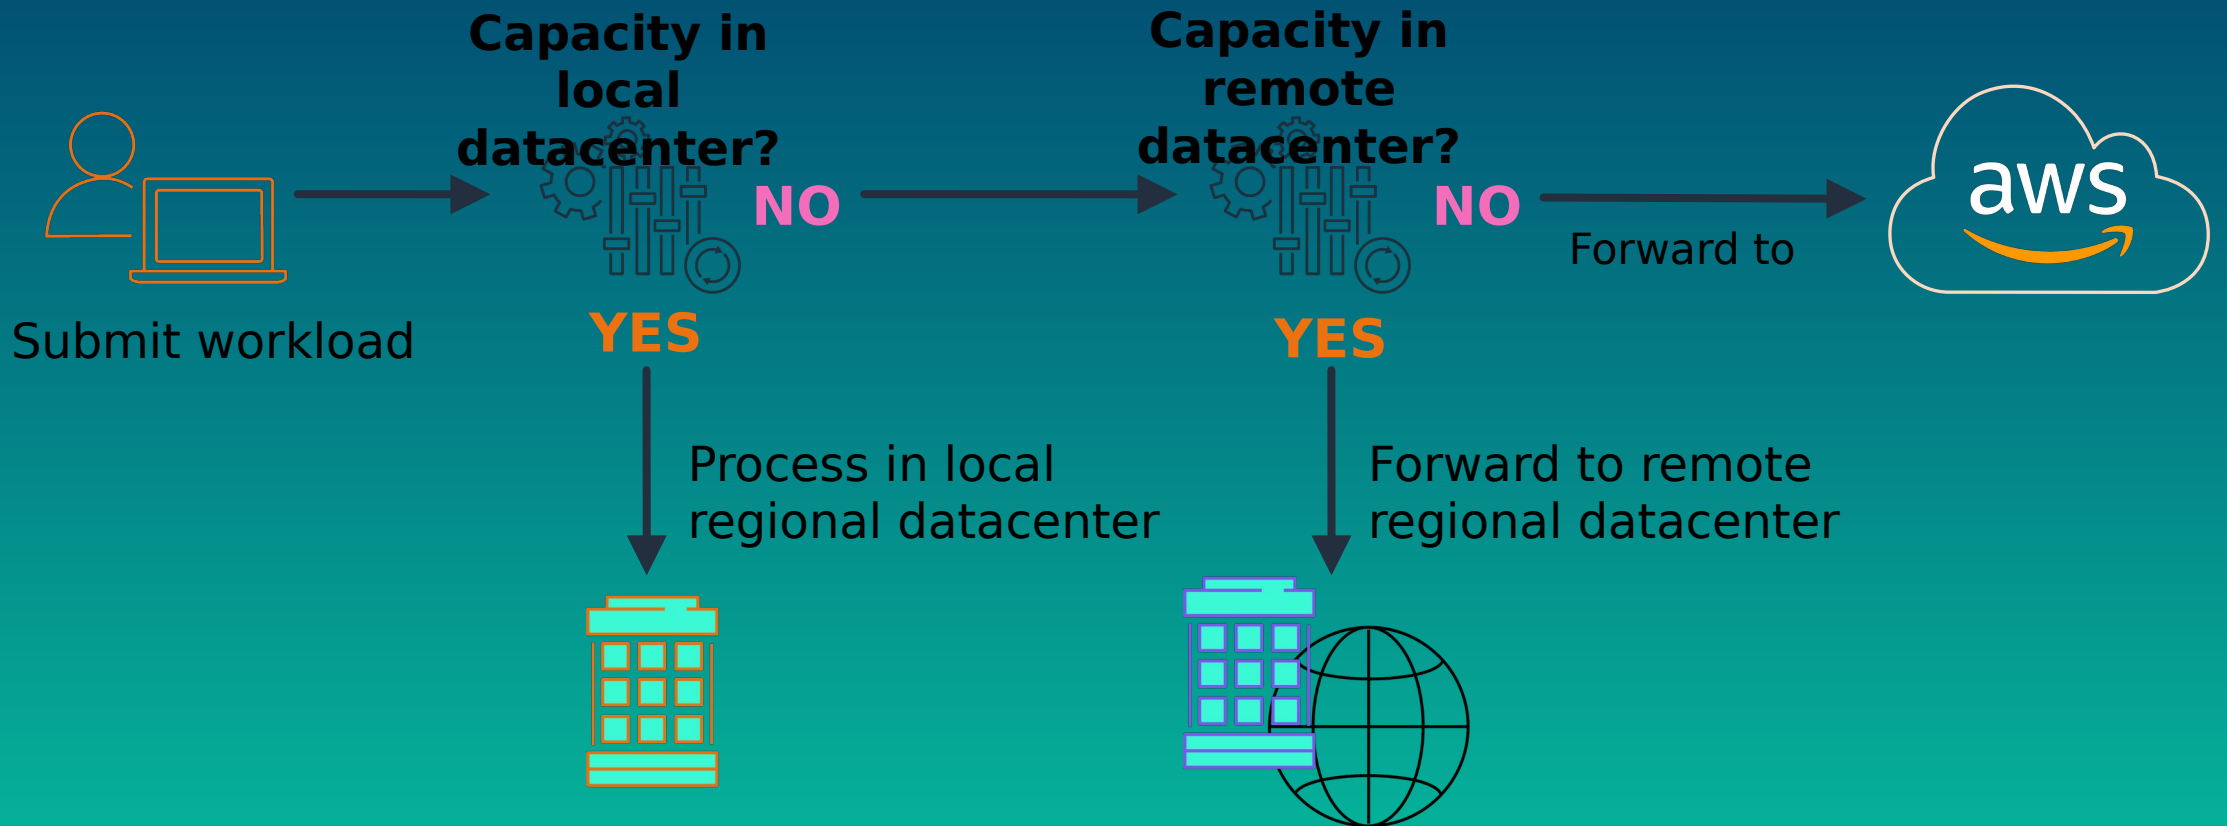
# Cloud Bursting

## CLOUD BURSTING ON AWS


Cloud Bursting

PoC and Pilot

All-in Migration



# Optimize for sustainability: Quick wins



Improve **energy efficiency** by switching to Graviton-based instances



Choose **serverless** when possible

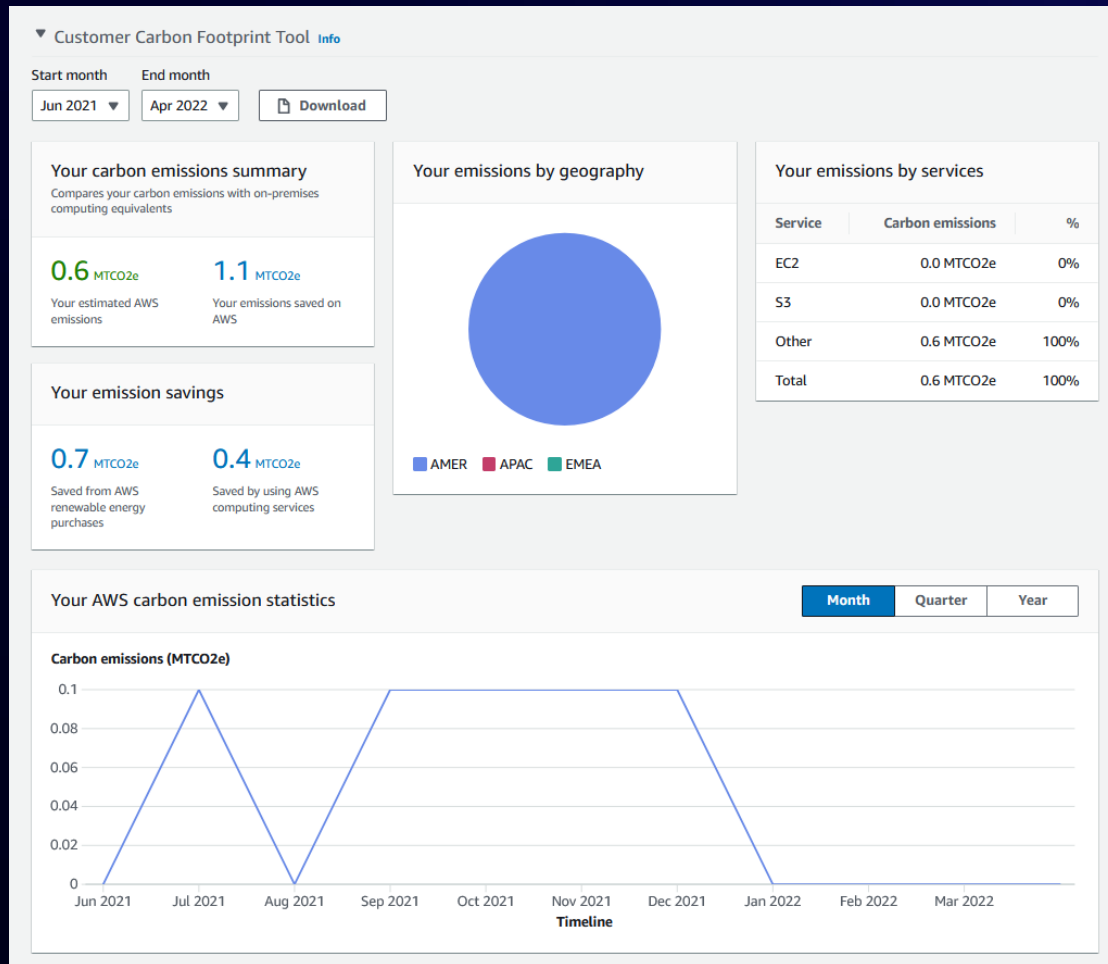


Integrate **AWS Instance Scheduler** to shut down and terminate when not in use



Use **AWS Compute Optimizer** for right-sizing recommendations of workloads

# AWS customer carbon footprint tool



**Calculate** carbon emissions generated from your AWS workloads

**Understand** historical carbon footprint and review changes in emissions over time

**Forecast** reflects Amazon's achievement of 100% renewable energy in 2023

# Sustainability **THROUGH** the Cloud

Leverage **AWS technologies and data services** to solve  
sustainability challenges



