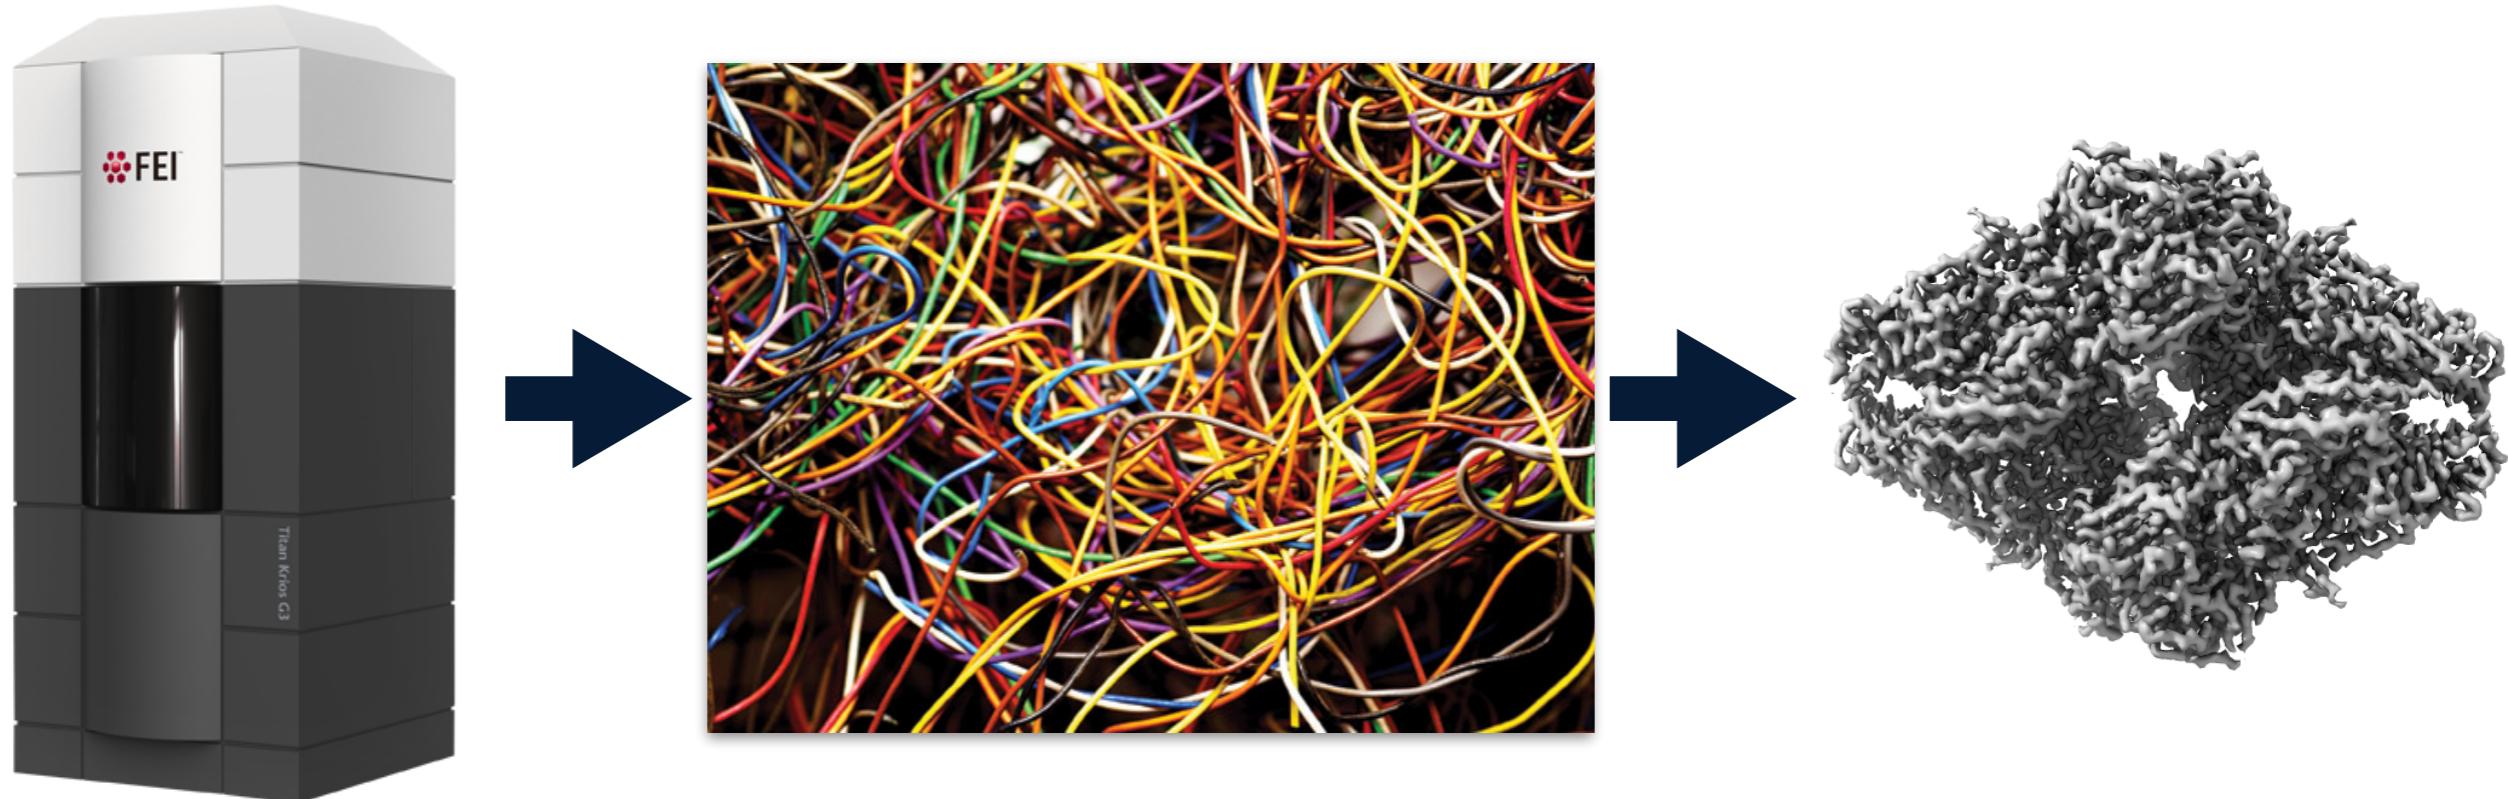


# Software tools to deploy and manage cryo-EM jobs in the cloud

---

Michael Cianfrocco  
Life Sciences Institute  
Department of Biological Chemistry  
University of Michigan

# 'State of the art' computing for cryo-EM

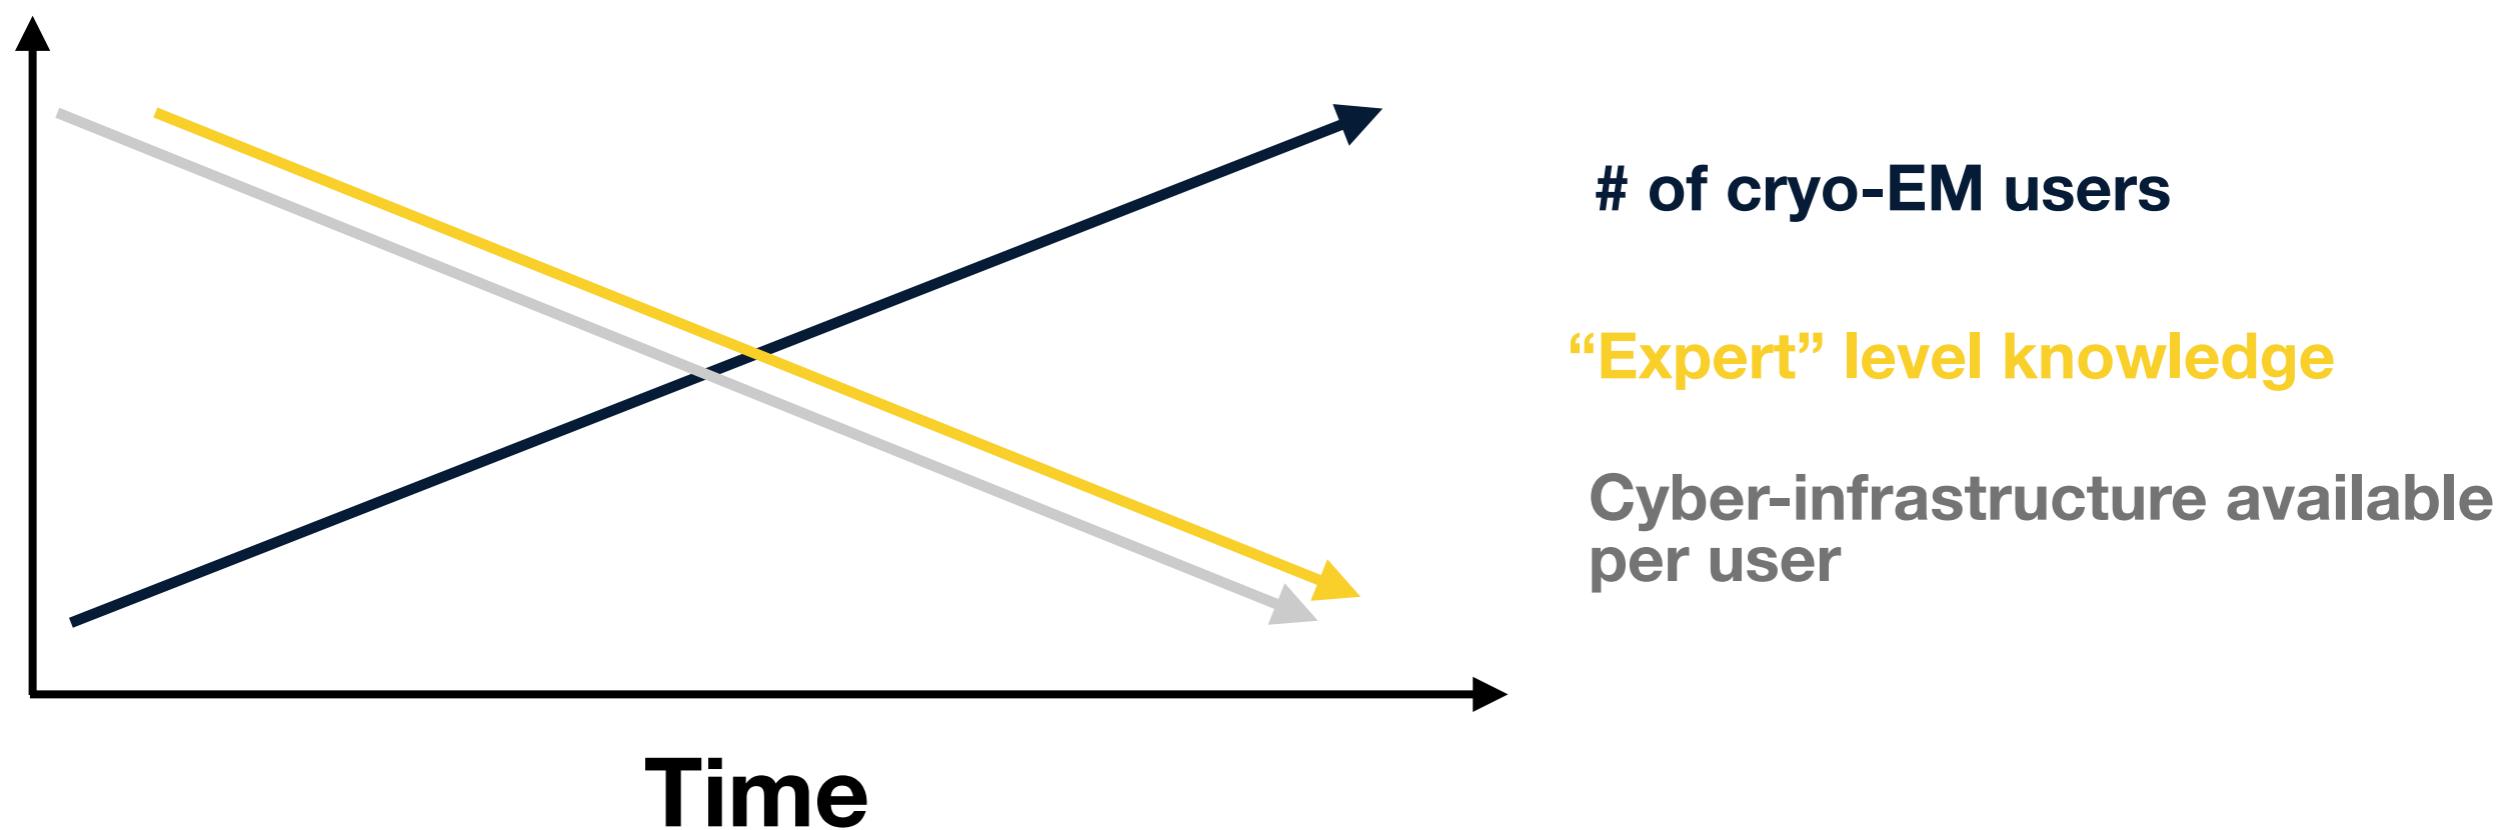


- ▶ Data processing occurring at a different site than data collection
- ▶ Data copying & moving; living at multiple locations
- ▶ New users left on their own to navigate the complex world of computing & storage

