



AWS for Research

Jessica Steed

Account Executive

Hatem Siyala

Solutions Architect

Why Use AWS for Research?



Simple for Researchers

Reducing user friction by providing powerful yet simple to access services



Security Controls & Roles

Intuitive to manage access to data



Budget & Cost Management

Cost transparency and simple estimation



Long-term Investment

We're here to stay with stable long-term technology



Open Innovation

Constantly defining new technologies and committed to providing tools to enable research



Culture & People

Shared mindset as we are constantly publishing, teaching, and engaging with academia

AWS Accelerates Research



Science, not Servers

Compute when you need it at any scale




Collaboration

Access datasets that span institutions and borders



Research Data Management

Storage, secure access, and management



Share & Reproduce Research

A common platform for reproducing scientific analyses



State-of-the-Art Analytics

Use data science methods in your research



Security

A collection of tools to protect data and privacy



Enterprise IT **and** Research

Enterprise IT cares about...



Stability & Operational
Excellence



Cost
Optimization



Security &
Compliance

Research cares about...



New Ideas &
Innovation



Speed &
Agility



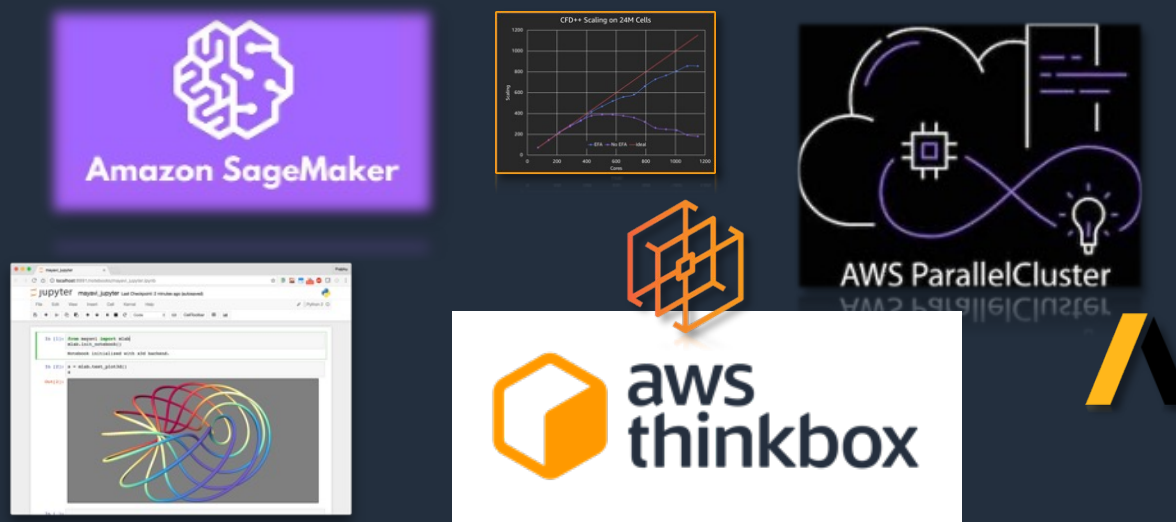
Value &
Results

Research is **not** one size fits all

Researcher Workspaces

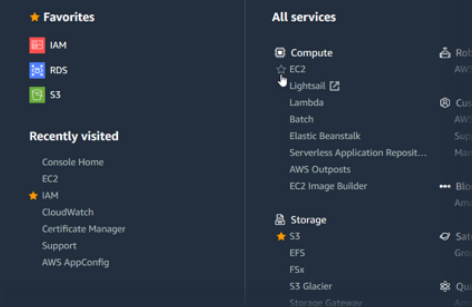


Solutions & Guidance

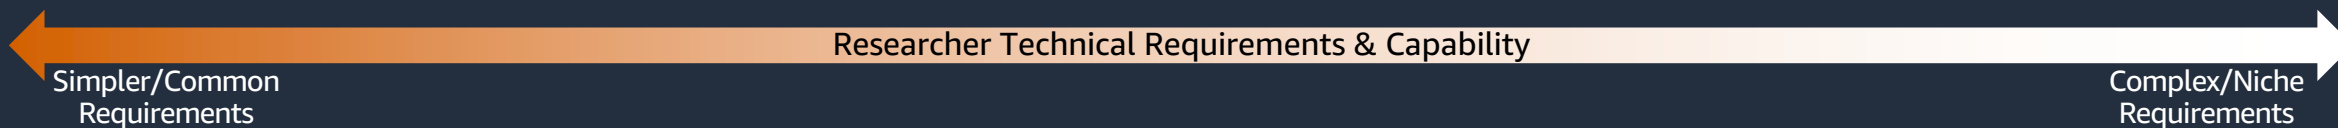


Native AWS

AWS Console, CLI, SDKs, etc.



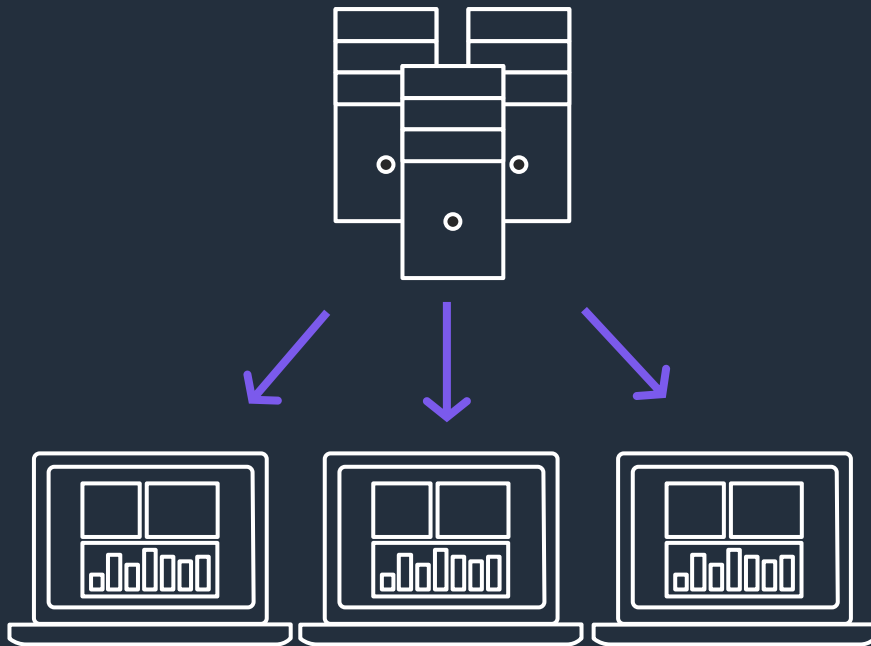
FSx Storage and data solutions to meet the specific researcher needs



A Cloud-Based Model Promotes Collaboration


Traditional Approach:

Move data to computing resources



Cloud Approach:

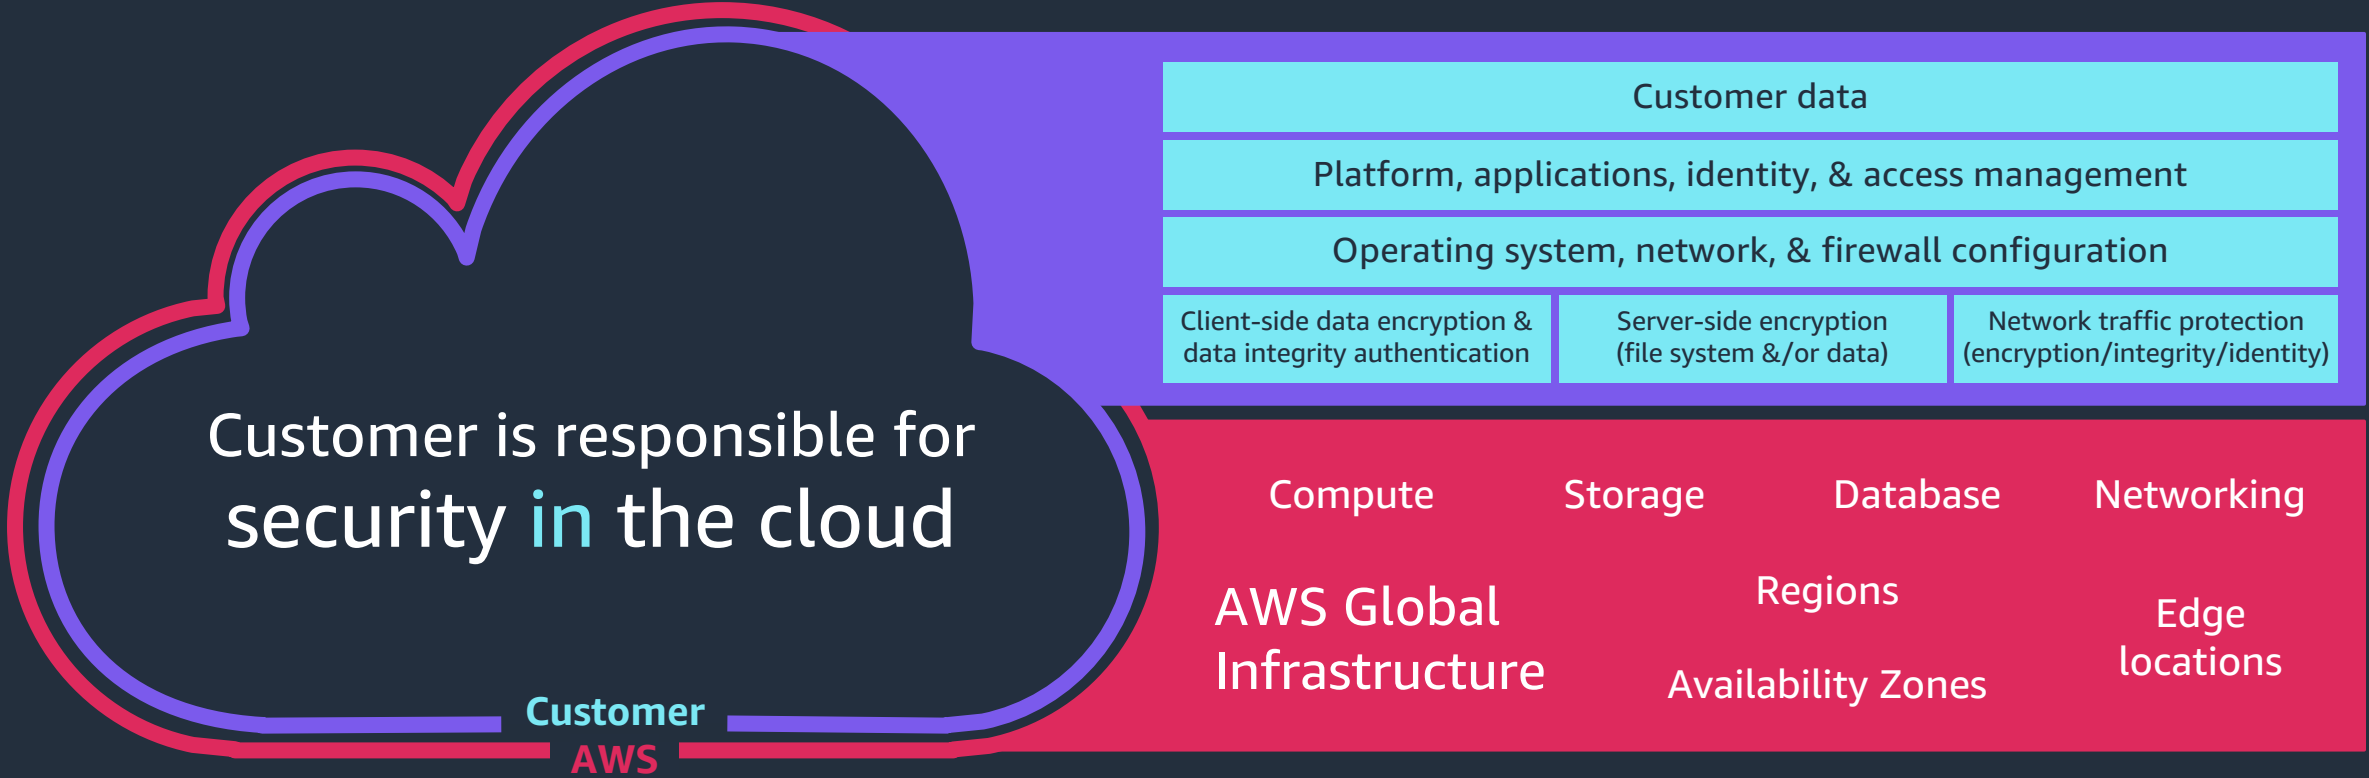
Move computing resources to data



Secure Research











Shared Security Responsibility Model










AWS is responsible for security **of** the cloud

Inherit Global Security & Compliance Controls





Certifications/Attestations


DoD CMMC		✓
FedRAMP		✓
FISMA		✓
HITRUST		✓
ISO 9001		✓
ISO 27001		✓
ISO 27017		✓
ISO 27018		✓

Laws/Regulations/Privacy

FERPA		✓
GDPR		✓
GLBA		✓
HIPAA		✓
IRS 1075		✓
ITAR		✓
PIPEDA [Canada]		✓

Alignments/Frameworks

CIS (Center for Internet Security)		✓
CJIS (US FBI)		✓
CSA (Cloud Security Alliance)		✓
FFIEC		✓
FIPS		✓
NIST		✓
PHR		✓

 = industry or global standard



AWS Solutions for Research



Introducing Amazon Lightsail for Research

- **For users with limited experience running large computation**
 - Ex: Researchers currently doing all analysis on their laptops and looking to quickly spin up an environment to test ideas
 - No need to configure or manage cloud infrastructure components
- **Access to applications through browser based, simple create wizard**
 - Pick from a fixed set of compute options matched to pre-installed application
 - Run scripts, analyze data, and install additional packages
 - Drag and drop web interface to upload datasets or download results
- **All-inclusive pricing**