# AWS Solutions Approach



## Customer Use Case Identification & Prioritization

**Innovation Programs - Think Big, Work Backwards**

**Prototype Builds & Technical Proof-of-Concept**

**Proof-of-Value & Strategic Business Advisory**



### **AWS Service**

**Purpose-built cloud products**

### **Partner Solution**

**Software, SaaS, or managed services from validated AWS partners**

### **AWS Solution**

**Ready-to-deploy solutions assembling AWS offerings, code, and configurations**

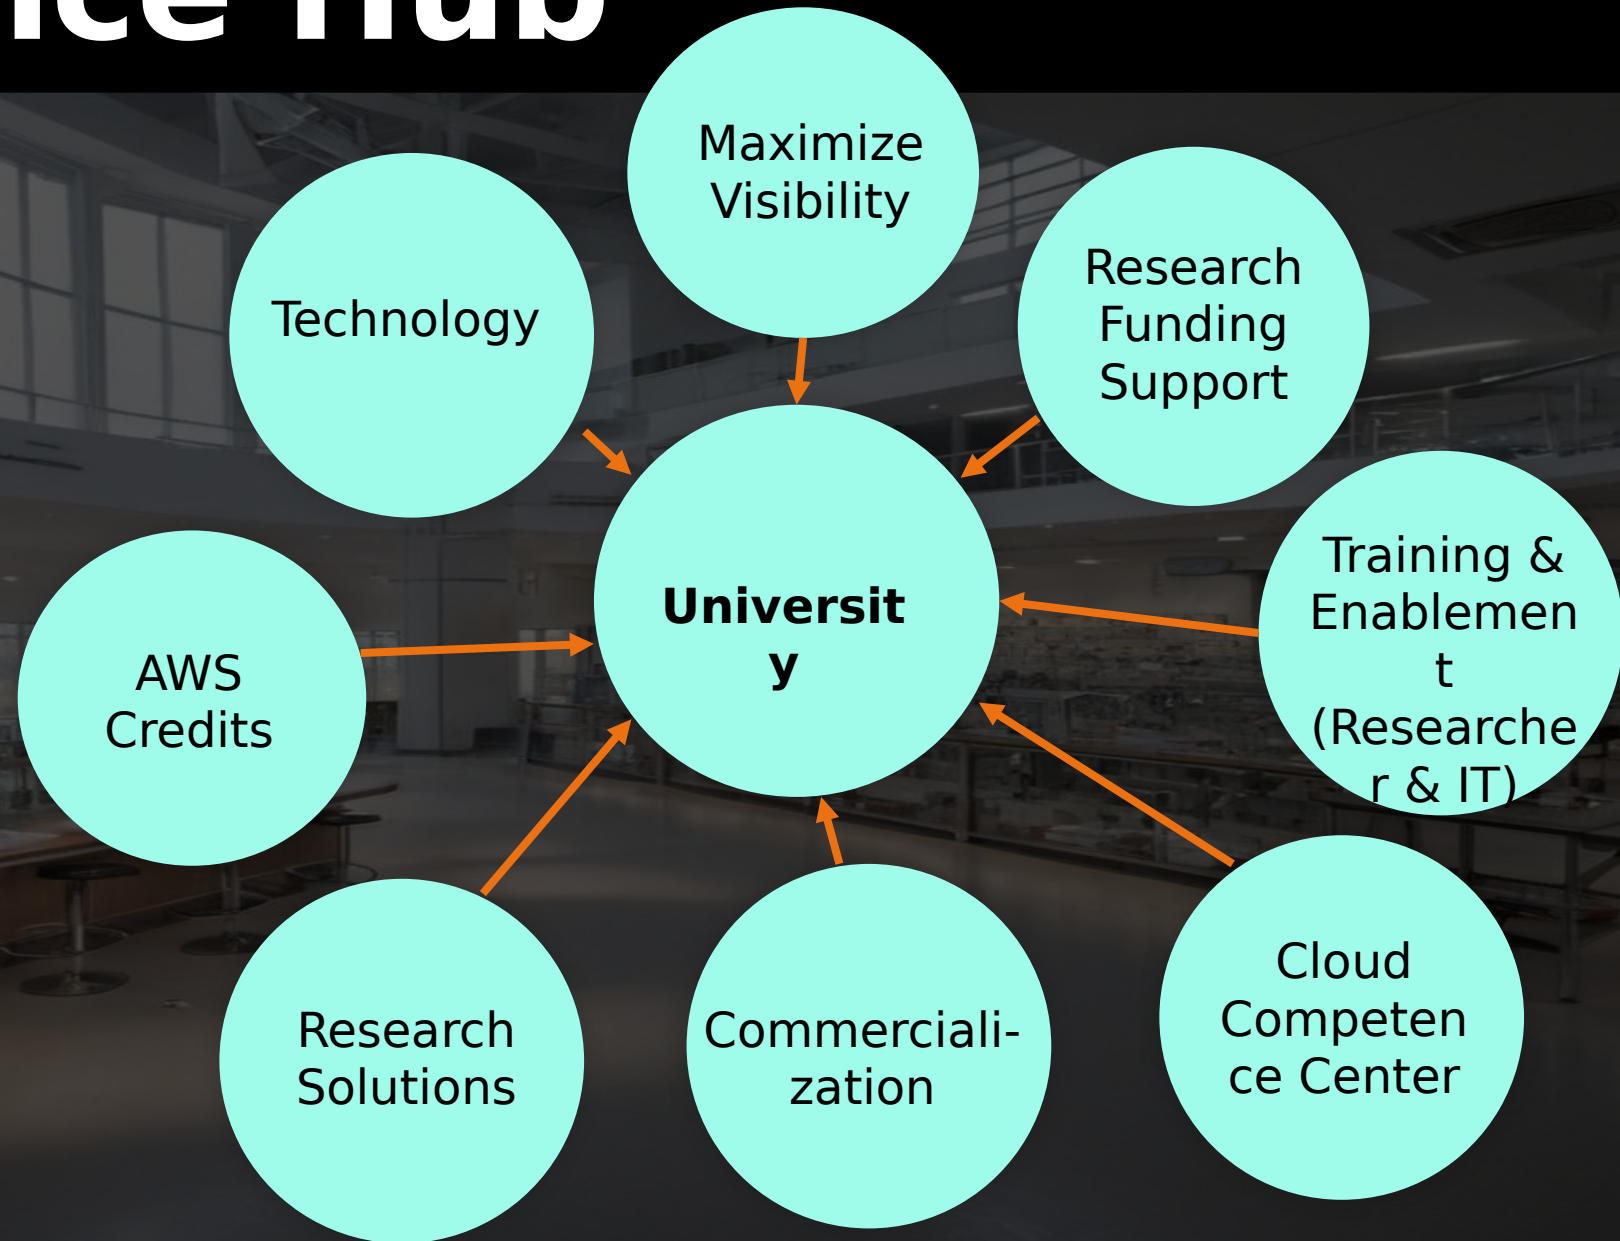
### **Guidance**

**Prescriptive architectural diagrams, sample code, and technical content**





# AWS Science Hub



# Solutions and services for the entire research process



- Research proposal support
- Letters of support
- Cloud economics
- Compliance documents

- OCRE Framework, Ubuntunet
- Deployment in minutes
- 200+ services
- Training
- Partner solutions
- Open source solutions (e.g. **Parallel Cluster**)
- Solution architects
- Professional services
  
- Single sign on
- Landing Zone

- AWS Snow Family
- AWS IoT for sensors
- AWS Ground Station
  
- Streaming data
- Event driven arch
  
- Databases
- Data lakes
- Data warehouses
  
- Bulk storage options

- **Interactive notebooks**
- **HPC (AWS Parallel Cluster)**
- **Machine learning**
- Containers
- Big data analytics
- Visualisation (NICE DCV)
- 400+ instance types
- GPUs, FPGAs, ARM, **Inferentia, Trainium**
- **Quantum computing**

- **Open Data Registry**
- Trusted Research Environments
- Content delivery network
- Long term low cost archival

- **AWS Marketplace**
- Building SaaS
- Partners
- **Proserve**



# Sustainability solutions are powered by data

Data is diverse, growing exponentially, and used by many applications. AWS storage and analytics services and data programs can help.