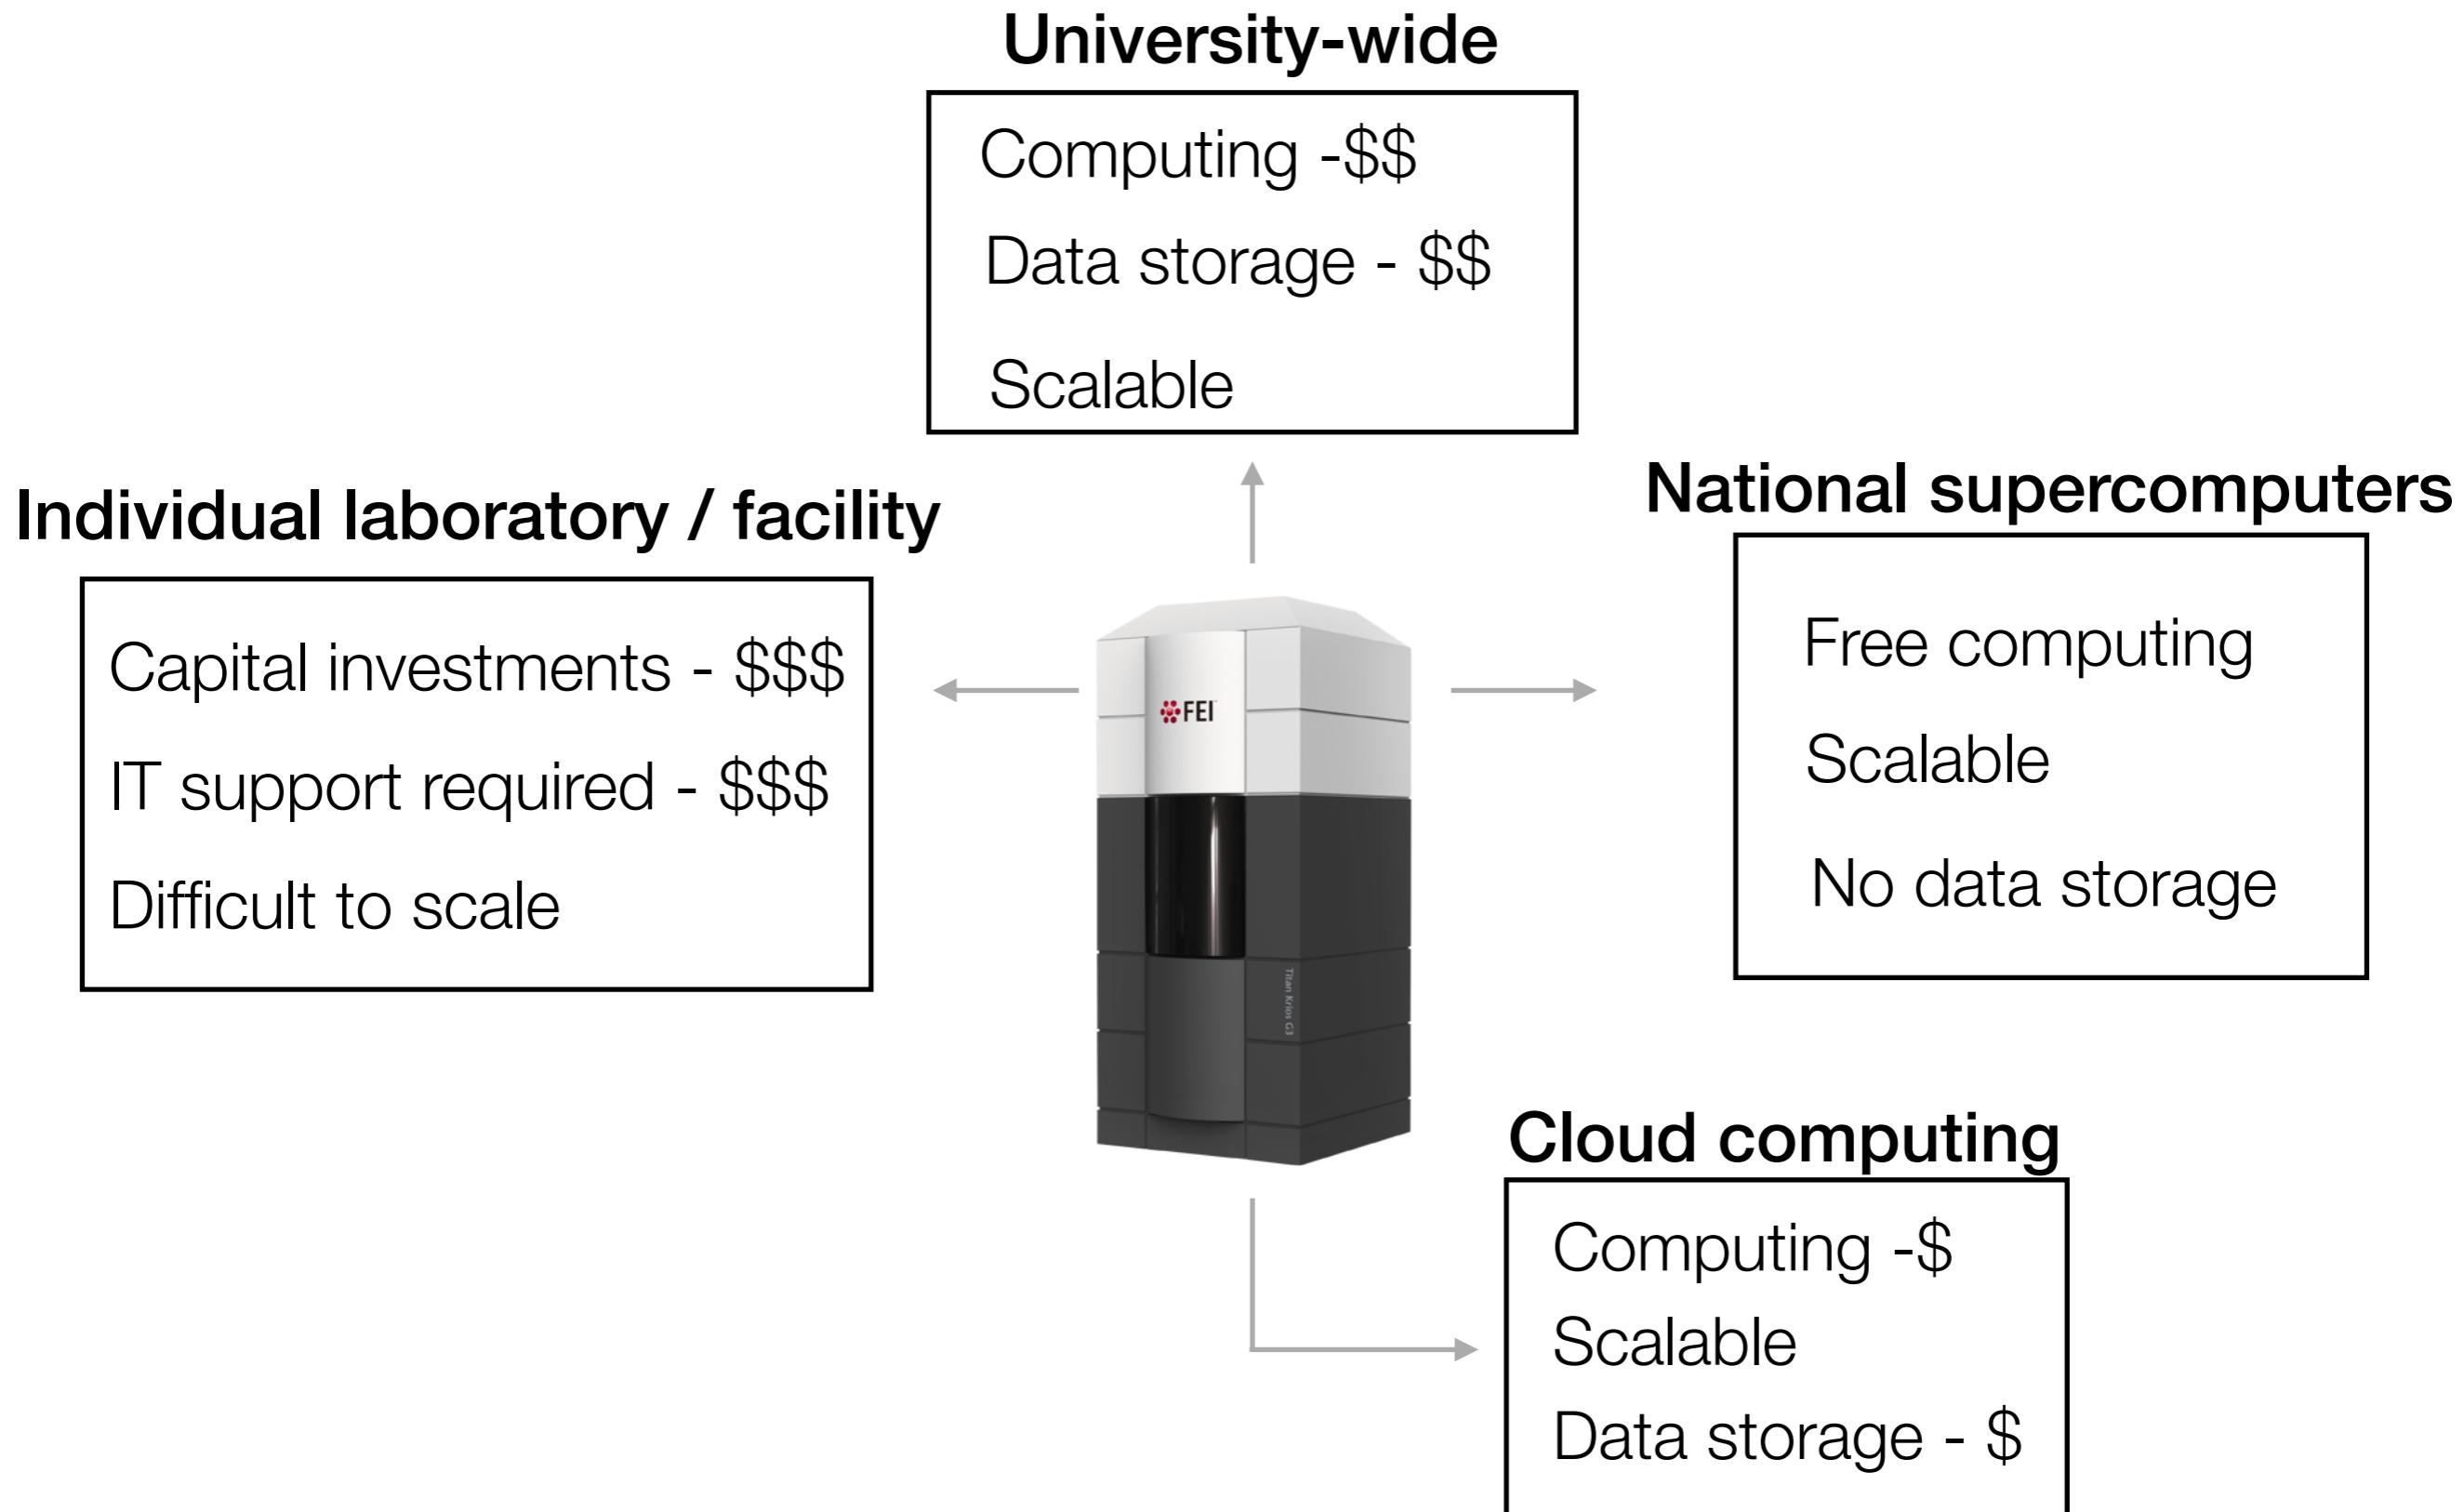


# Current state of computing in cryo-EM

---

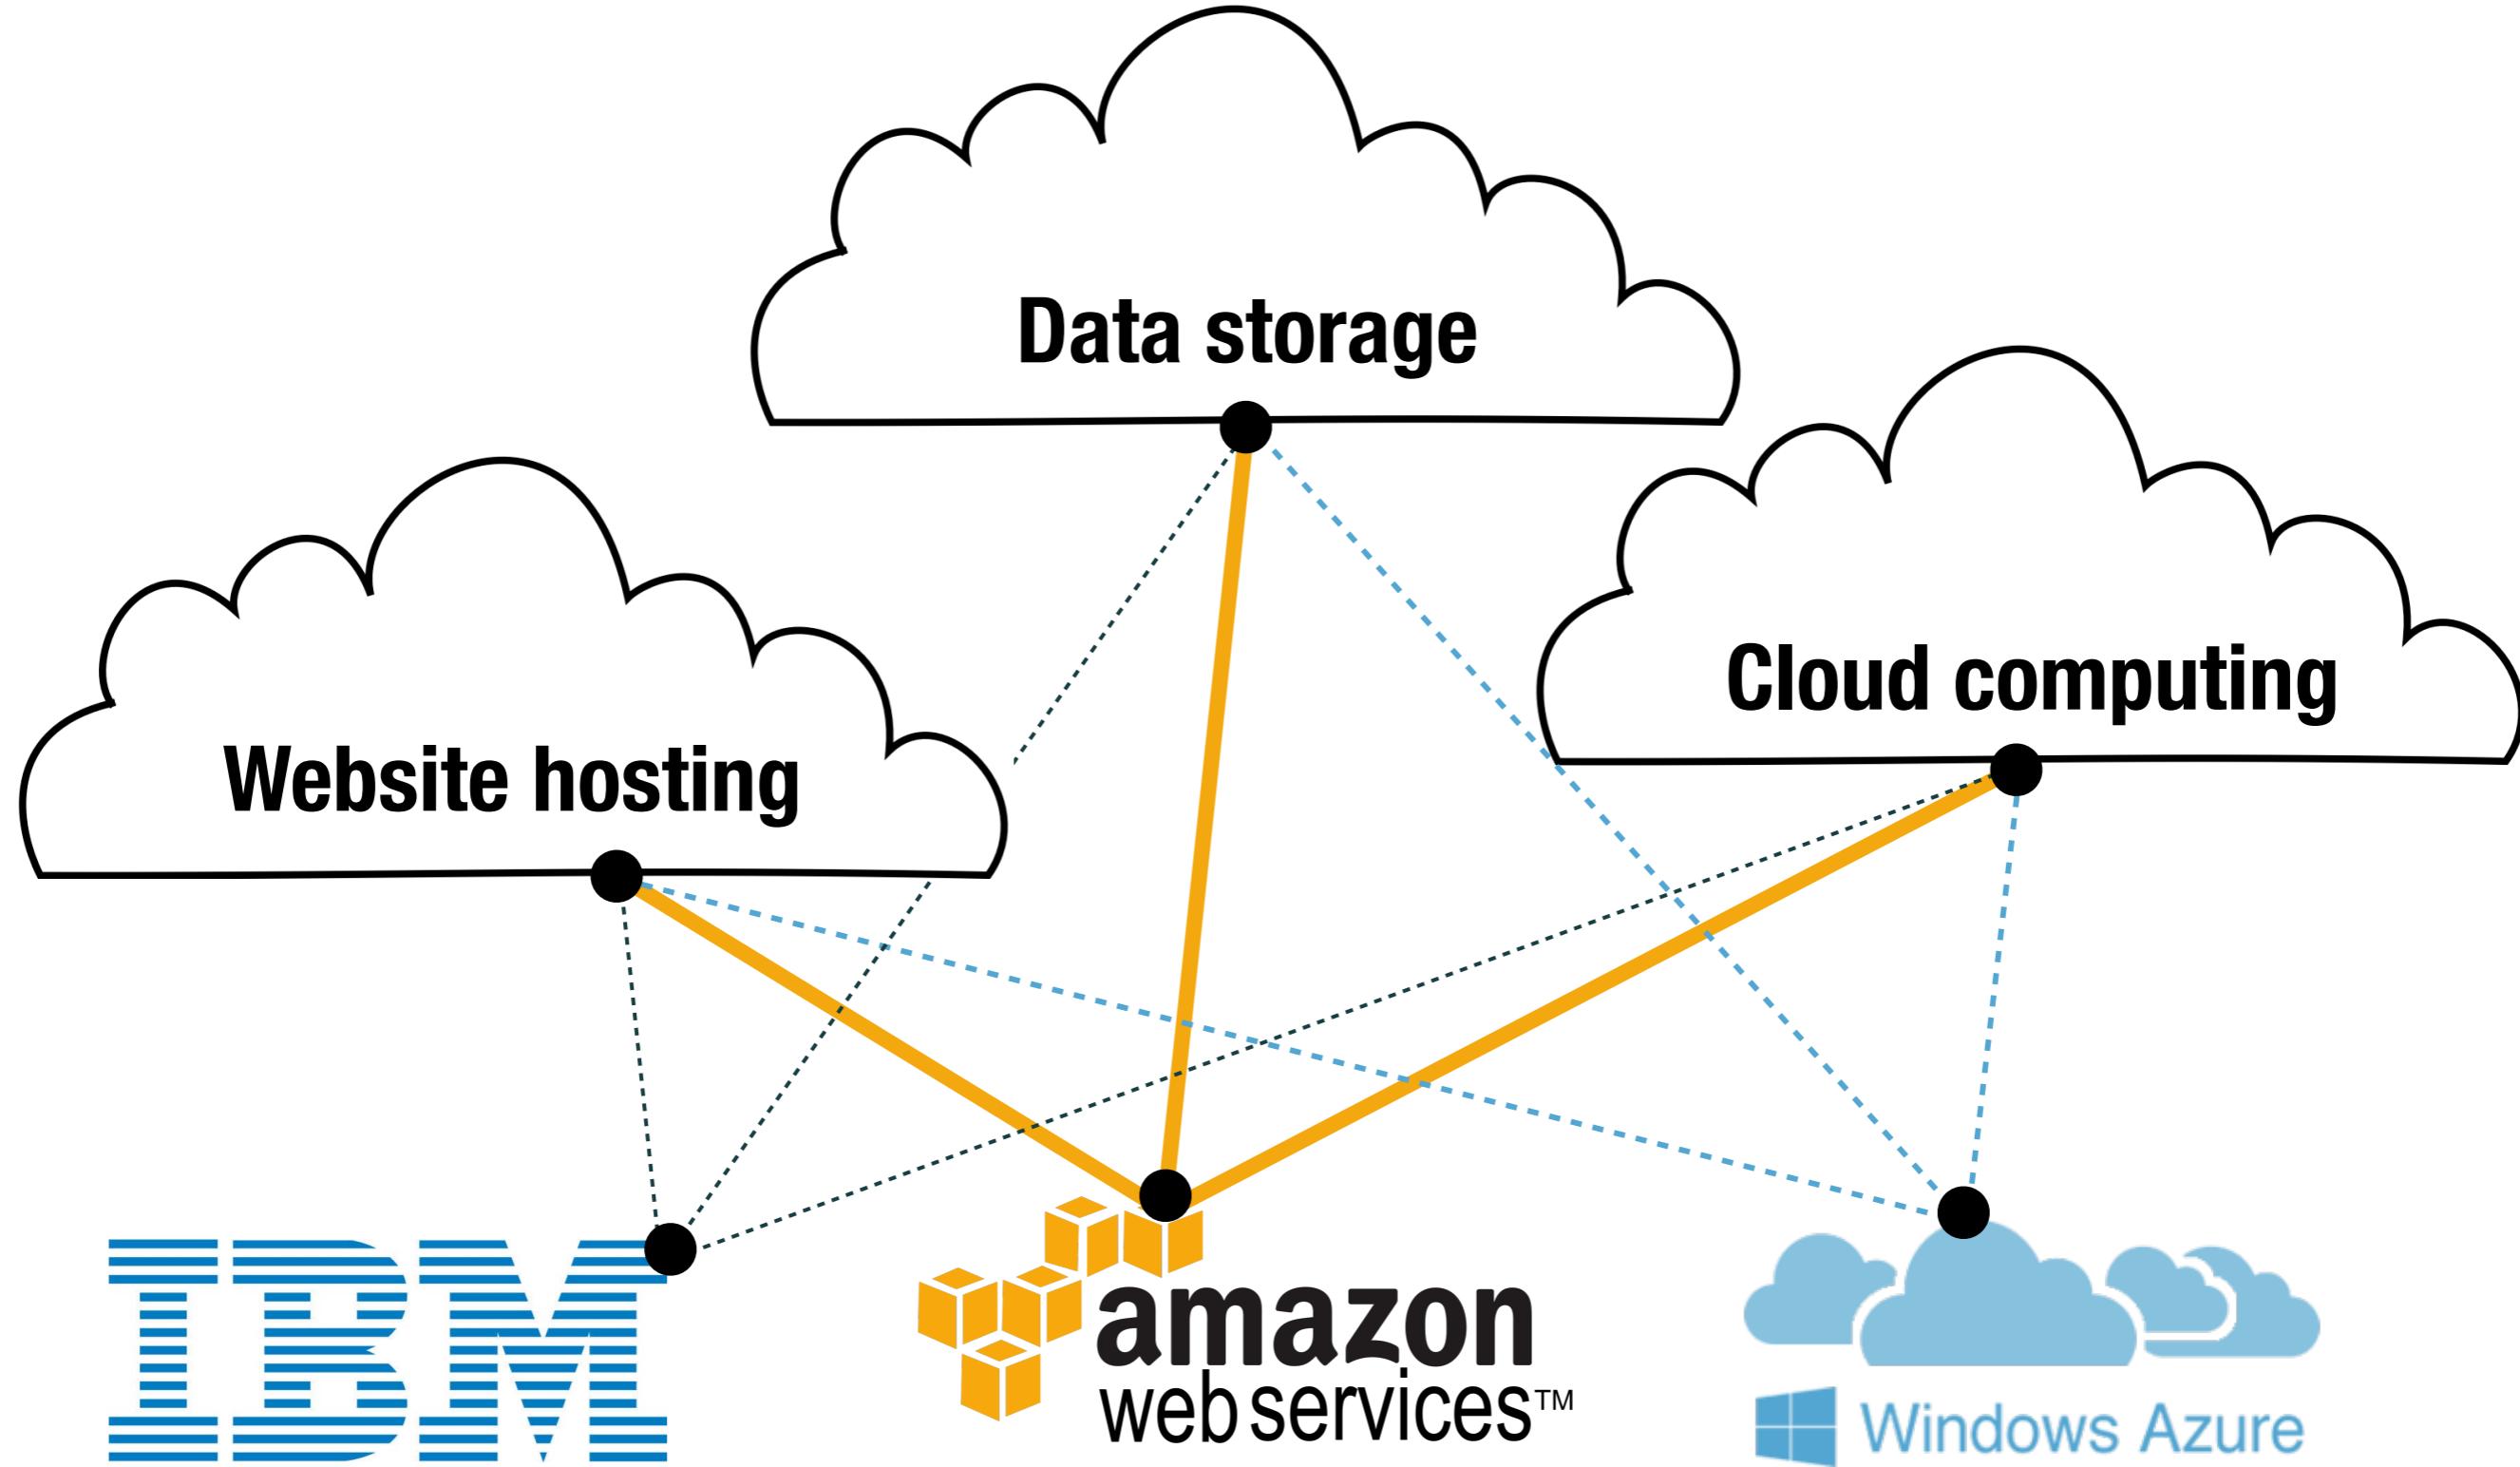


# How do we scale computing and storage for cryo-EM?



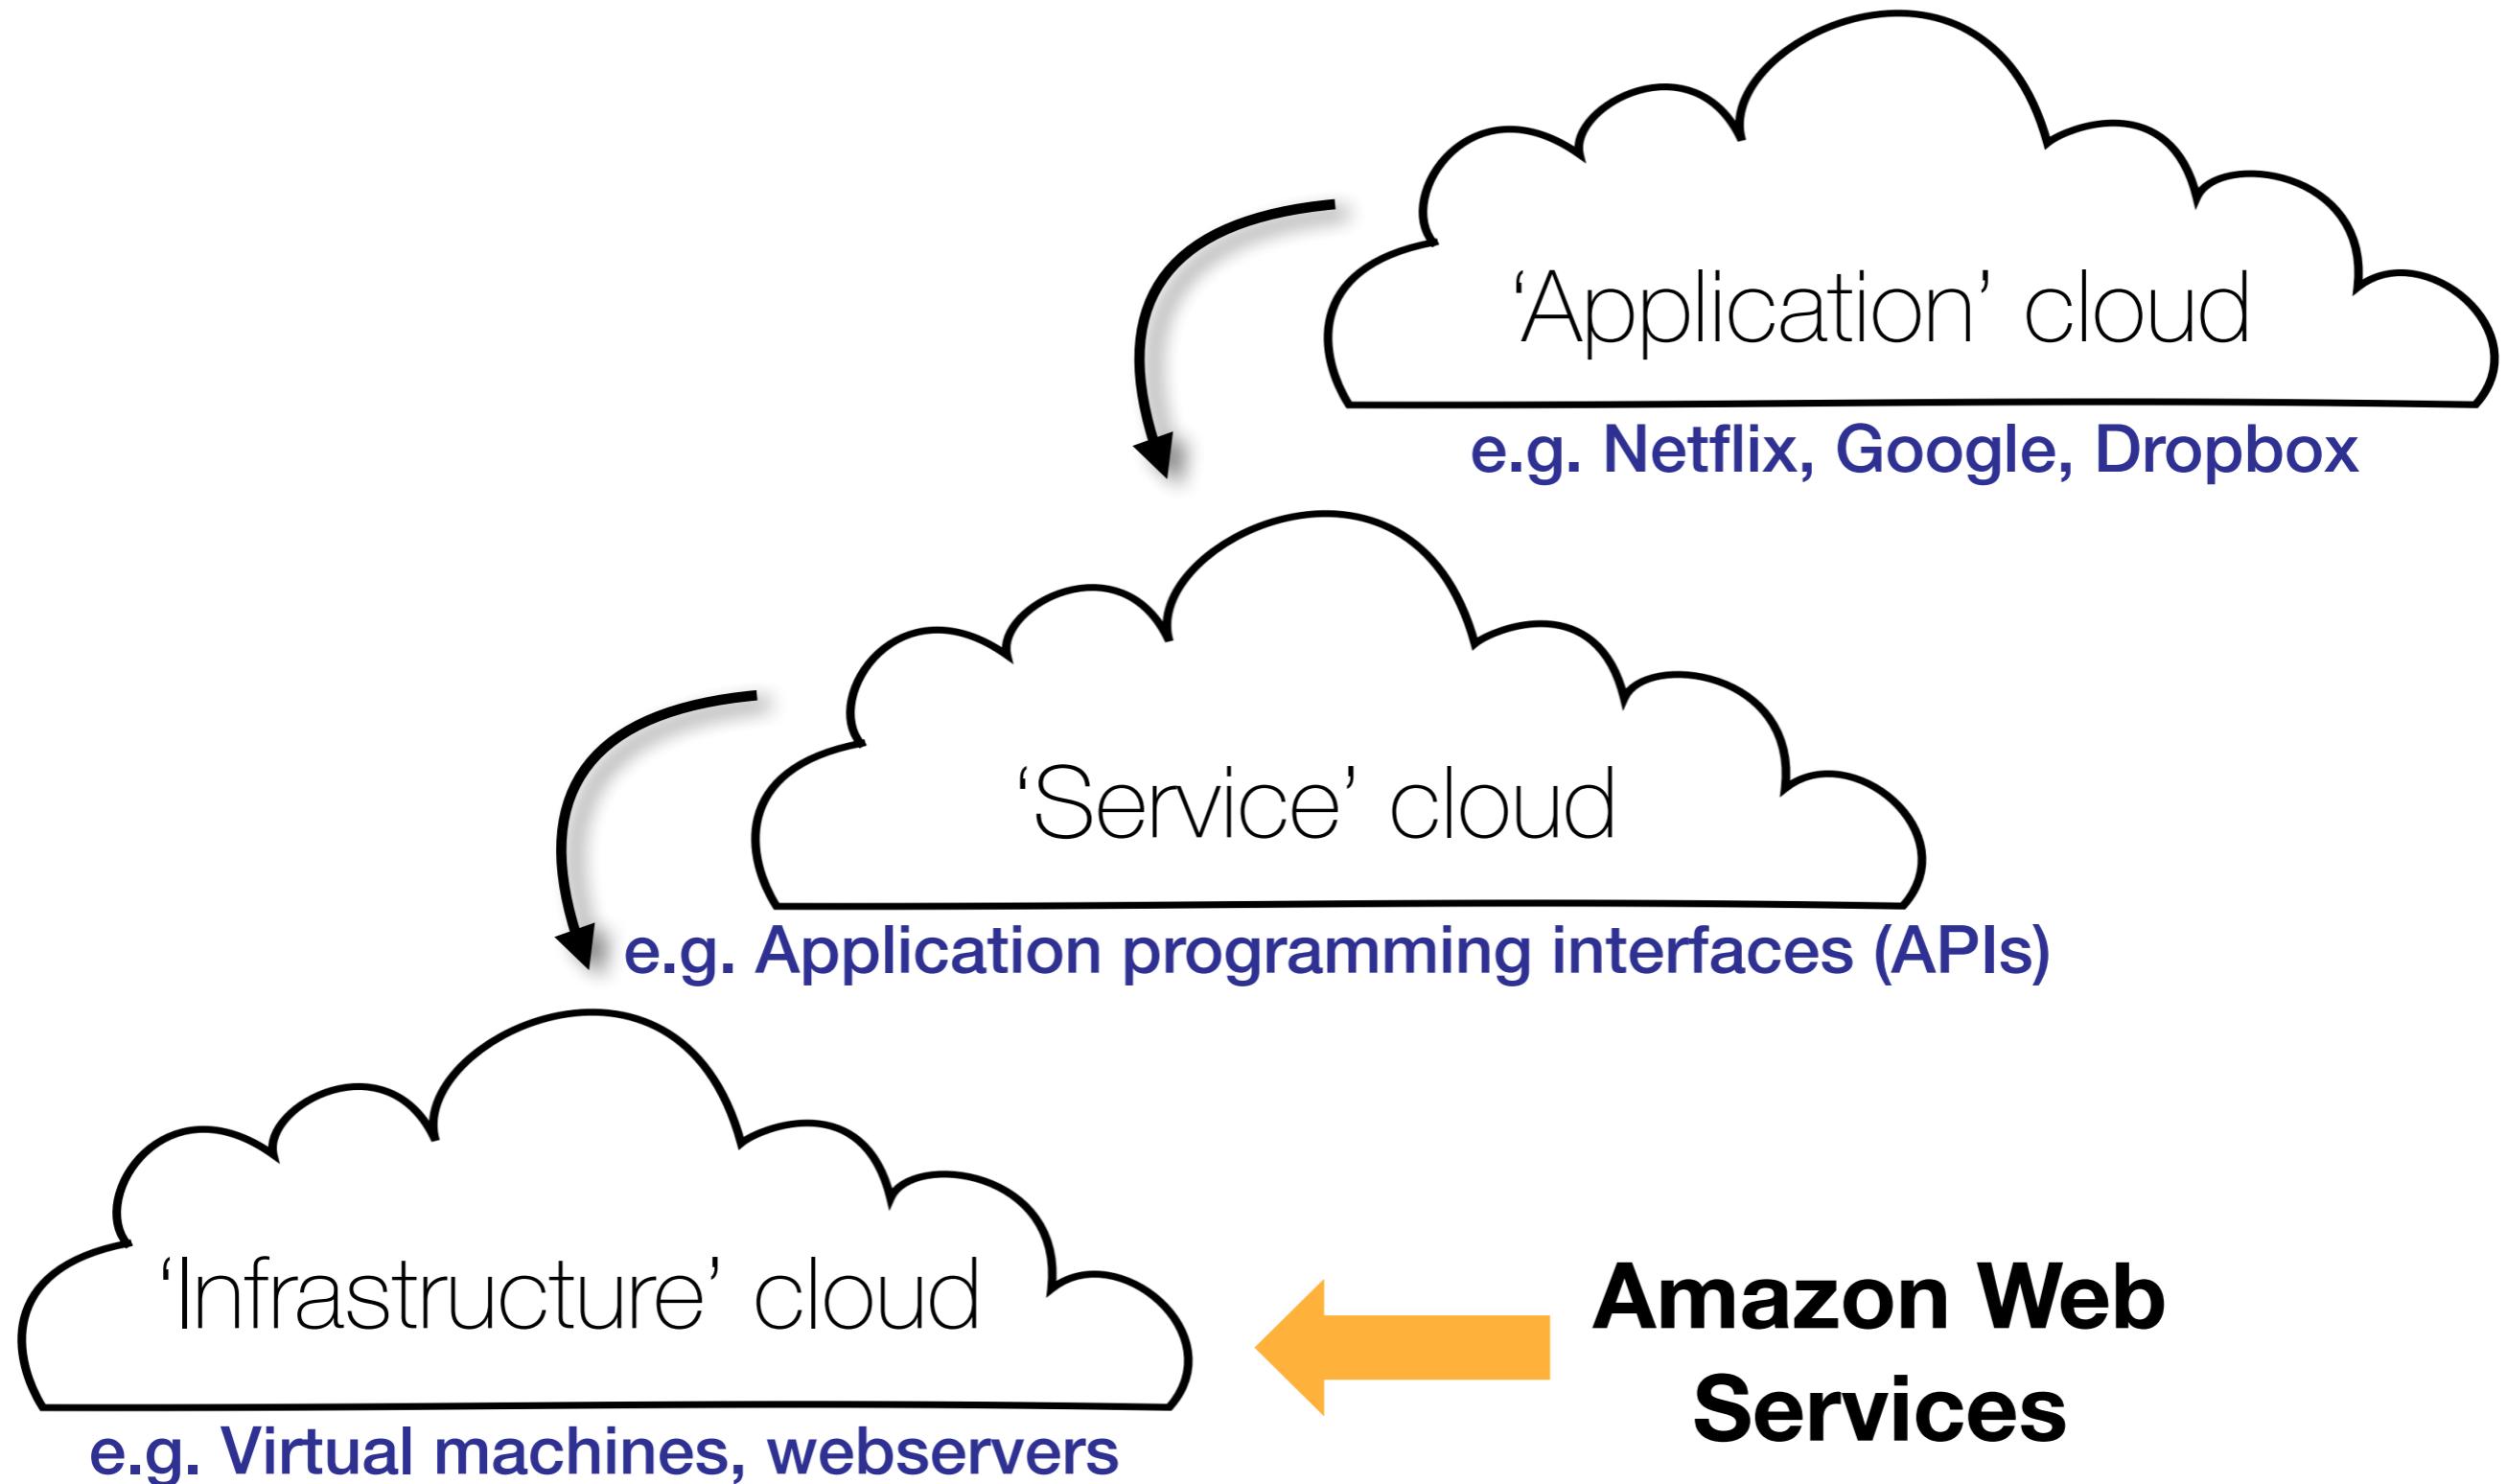
# The cloud is a ubiquitous feature of our daily lives

---



# The cloud is a ubiquitous feature of our daily lives

---



# Key principles of cloud computing

---

## 1. Cost effective

- Economies of scale
- Someone else takes care of IT hardware support
- *Pay as you go (pro-rated to the minute)*

## 2. Reliable data storage

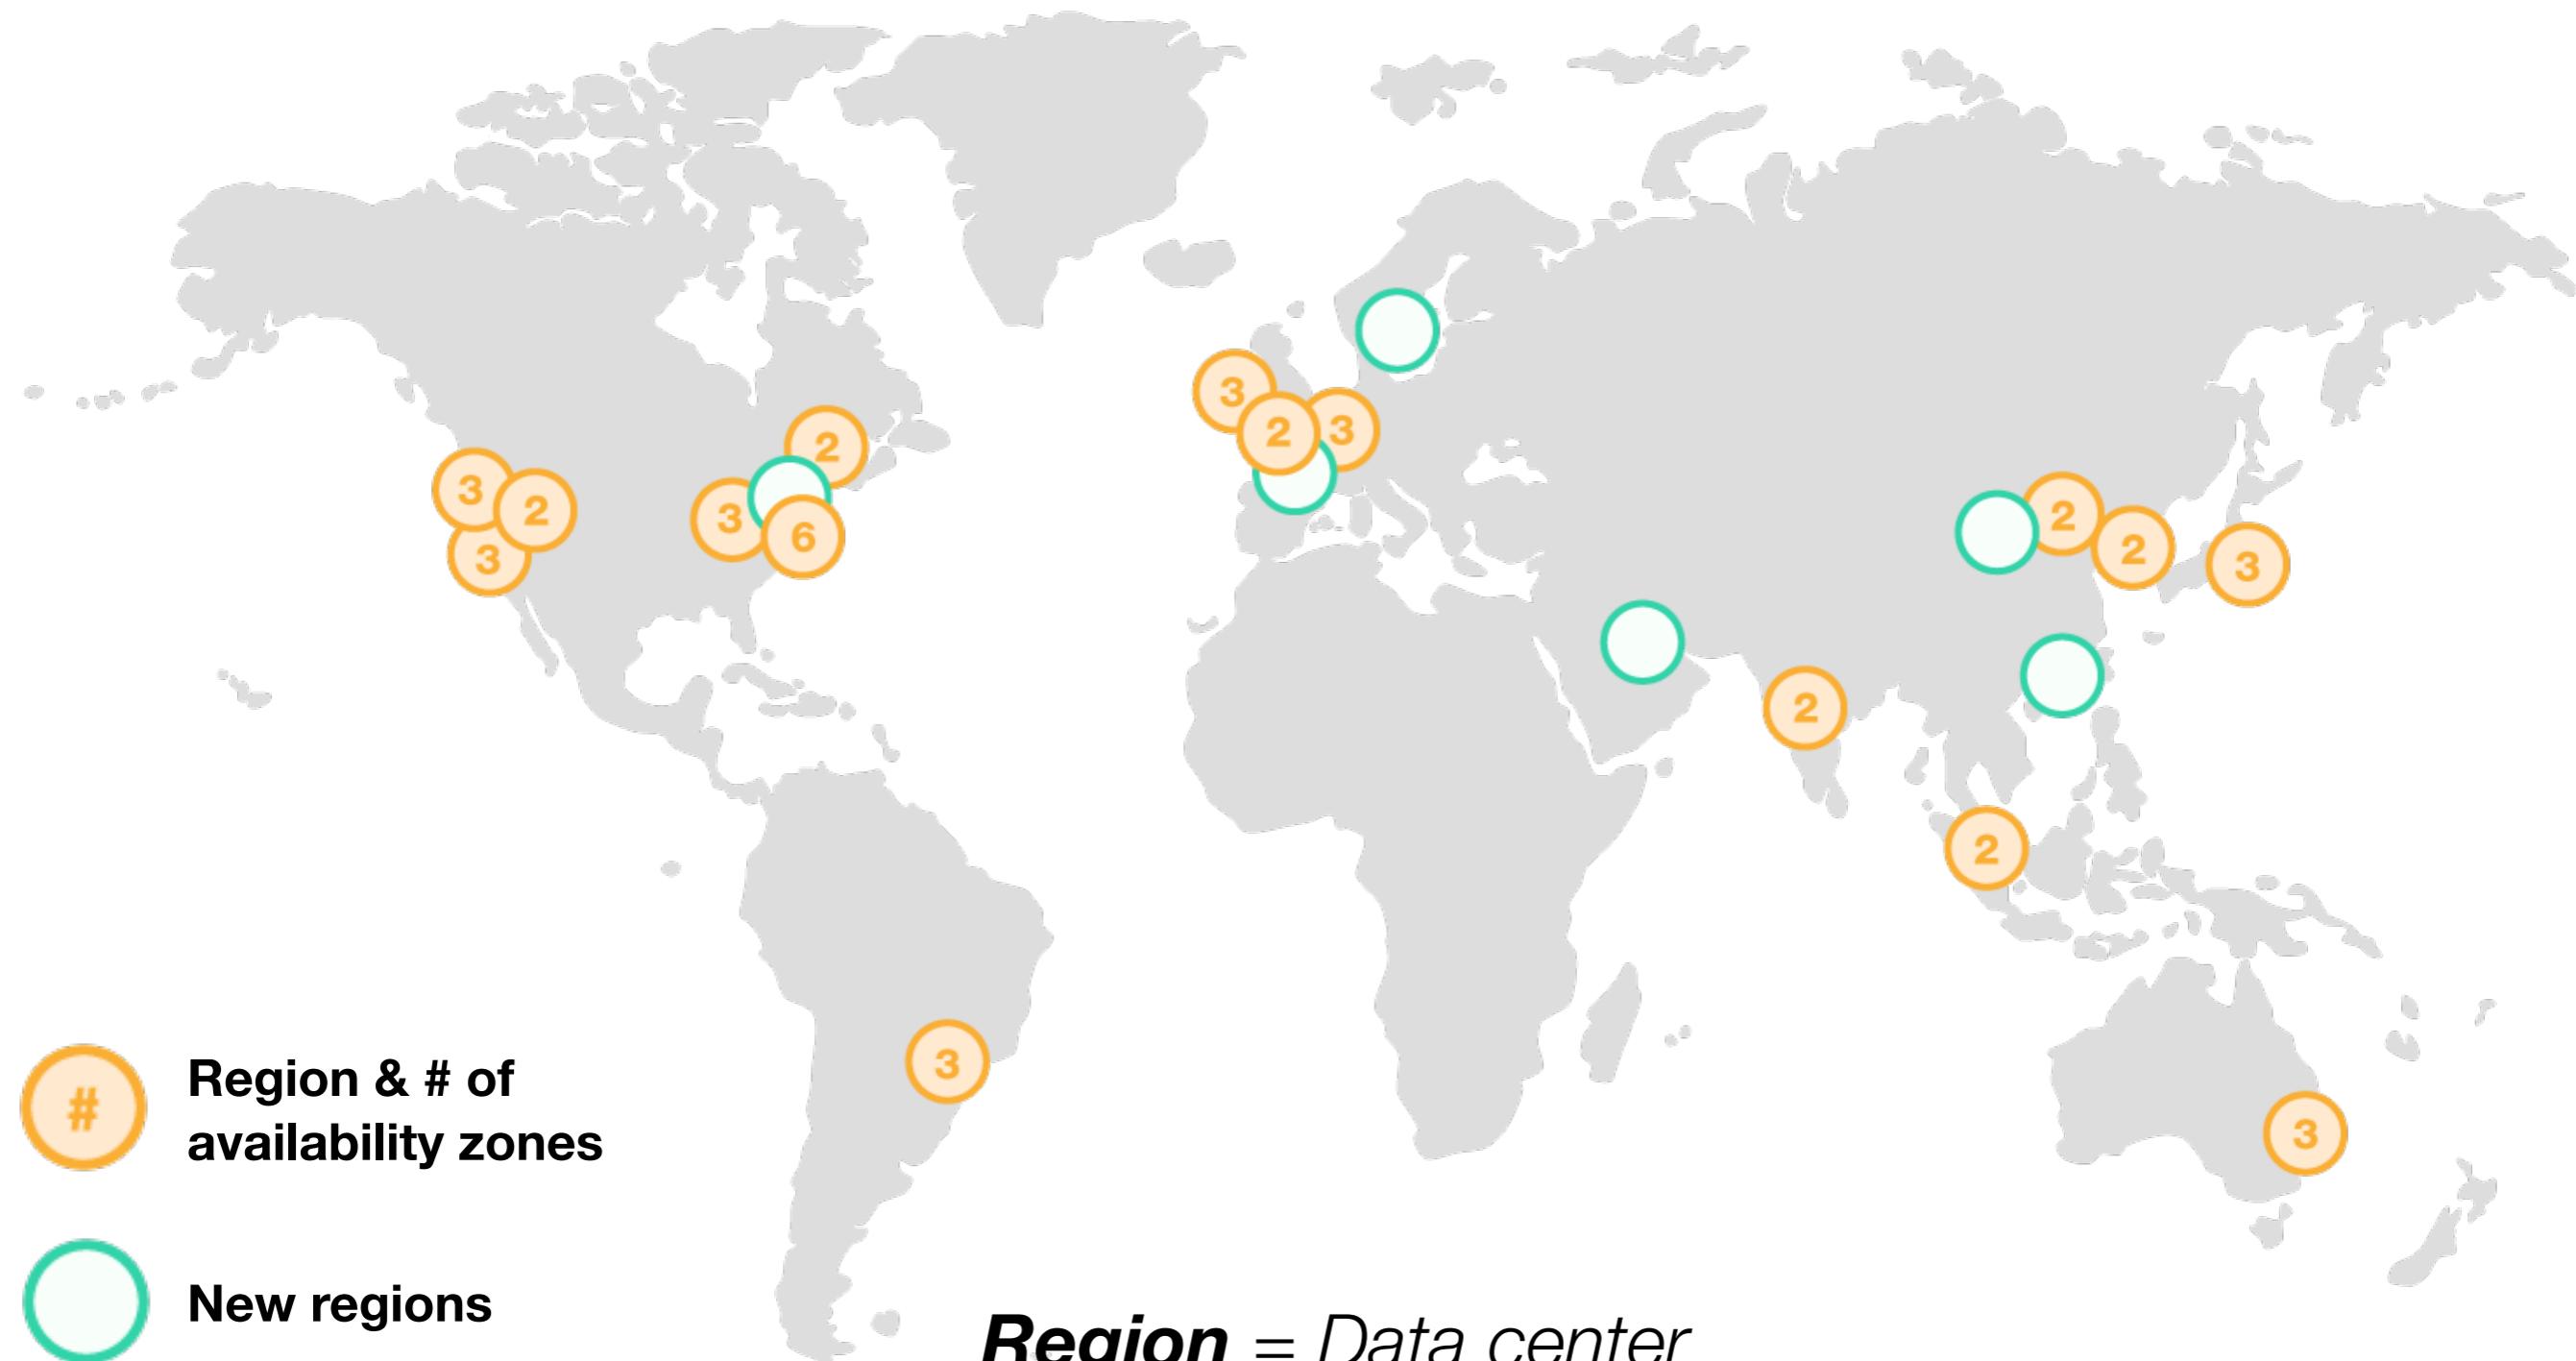
- Backed up offsite, multiple locations

## 3. Flexibility

- Global footprint
- As more computing power is needed, easy to add more
- No need to ‘choose’ between CPU vs GPU resources
- Instant access to hundreds of CPUs