 - Includes compute, storage, and data transfer only for duration
 - Data can persist after power down
 - Budgeting tools to manage costs, power down idle research computers, etc.



Amazon Lightsail for Research

The image shows two overlapping screenshots of the Amazon Lightsail console. The top screenshot is the 'Choose application' step, and the bottom screenshot is the 'Select hardware bundle' step.

Choose application

Get up and running quickly

Create your virtual computer

Start your work

Choose an application to install on your computer.

- Ubuntu 20.04 LTS** - Choose
- RStudio 1.0** - Choose
- VSCodium 1.0** - Choose

Select hardware bundle

Select a hardware bundle

Bundle	Price	vCPUs	Memory	Storage	Monthly Data Transfer
Standard XL	\$0.90 USD per hour	4	8 GB	50 GB	1 TB
Standard 2XL	\$1.11 USD per hour	8	16 GB	50 GB	1 TB
Standard 4XL	\$1.53 USD per hour	16	32 GB	50 GB	1 TB
GPU XL	\$2.37 USD per hour	4	16 GB	50 GB	1 TB
GPU 2XL	\$2.64 USD per hour	8	32 GB	50 GB	1 TB
GPU 4XL	\$3.18 USD per hour	16	64 GB	50 GB	1 TB

Summary

AWS Region: us-east-2 (Ohio)

Application: RStudio 1.0 on Unix

Hardware bundle: Standard XL

- 4 vCPUs
- 8 GB memory
- 50 GB storage
- 1 TB monthly data transfer

Hourly total: \$0.90 USD

Create virtual computer Cancel

By using this software, you agree to the [software's end user license agreement](#).