**Open Data  
Sponsorship  
Program**

**Amazon Sustainability  
Data Initiative (ASDI)**

**AWS  
Data Exchange**

*Image from Landsat 8 satellite, courtesy of the U.S. Geological Survey*





# Open data on AWS

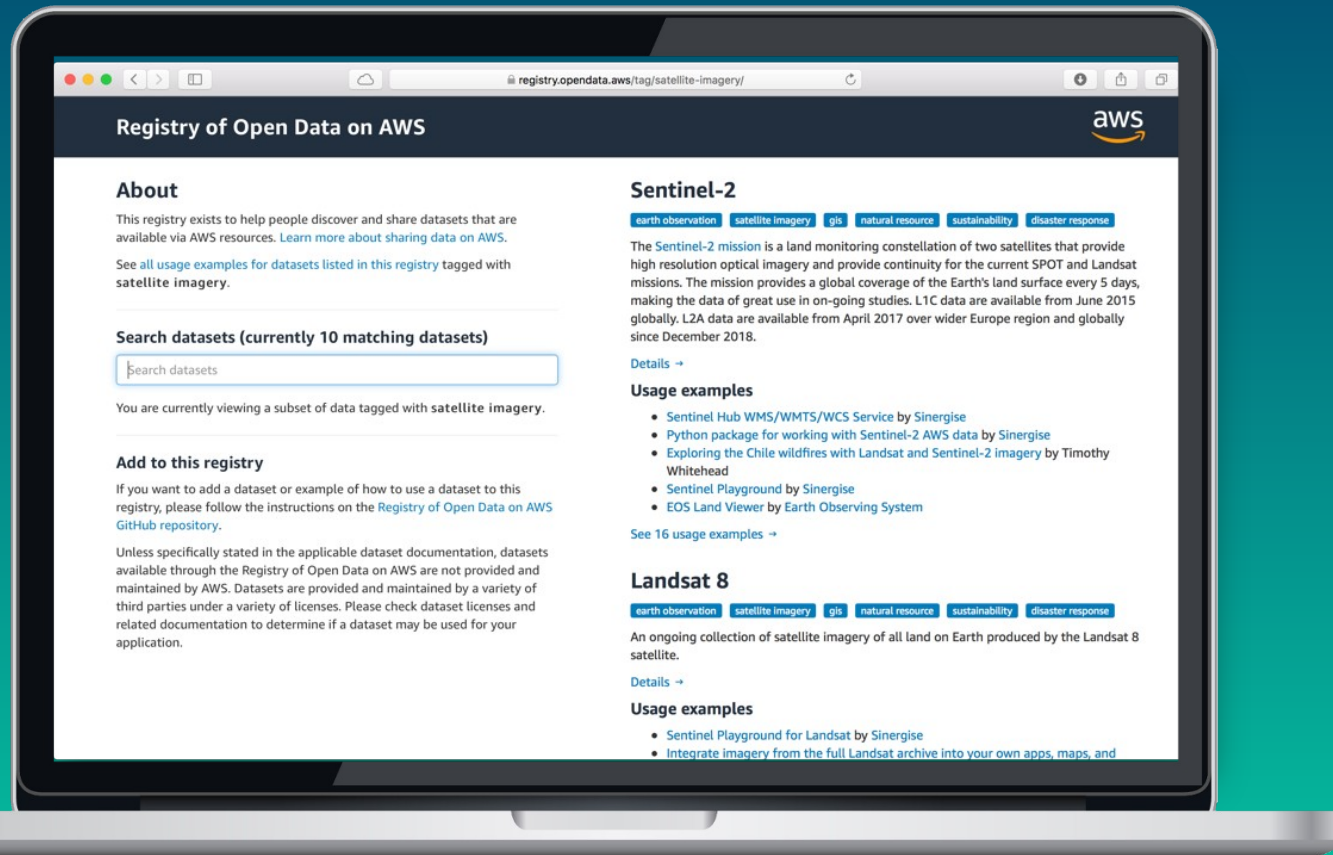
AWS hosts a variety of public datasets to lower the cost and improve the speed of research.

<https://registry.opendata.aws/>

## Examples

- 1000 Genomes Project
- The Cancer Genome Atlas
- International Cancer Genome Consortium
- Landsat 8
- Common Crawl
- SpaceNet
- OpenStreetMaps


... Regularly updated




# FastMRI: Accelerating MRI Scans with AWS and AI

- Dataset: 7 million+ brain and knee MRI scans
- Global Impact: Users from 79 countries, 11 publications
- Results: Up to 98% similarity to full scans in 15% of the time

Fully sampled



Undersampled



<https://aws.amazon.com/blogs/publicsector/nyu-langone-center-increases-mri-accessibility-cooperative-data-sharing-research/>



[Amazon Open Data Sponsorship Program](#)



© 2024, Amazon Web Services, Inc. or its affiliates.

# AWS investments supporting research and education




## TRAININGS

AWS Training Academy  
Demo/Immersion days



## SUPPORT

Technical Support  
AWS Research Scientists  
Grant Proposals



## INCENTIVES

Proof of concepts  
AWS Cloud Credits for Research  
Amazon Open Data Sponsorship  
Program  
Data Egress Waiver





# Voice of the researcher: requirements



## Reduce time to run

Access research environments in minutes



## Security

Maintain consistent security, compliance, and governance



## Resilient data ingestion

Ingest datasets at scale



## Integrated visualization

Post-analysis visualization and insights



## Spend controls

Cost visibility, centralized budgeting, and chargeback management



## Universally accessible

Collaborate from any location



## End-to-end resource provisioning

Storage, compute (AWS ParallelCluster), and visualization from single pane of glass



# Research Gateway

Secure, performant, and scalable research workbench that . . .

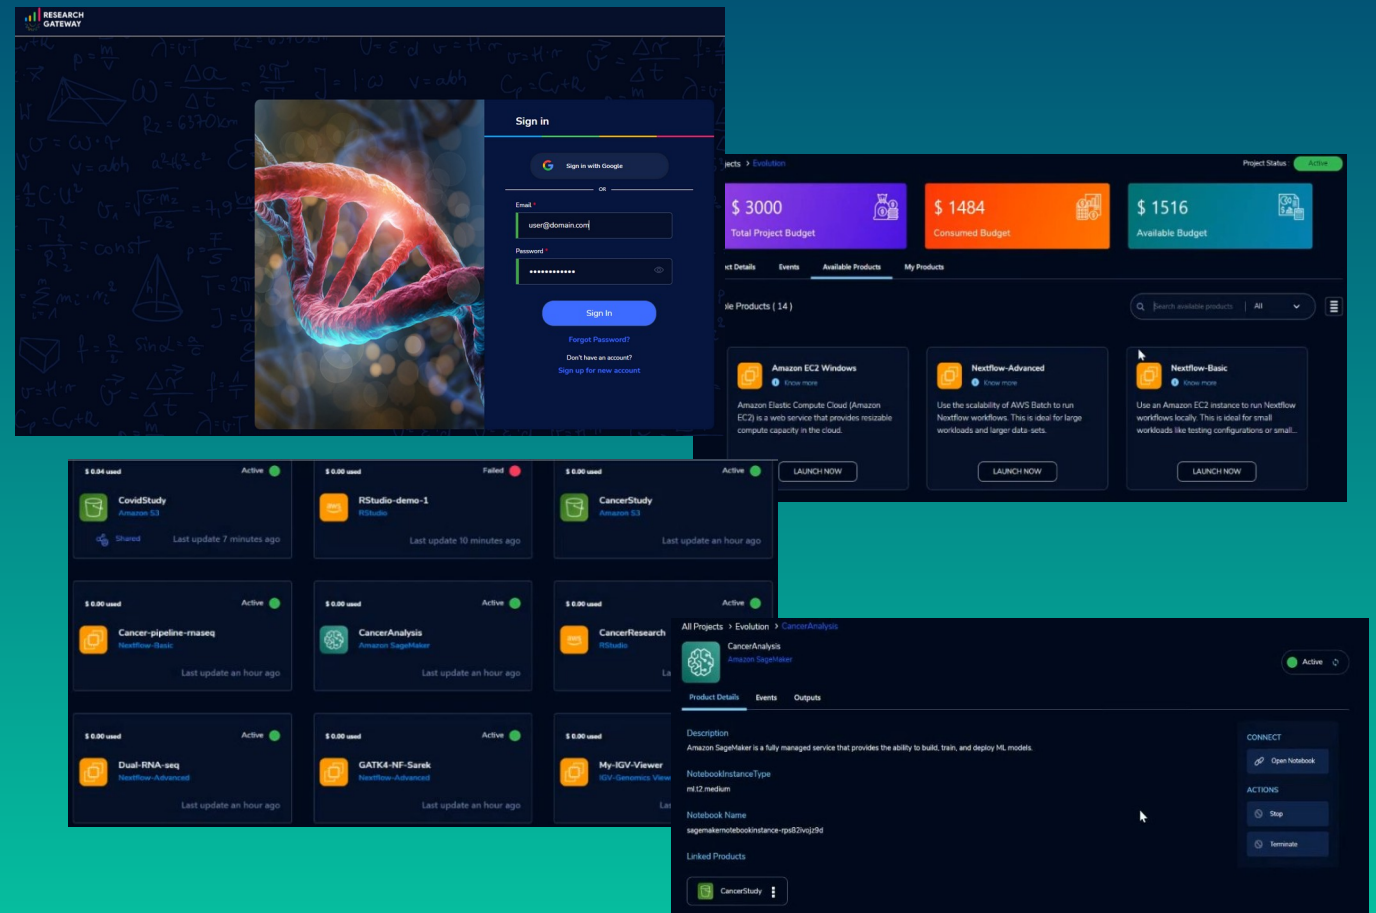
. . . can be provisioned in minutes

. . . is backed by secure and resilient data ingestion framework

. . . is delivered with wide selection of compute resources, including HPC

. . . is delivered with role-based access to workbench and data

. . . has clear budget and consumption costs breakdown by project, user, and workload