# Amazon Web Services (AWS) is the world's largest cloud provider



**Region** = Data center

**Availability Zone** = Building within data center



# Computing options on AWS

**Instance** = virtual machine on AWS

**vCPU** = hyper thread on CPU core

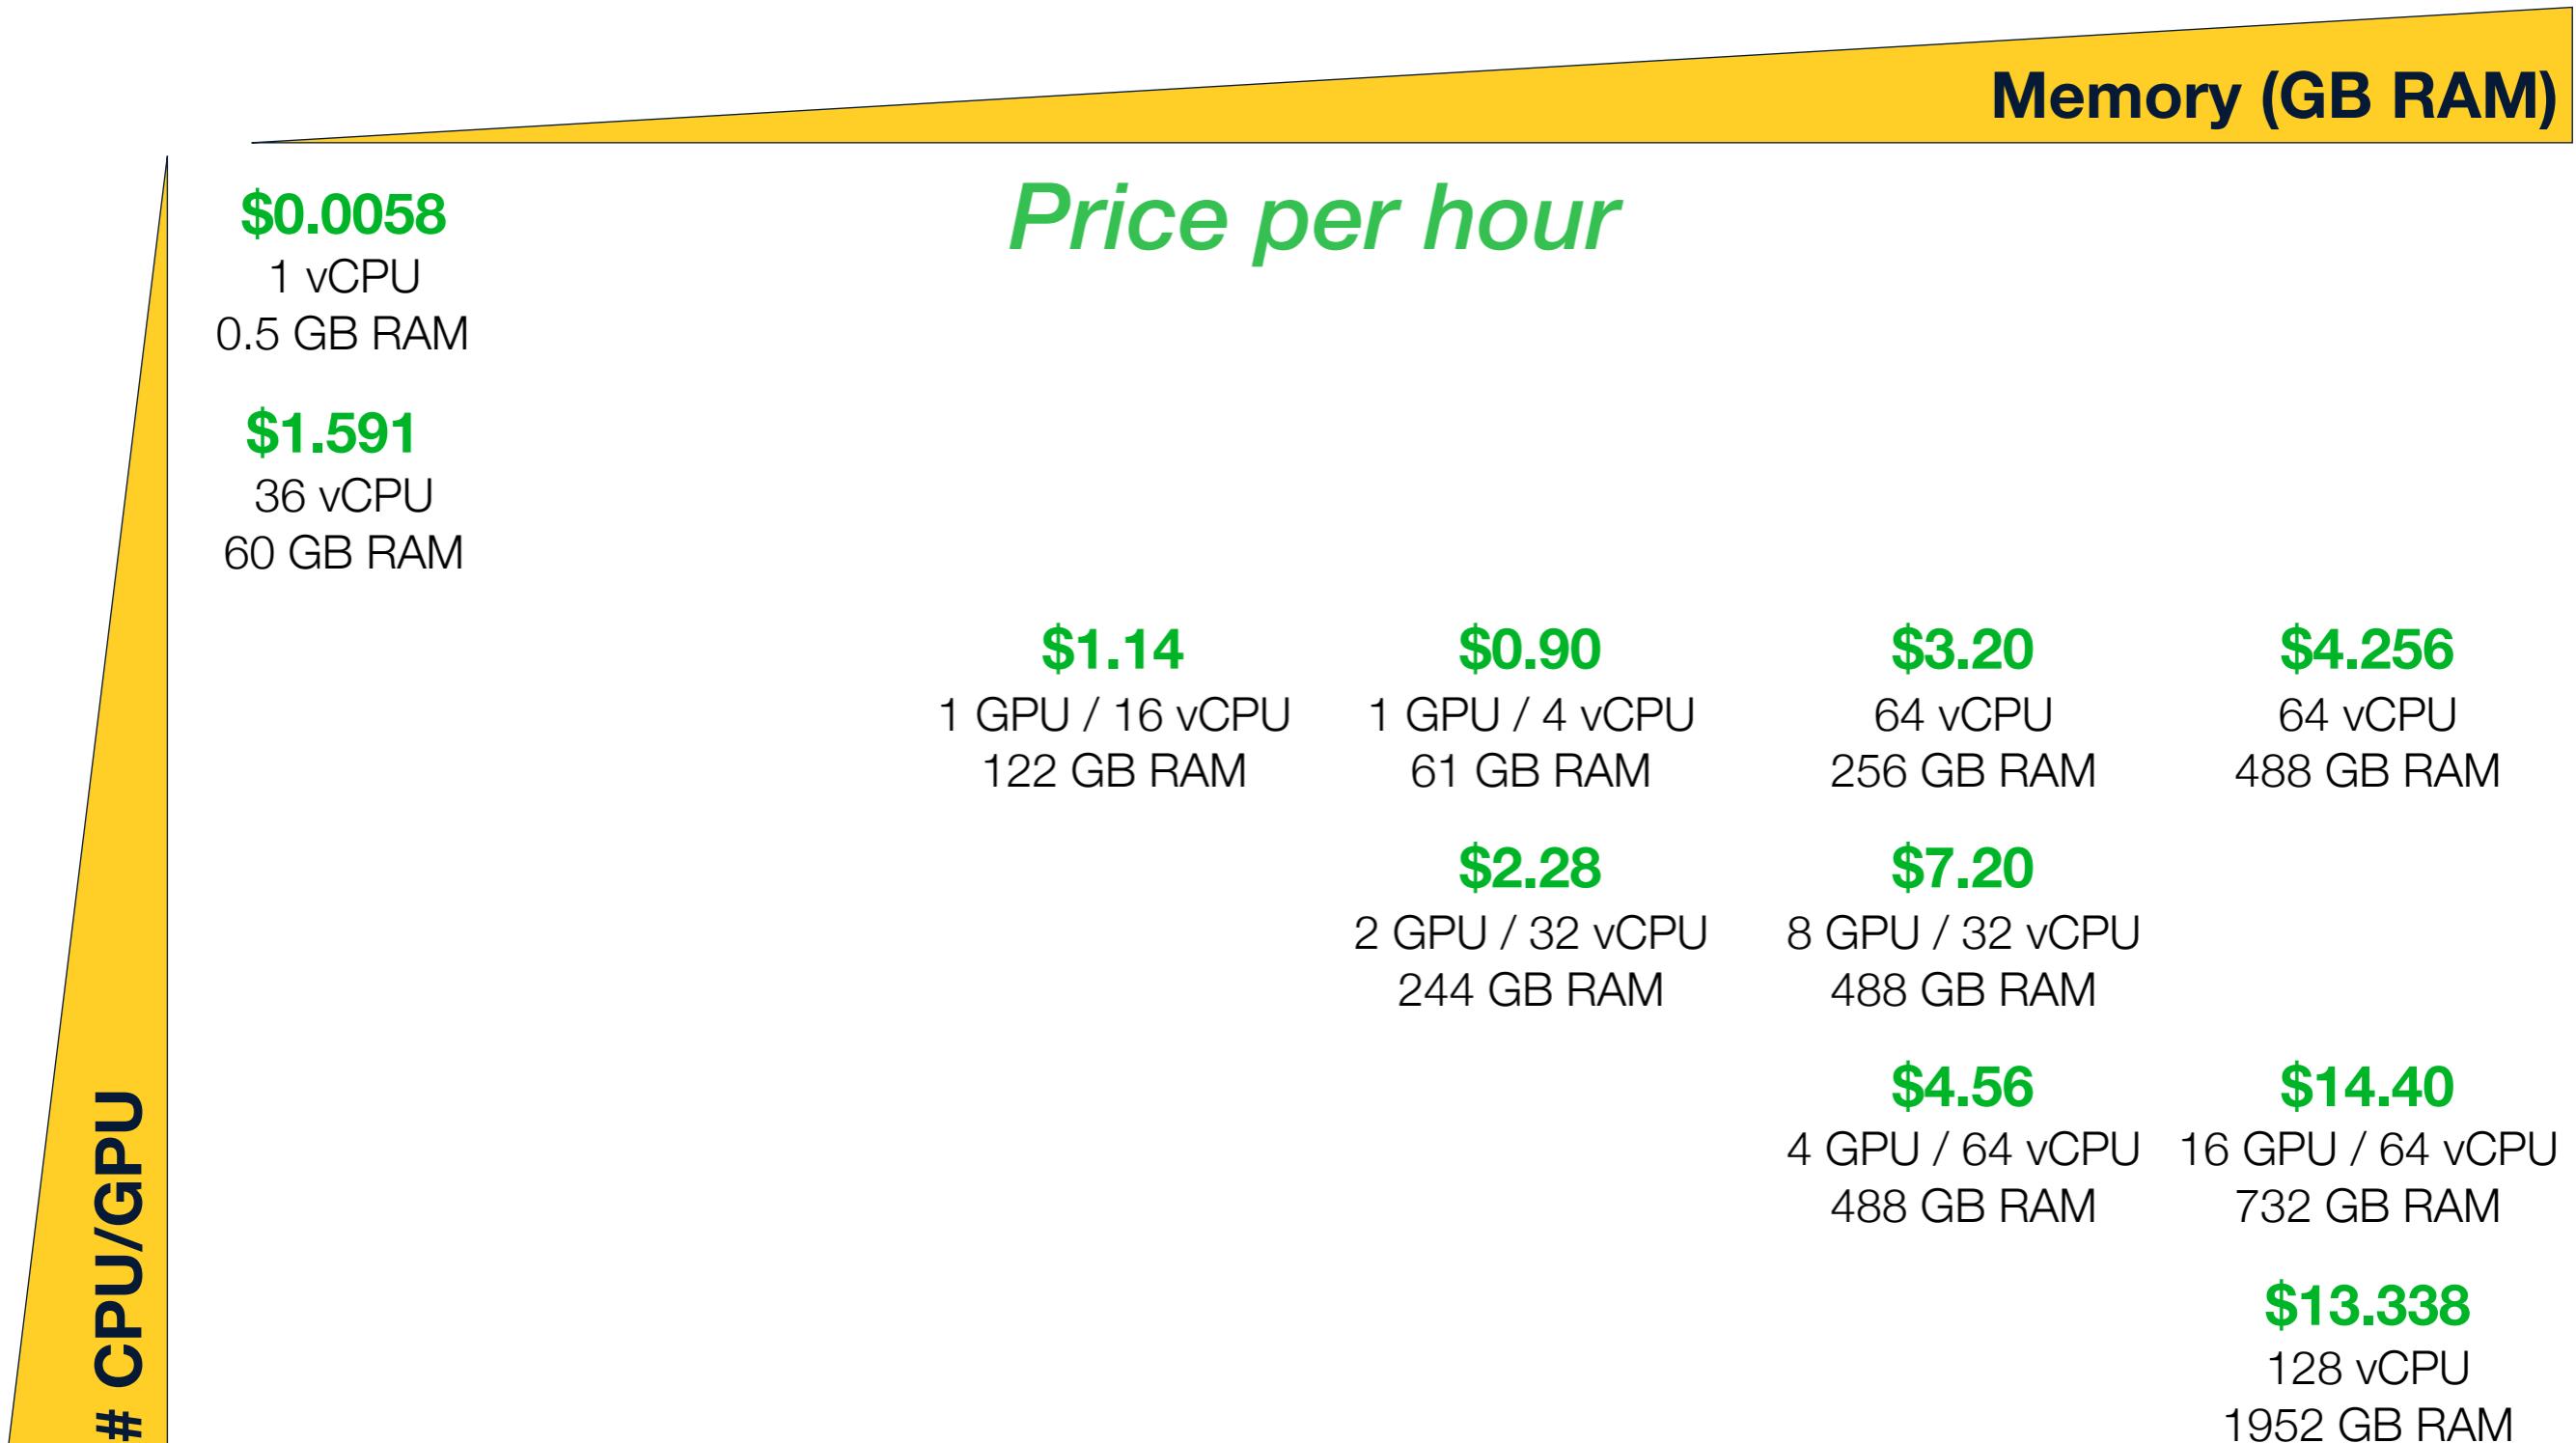
# CPU/GPU	Memory (GB RAM)			
<b>t2.nano</b>	1 vCPU	0.5 GB RAM		
<b>c4.8xlarge</b>	36 vCPU	60 GB RAM		
	<b>g3.4xlarge</b>	<b>p2.xlarge</b>	<b>m4.16xlarge</b>	<b>r4.16xlarge</b>
	1 GPU / 16 vCPU 122 GB RAM	1 GPU / 4 vCPU 61 GB RAM	64 vCPU 256 GB RAM	64 vCPU 488 GB RAM
	<b>g3.8xlarge</b>	<b>p2.8xlarge</b>		
	2 GPU / 32 vCPU 244 GB RAM	8 GPU / 32 vCPU 488 GB RAM		
	<b>g3.16xlarge</b>	<b>p2.16xlarge</b>		
	4 GPU / 64 vCPU 488 GB RAM	16 GPU / 64 vCPU 732 GB RAM		
	<b>x1.32xlarge</b>			
			128 vCPU 1952 GB RAM	
<b>...59 total instances (48 more than shown here)</b>				



# Computing options on AWS

**Instance** = virtual machine on AWS

**vCPU** = hyper thread on CPU core



# Reserving instances on AWS

*Pay-per-hour*

## On-demand

Pay ‘standard’ price  
Instance is 100% yours

## Reserved

Pay 1-3 yr upfront to reserve instance  
Instance is 100% yours

## Spot

Bid on marketplace  
Instance is yours until somehow outbids you

Cheaper

Cheaper



# Data storage on AWS

---

## Local SSDs

Directly attached to instance

No cost

## Elastic block storage (EBS)

SSDs that can be attached/detached

\$0.10/GB/mo.

***Object storage:*** Scalable storage (active & cold)

## Simple storage service (S3)

\$0.023/GB/mo.

## Glacier

\$0.004/GB/mo.



# Previously...

## Problems:

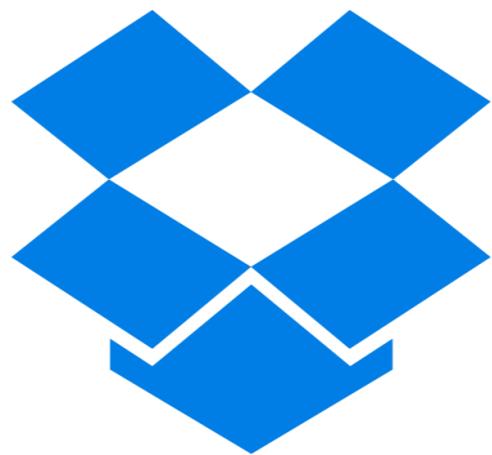
- ▶ **Manual** deployment & management of AWS resources
- ▶ **Cumbersome** data movement
- ▶ Medium/expert-level **linux experience required** to interface

Pay < \$10/hr for 150+ CPUs with spot pricing

# What is the best way to use the cloud?

---

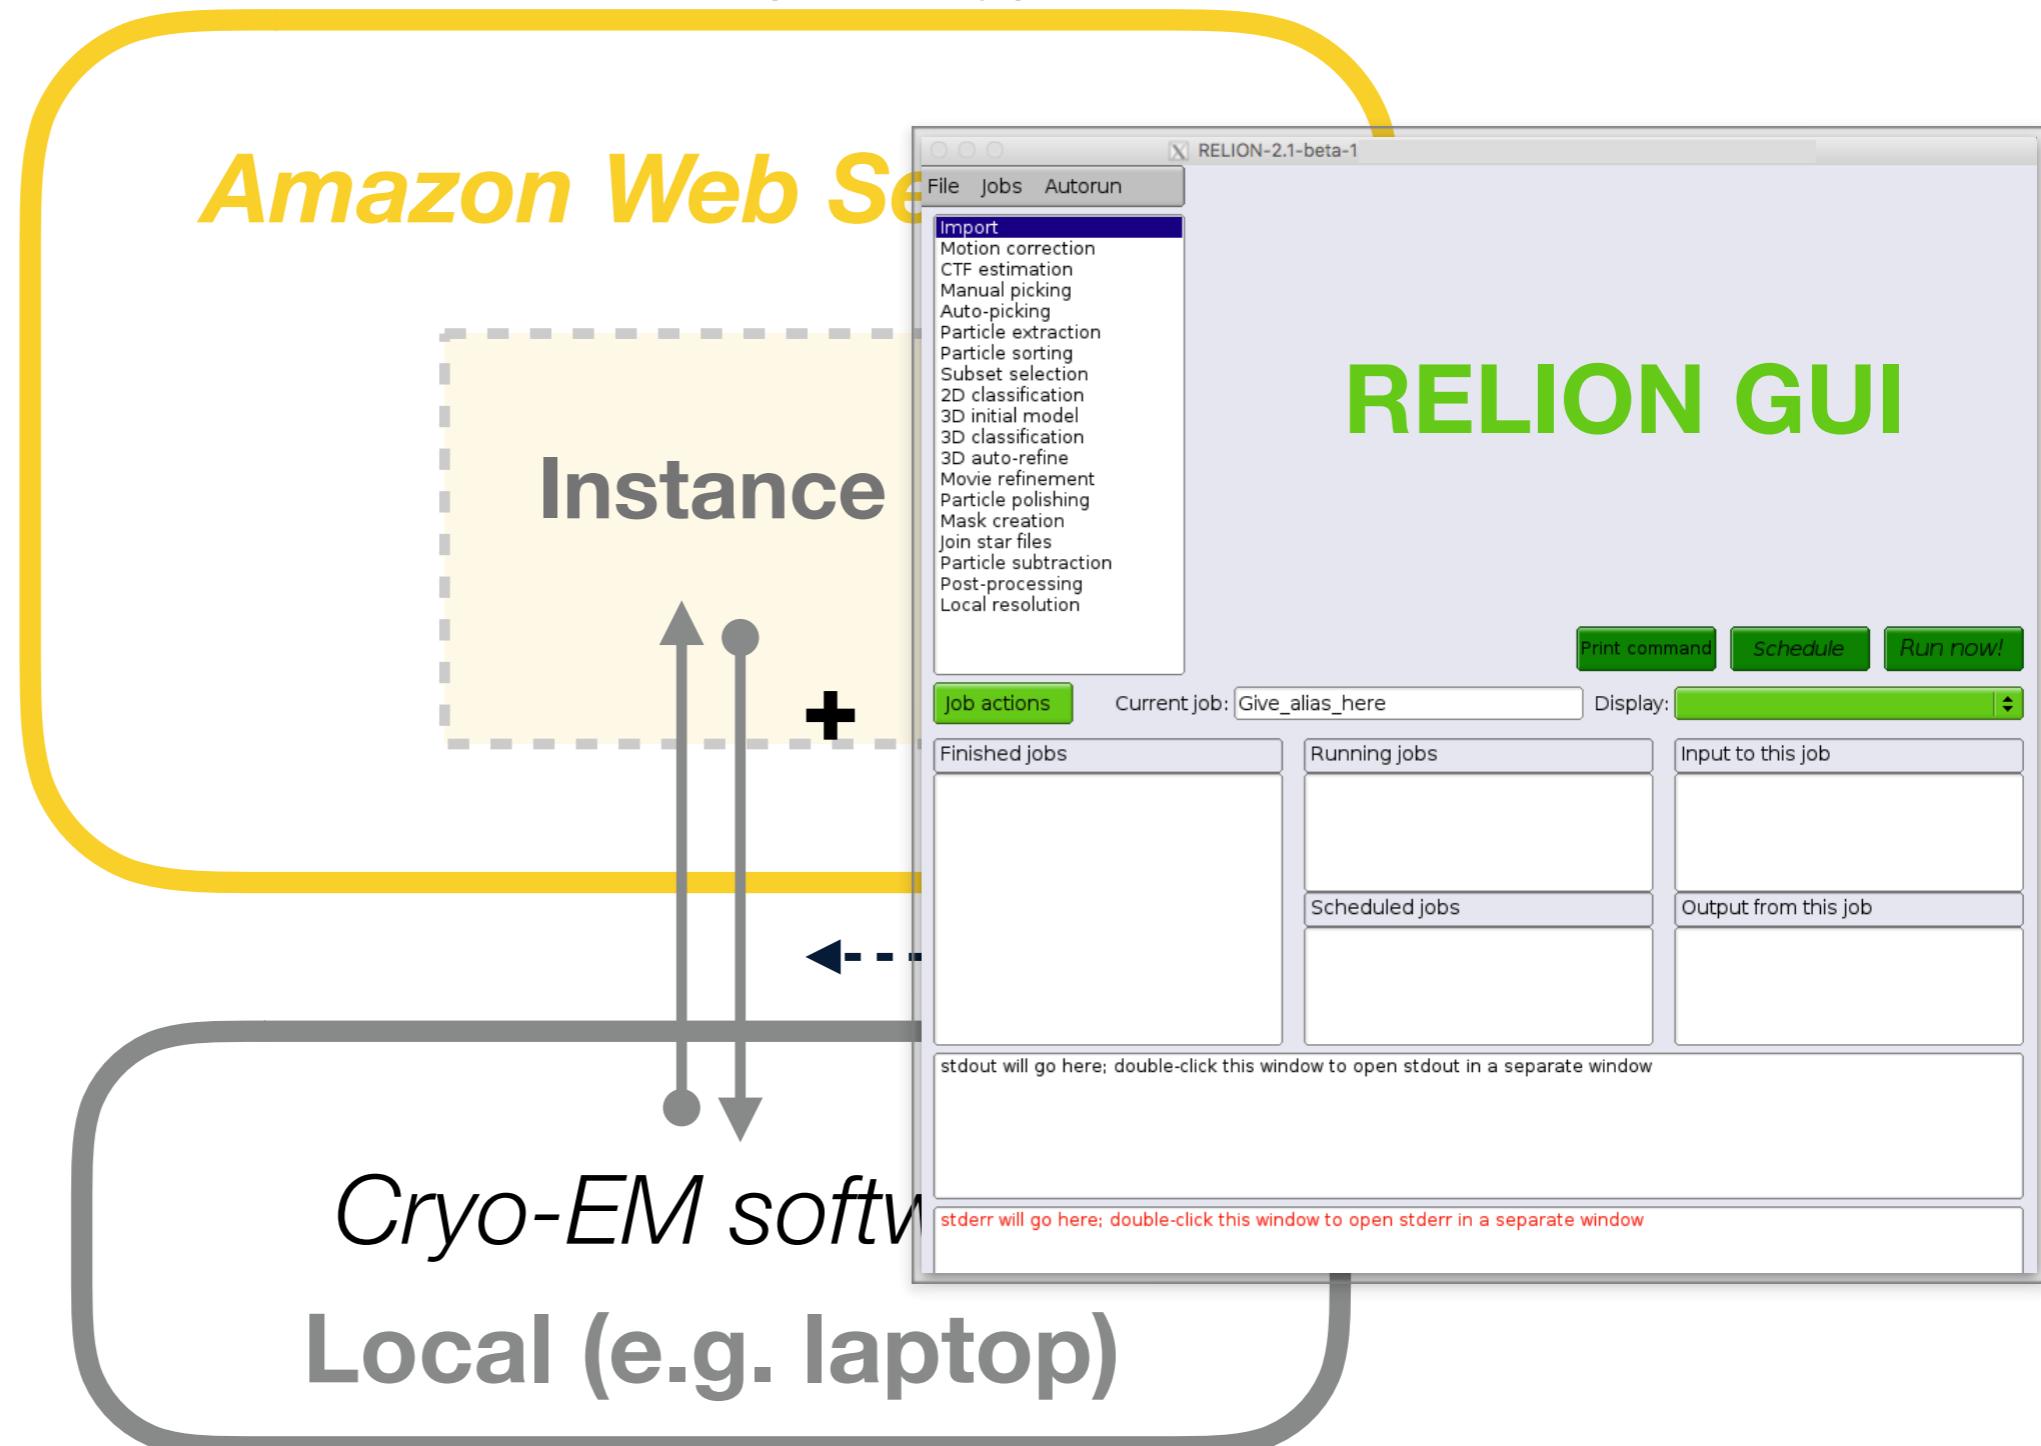
*“Hybrid cloud architecture is the integration of on-premises resources with cloud resources.”*



e.g. **Dropbox**

# Building a hybrid cloud infrastructure for cryo-EM

*RELION as a prototype*



# Building a hybrid cloud infrastructure for cryo-EM

*RELION as a prototype*

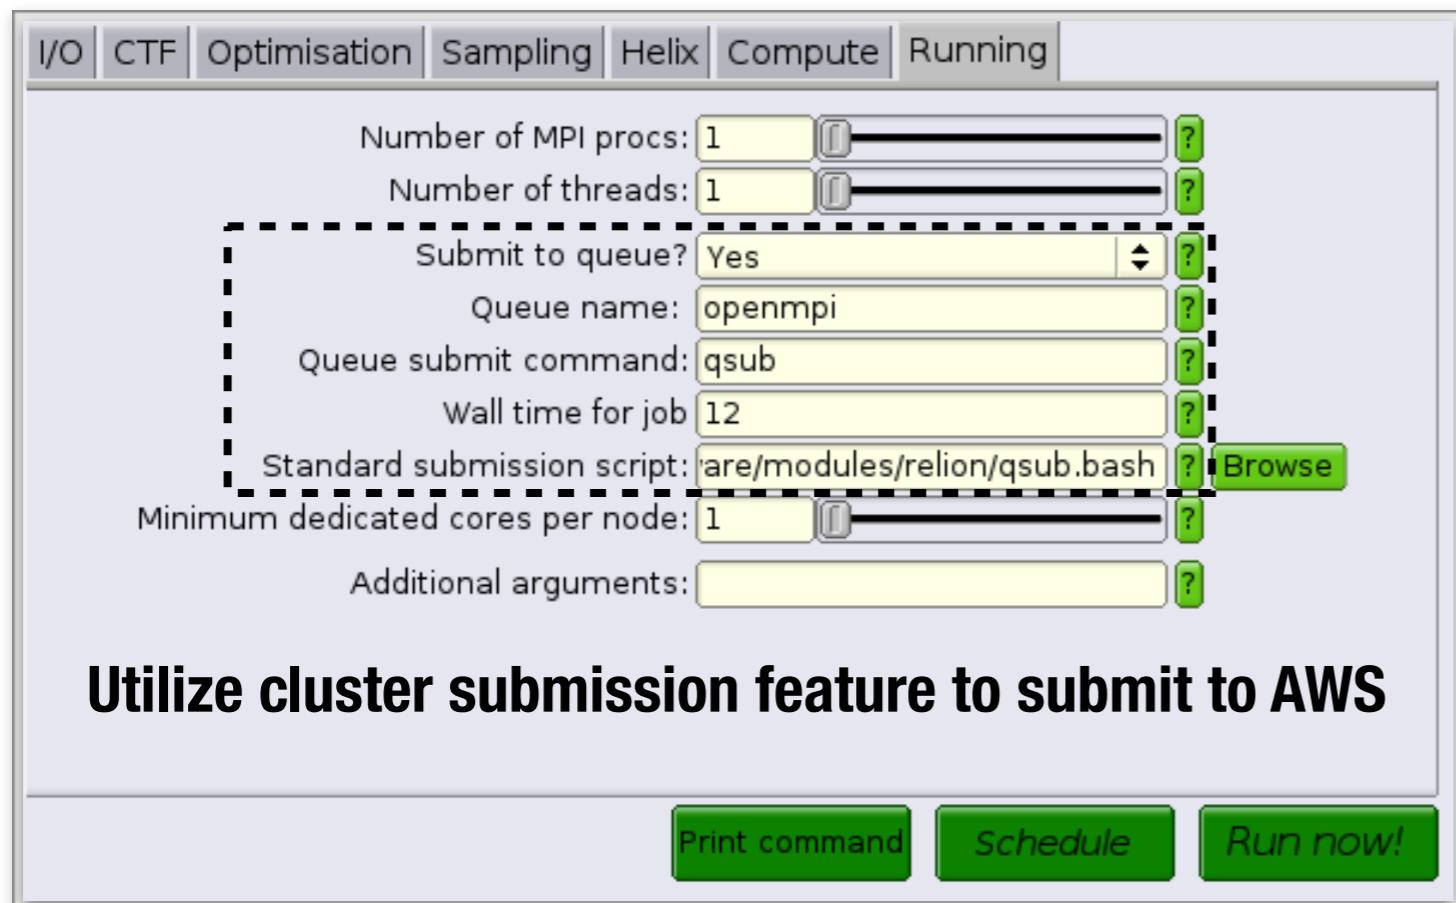
I/O | CTF | Optimisation | Sampling | Helix | Compute | Running

Number of MPI procs: 1  
Number of threads: 1

Submit to queue? Yes  
Queue name: openmpi  
Queue submit command: qsub  
Wall time for job: 12  
Standard submission script: /are/modules/relion/qsub.bash  
Minimum dedicated cores per node: 1  
Additional arguments:

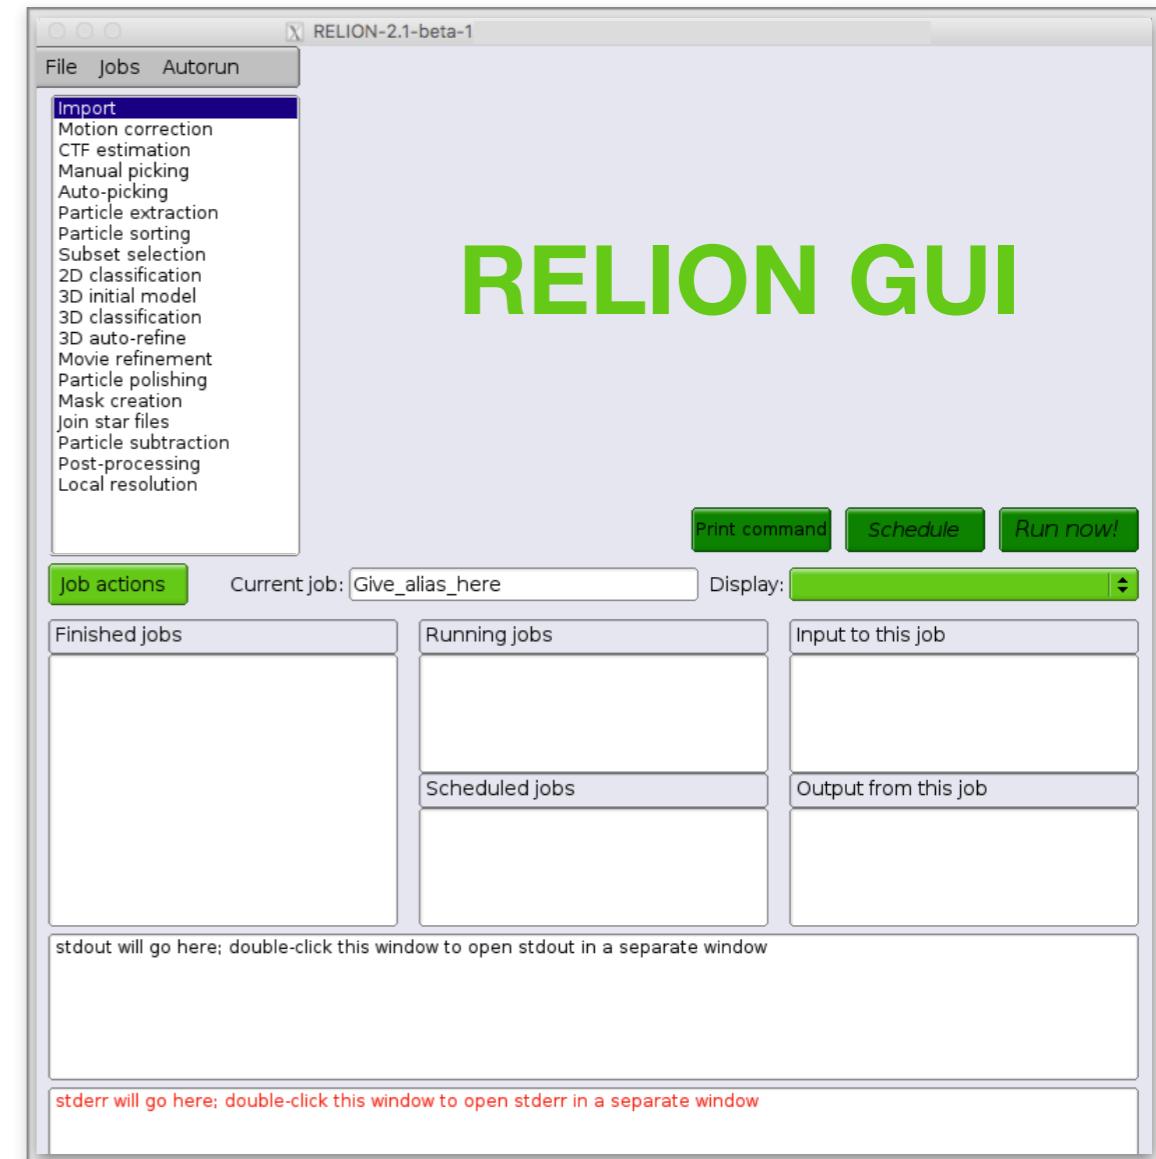
**Utilize cluster submission feature to submit to AWS**