Your use of AWS services is subject to the [AWS Customer Agreement](#).

Name your virtual computer

Virtual computer names must be unique within each AWS Region in your Lightsail account.

noaa-gsod-analysis

Valid characters: A-Z, a-z, 0-9, _ (underscore), and - (hyphen)



AWS ParallelCluster

Cluster properties [Info](#)

Name

dem

The name of the cluster.
Valid characters are lowercase letters, numbers, and hyphens.

Region

The AWS Region where the cluster is created.

US-V

Operating System

The operating system of the cluster.

Queue 1 - on-demand

Queue name

on-demand-queue

Subnet IDs

Select a subnet

subnet-04e5d7521c24f6a
us-west-2b - usw2-az2 (demo)

Purchase type [Info](#)

On-Demand

Compute resources

Configure up to 50 compute resources.

Resource 1 - on-demand

Static nodes

1

Storage settings

Add up to [20 file systems](#) for shared storage.

 No shared storage options selected.

Storage types

Select file system types

- Amazon FSx for Lustre (FSX)
- Amazon FSx for NetApp ONTAP (FSX)
- Amazon FSx for OpenZFS (FSX)
- Amazon File Cache
- Amazon Elastic File System (EFS)
- Amazon Elastic Block Store (EBS)

Add storage

AWS ParallelCluster

```
1 ▼ HeadNode:
2     InstanceType: c5n.large
3 ▼ Ssh:
4     KeyName: KP - PDX
5 ▼ LocalStorage:
6 ▼     RootVolume:
7     VolumeType: gp3
8 ▼ Networking:
9     SubnetId: subnet - 04e5 d7521c24f6a7e
10 ▼ Scheduling:
11     Scheduler: slurm
12 ▼ SlurmQueues:
13 ▼     -Name: on-demand - queue
14 AllocationStrategy: lowest - price
15 ▼ ComputeResources:
16 ▼     -Name: on - demand - queue - cr - 0
17 ▼ Instances:
18 ▼     -InstanceType: c5n.large
19 MinCount: 1
20 MaxCount: 40
```

SageMaker

Features

What	AWS service
Use cases	Development of machine learning (ML) models
Configuration	Configurable
Setup persona	Research IT / Partner / ProServe
Setup skills	Advanced
User skills	Advanced
Compliance	Depends on Setup

Pricing & Billing

Egress traffic	Per service usage
Pricing	Metered service
Billing	As per services used

Links

[Public service page](#)