# What are Research Gateway use cases?



**Analytics**

nextflow



**Genomics workflow**



**High performance computing**




Amazon SageMaker  
**Machine learning and analytics**



**Customized AMIs**



**Research budget tracking**



**Bioinformatics toolkit**

egress

**Secure data egress**



**Study open datasets**



**Standard product catalog and marketplace access**



**Secure research workspaces**



**AWS data lake and secure ingress**



**BioContainers and researcher tools**



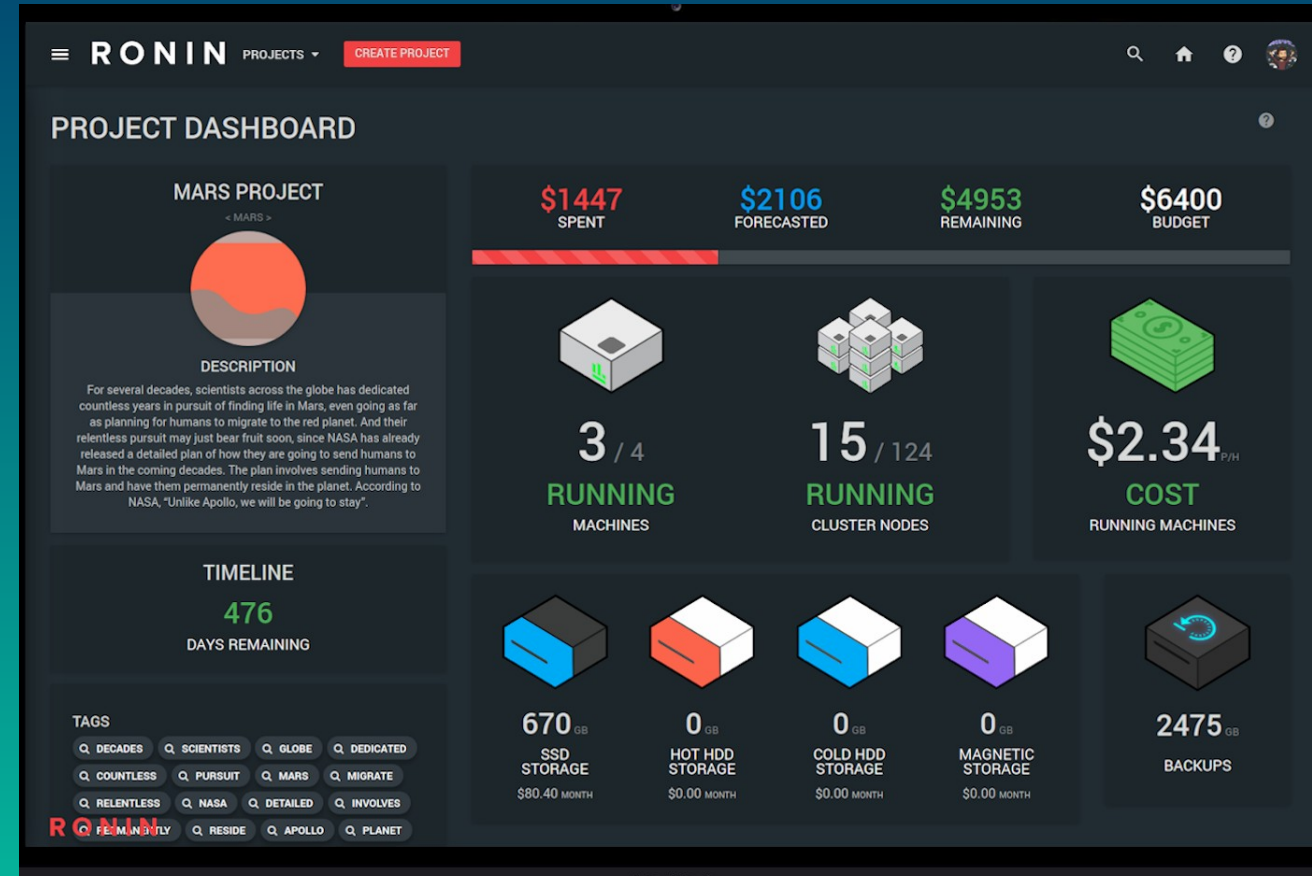
**HIPAA regulations and trusted research**

InCommon  
TRUSTED ACCESS  
PLATFORM

**Research community collaboration**

# RONIN is a cloud orchestration and collaboration platform, lowering the entry level to using the cloud for researchers and research IT.

- Deployed within one AWS account, serving many researchers
- Enables researchers self-service access to AWS resources
- Enforces an institution's security policy



# Research and Engineering Studio on AWS

Open source, easy-to-use web-based portal for administrators to create and manage secure cloud-based research and engineering environments.

## Benefits:

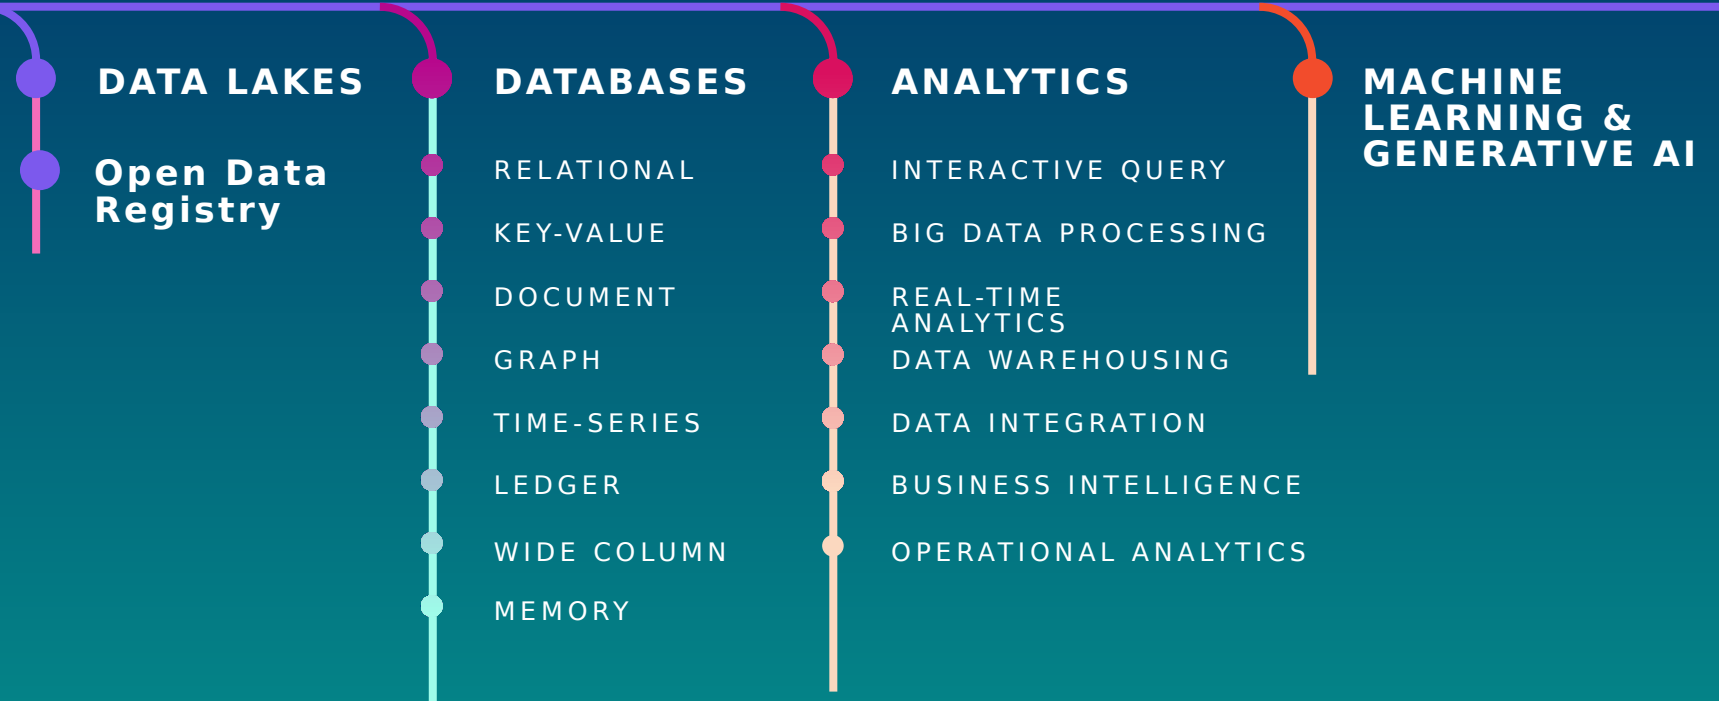
- Minimize administrative overhead
- No cloud expertise required
- Flexible access to services