**Print command** **Schedule** **Run now!**



# Building a hybrid cloud infrastructure for cryo-EM

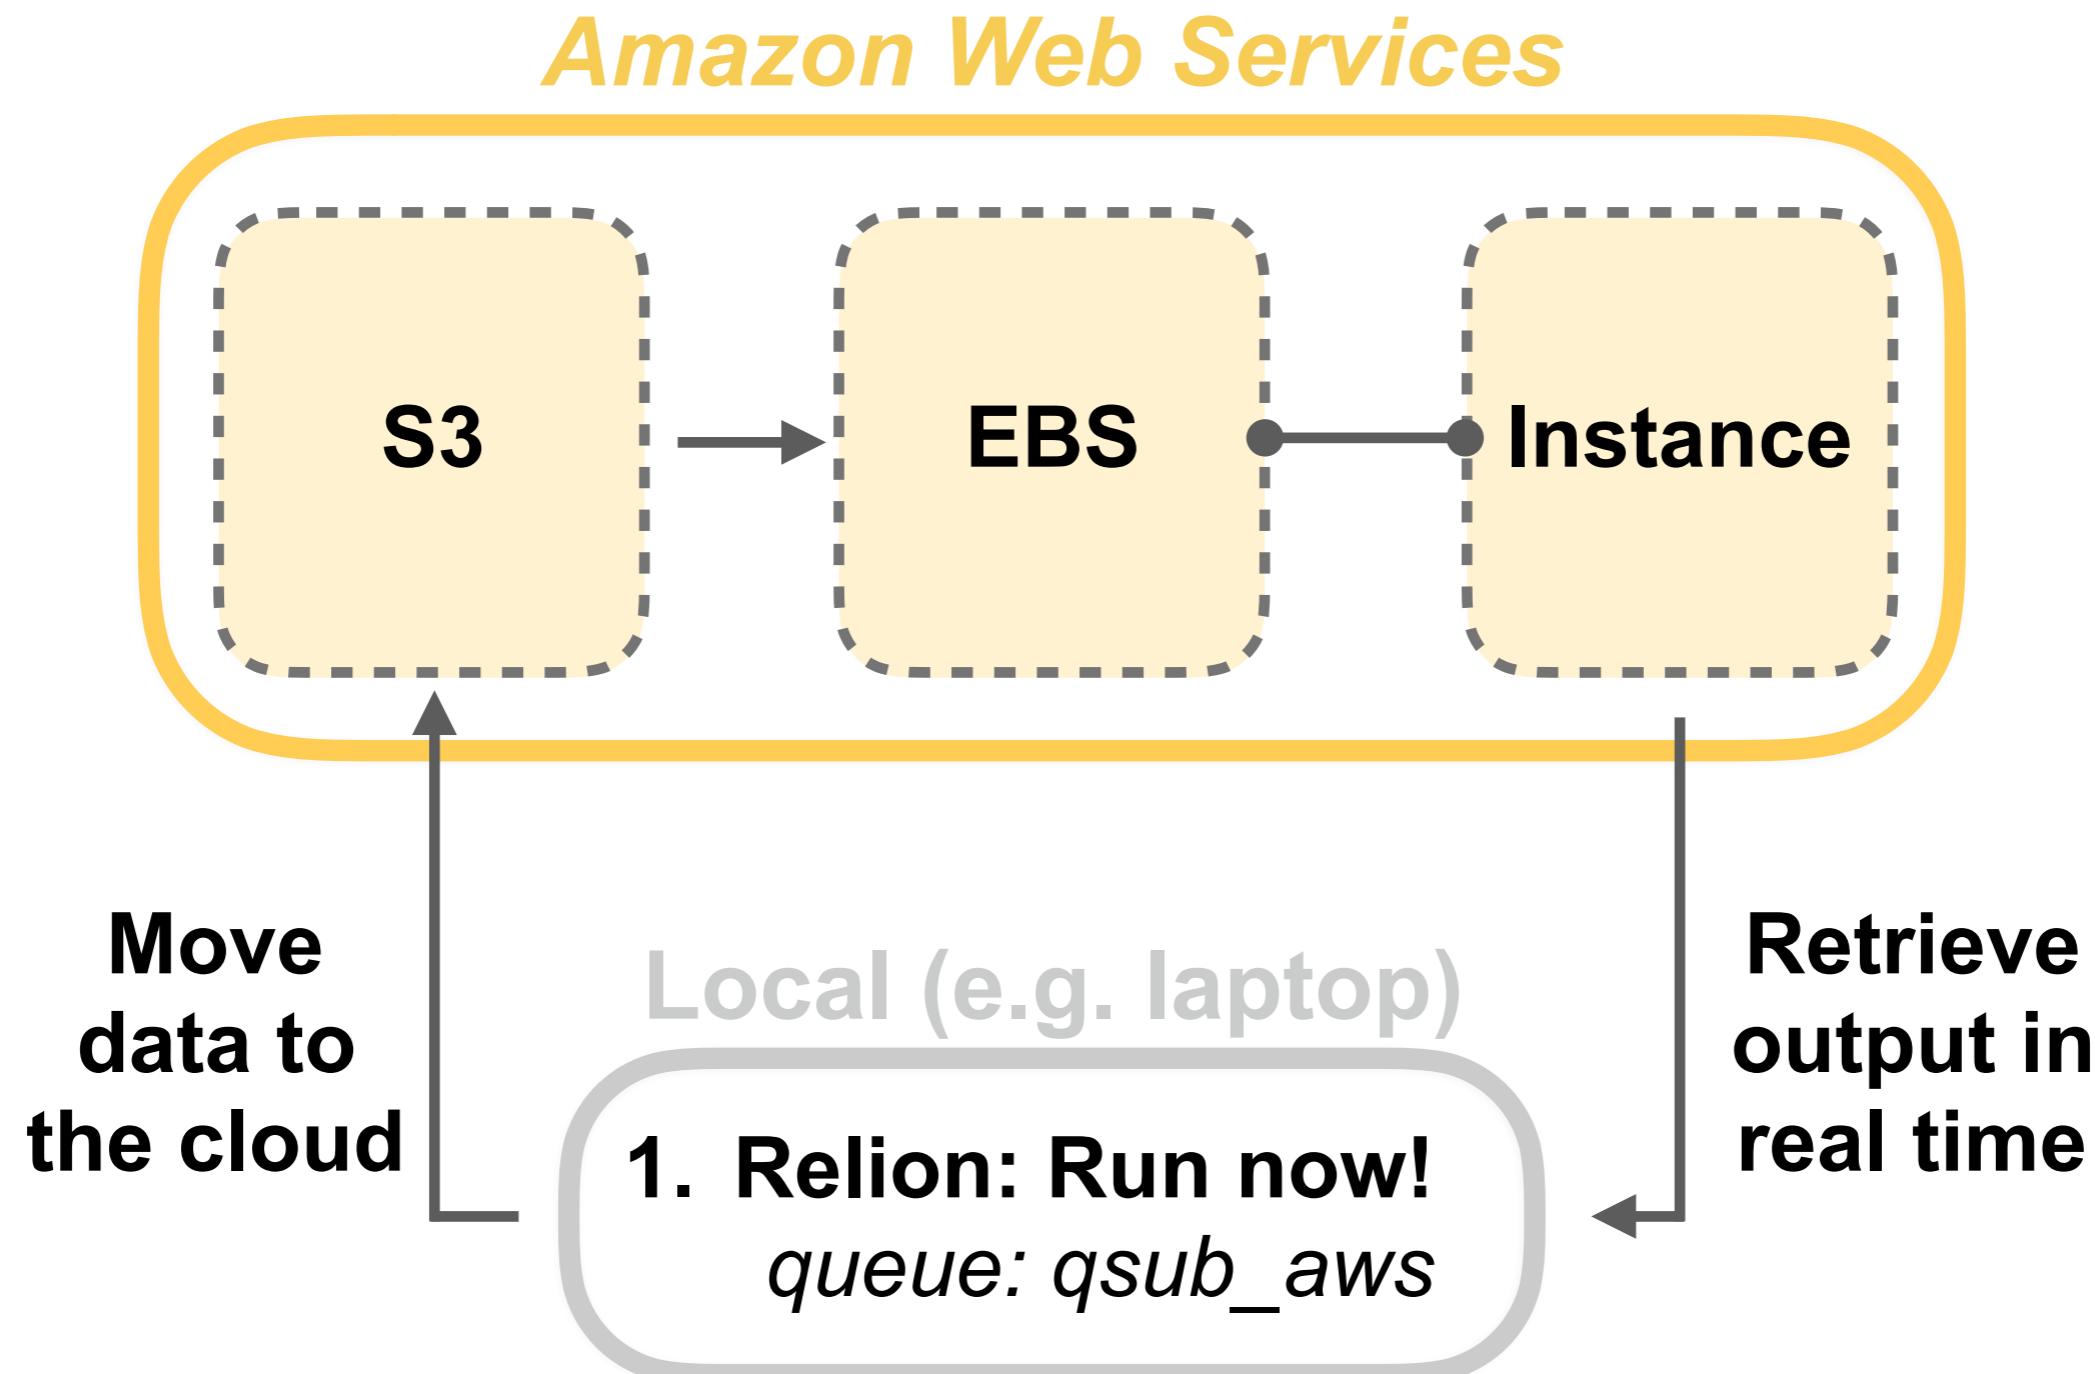
*RELION as a prototype*

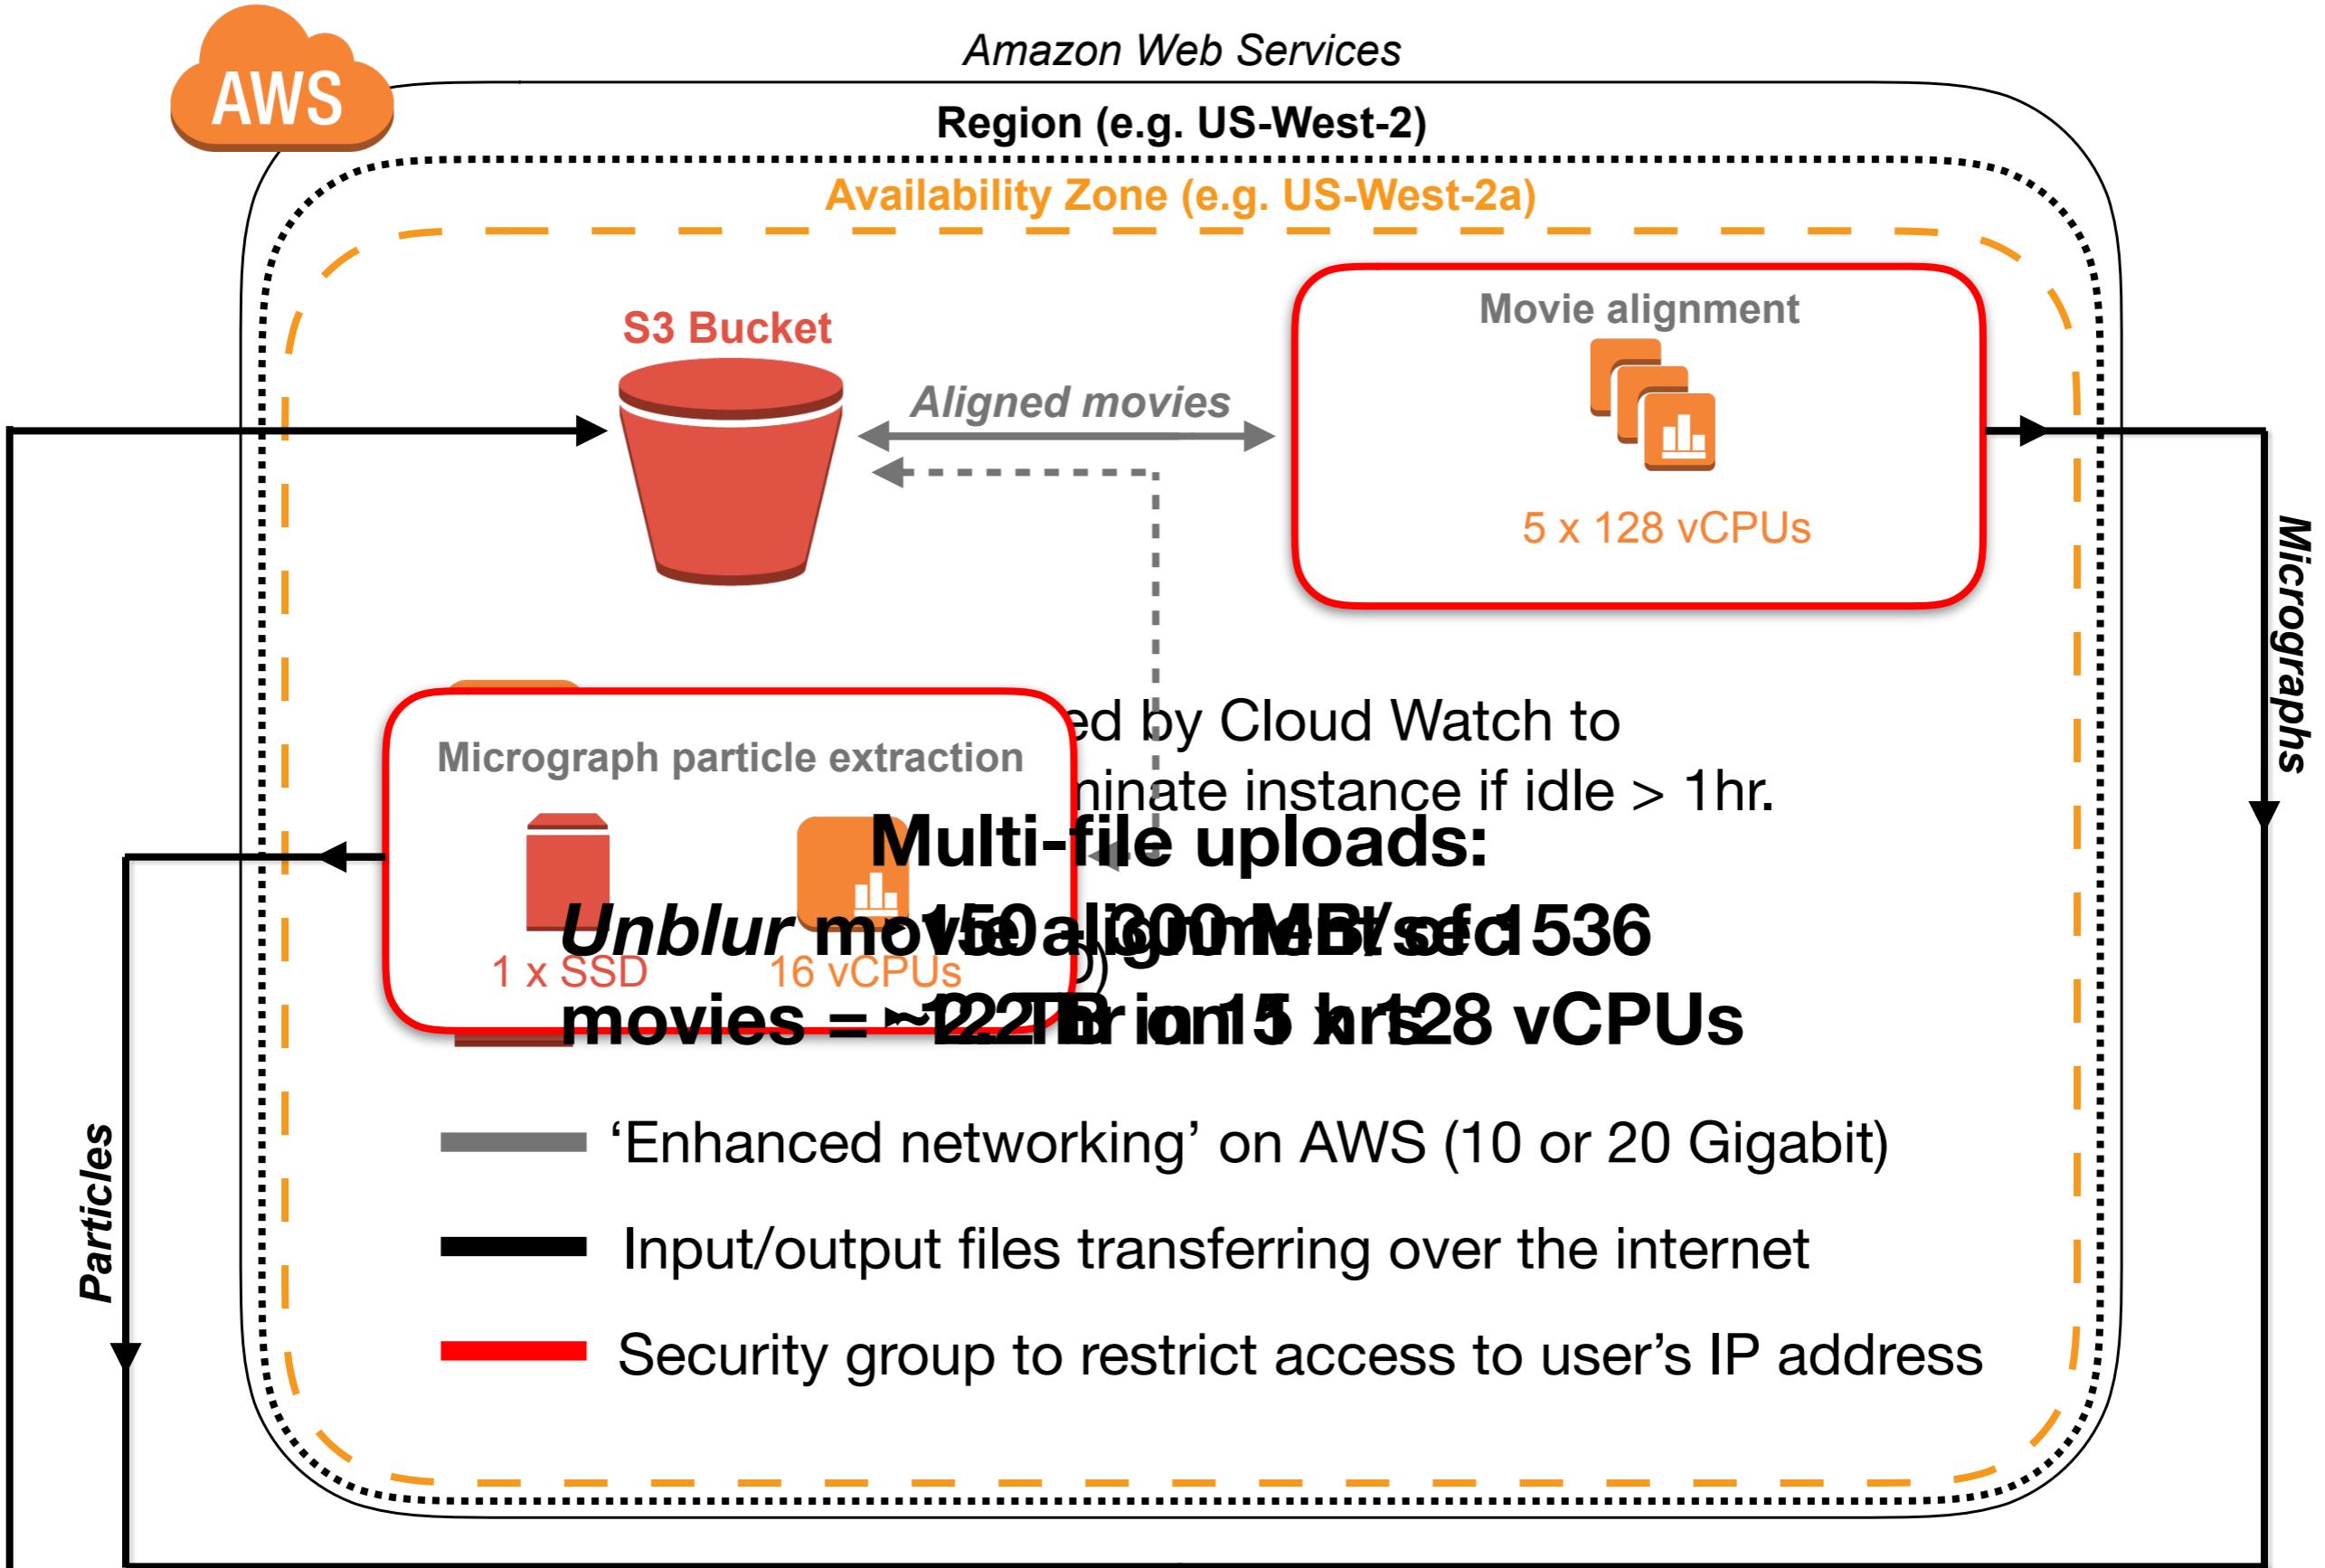


**Implement for all jobtypes**



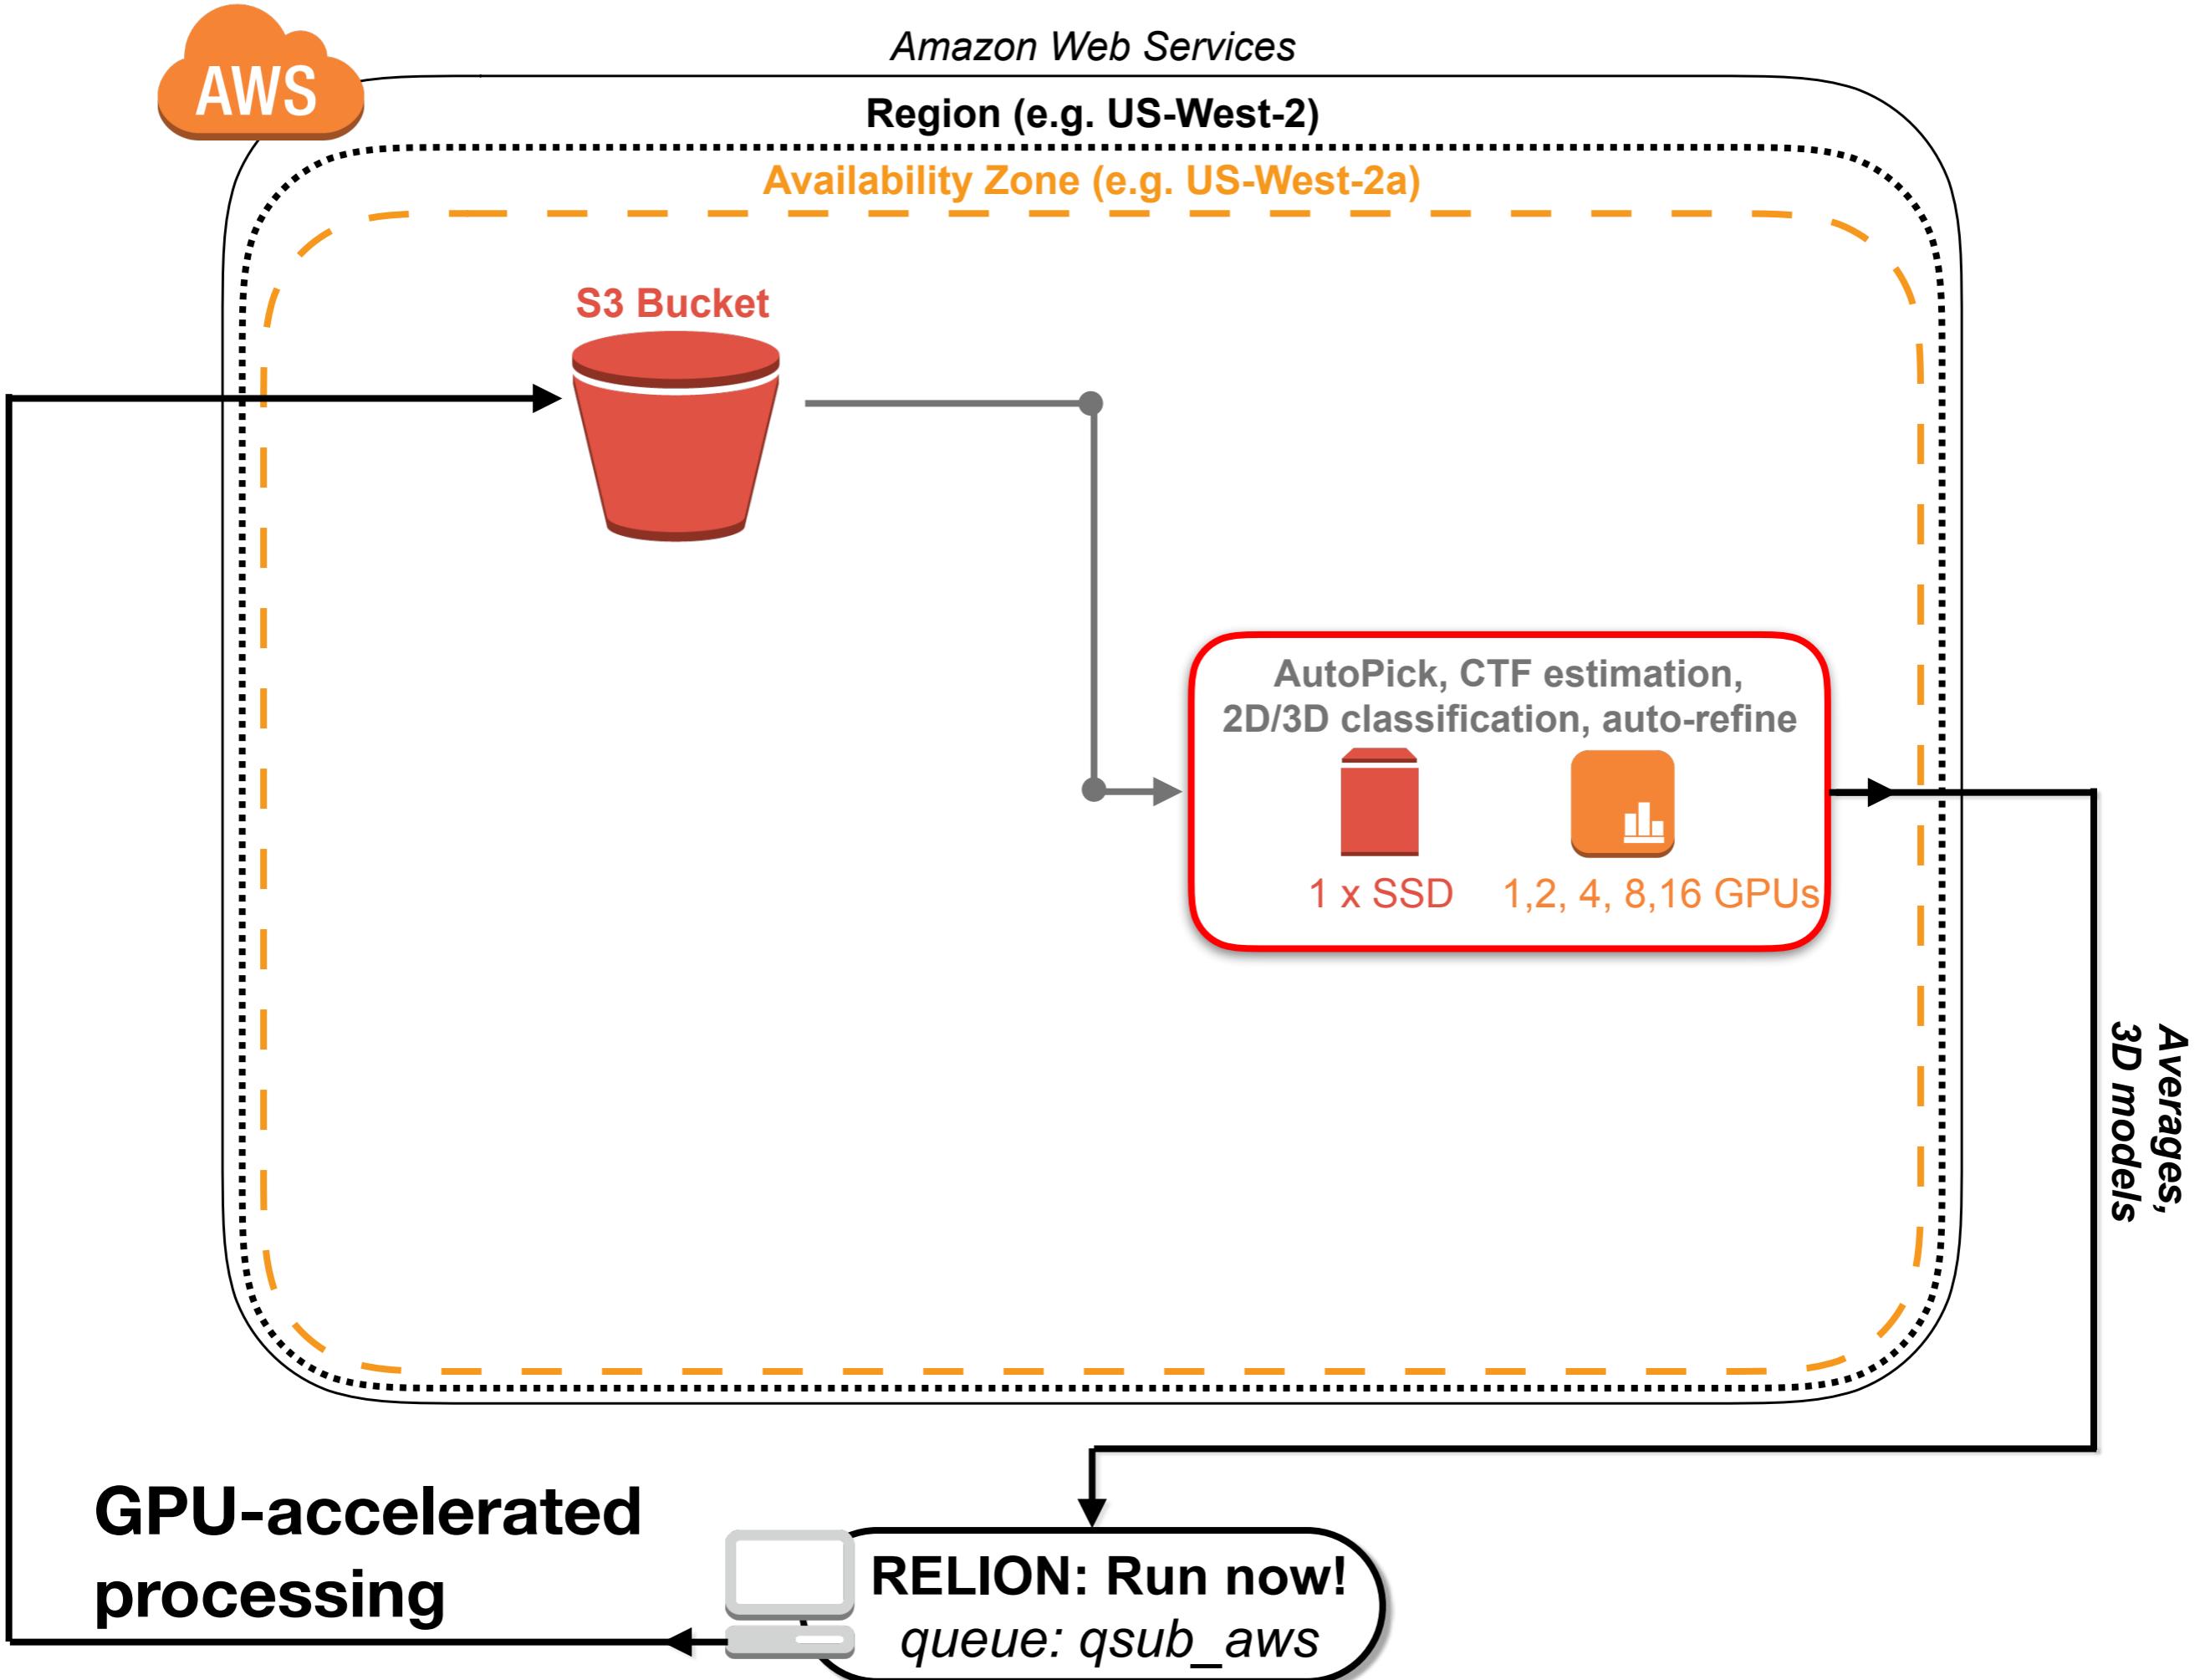
# *cryoem-cloud-tools*: software to run cryo-EM on the cloud