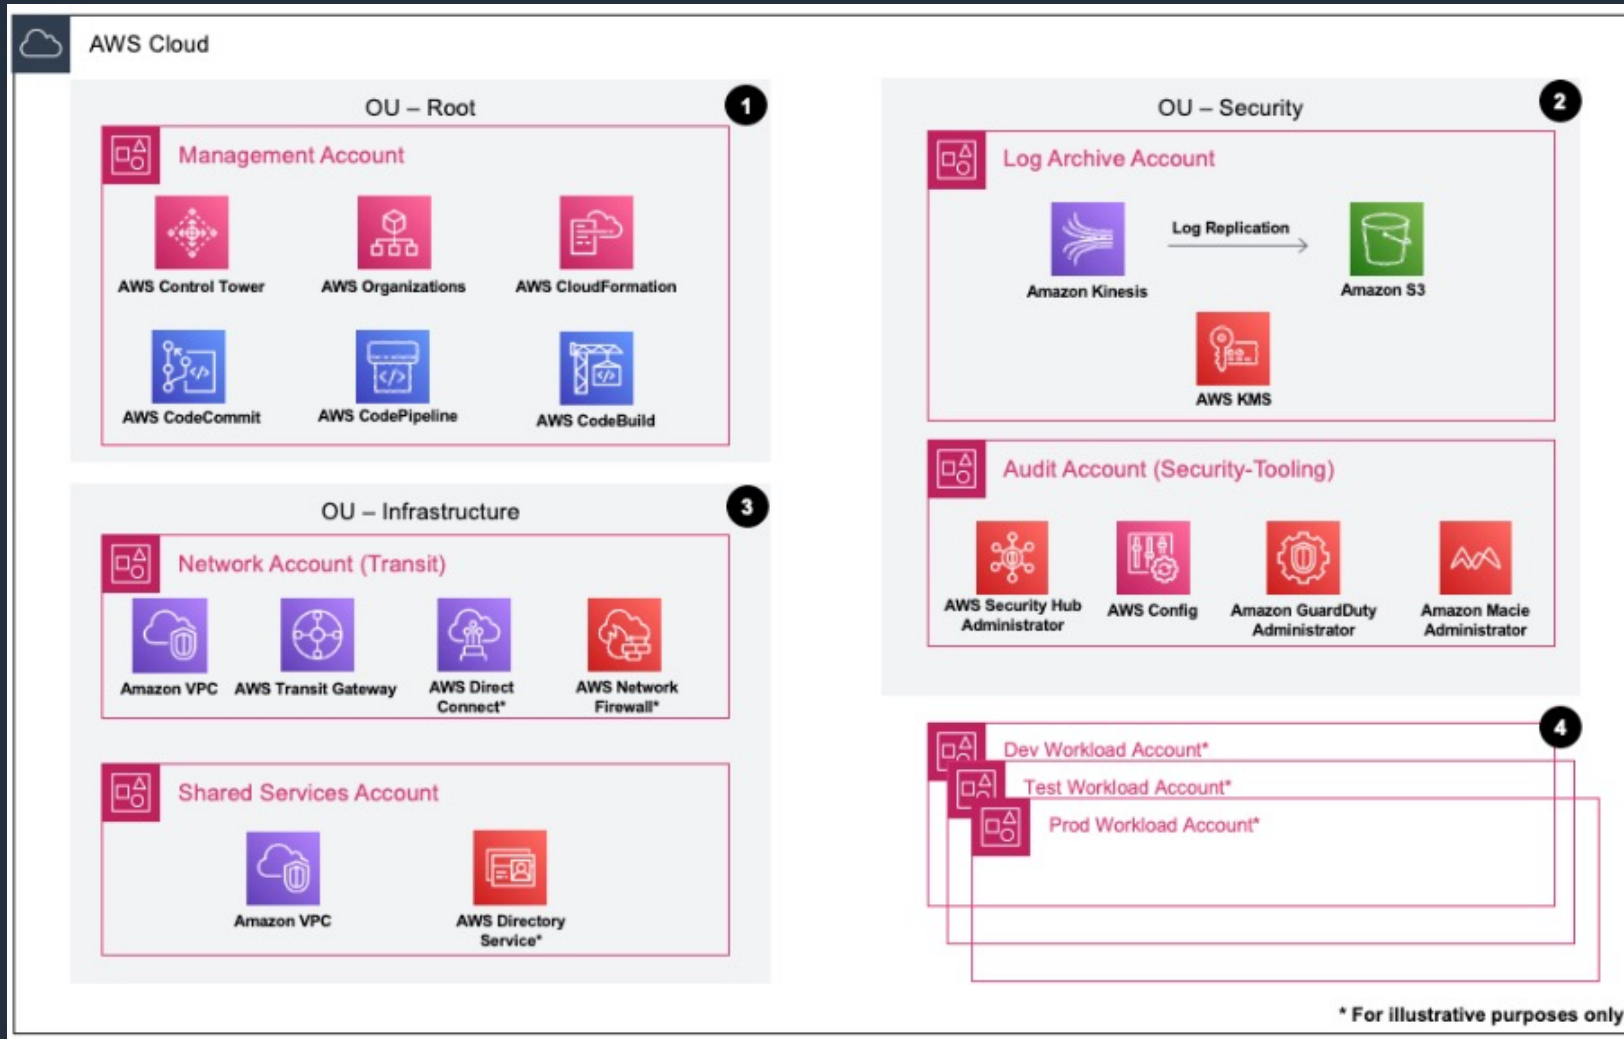
Solution Data

Description	Build, train, and deploy ML models for any use case with fully managed infrastructure, tools, and workflows
Key points	<ul style="list-style-type: none">• Access, label, and process large amounts of structured data (tabular data) and unstructured data (photo, video, geospatial, and audio) for ML• Reduce training time from hours to minutes with optimized infrastructure. Boost team productivity up to 10 times with purpose-built tools• Enable more people to innovate with ML through a choice of tools—IDEs for data scientists and no-code interface for business analysts• Automate and standardize MLOps practices and governance across your organization to support transparency and auditability
Category	Research tools