## Use cases:

- Collaborate using shared research and engineering environments
- Define and manage projects
- Enable access to AWS without creating individual accounts

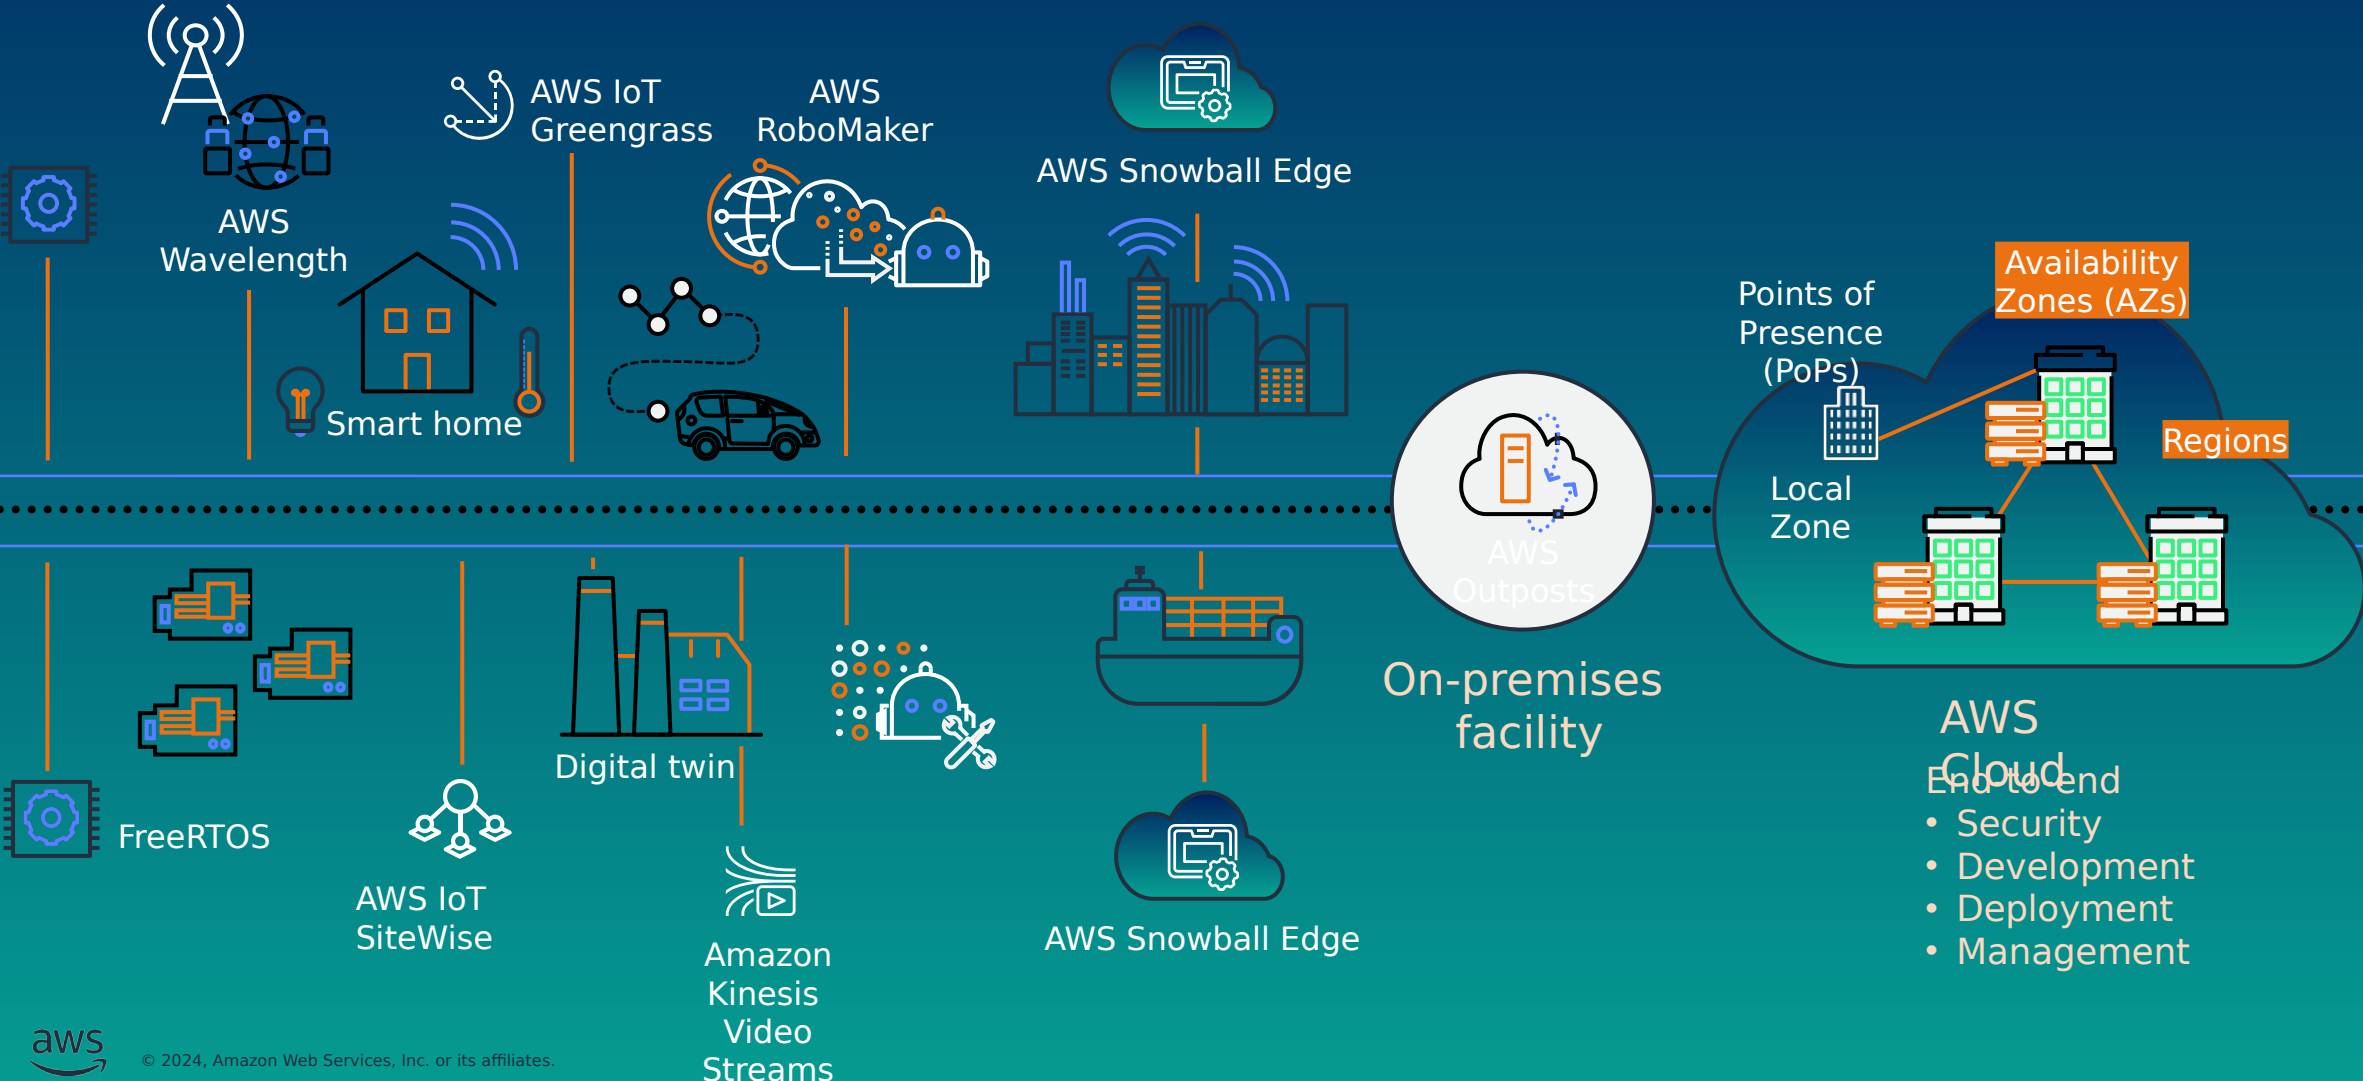


# The most comprehensive set of data services


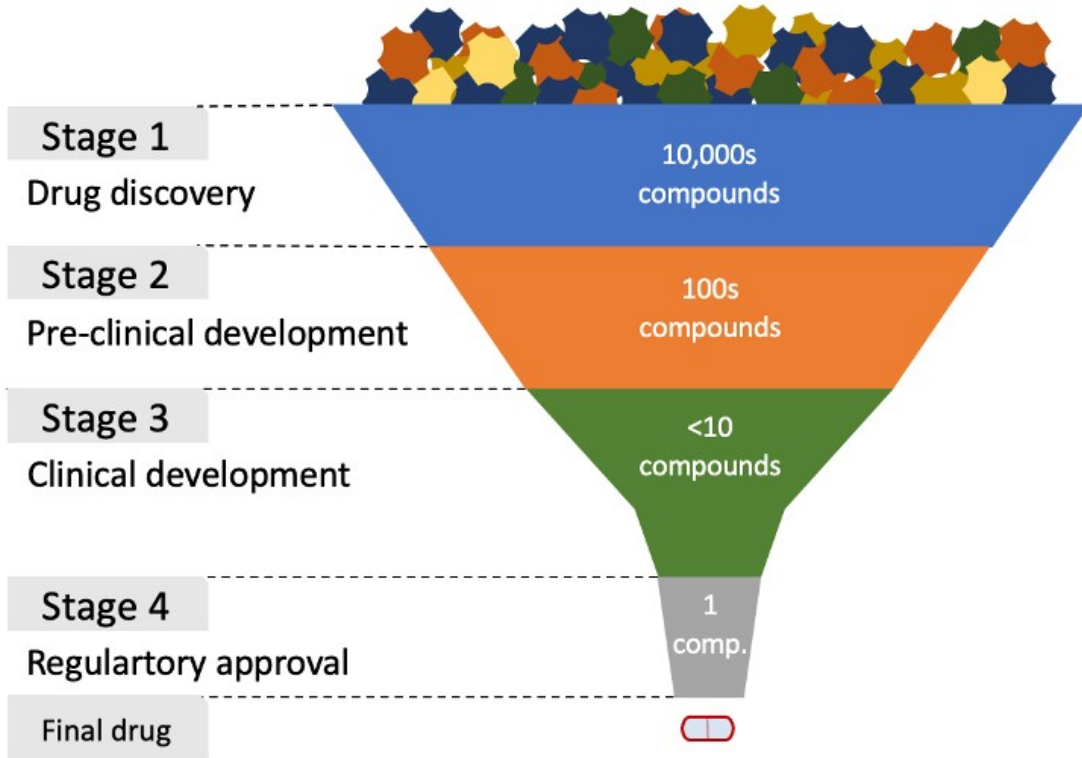




# A complete platform for building and deploying edge applications




# Agility and productivity - GROMACS: Max Planck Institute



<https://aws.amazon.com/blogs/hpc/running-20k-simulations-in-3-days-with-aws-batch/>  
<https://pubs.acs.org/doi/10.1021/acs.jcim.2c00044#>



© 2023, Amazon Web Services, Inc. or its affiliates.



Amazon S3



NICE DCV



Amazon EC2  
Auto Scaling




AWS ParallelCluster

# Clemson University - Natural Language Processing

<https://aws.amazon.com/blogs/aws/natural-language-processing-at-clemson-university-1-1-million-vcpu-ec2-spot-instances/>


The 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





# NATIONAL INSTITUTE OF NUCLEAR PHYSICS AND AWS WORK TOGETHER TO ACCELERATE QUANTUM COMPUTING RESEARCH

“ We are pleased to partner with AWS in what is an important element of our global strategy in quantum computing research.”

Marco Pallavicini, executive board member of INFN.



© 2024, Amazon Web Services, Inc. or its affiliates.



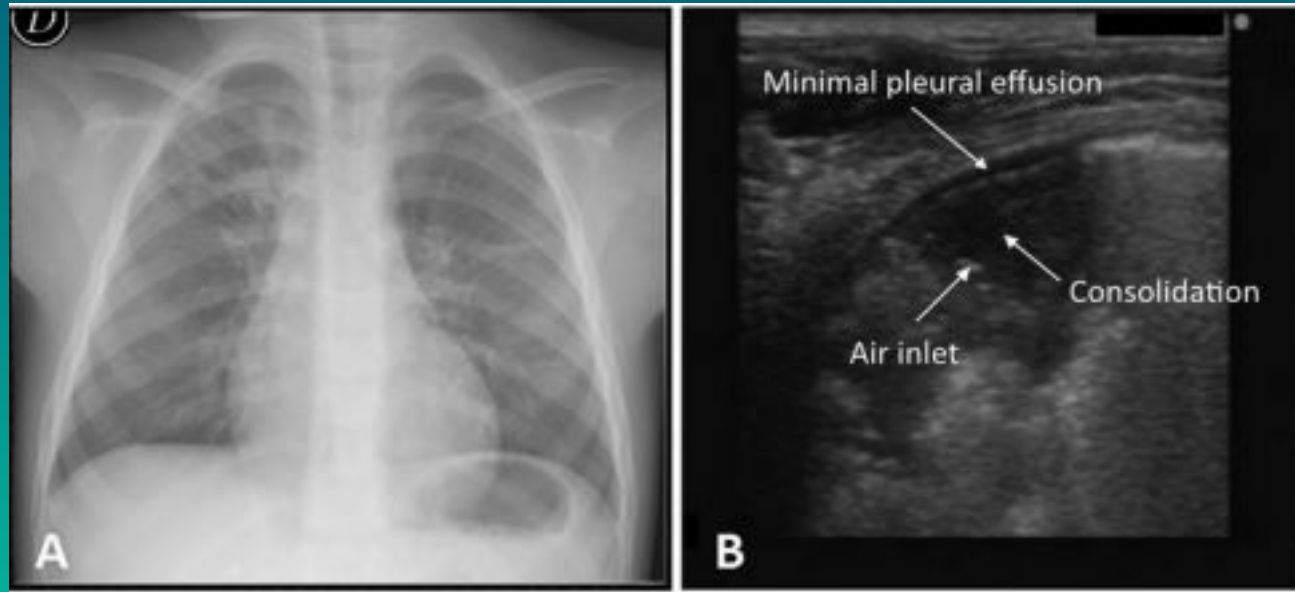
Amazon Braket



# Sharing & collaborating UniTrento - ICLUS

SOLUTION

University of Trento's Department of Information Engineering and Computer Science used AWS to run a International Project (<https://www.disi.unitn.it/iclus>) where AI algorithms were used to analyze ultrasound images of lungs to determine possible Covid infections.



More than 60 institutes involved

## Italy

- 118 Castelnuovo, Garfagnana (LU)
- APSS Trento, Trento (TN)
- Ausl Romagna - Cesena, Cesena (NULL)
- Azienda Ospedaliera Università di Padova, Padova (PD)