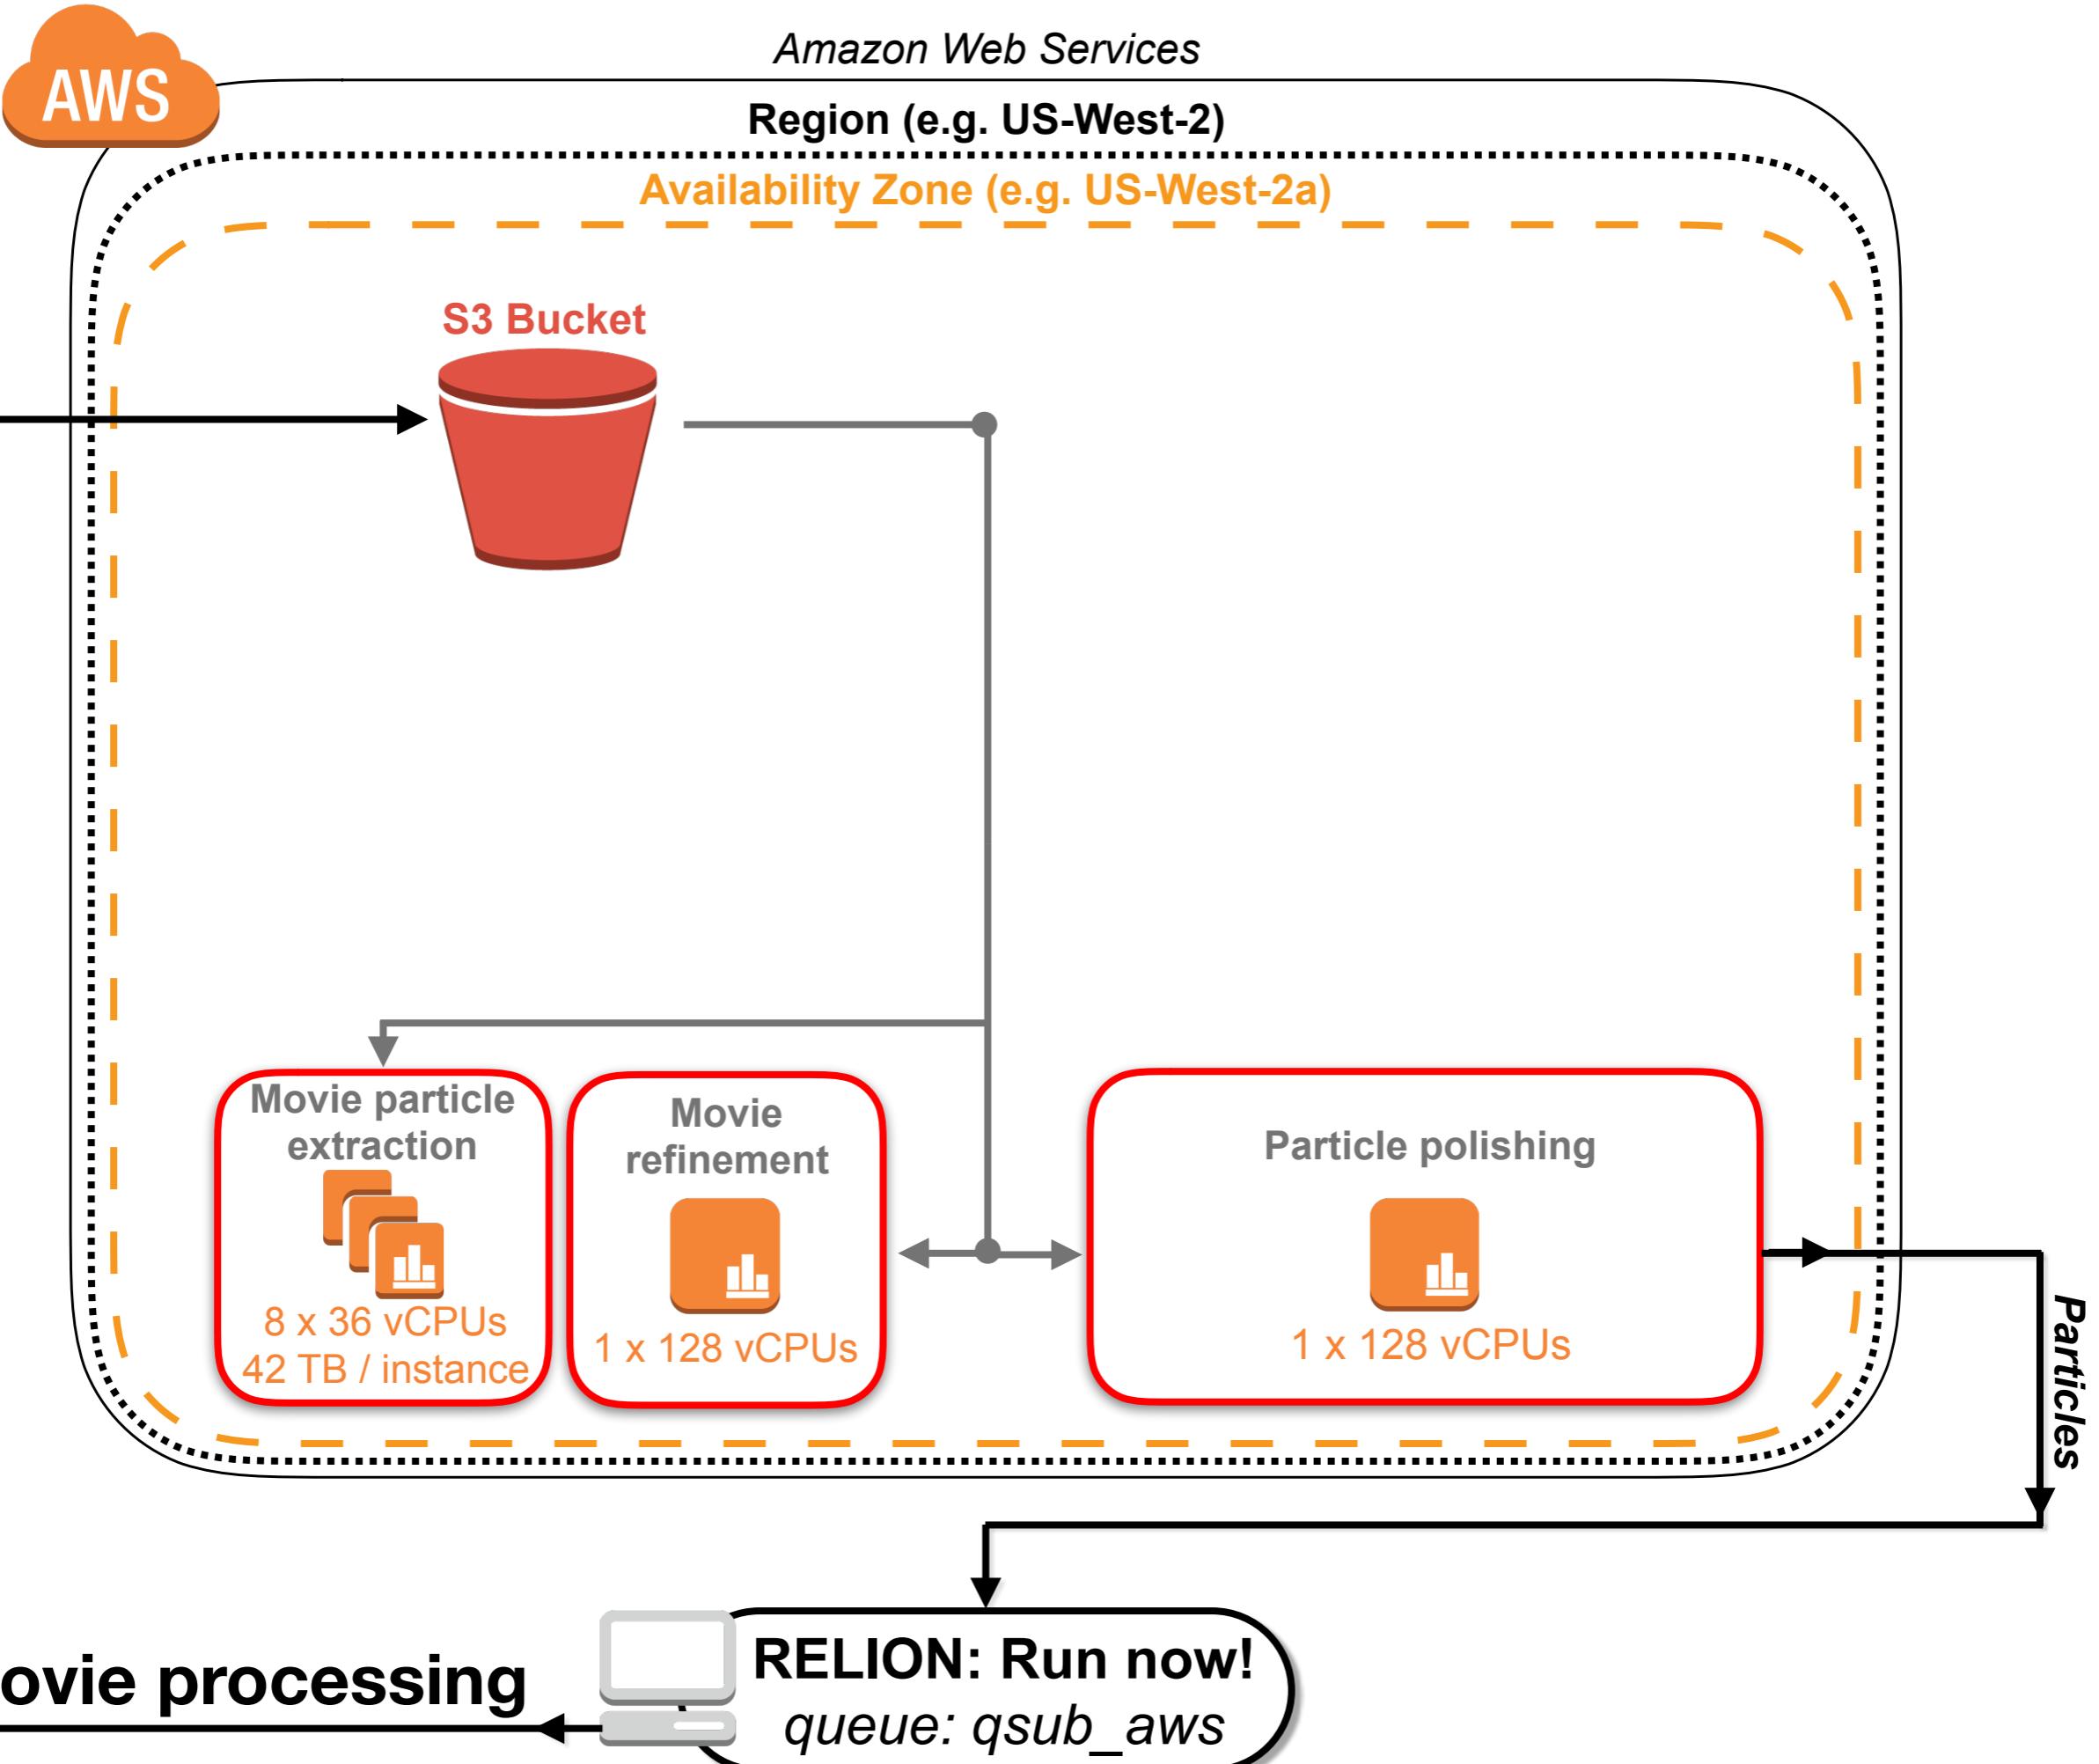


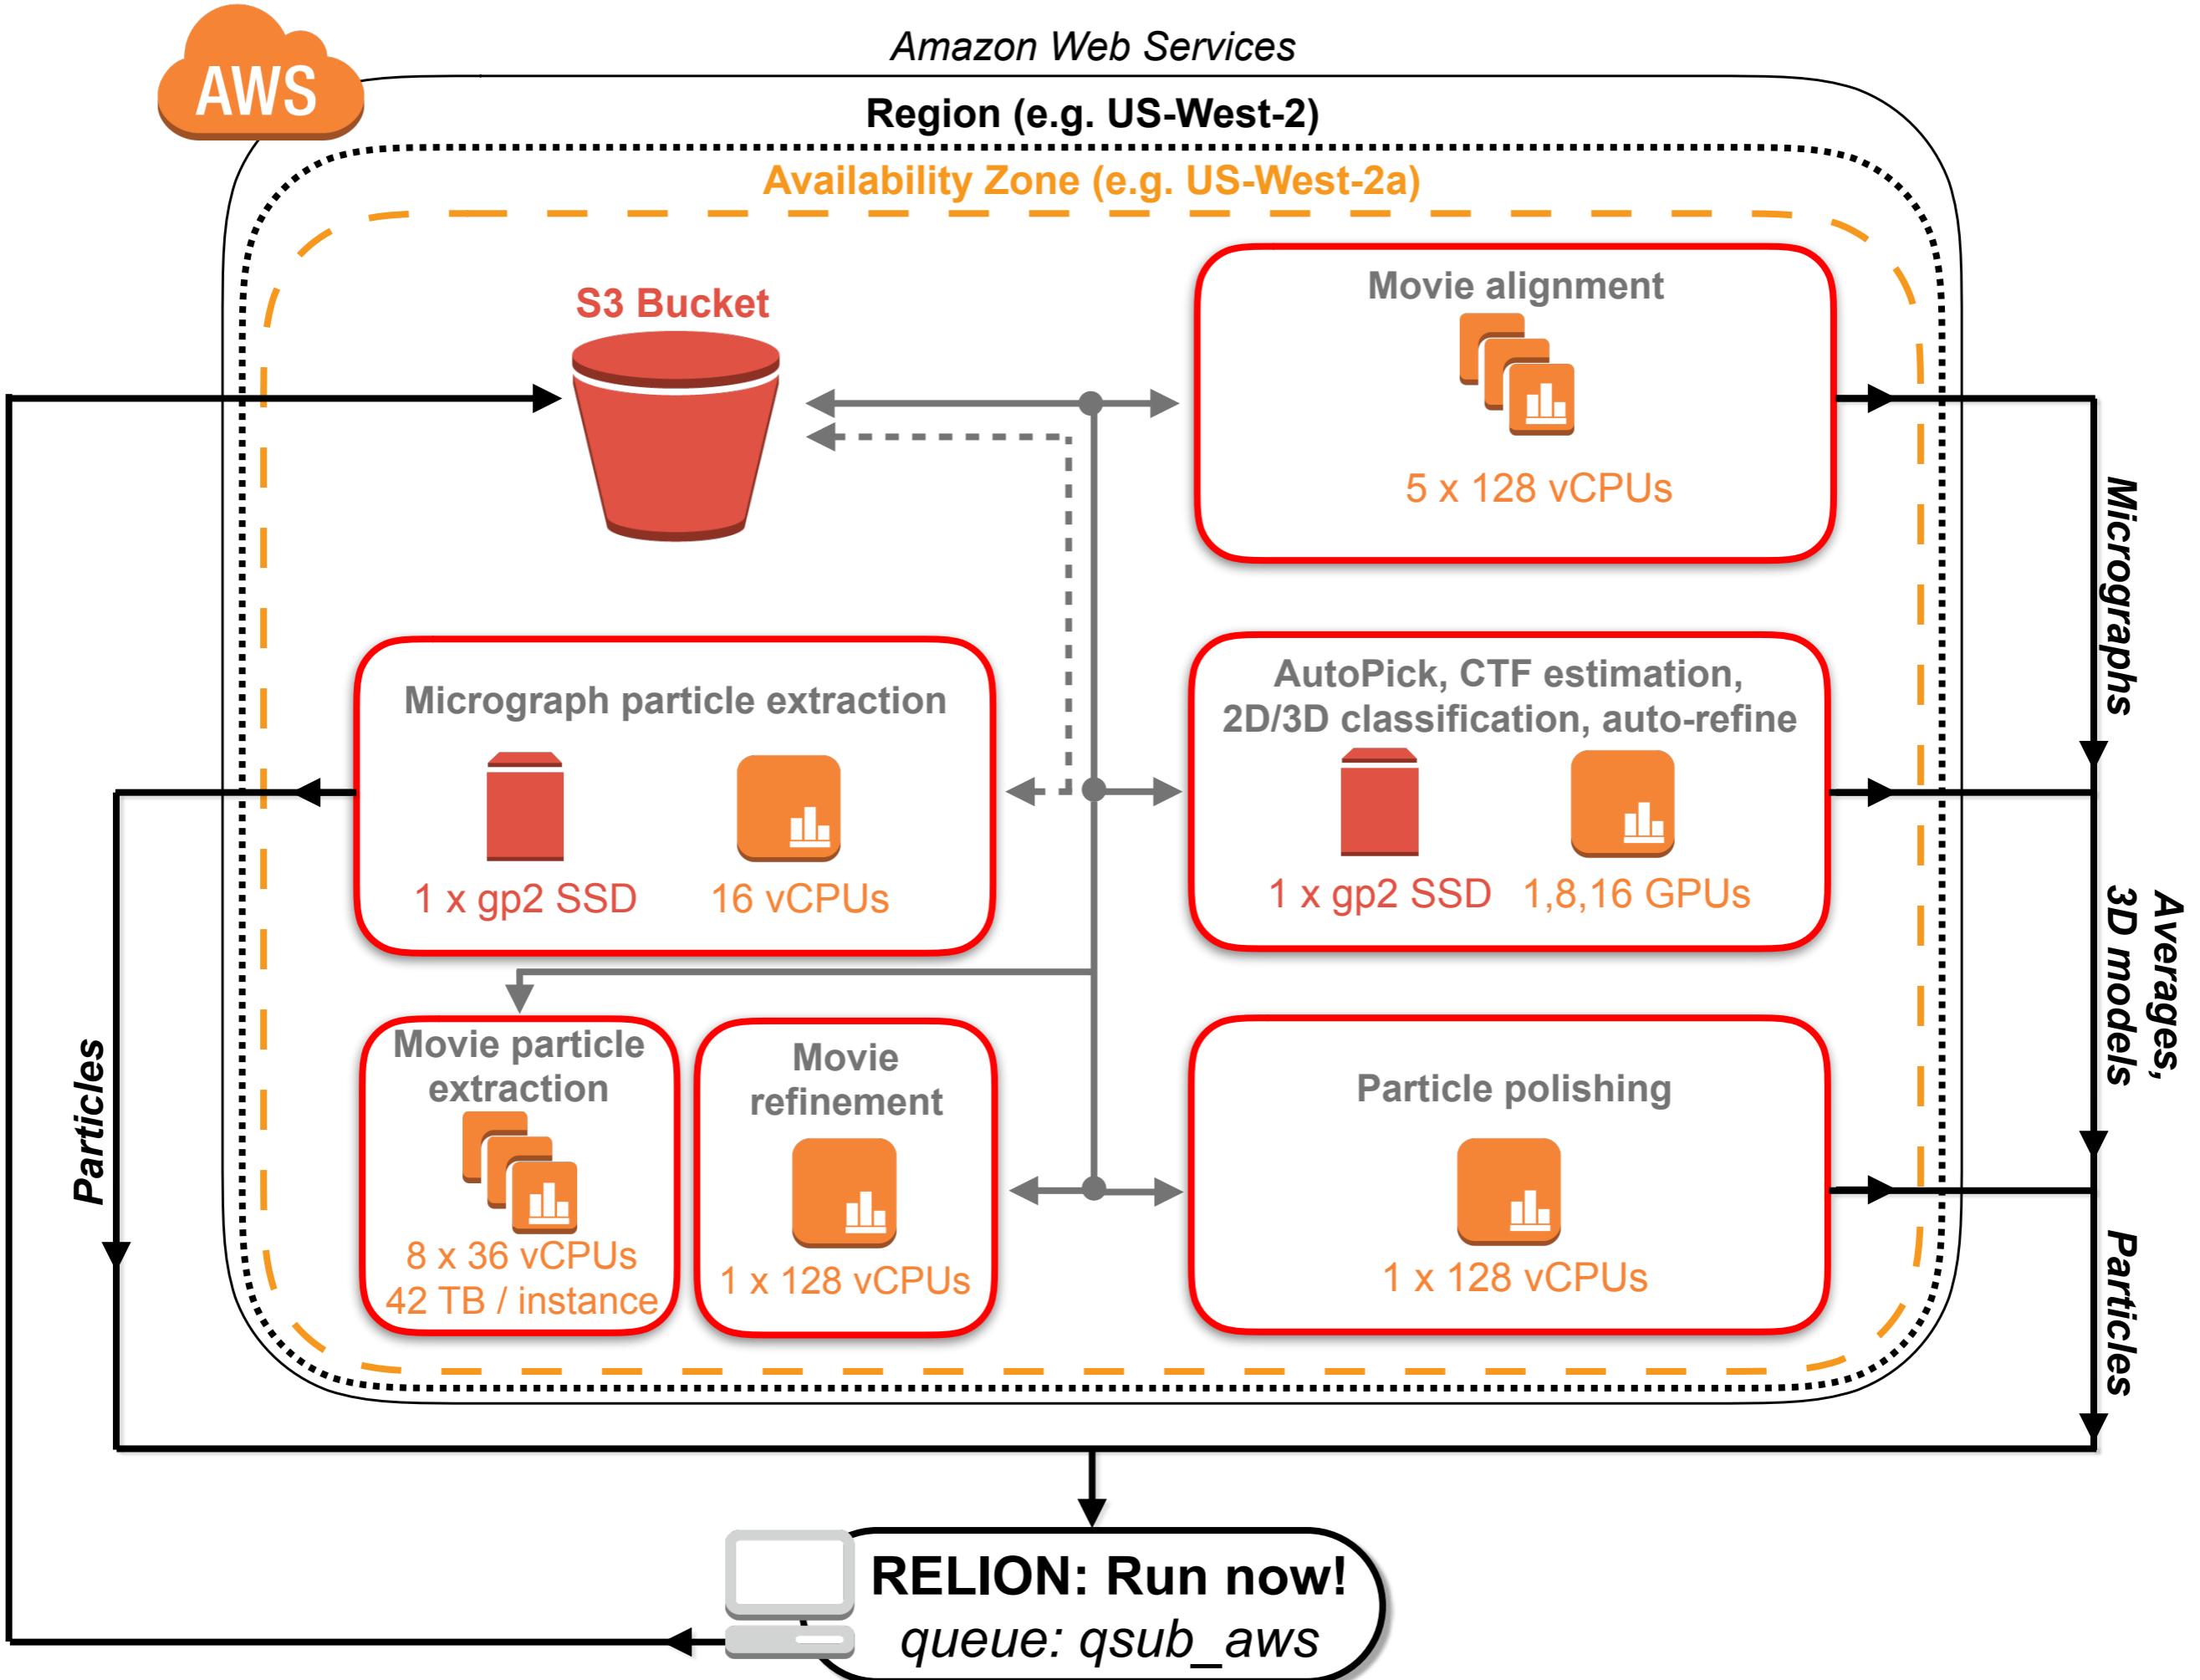


**Movie realignment**

**RELION: Run now!**  
queue: qsub\_aws

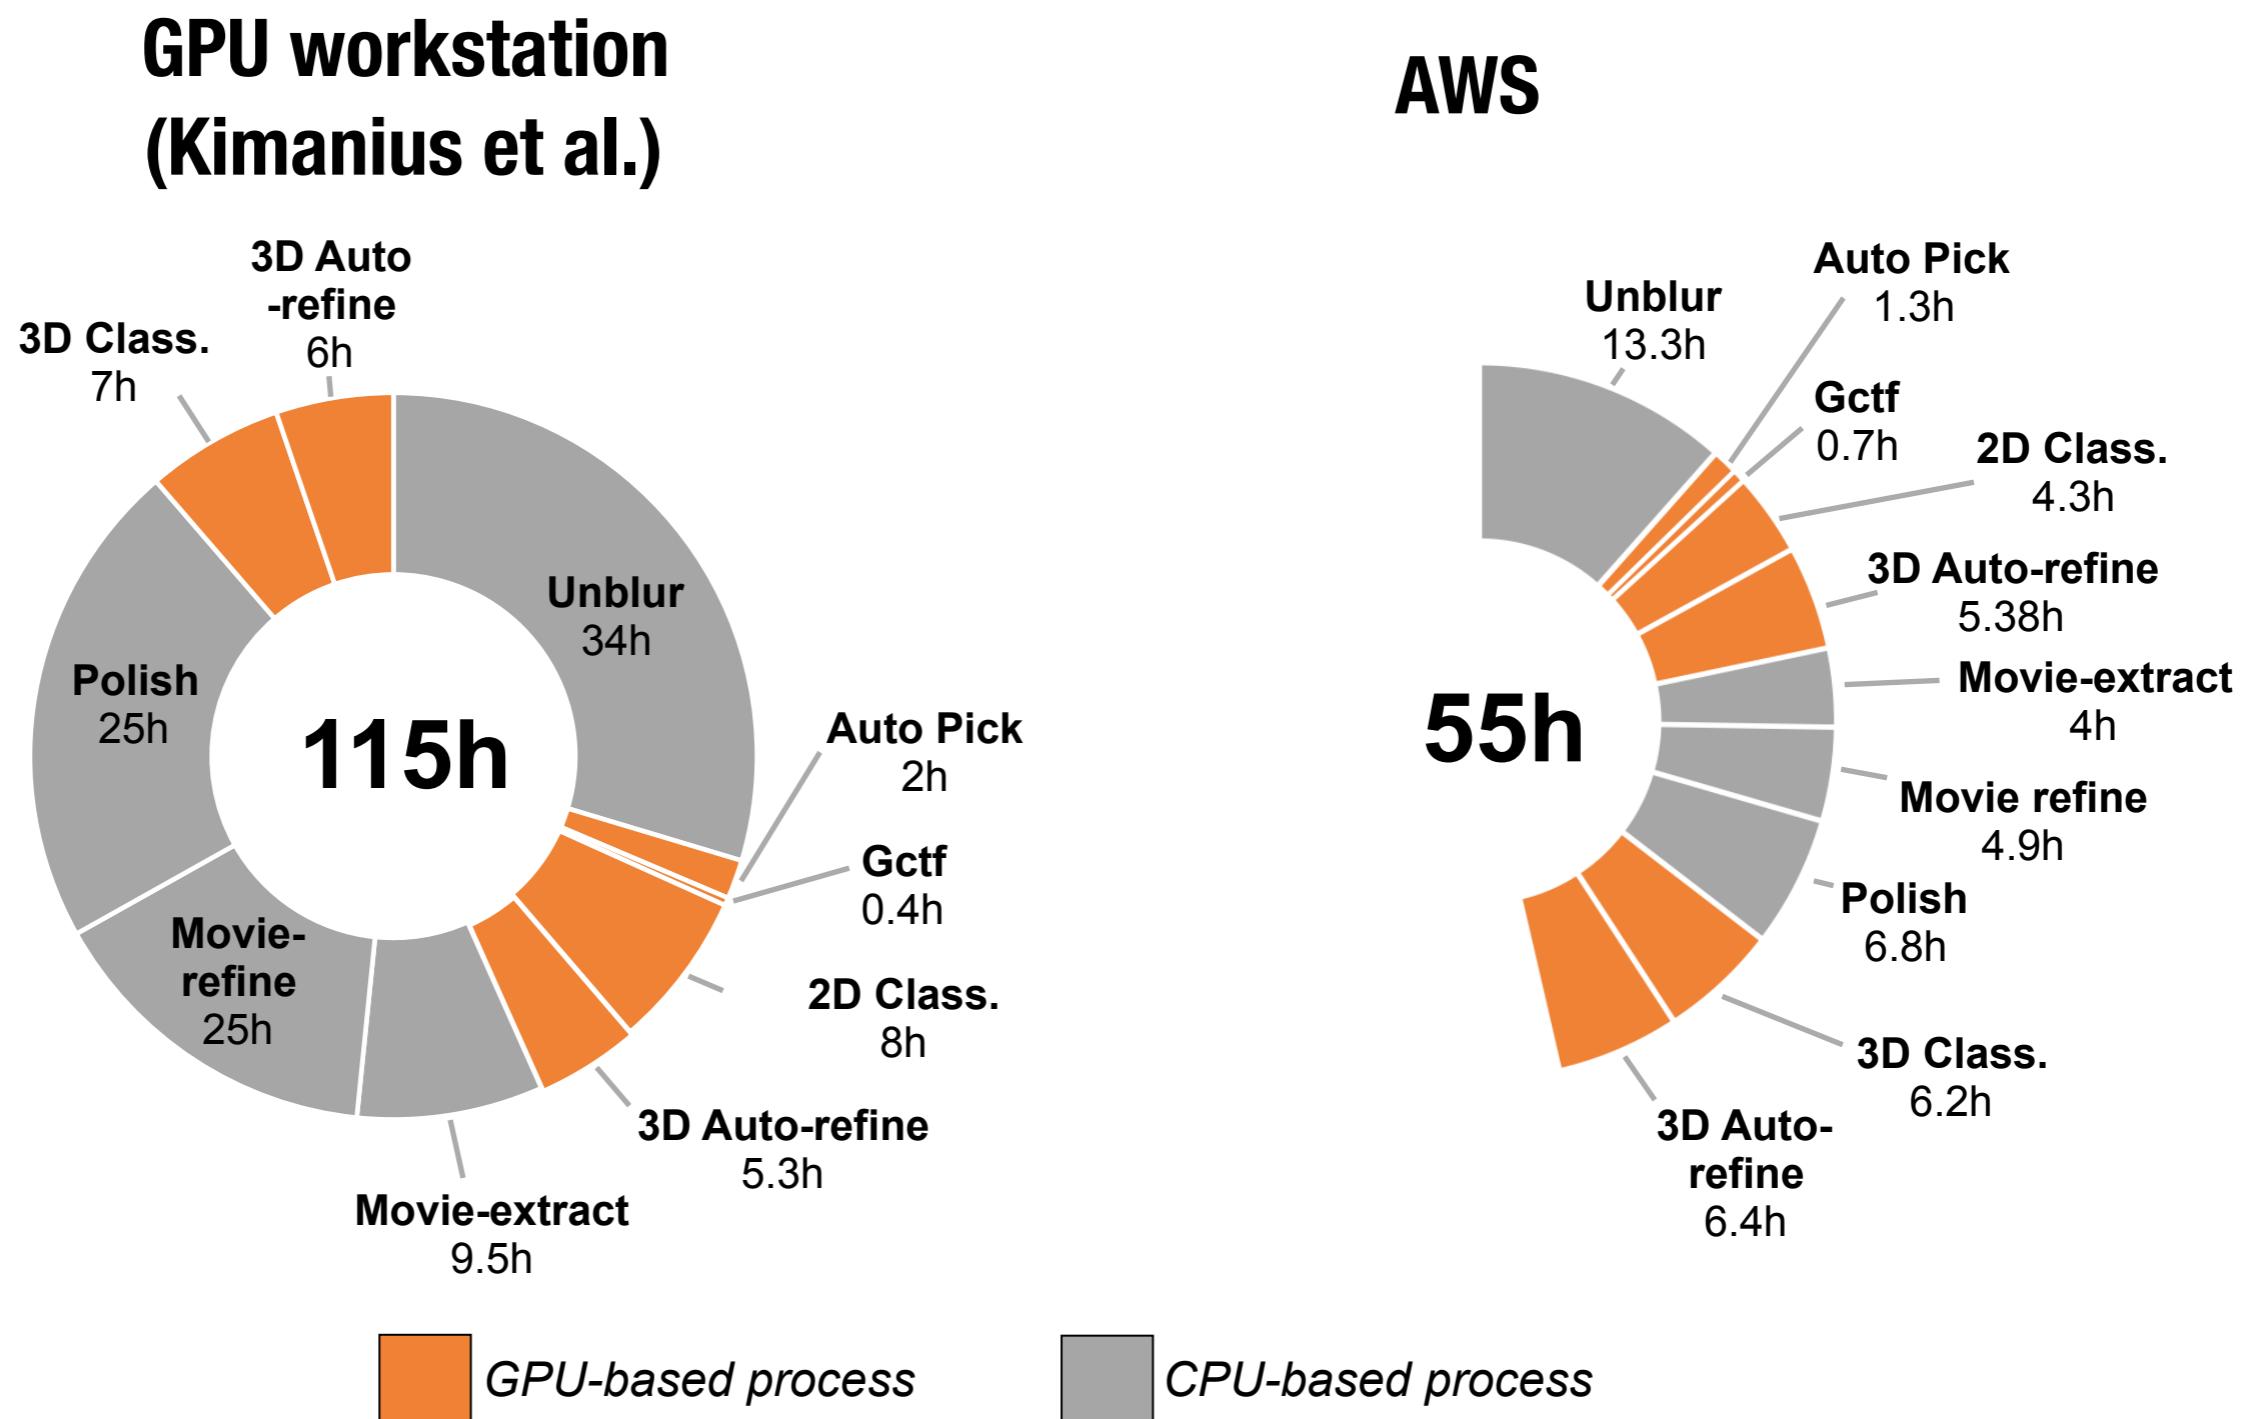




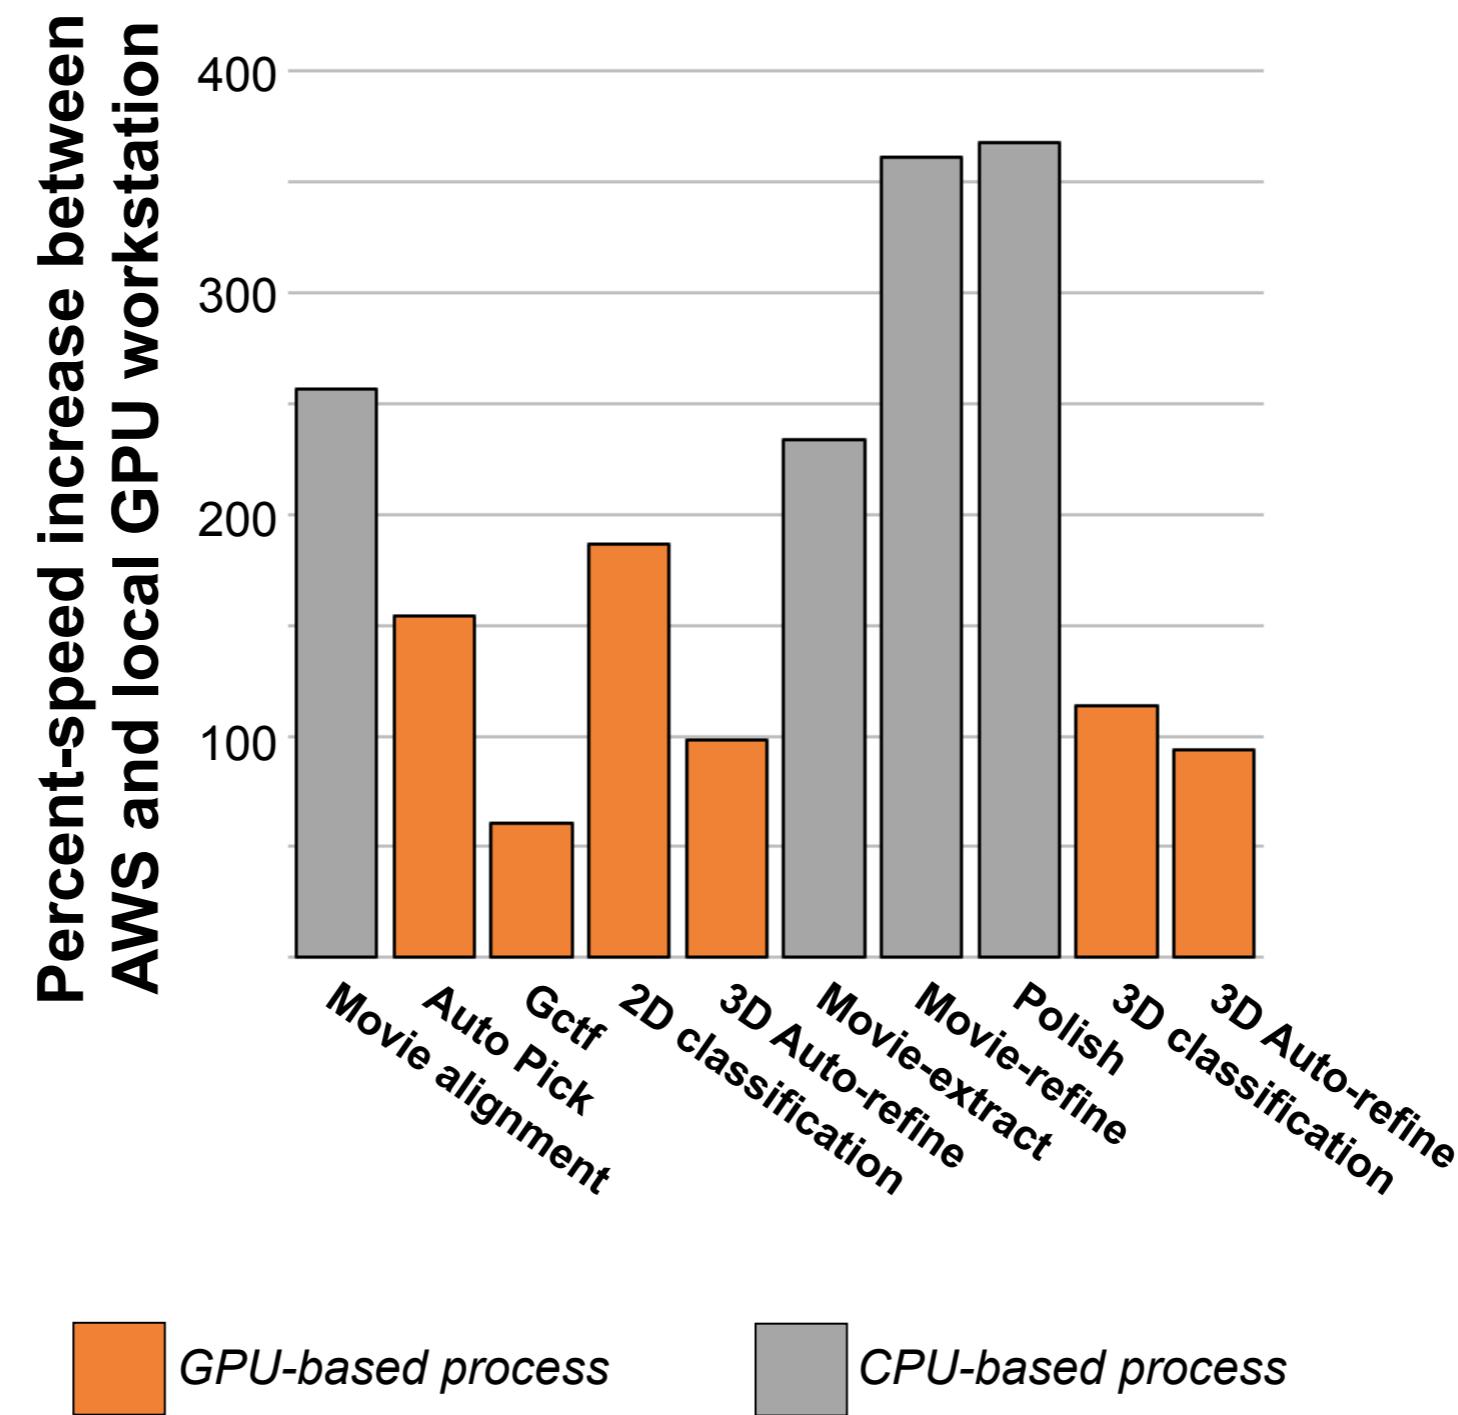


# Comparison of AWS vs. local GPU workstation for RELION2

Comparing 2.2 Å beta-galactosidase dataset (*EMPIAR 10061*)

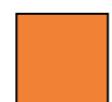
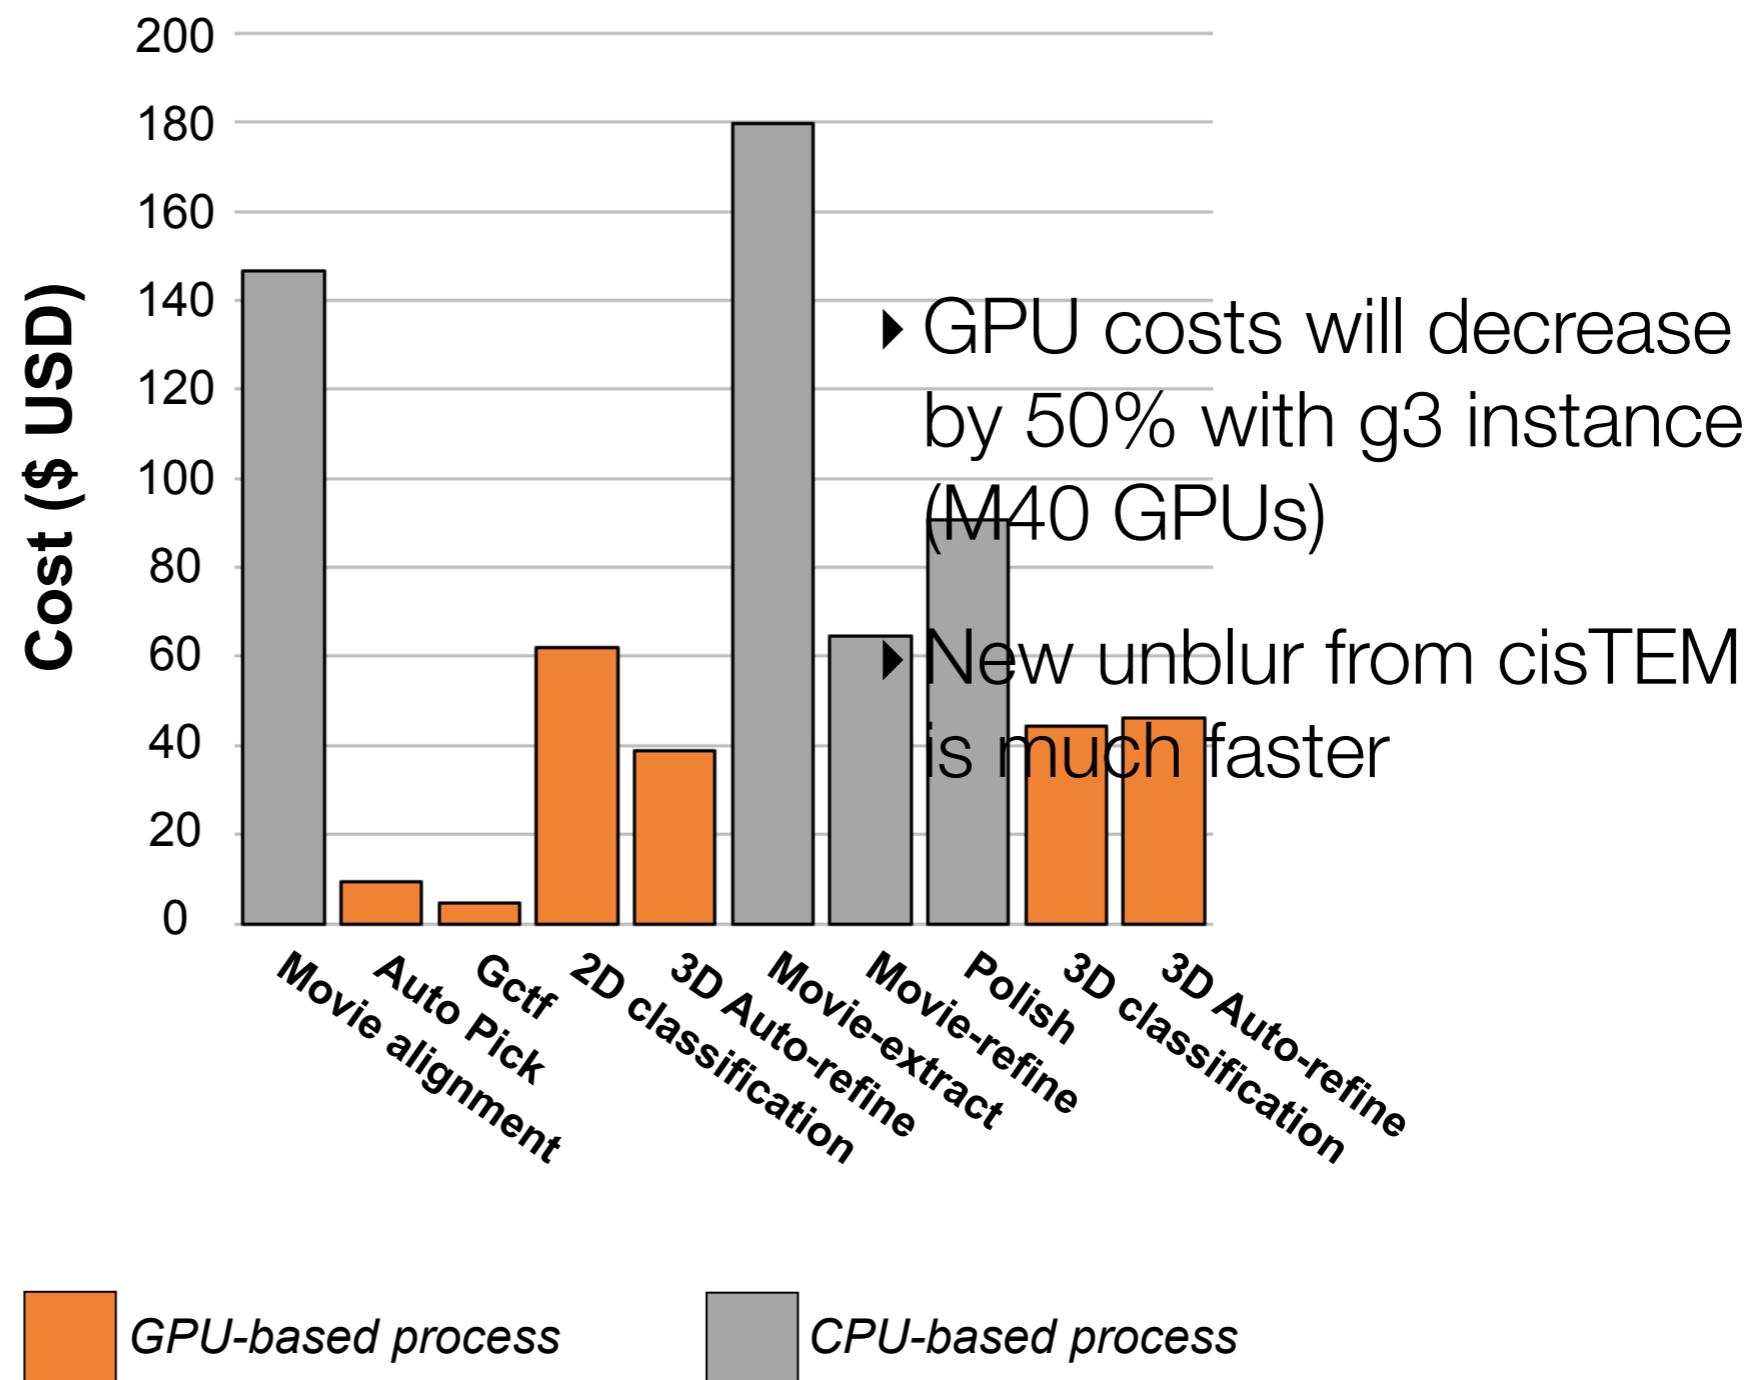