Other Research Resources



Industry Landing Zones

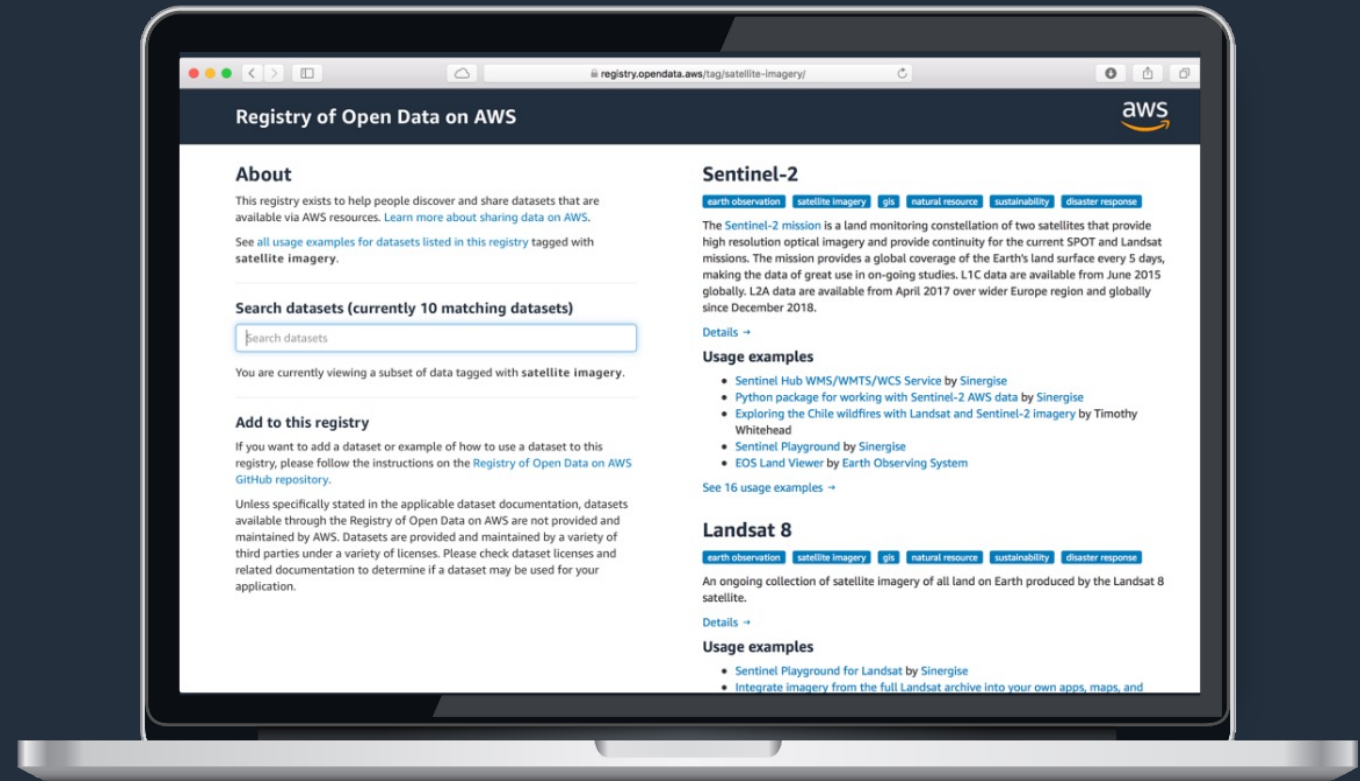


Open Data on AWS

AWS HOSTS PUBLIC DATASETS TO LOWER THE COST AND IMPROVE THE SPEED OF RESEARCH.

Examples include:

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



Global Data Egress Waiver

Why

Researchers need predictable budgets

Who

Available to degree-granting and research institutions

Must use a research network (e.g. Internet2) or AWS Direct Connect

What