- Azienda Ospedaliera Universitaria Federico II, Napoli (NA)
- Azienda Ospedaliera Universitaria Policlinico Vittorio Emanuele, Catania (CT)
- Bresciamed, Brescia (BS)
- Cardiologia - Ospedale Policlinico San Martino, Genova (GE)
- Emergency Department of Arzignano Hospital - AULSS8 Berica, Vicenza (VI)
- Fondazione Policlinico San Matteo IRCCS, Pavia (PV)
- Fondazione Policlinico Universitario A. Gemelli IRCCS, Roma (RM)
- Mater Olbia Hospital, Olbia (SS)
- Ospedale Civile di Voghera, Voghera (PV)
- Ospedale dei Bambini Vittore Buzzi, Milano (MI)
- Ospedale di Sanremo Asl1 Imperiese, Sanremo (IM)
- Ospedale di Tione, Tione (TN)

## Other countries

- Augusta University - Department of Emergency Medicine, Augusta, United States of America
- Clinic of thoracic and vascular surgery, Gera, Germany
- Contra Costa Regional Medical Center, Martinez, California, USA
- DeepMed I/O, Manchester, United Kingdom
- Department of Obstetrics and Gynecology - University Hospitals Leuven, Leuven, Belgium
- DSP Medea, Medea, Algeria
- Eindhoven University of Technology, Eindhoven, The Netherlands
- Hospital das Clínicas da Universidade de São Paulo, São Paulo, Brazil
- Hospital General de Catalunya, Barcelona, Spain
- IFC - CNR LECCE, Lecce, ITALIA
- Indian Institute of Technology, Kharagpur, India
- Indian Institute of Technology, Jodhpur, India
- Indian Institute of Technology Patna, Patna, India
- Institute of Biomedical Engineering, University of Oxford, Oxford, United Kingdom
- Intelligent Ultrasound, Cardiff, United Kingdom
- Klinikum rechts der Isar der TU, München, Germany
- Leitat Technological Center, Barcelona, Spain
- Massachusetts General Hospital, Boston, USA
- Michigan State University, Lansing, USA
- North Carolina State University, Raleigh, United States of America

More than 29 joint publications



© 2024, Amazon Web Services, Inc. or its affiliates.