# Comparison of AWS vs. local GPU workstation for RELION2



# Cost breakdown of AWS

*Total computing cost: \$688.15*



GPU-based process



CPU-based process



# Cost comparison between AWS and local GPU workstation

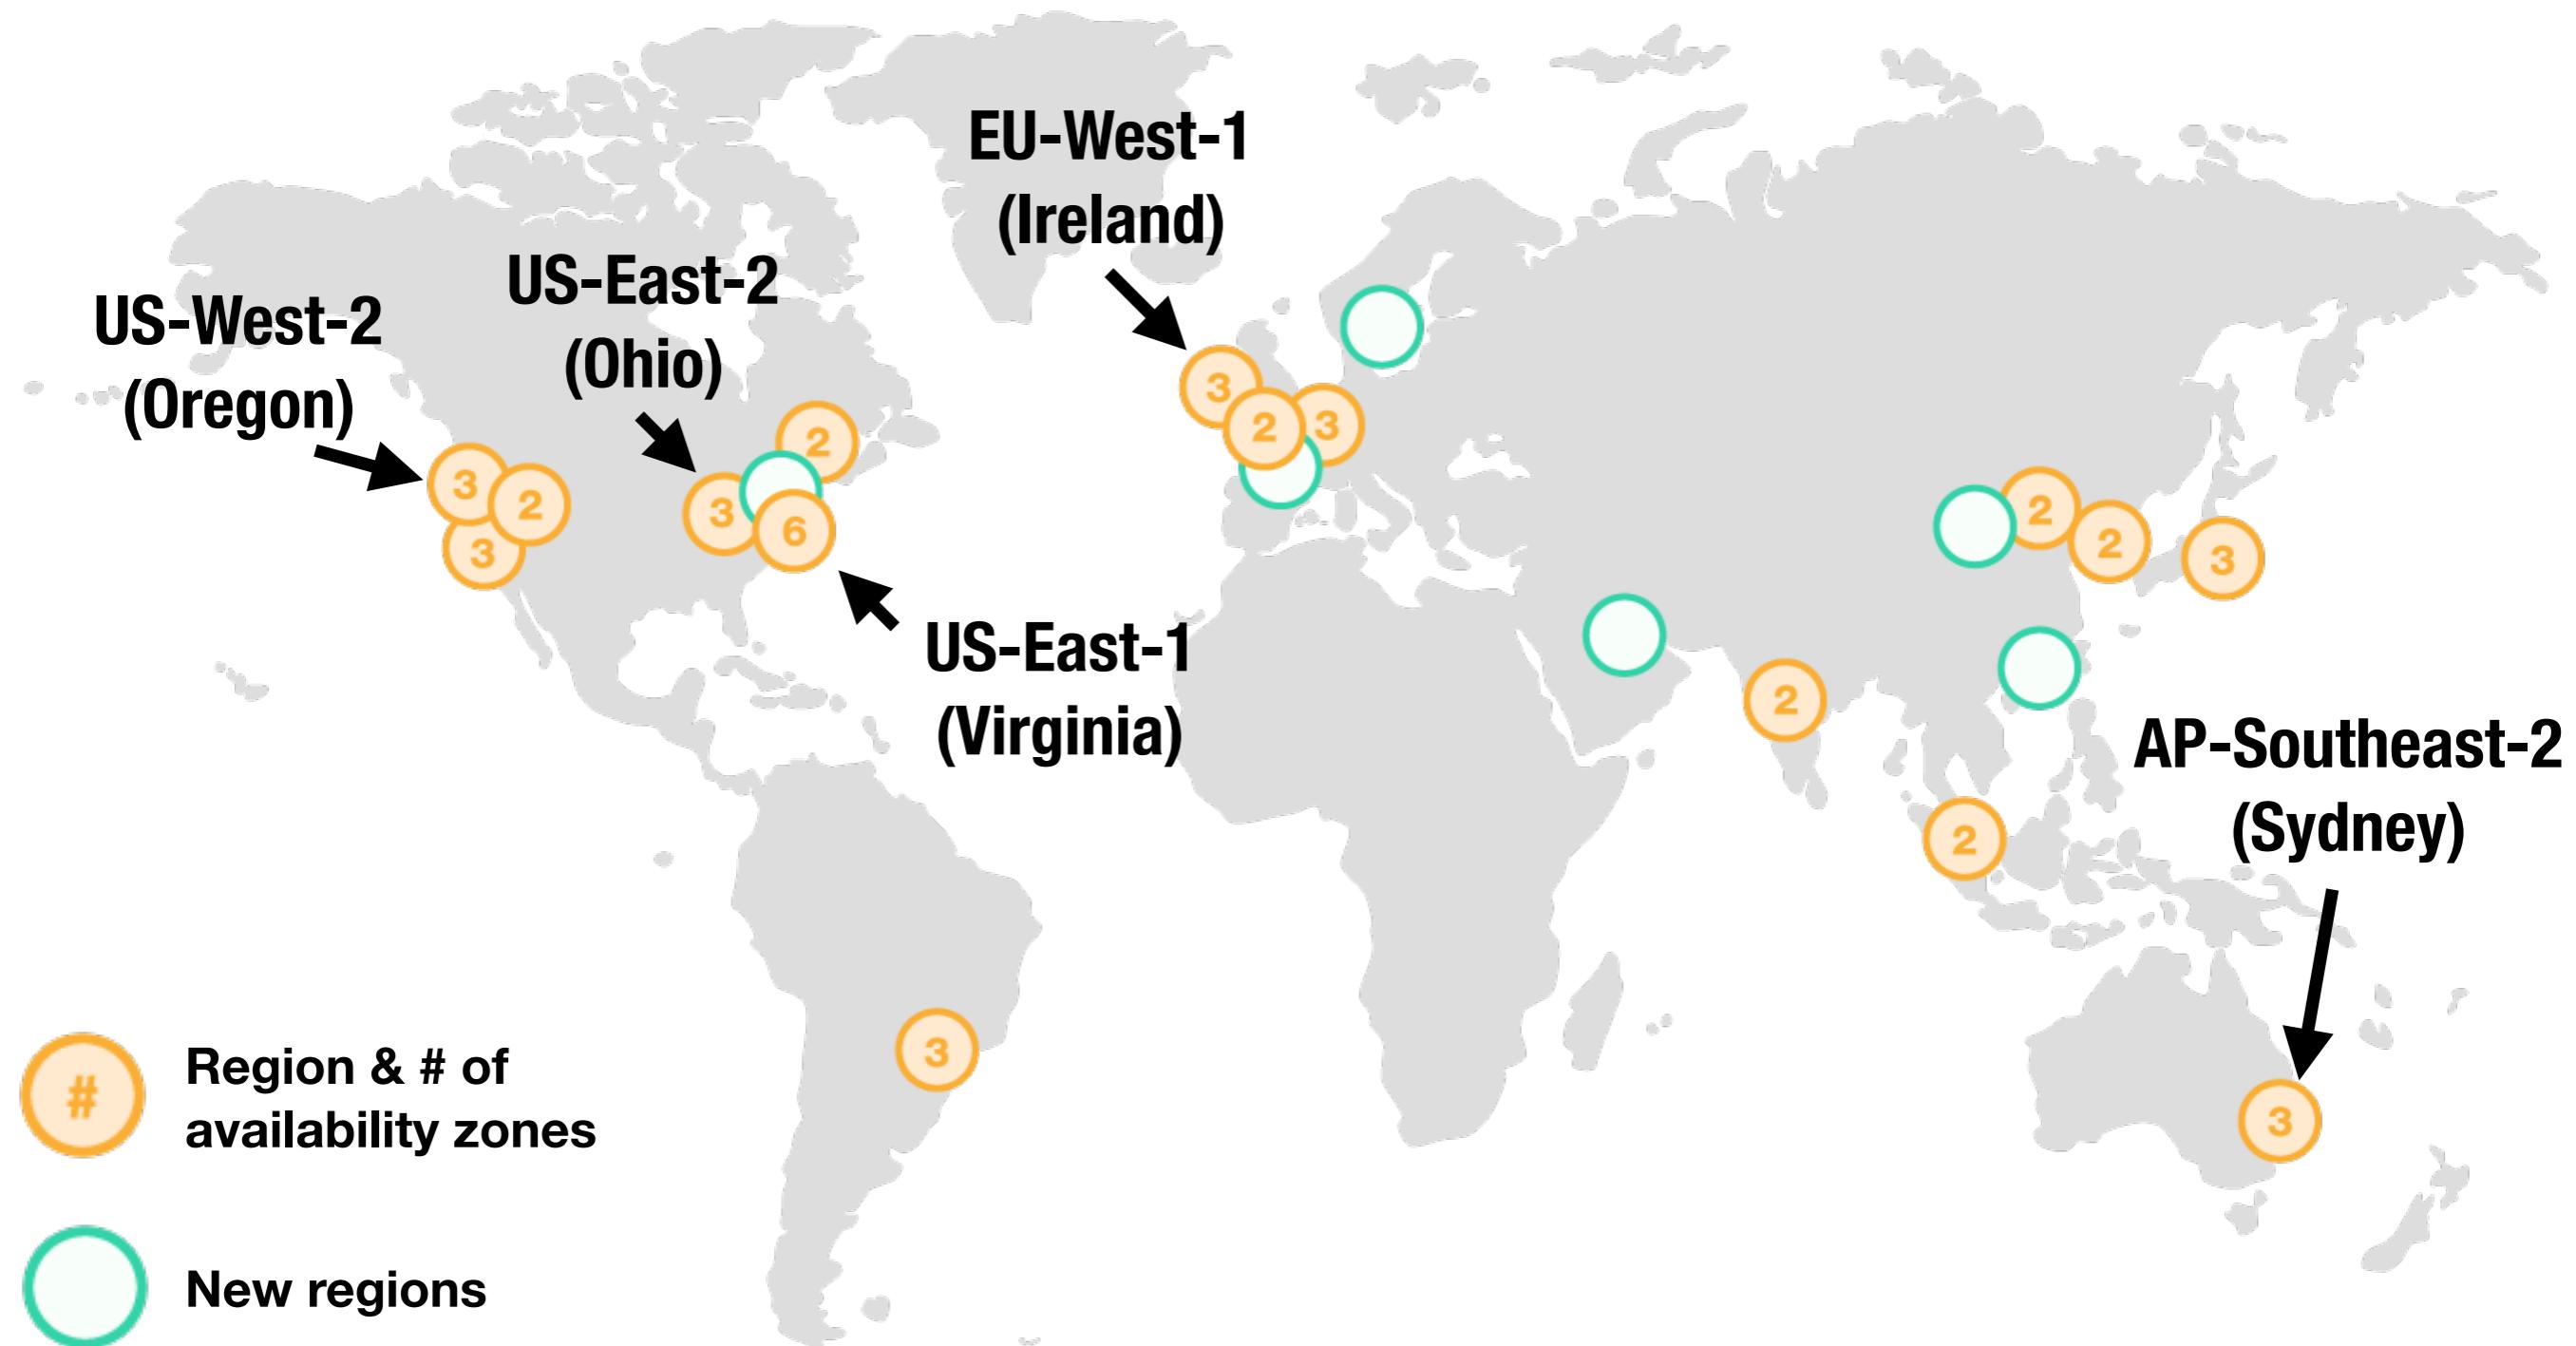
---

	AWS	Local GPU workstation
Computing costs*	\$688.17	\$72.45
Processing time (h)	55	115
Number of working days	5	10

\*For local GPU workstation, \$10,000 workstation cost = \$0.63/hr (60% utilization).