Waives charges for data downloads

Waives up to 15% of the customer total account bill

Data uploads are no-charge

How

Contact your AWS account manager

Or via research portals (e.g. Arcus, GEANT)

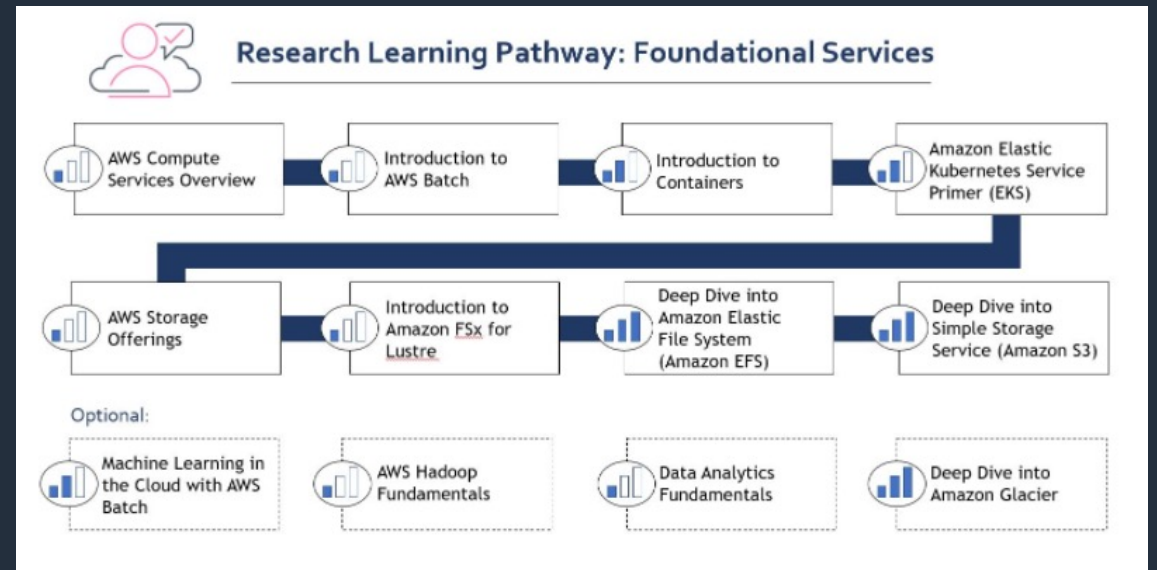
Or via certain resellers (e.g. DLT/I2 NET+, VMWare on AWS)



No-cost Research Learning Pathway

Provides a structured, curated list of online courses tailored to researchers' needs

AWS Skill Builder provides 500+ no-cost digital courses



Q&A