AWS RESEARCH –17<sup>TH</sup> SEPTEMBER 2024

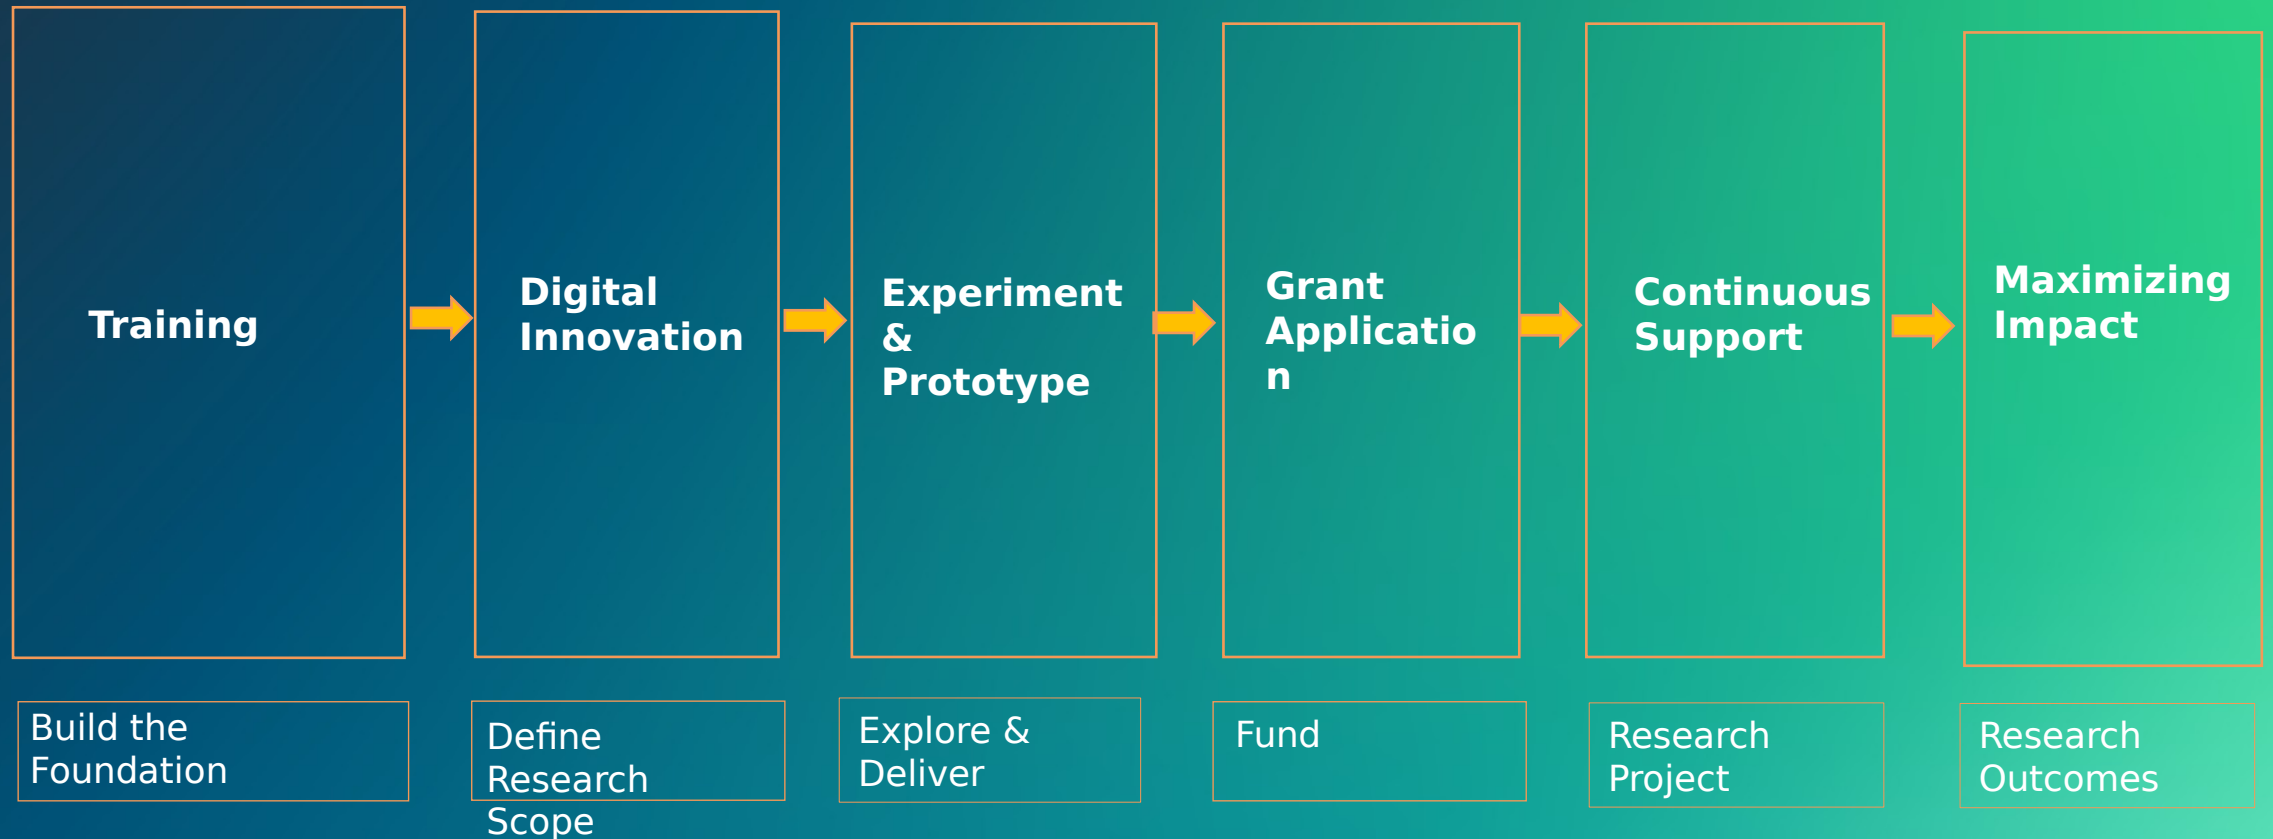
# Services and case studies: cloud in action

**AWS Education & Research  
Team**

Roberta Piscitelli, PhD, MBA - [piscitr@amazon.com](mailto:piscitr@amazon.com)