# Global footprint of *cryoem-cloud-tools*

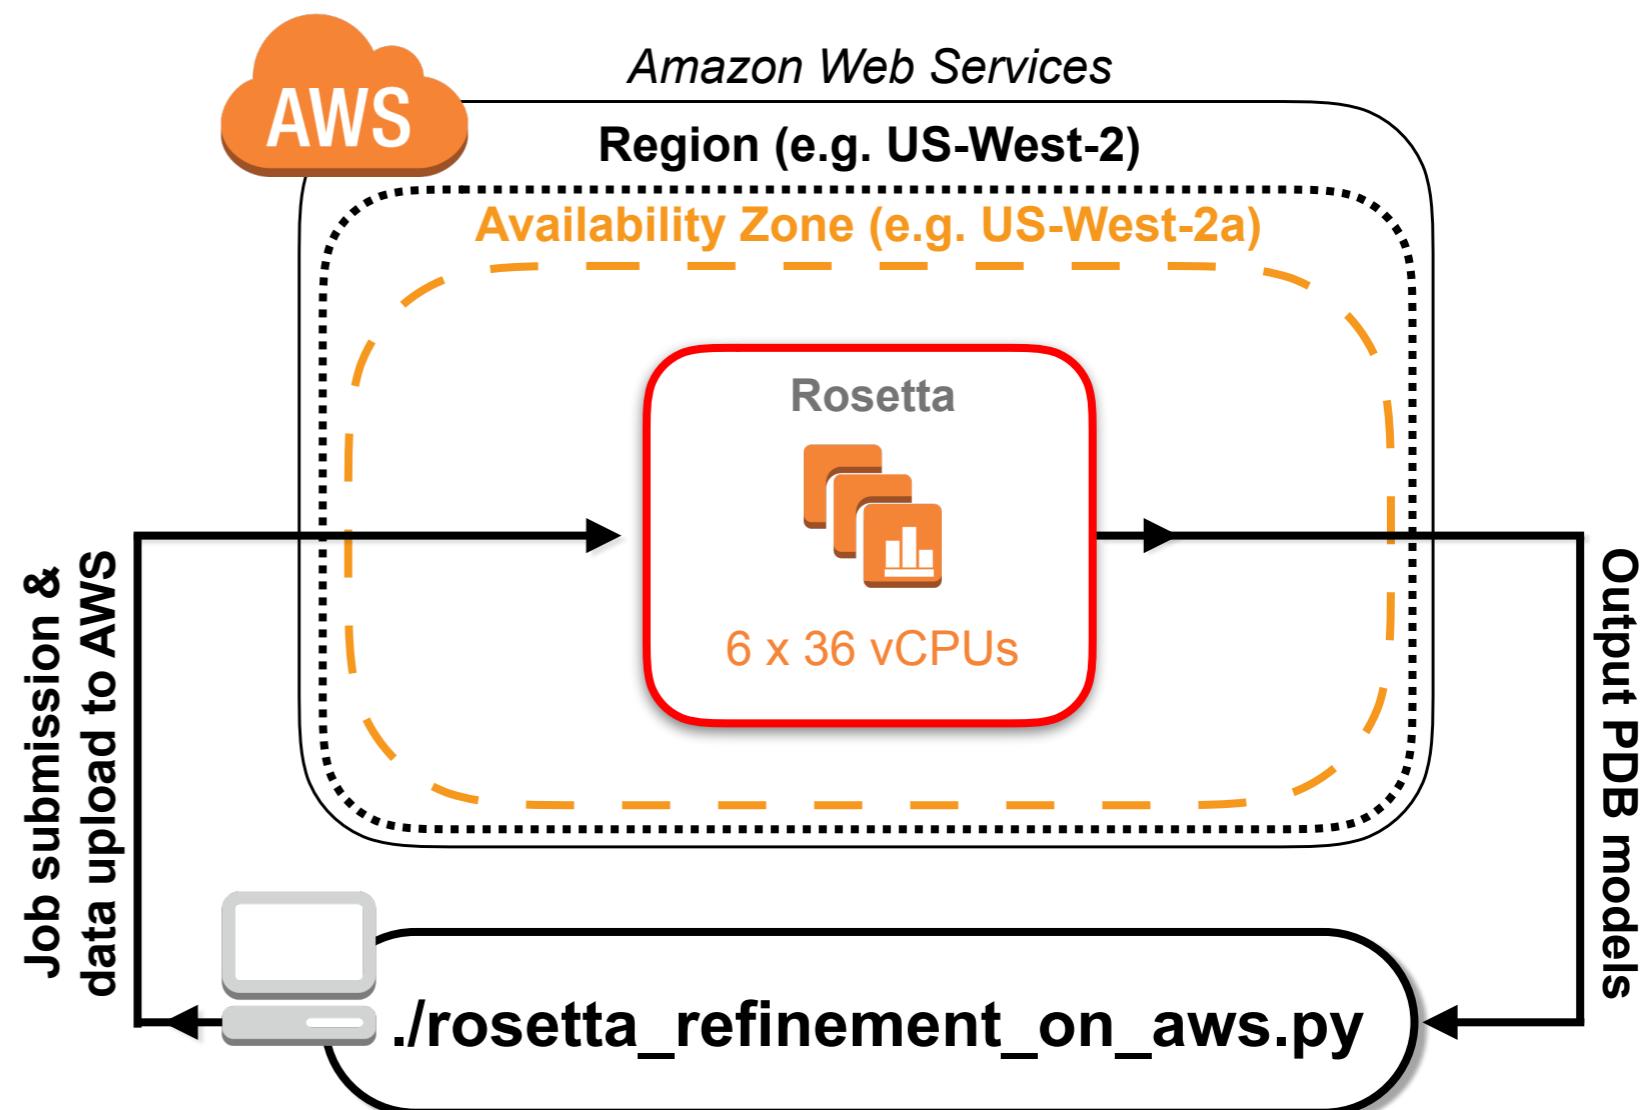
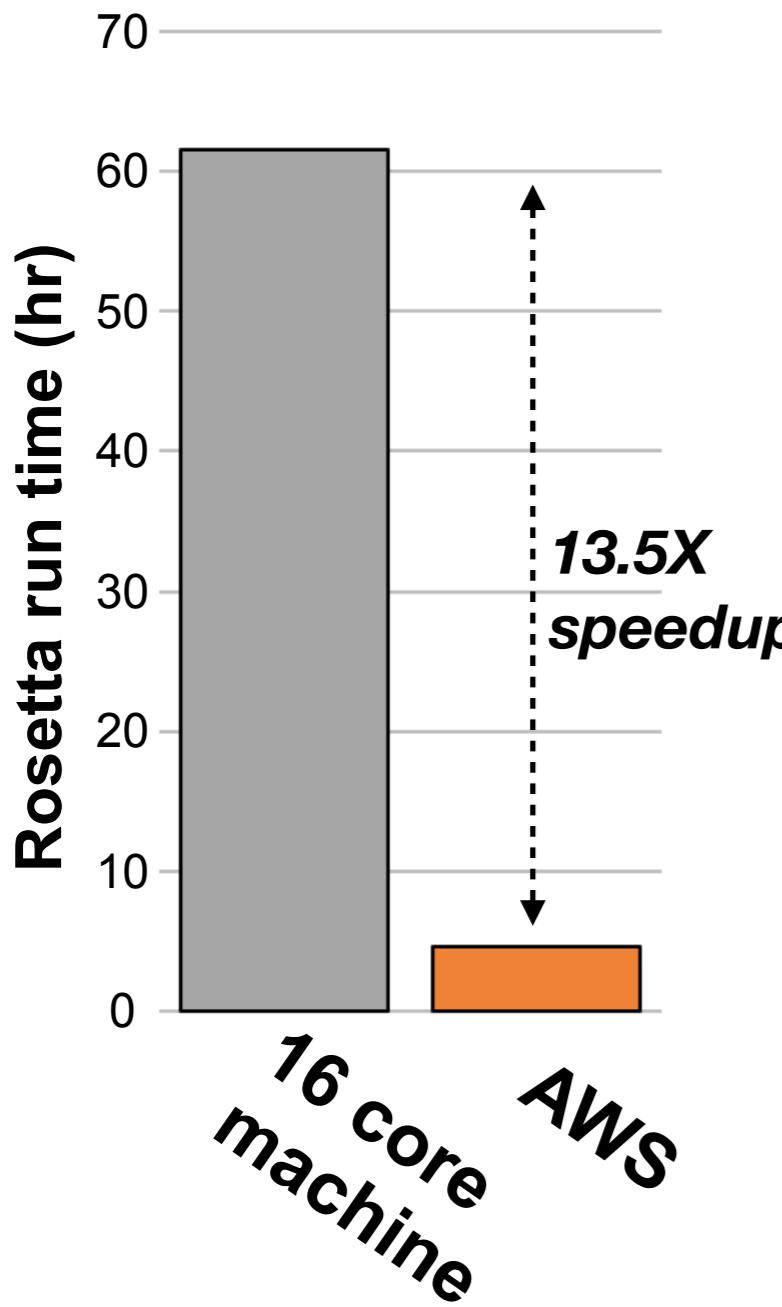


*Tracking with expanding GPU availability on AWS*



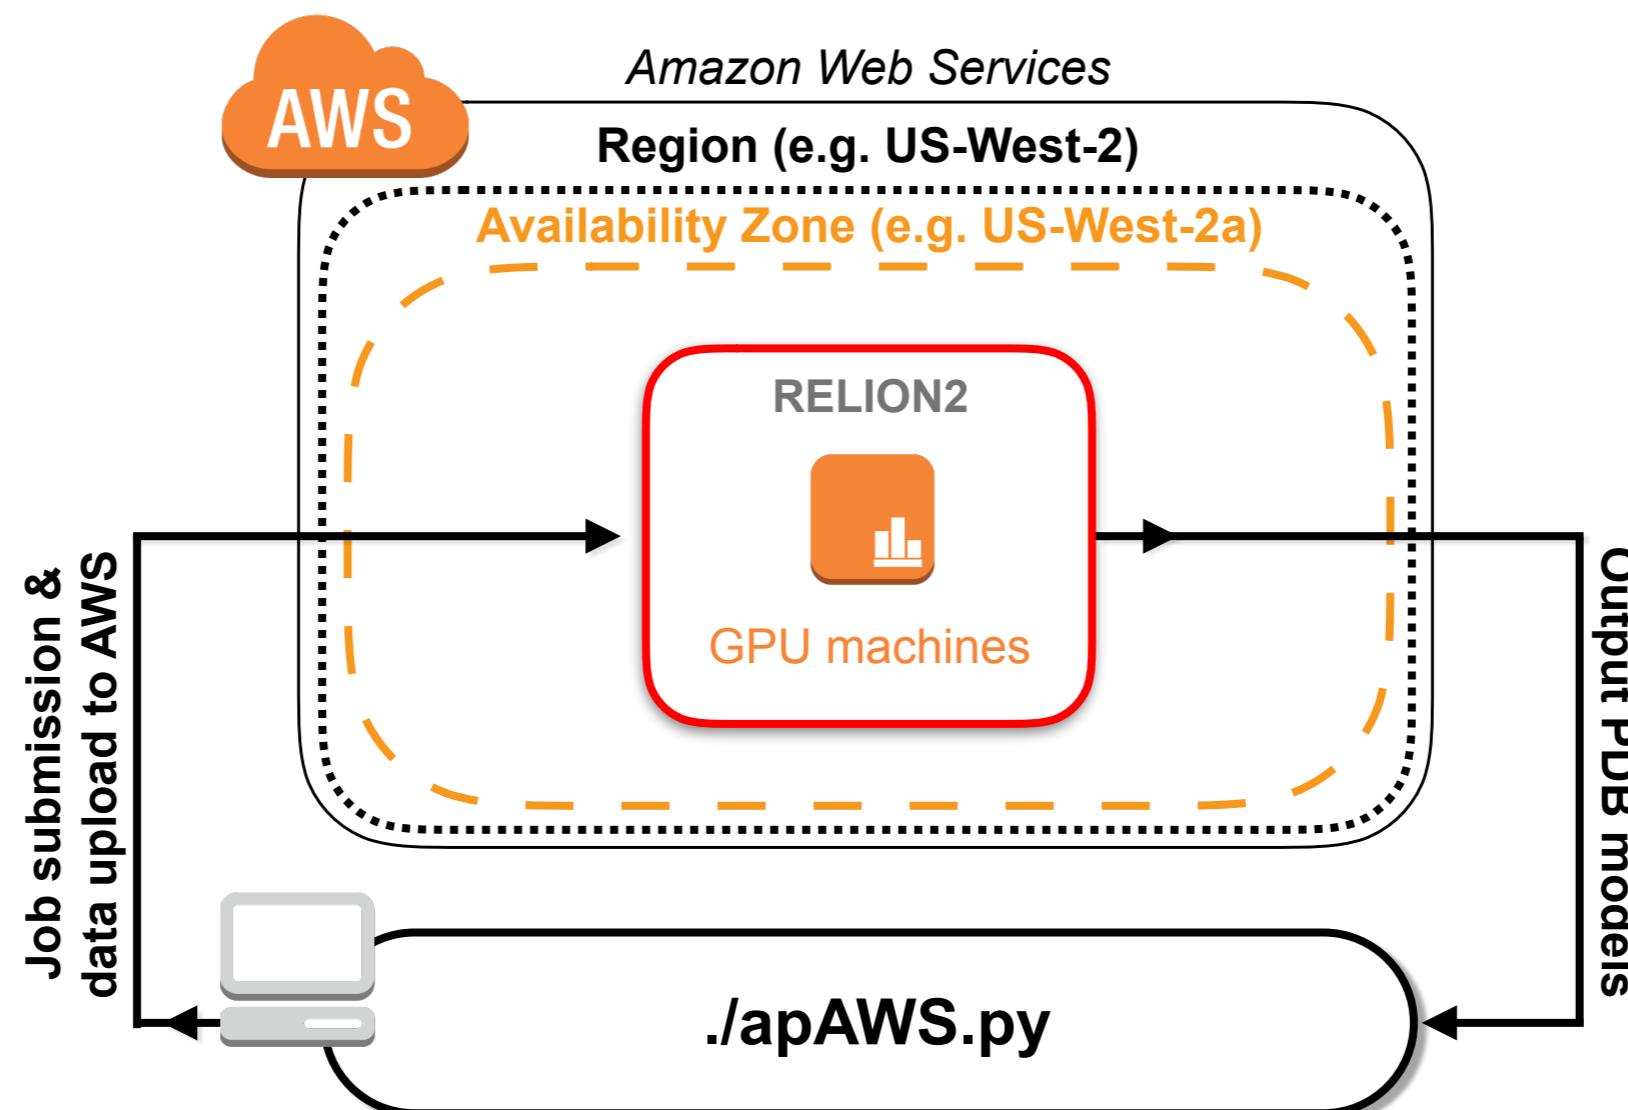
# Extending *cryoem-cloud-tools*: Rosetta on AWS

216 structures for each CM & Relax into 2.2 Å beta-galactosidase



Total AWS Cost: \$43.60

# Extending *cryoem-cloud-tools*: Appion on AWS



# Extending *cryoem-cloud-tools*: Appion on AWS

 Relion 2D Maximum Likelihood Alignment in AWS

Project: SEMC - Aldolase (381)  
Session: 17sep21j - aldolase 1.2/1.3 ultrAufoil (130K image shift, krios2 k2)  
Image Path: /gpfs/leginon/lkim/17sep21j/rawdata

Hide | Expand | Contract

Object Selection : 3

Select Particle Picker...  
3 complete

Repeat an image loop run

CTF Estimation

Estimate the CTF...  
3 complete  
1 running

Transfer results to another preset

Repeat an image loop run

Stacks : 18

Stack Creation...  
23 complete

Particle Alignment

Select Particle Alignment...  
11 complete

Run Feature Analysis...

Run Particle Clustering...

Run MaskItOn

Template Stacks

Ab Initio Reconstruction

EMAN Common Lines

SIMPLE Common Lines

Refine Reconstruction

Run name: rmaxlike40

Output directory: /gpfs/appion/lkim/17sep21j/align/

Description:

Stack Selection

Select a stack of particles to use: qstack10 ID: 98 (57,644 particles, 1.70 Å/pix, 96x96) (Relion Stack)

Select AWS Instance type

Amazon GPU Instance types

Select instance type: p2.xlarge (1 GPU, \$0.90/hour)

p2.8xlarge (8 GPU's,\$7.20/hour)

p2.16xlarge (16 GPU's, \$14.40/hour)

g3.8xlarge (2 GPU's, \$2.28/hour)

g3.16xlarge (4 GPU's, \$4.56/hour)

p3.2xlarge (1 GPU, \$3.06/hour)

p3.8xlarge (4 GPU, \$12.24/hour)

p3.16xlarge (8 GPU, \$24.48/hour)

Unbinned Clip Diameter (pixels): 96

Particle Binning: 1

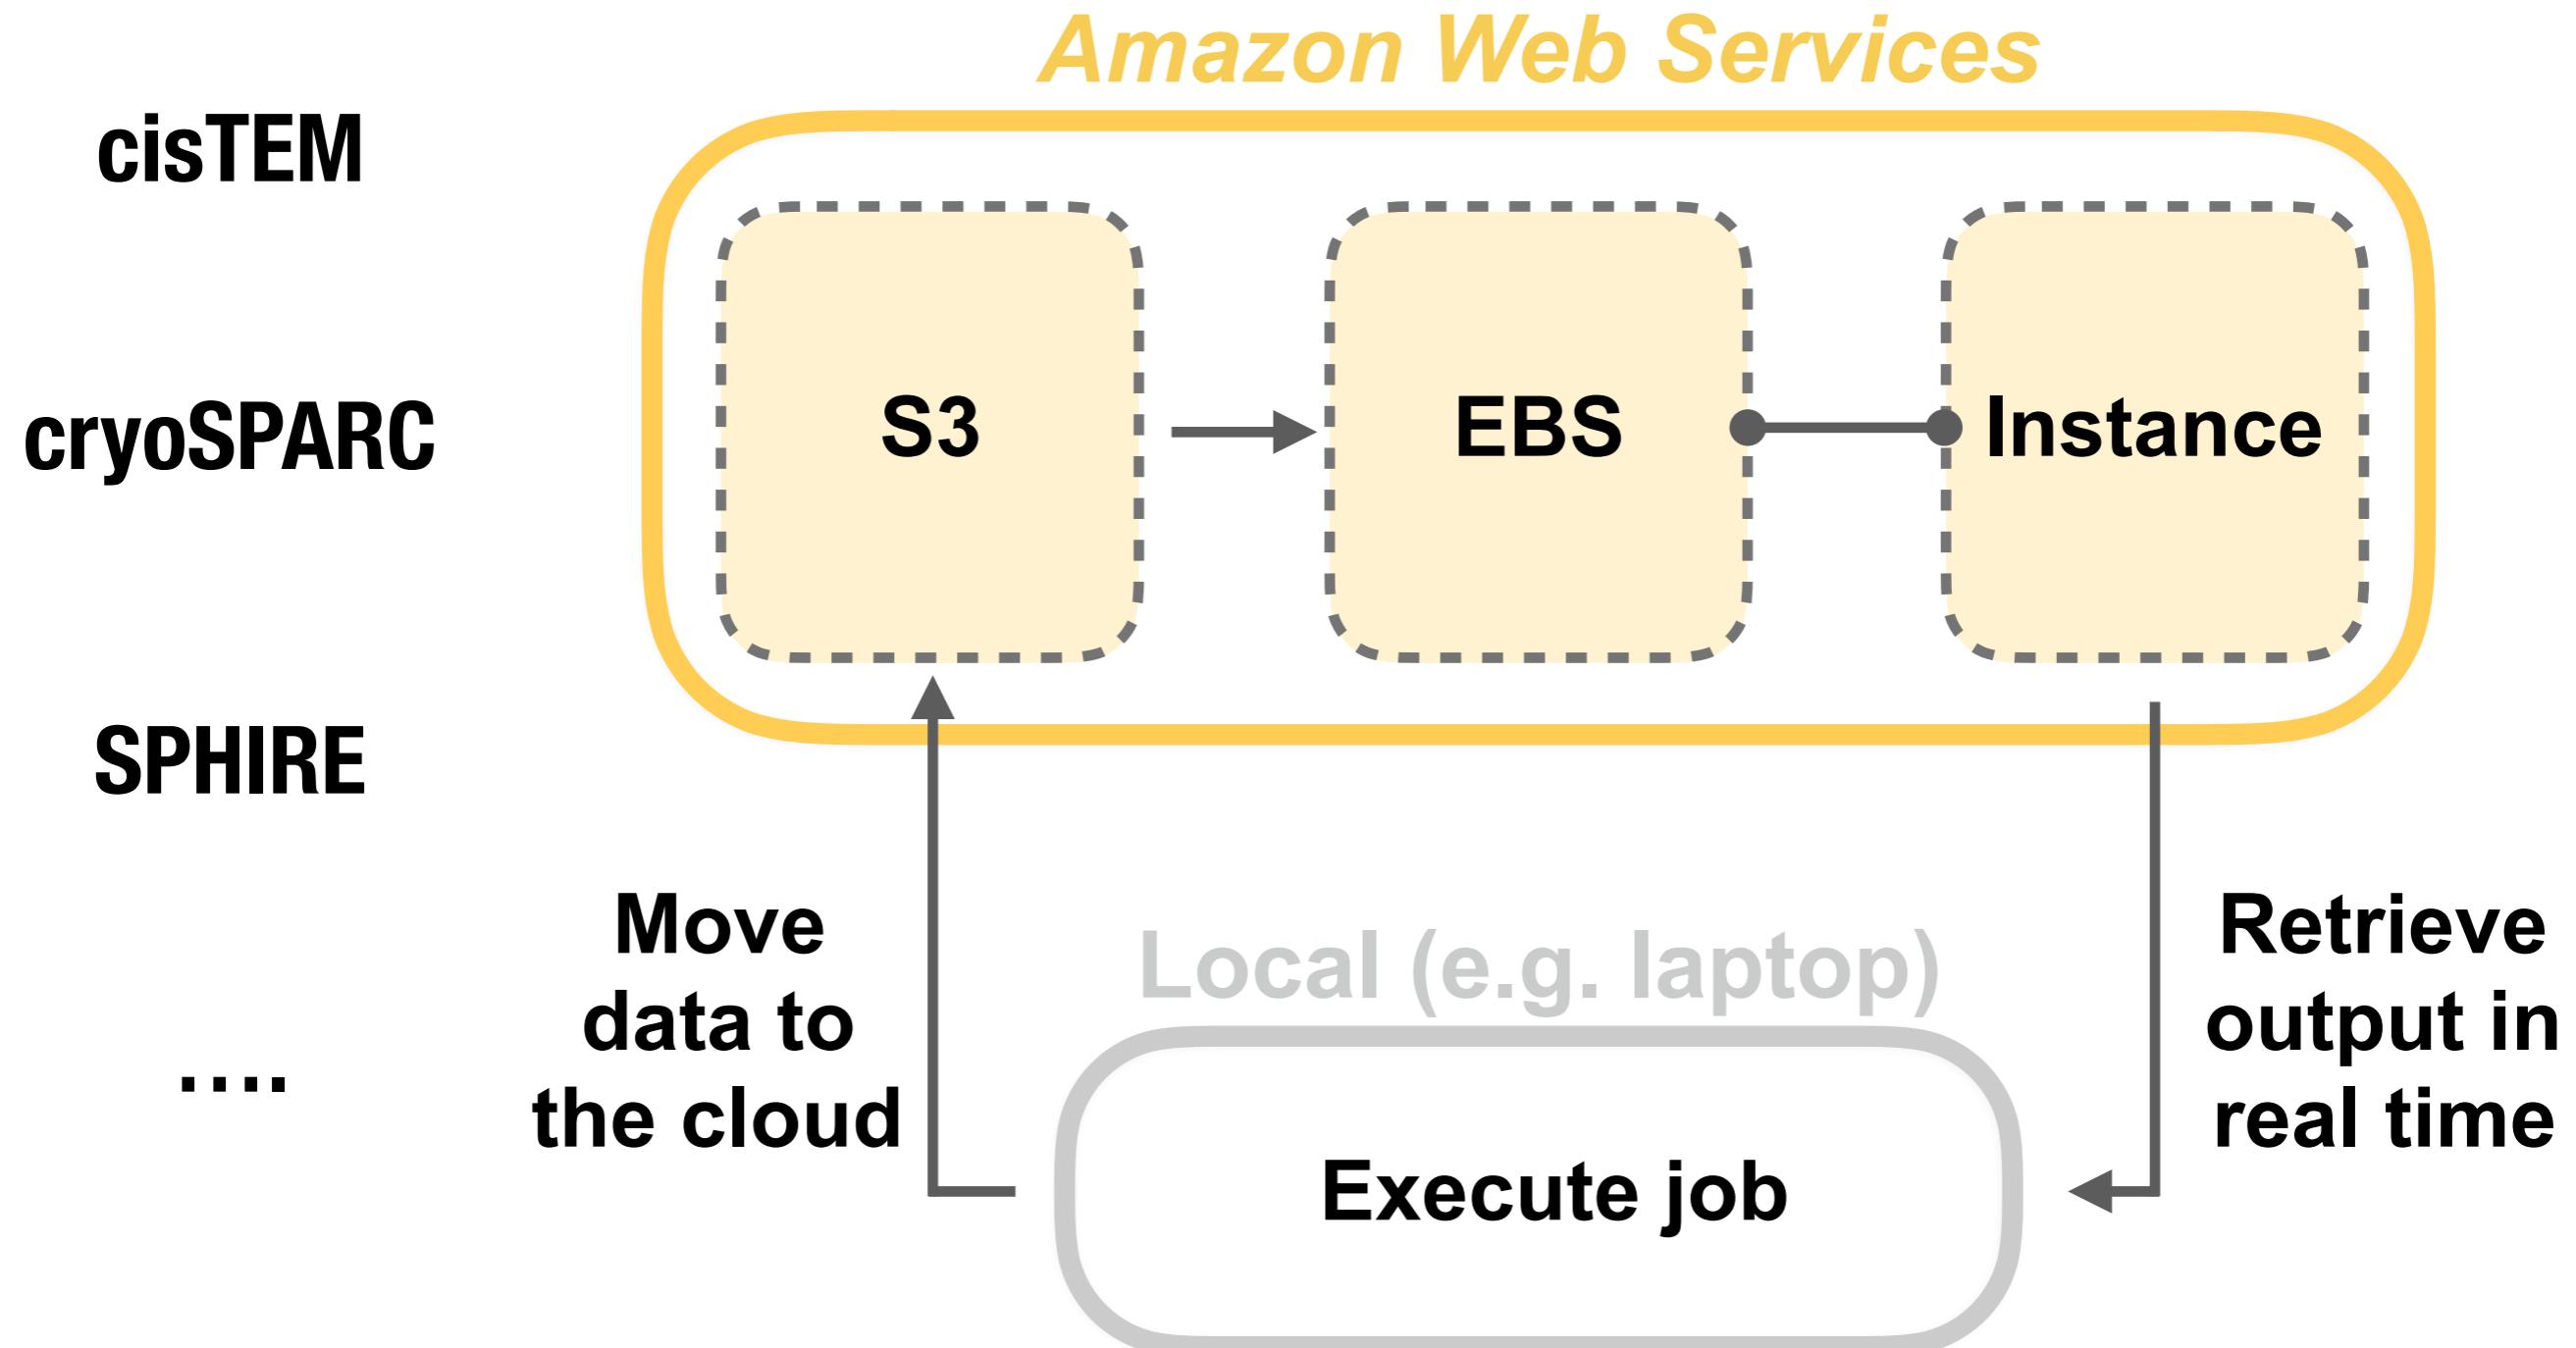
Number of Particles: 5764

Filters

(Ignore these parameters when using a Relion Stack)



# Extending *cryoem-cloud-tools*: Future directions



# Take-home points

---

## **Is the cloud useful for cryo-EM?**

Yes - the cloud offers the following features:

- 1) Outsourcing of IT administration
- 2) Large choice of computing configurations
- 3) Active & cold data storage options

## **What types of resources are at the disposal of users?**

Small to large CPU/GPU systems.

## **What are the appropriate workflows and benchmark comparisons?**

Benchmarking can be found online at [cryoem-tools.cloud](http://cryoem-tools.cloud) under the RELION2 section of AWS

*Punchline:* There are many types of GPU machines to satisfy all types of RELION2 computing tasks



# Take-home points

---

## How does the cost compare to local infrastructure?

Per hour - these machines are more expensive than simply buying a standalone workstation. But, standalone workstations don't come with back up storage.

Example storage quote:

\$200/TB/year -> \$0.0167/GB/mo. (S3: \$0.023/GB/mo.) WITHOUT IT support

## When should someone use the cloud?

- 1) Whenever someone is waiting for computing time. For RELION, 4xGPU machines cost \$4.56/hr to run, which is much less than the hour wage of individuals working in the lab
- 2) Large facilities - the cloud lets users launch jobs from a single point of contact and can expand (seemingly) indefinitely
- 3) Need to solve successive structures under time limit



# Questions?

---

**[cryoem-tools.cloud](#)**

***Demo of cryoem-cloud-tools @  
Thursday breakout session***





# Take-home points

---

Cloud for software dissemination

No need to be pigeon-holed into a single setup

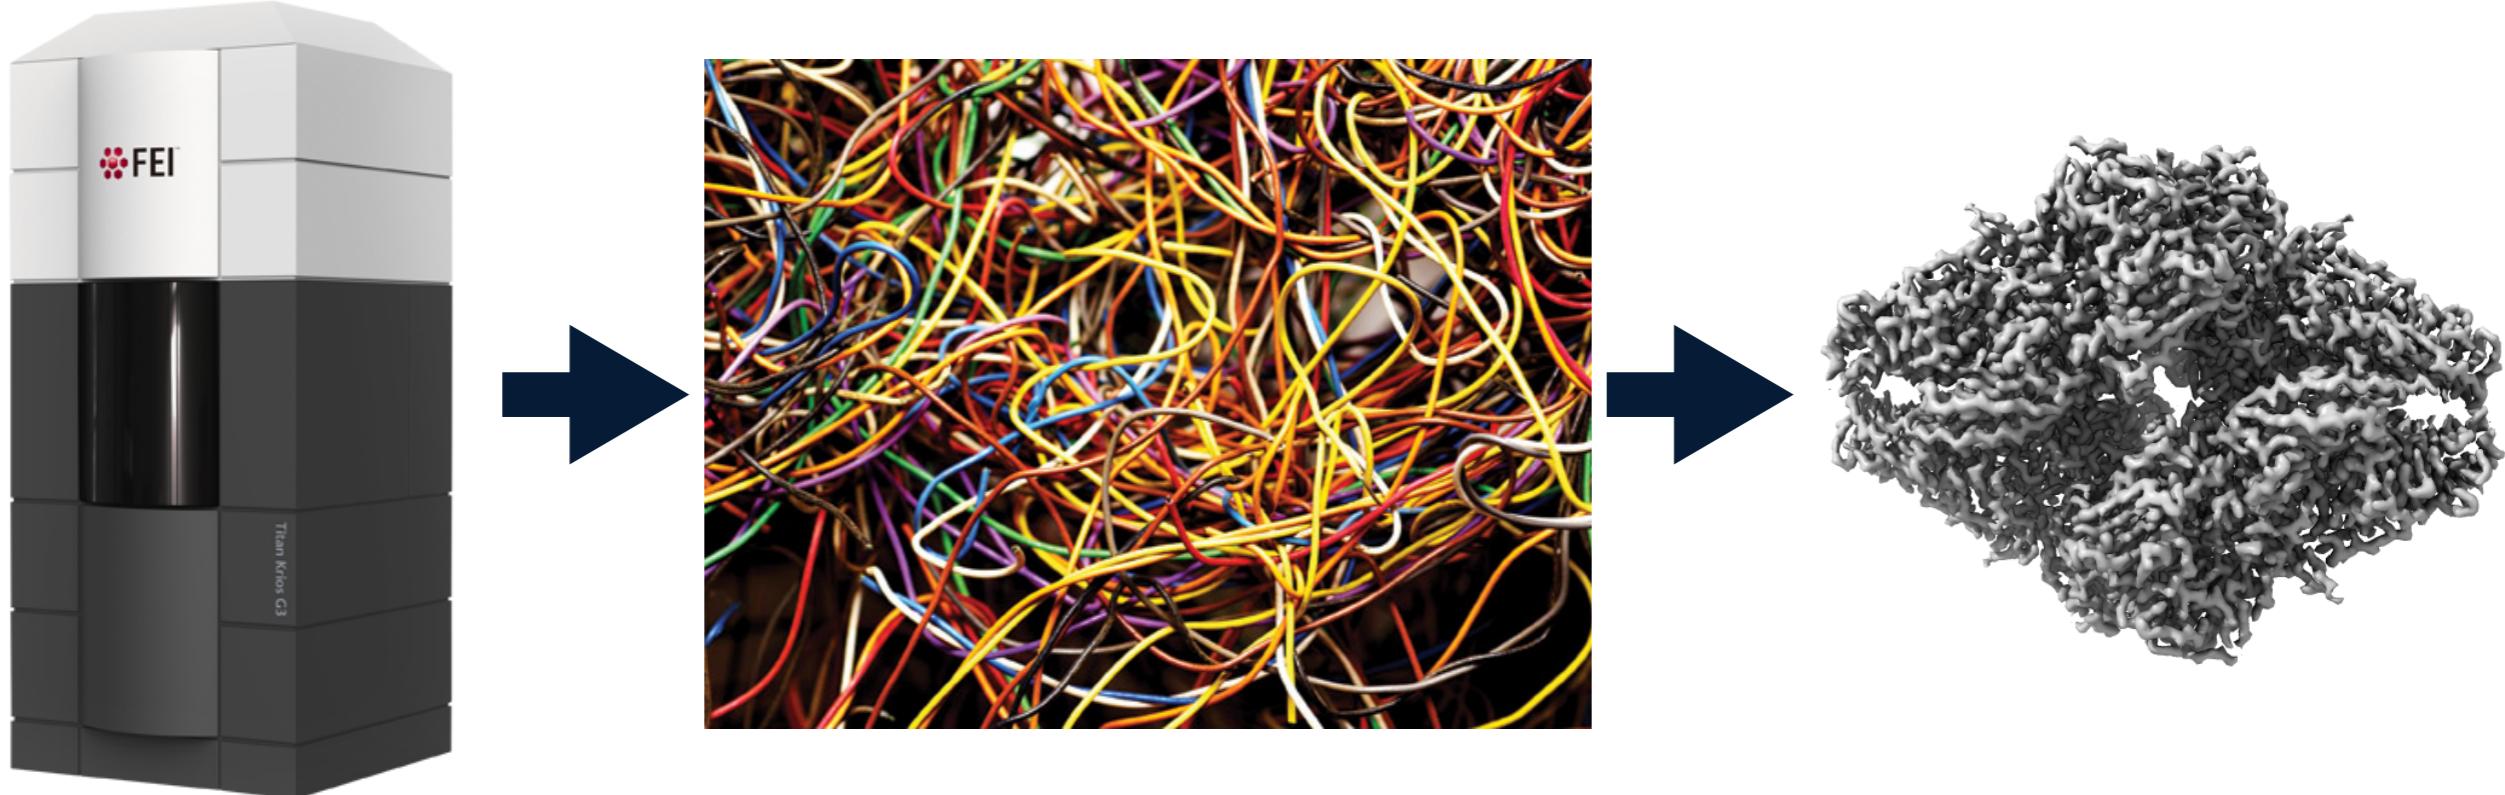
‘Hybrid’ architectures will have the biggest impact on the field - end users don't need to see the backend of AWS

Platform for plugging in new analysis (e.g. ML)

Underlying data movement tools > agnostic to cloud; could be adapted for anywhere



# What do we need?



On-the-fly data movement, analysis, and storage as movies are collected

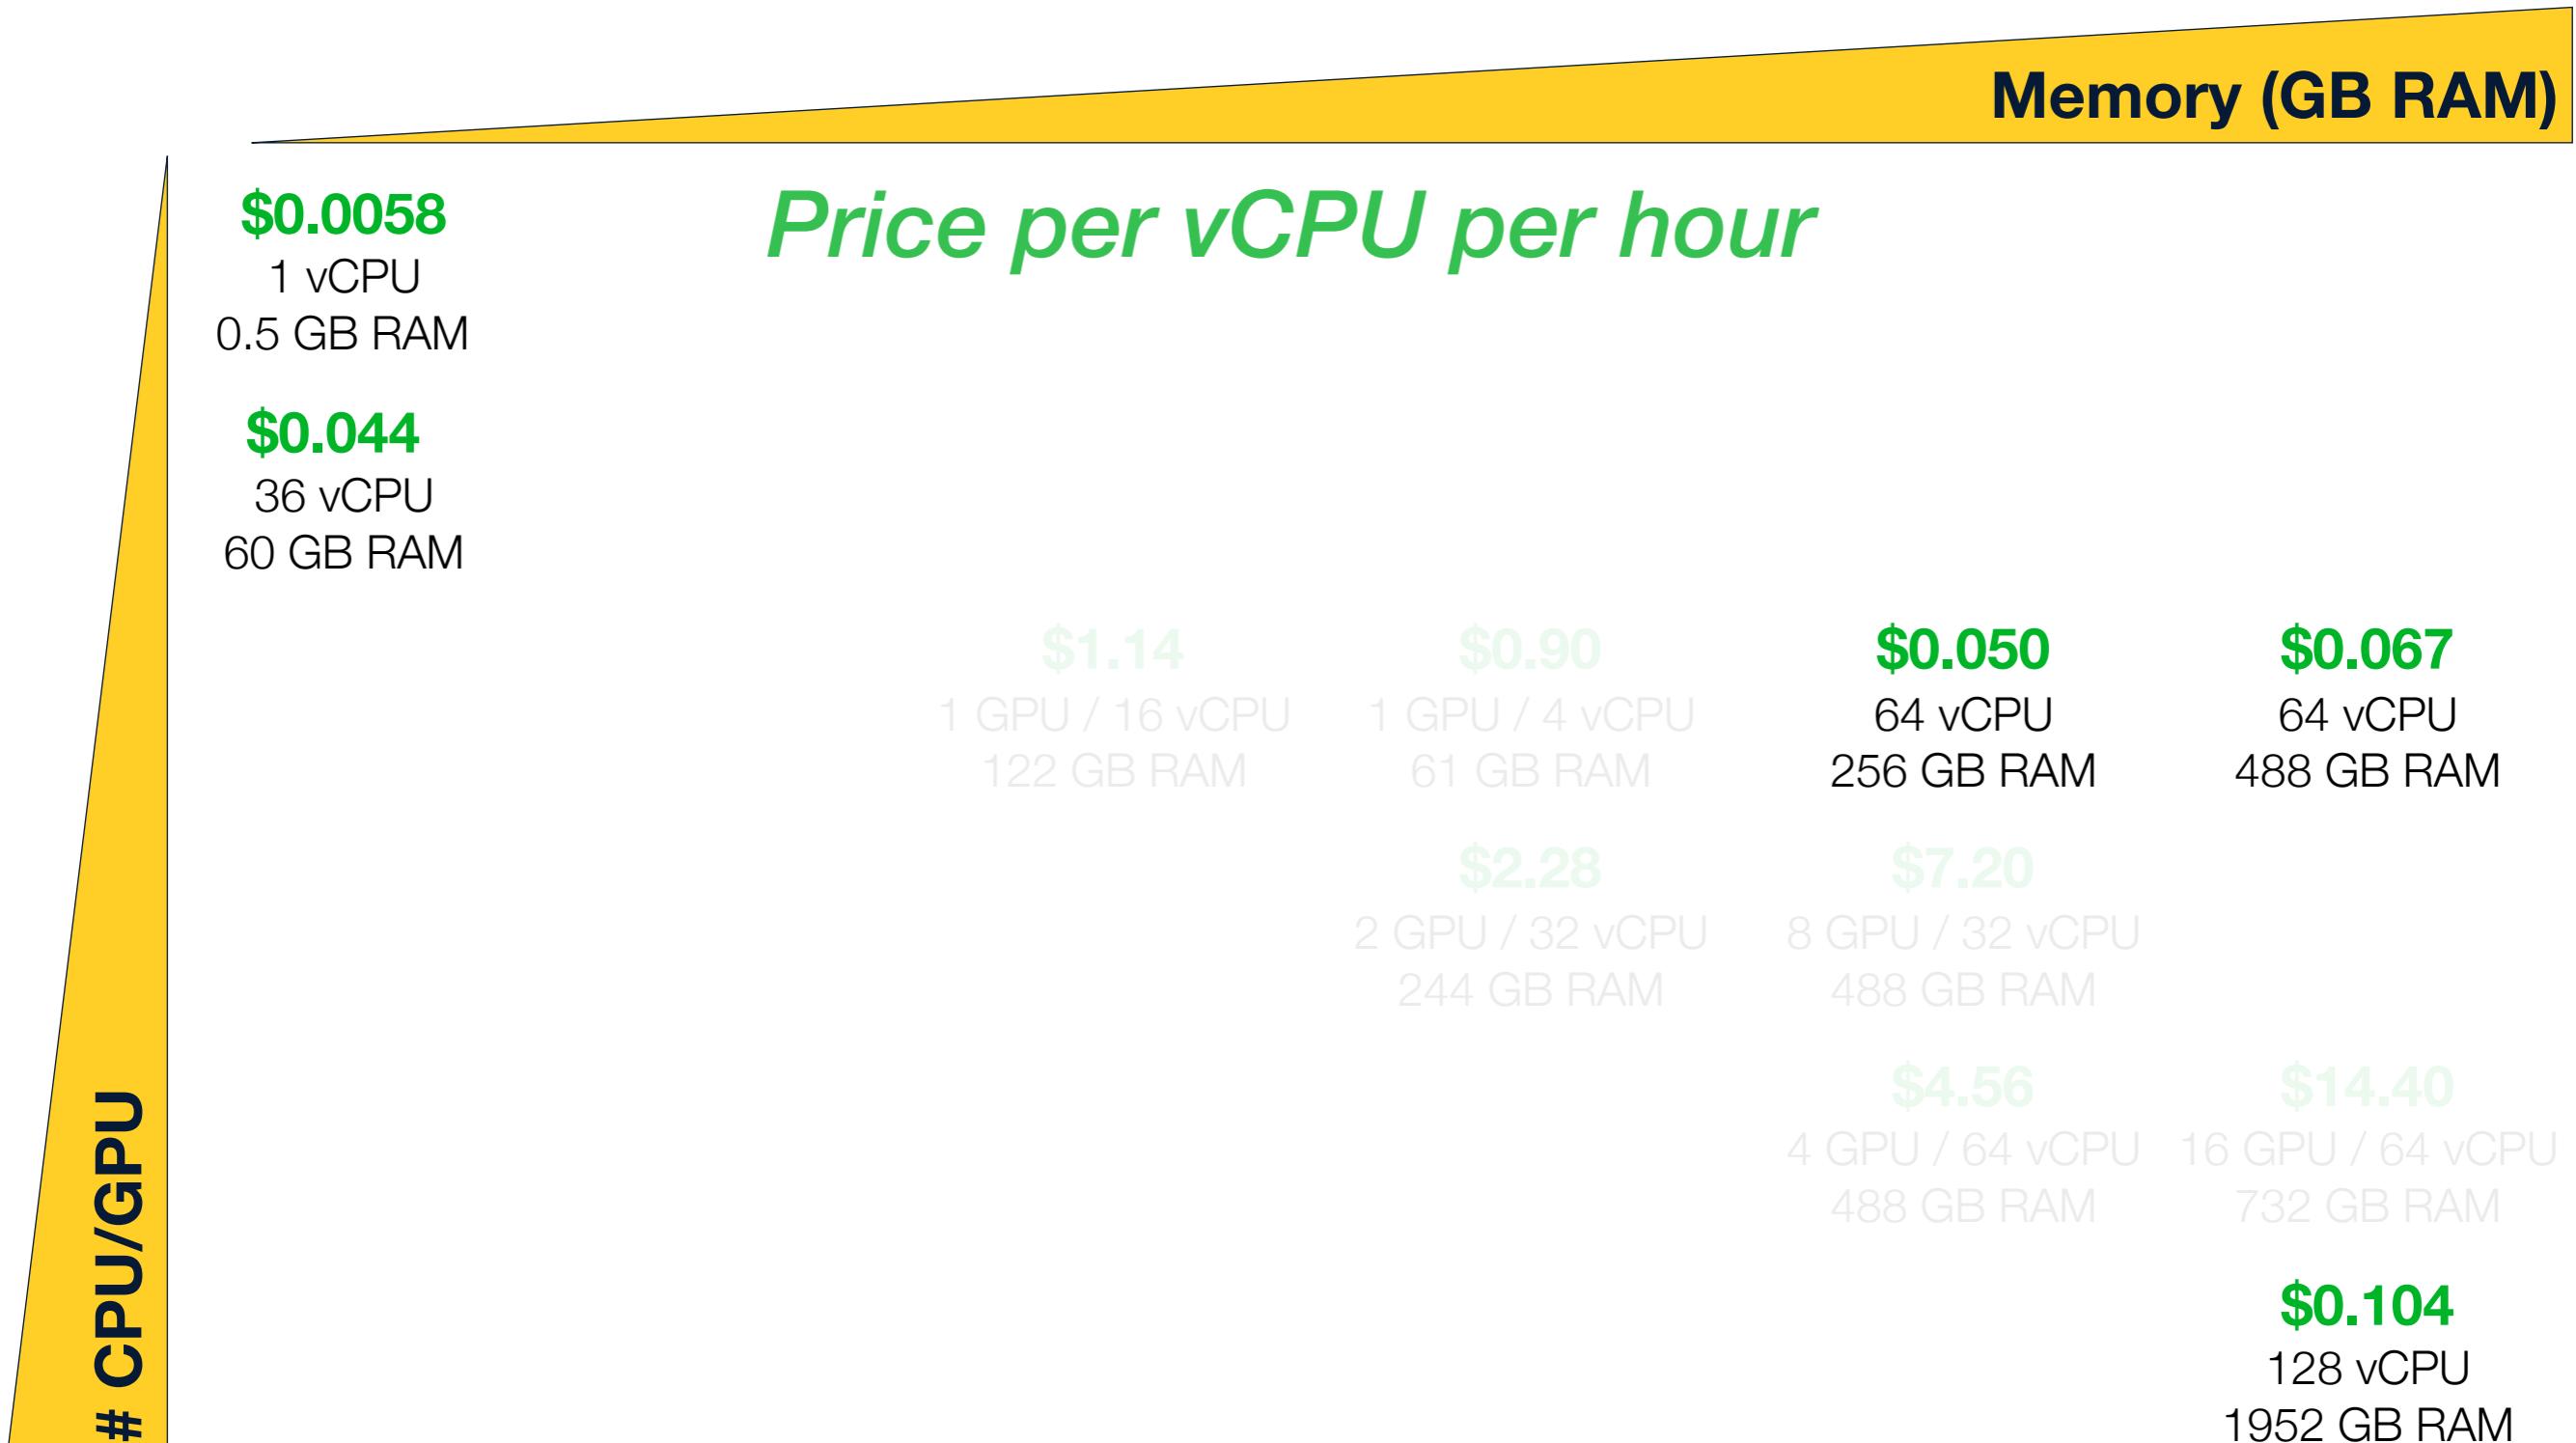
- ✓ For *all* users who use facility



# Computing options on AWS

**Instance** = virtual machine on AWS

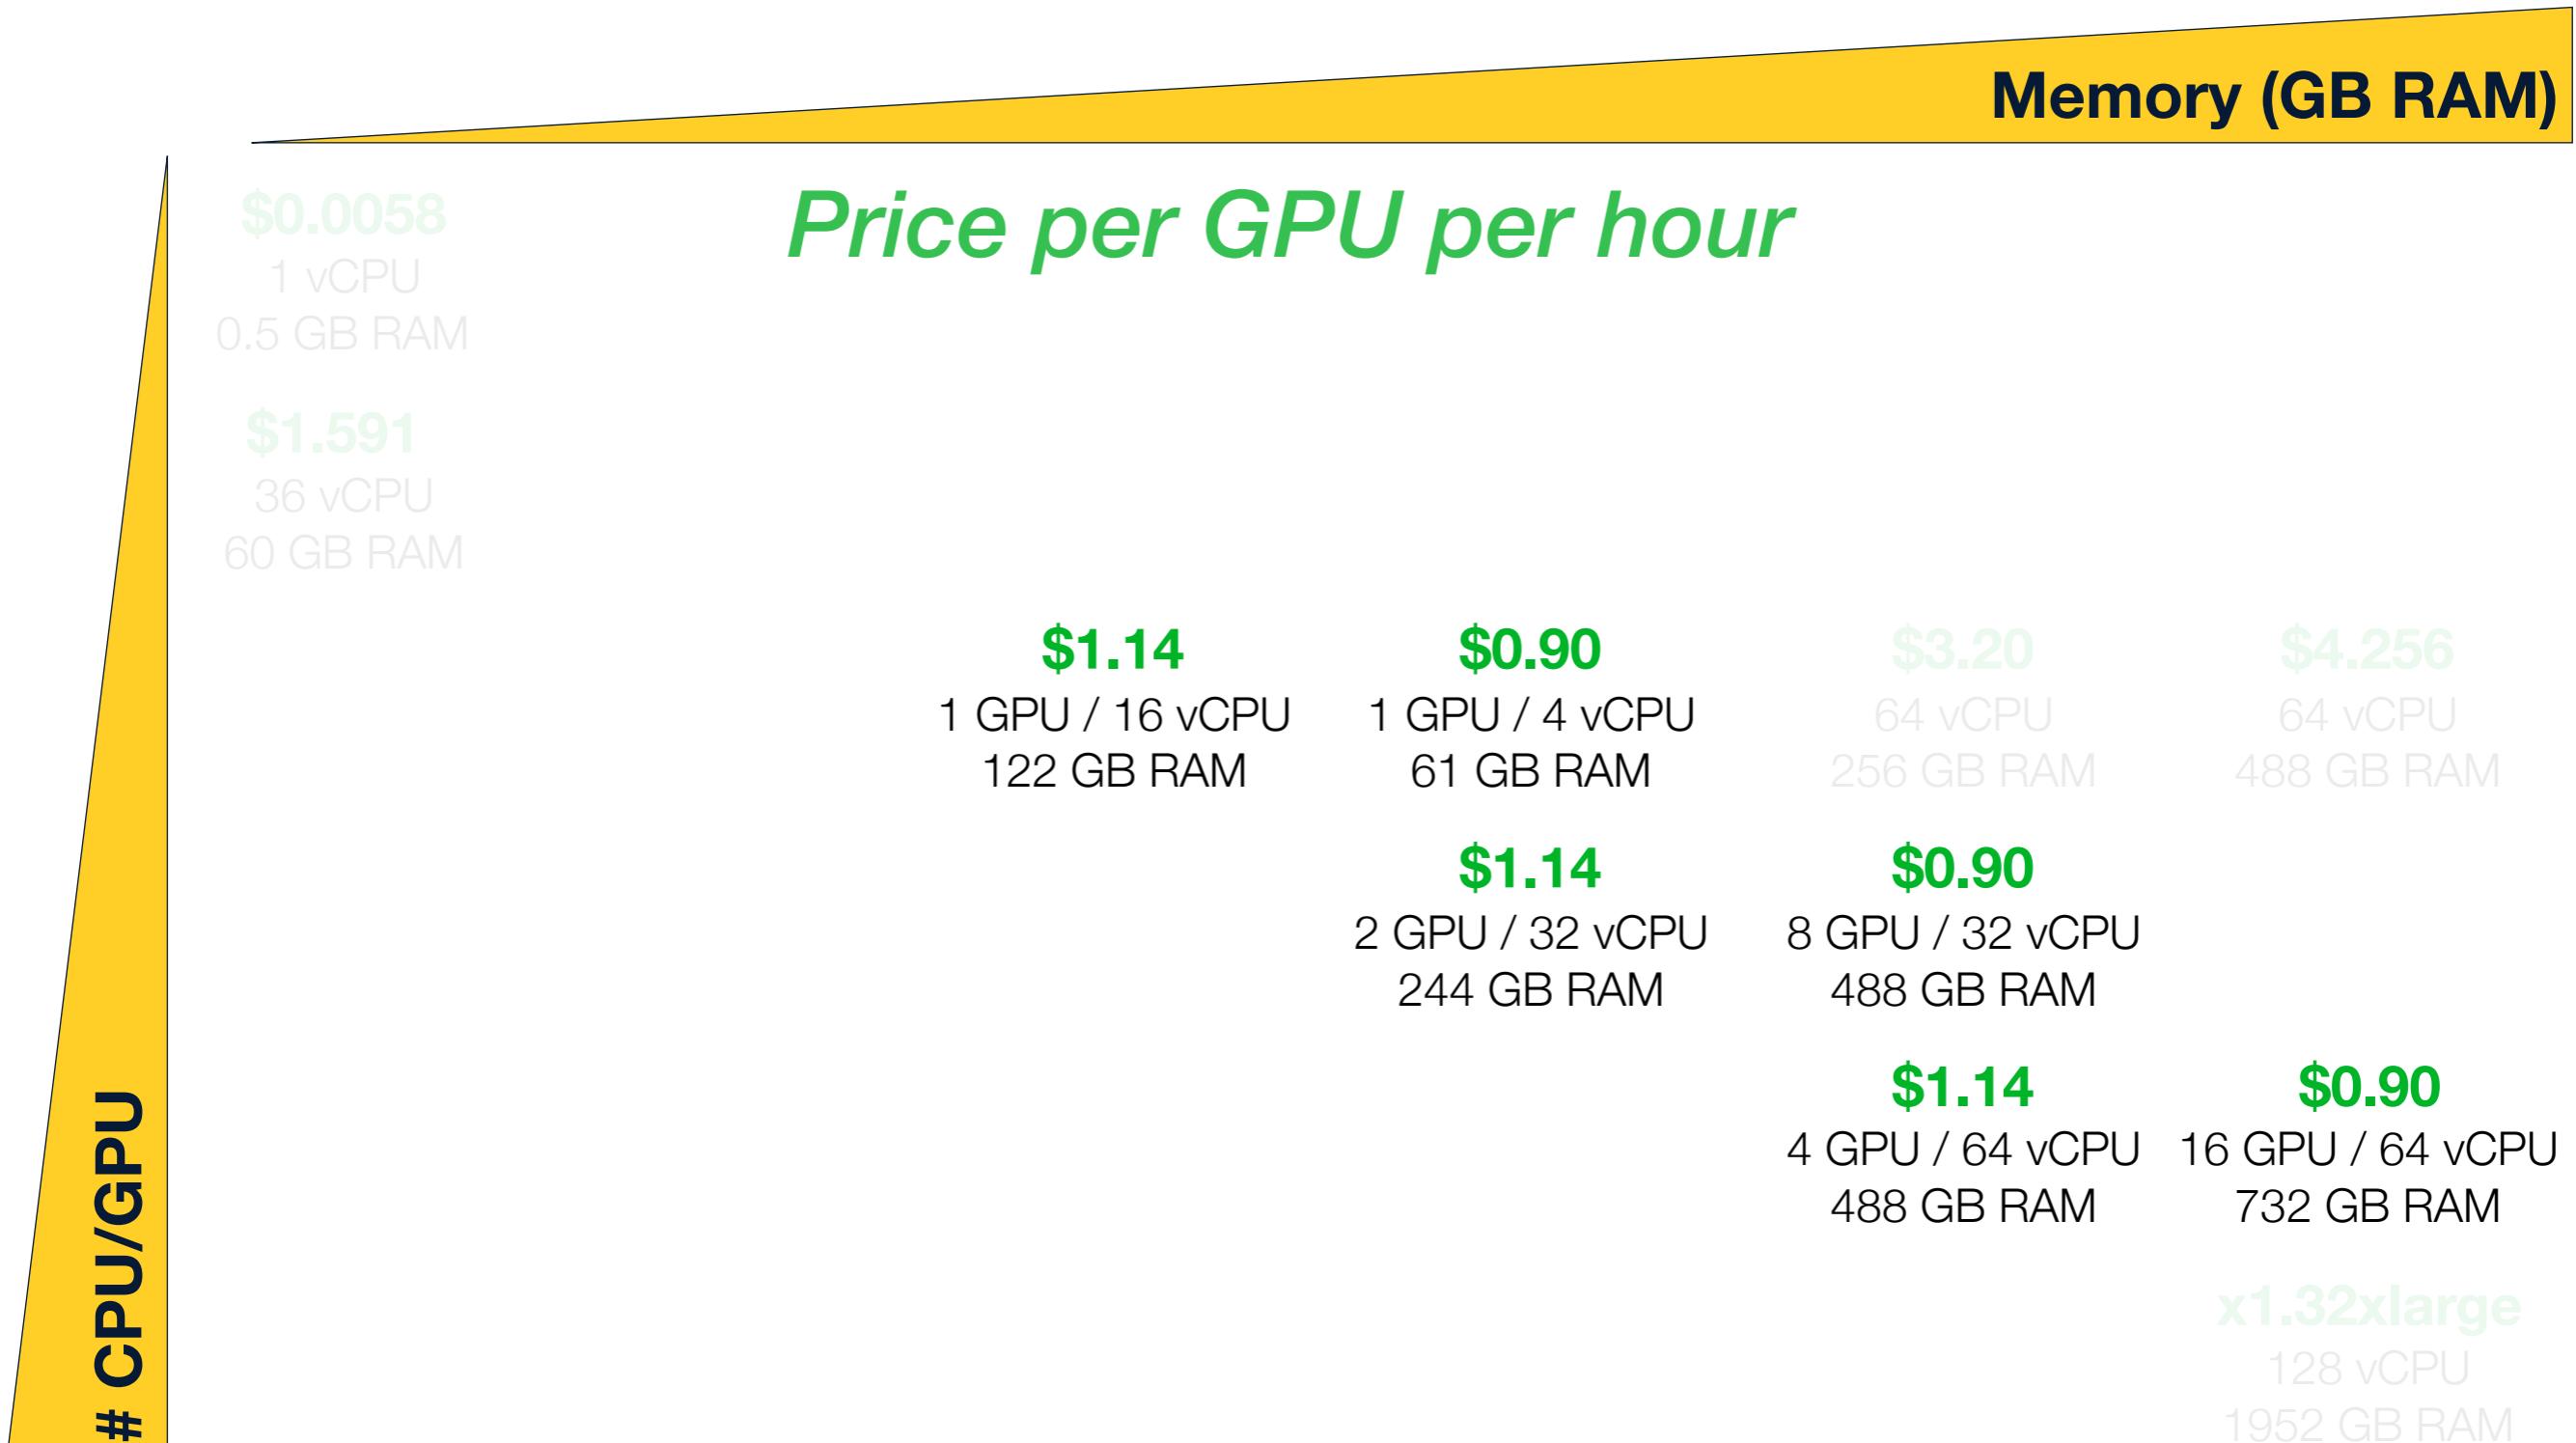
**vCPU** = hyper thread on CPU core



# Computing options on AWS

**Instance** = virtual machine on AWS

**vCPU** = hyper thread on CPU core



	GPU Workstation					Amazon Web Services							
	Total processing time (hr)	Estimated elapsed time (working days) <sup>^</sup>	Number of GPUs	Number of CPU cores	Amount of RAM (GB)	Total processing time (hr)	Estimated elapsed time (working days) <sup>^</sup>	Data moving time (hr)	Job running time (hr)	EC2 Cost (USD)*	Number of VMs	VM type	VM cost per hour (USD)
<b>Movie alignment</b>	34.0	1	4	8	64	13.26	1	11.06	2.20	\$146.72	5	x1.32xlarge	\$13.34
<b>Autopick</b>	2.0	3	4	8	64	1.30	2	1.10	0.20	\$9.36	1	p2.8xlarge	\$7.20
<b>CTF Estimation</b>	0.4	3	4	8	64	0.66	2	0.55	0.11	\$4.75	1	p2.8xlarge	\$7.20
<b>Extract</b>	-	3	4	8	64	1.30	2	0.68	0.62	\$1.04	1	m4.4xlarge	\$0.80
<b>2D classification</b>	8.0	3	4	8	64	4.30	2	0.17	4.13	\$61.92	1	p2.16xlarge	\$14.40
<b>3D autorefine</b>	5.3	4	4	8	64	5.38	2	0.08	5.30	\$38.74	1	p2.8xlarge	\$7.20
<b>Movie extract</b>	9.5	4	4	8	64	4.07	3	3.35	0.72	\$179.73	8	d2.8xlarge	\$5.52
<b>Movie-refine</b>	17.5	5	4	8	64	4.85	3	0.75	4.10	\$64.69	1	x1.32xlarge	\$13.34
<b>Polish</b>	25.0	7	4	8	64	6.80	3	1.40	5.40	\$90.70	1	x1.32xlarge	\$13.34
<b>3D classification</b>	7.0	8	4	8	64	6.17	4	0.17	6.00	\$44.42	1	p2.8xlarge	\$7.20
<b>3D autorefine</b>	6.0	8	4	8	64	6.40	4	0.10	6.30	\$46.08	1	p2.8xlarge	\$7.20
<b>Rosetta</b>	N/A	9	4	8	64		5				6	c4.8xlarge	\$1.59
<b>S3 storage cost#</b>	-	-	-	-	-	-	-	-	-	\$690.00	-	-	-
<b>Totals</b>	114.7	10	-	-	-	54.5	5	19.4	35.1	\$1,378.15	-	-	-