# AWS grant acceleration support package

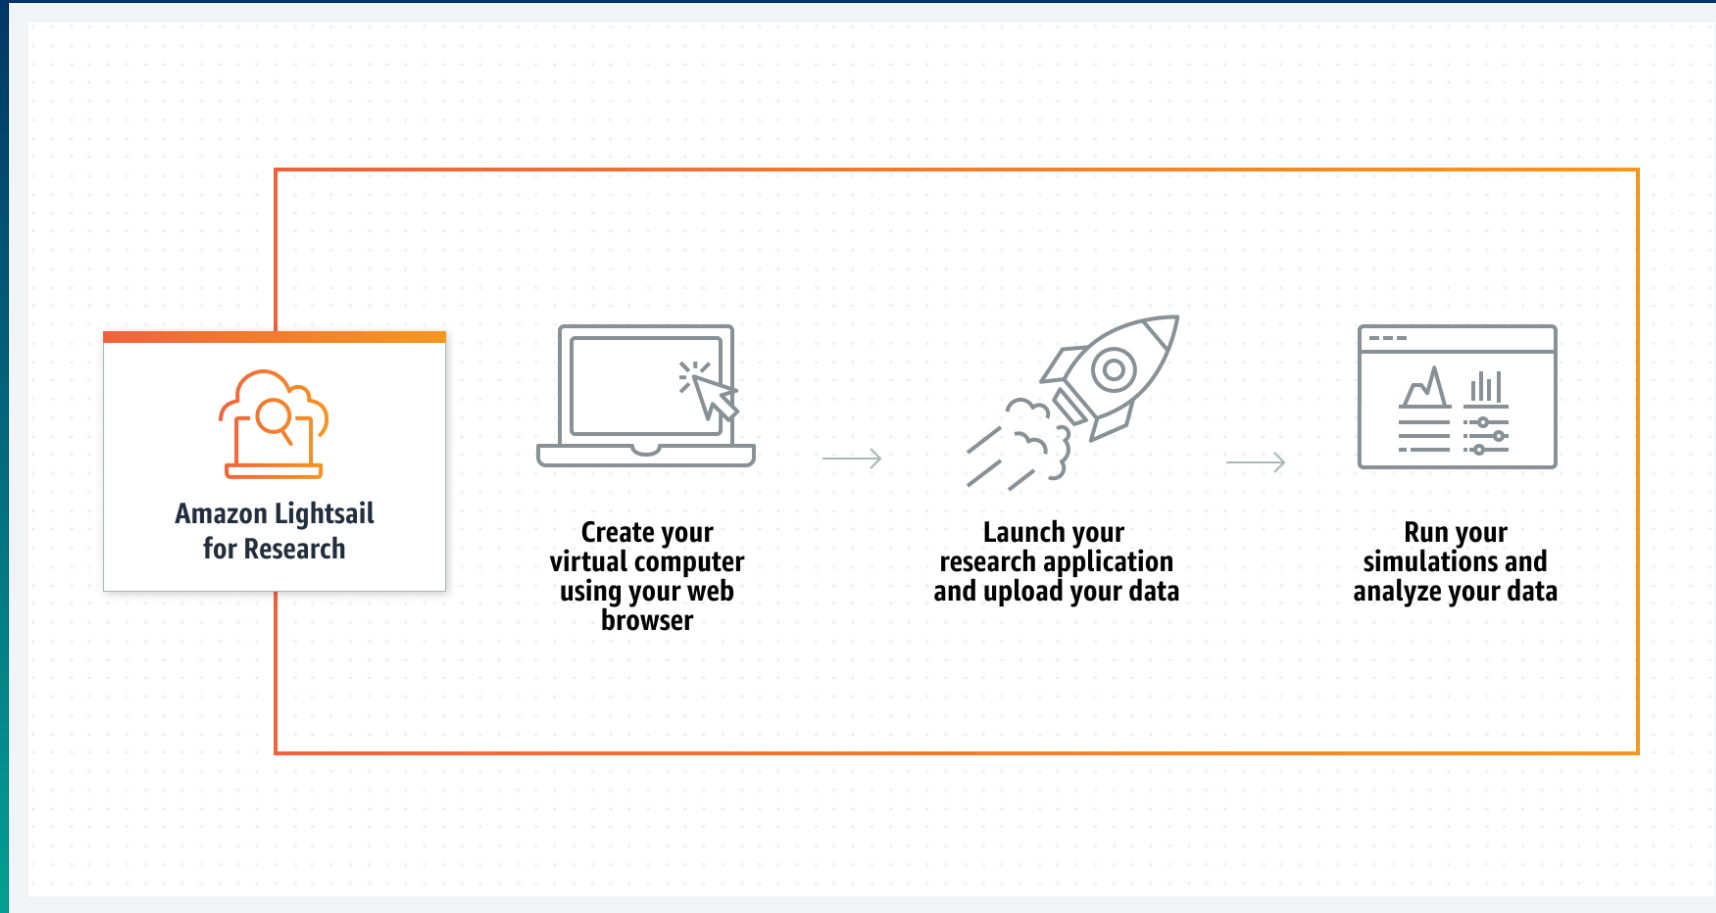


# Build research



# Lightsail for Research

<https://aws.amazon.com/lightsail/research/>



# Acquiring data




# Data analysis





# Multimodal analytics

PURPOSE-BUILT SERVICES FOR HEALTHCARE AND LIFE SCIENCES



## Amazon Omics

Transform genomic, transcriptomic, and other omics data into insights



## Amazon HealthLake Imaging and Analytics

Provide a complete view of individual or patient population health data



## Amazon Comprehend Medical

Understand medical context using natural language processing



## Amazon Transcribe Medical

Automatically convert medical speech to text

# The AWS ML Stack

Broadest and most complete set of machine learning capabilities

## AI SERVICES

### HEALTH AI



NEW

Amazon HealthLake



Amazon Transcribe Medical



Amazon Comprehend Medical

### INDUSTRIAL AI



NEW

AWS Panorama + Appliance



NEW

Amazon Monitron



NEW

Amazon Lookout for Equipment



NEW

Amazon Lookout for Vision

### ANOMALY DETECTION



NEW

Amazon Lookout for Metrics



NEW

Amazon DevOps Guru



Amazon CodeGuru

### CODE AND DEVOPS

### VISION



Amazon Rekognition

### SPEECH



Amazon Polly



Amazon Transcribe  
+Medical

### TEXT



Amazon Comprehend  
+Medical



Amazon Translate



Amazon Textract

### SEARCH



Amazon Kendra

### CHATBOTS



Amazon Lex

### PERSONALIZATION



Amazon Personalize

### FORECASTING



Amazon Forecast

### FRAUD



Amazon Fraud Detector

### CONTACT CENTERS



Contact Lens

Voice ID

For Amazon Connect

## ML SERVICES



Amazon SageMaker

Label data

NEW

Aggregate & prepare data

NEW

Store & share features

Auto ML

Spark/R

NEW

Detect bias

Visualize in notebooks

Pick algorithm

Train models

Tune parameters

NEW

Debug & profile

Deploy in production

Manage & monitor

NEW  
CI/CD

Human review

### SAGEMAKER STUDIO IDE

NEW: SageMaker JumpStart

NEW: Model management for edge devices

## FRAMEWORKS & INFRASTRUCTURE

TensorFlow

mxnet

PyTorch

GLUON

Keras

skit-learn



DeepGraphLibrary

Deep Learning AMIs & Containers

GPUs & CPUs

Elastic Inference

Trainium

Inferentia

FPGA



© 2024, Amazon Web Services, Inc. or its affiliates.



# Amazon

# SageMaker

Build, train, and deploy ML models at scale

Automatic model fine-tuning & distributed training

Flexible model deployment options


Tools for ML operations

Built-in features for responsible AI



# University of Oxford Introduces a Sector-Leading Image Recognition ML Prototype to Augment Digitization in Numismatics

<https://aws.amazon.com/solutions/case-studies/oxford-case-study/>



“ I thought this project would be complex and time consuming, but **using AWS made it easy.** ”


Anjanesh Babu  
Systems architect and network manager, Gardens and Museums IT, University of Oxford's  
Gardens, Libraries & Museums




© 2024, Amazon Web Services, Inc. or its affiliates.



Amazon SageMaker



Amazon EC2



Amazon S3

# Large Models built on AWS



Clibrain develops Lince Zero—the first large language model (LLM) optimized for Spanish using Amazon SageMaker



Stability AI will build AI models on compute clusters with thousands of GPU or AWS Trainium chips, reducing training time and cost by 58%



BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance, utilizing AWS infrastructure services



# New LLM development and refinement UAE Technology Innovation Institute



## CHALLENGE

TII researchers needed an easy way to develop, iterate, and distribute Falcon-40B LLM. Experimenting at such massive scale was slowed by limited computational resources and complex on-prem infrastructure. Researchers sought more flexibility to efficiently enhance, evaluate, and provide access to Falcon under dynamic capacity requirements to empower their team's foundation model research.

## SOLUTION

TII utilized SageMaker's managed machine learning infrastructure and tools to train, host, and deploy versions of Falcon with greater speed and experimentation. SageMaker provided on-demand access to training clusters along with modular deployment options for easy distribution to varying client workloads—key for customizing model access.

## OUTCOME

- ✓ TII successfully implemented Falcon-180B by using SageMaker and custom innovation. This allows TII's contribution within UAE's 2031 National AI Strategy, fostering economic growth and social progress. Releasing UAE's Falcon 180B, World's Top-Ranked Open Source AI Model will further encourage AI academic research and scientific collaboration.

# Innovating at the silicon level


**AWS Trainium 2**



**AWS Inferentia 2**



# Graviton, Inferentia, and Trainium



Graviton-based Elastic Compute Cloud instances use up to **60% less energy** than comparable Amazon EC2 instances

---

Models built on Trainium result in **energy-consumption reductions of up to 25%** versus comparable instances

---

Our Inferentia2 machine learning accelerator delivers up to **50% higher performance per watt**





# Amazon Q Developer

AI-powered code suggestions  
in the IDE and the command  
line

A screenshot of a code editor interface. At the top, there is a tab labeled 'main.js' with a small icon to its left. Below the tab, the editor area is dark with white text. On the left side of the editor, there is a vertical list of line numbers from 1 to 21. The rest of the editor area is currently blank.

# Amazon Bedrock


simplifies



Choice



Customization



Integration

AI21labs

amazon

ANTHROPIC

cohere

Meta

stability.ai

JURASSIC-2

AMAZON TITAN

CLAUDE

COMMAND + EMBED

LLAMA 2

STABLE DIFFUSION XL




# How AWS supports Gen AI research



Use familiar tools  
that meet your  
requirements

Pre-trained FM



Customization

Fine-tuned FM  
Industries (Legal,  
Telco, Financial,  
Education...)



FM-as-a-service

FM  
monetization  
Licensing



FM-powered  
applications

Start-up  
ecosystem  
Public  
administration  
Industry &  
Academia  
Citizen services






# Publishing and sharing research data



# Reach out to your contacts at Amazon Web Services



**Roberta Piscitelli**

Tech BDM Research | EMEA

[piscitr@amazon.com](mailto:piscitr@amazon.com) | +44 793 3576  
787



**Andreas Wagner**

Sr. Account Manager | Research &  
Nonprofits

[wagneand@amazon.com](mailto:wagneand@amazon.com) | +49 151 0913  
5033





# Thank you!

**AWS Education & Research**

**Team**

Roberta Piscitelli - [piscitr@amazon.com](mailto:piscitr@amazon.com)



A night sky filled with stars and constellations, with the silhouettes of trees at the bottom. The text "Responsible AI" is centered in the middle of the image.

# Responsible AI



# Key dimensions

## Fairness

- How a system impacts different subpopulations of users (e.g. by gender, ethnicity)

## Explainability

- Mechanisms for understanding and evaluating the outputs of an AI system

## Privacy and Security

- Data protected from theft and disclosure

## Robustness

- Mechanisms to ensure reliable operability of an AI system

## Governance

- Processes for defining, implementing and enforcing responsible AI practices within an organization

## Transparency

- Communicating information about an AI system so that stakeholders can make informed and ethical choices about its use

# AWS announces new tools for responsible AI innovation across the generative AI lifecycle

## Implementation of guardrails for Amazon Bedrock

Implements predefined safeguards and filters based on customer policies to avoid harmful content

## Model Evaluation for Bedrock & SageMaker

Compares foundation models on custom metrics like accuracy, safety, and subject criteria

## Fight misinformation

Invisible watermarks on AI-generated images to combat disinformation in amazon titan

## Building Trust & Transparency

IP Indemnification - Protects customers from infringement lawsuits related to outputs of generative AI services

## AWS's Commitment to Responsible AI

Collaborating with policymakers and the AI community on responsible practices  
Investing in purpose-built tools like SageMaker Clarify and ML Governance to operationalize responsible AI

<https://aws.amazon.com/blogs/machine-learning/announcing-new-tools-and-capabilities-to-enable-responsible-ai-innovation/>



© 2023, Amazon Web Services, Inc. or its affiliates.

