
aws
SUMMIT
SAN FRANCISCO | APRIL 20–21, 2022

CMP208

The future of computational fluid dynamics in the cloud

Sean Smith
Senior HPC Solutions Architect
AWS

Karthik Raman
Principal HPC Solutions Architect
AWS

Agenda

Introduction

What is CFD, and why does it matter?

1. Compute selection

2. Storage

3. Visualization

4. Cost optimization

5. Innovating with ML

What is CFD, and why does it matter?

Companies need a way to assess the performance of new designs

The common goal is to go from conceptual design to full design in the shortest possible time with the least expense

Physical tests are difficult; **specifically, they are expensive and time-consuming**

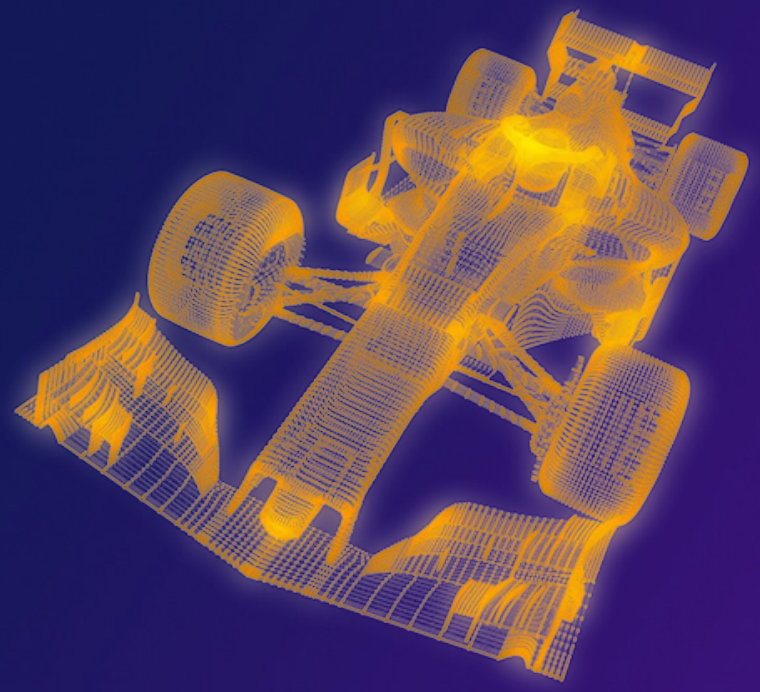

Formula 1 on AWS

Formula 1 designed the 2022 car using CFD models on AWS; they faced several challenges with traditional methods

- **No access to wind tunnels** for 12 to 24 months, only CFD
- Reduced average simulation time from **60 hours** down to **10 hours** using AWS

Agility = the right tool for the job

Company	Simulator	Hardware
Ansys	Fluent	CPU, GPU
Dassault	Abaqus	CPU, GPU
Siemens	STAR-CCM+	CPU
LSTC/Ansys	LS-DYNA	CPU, GPU
OpenFOAM Foundation	OpenFOAM	CPU, GPU
COMSOL	Multiphysics	CPU
Altair	AcuSolve	CPU

Reference architecture

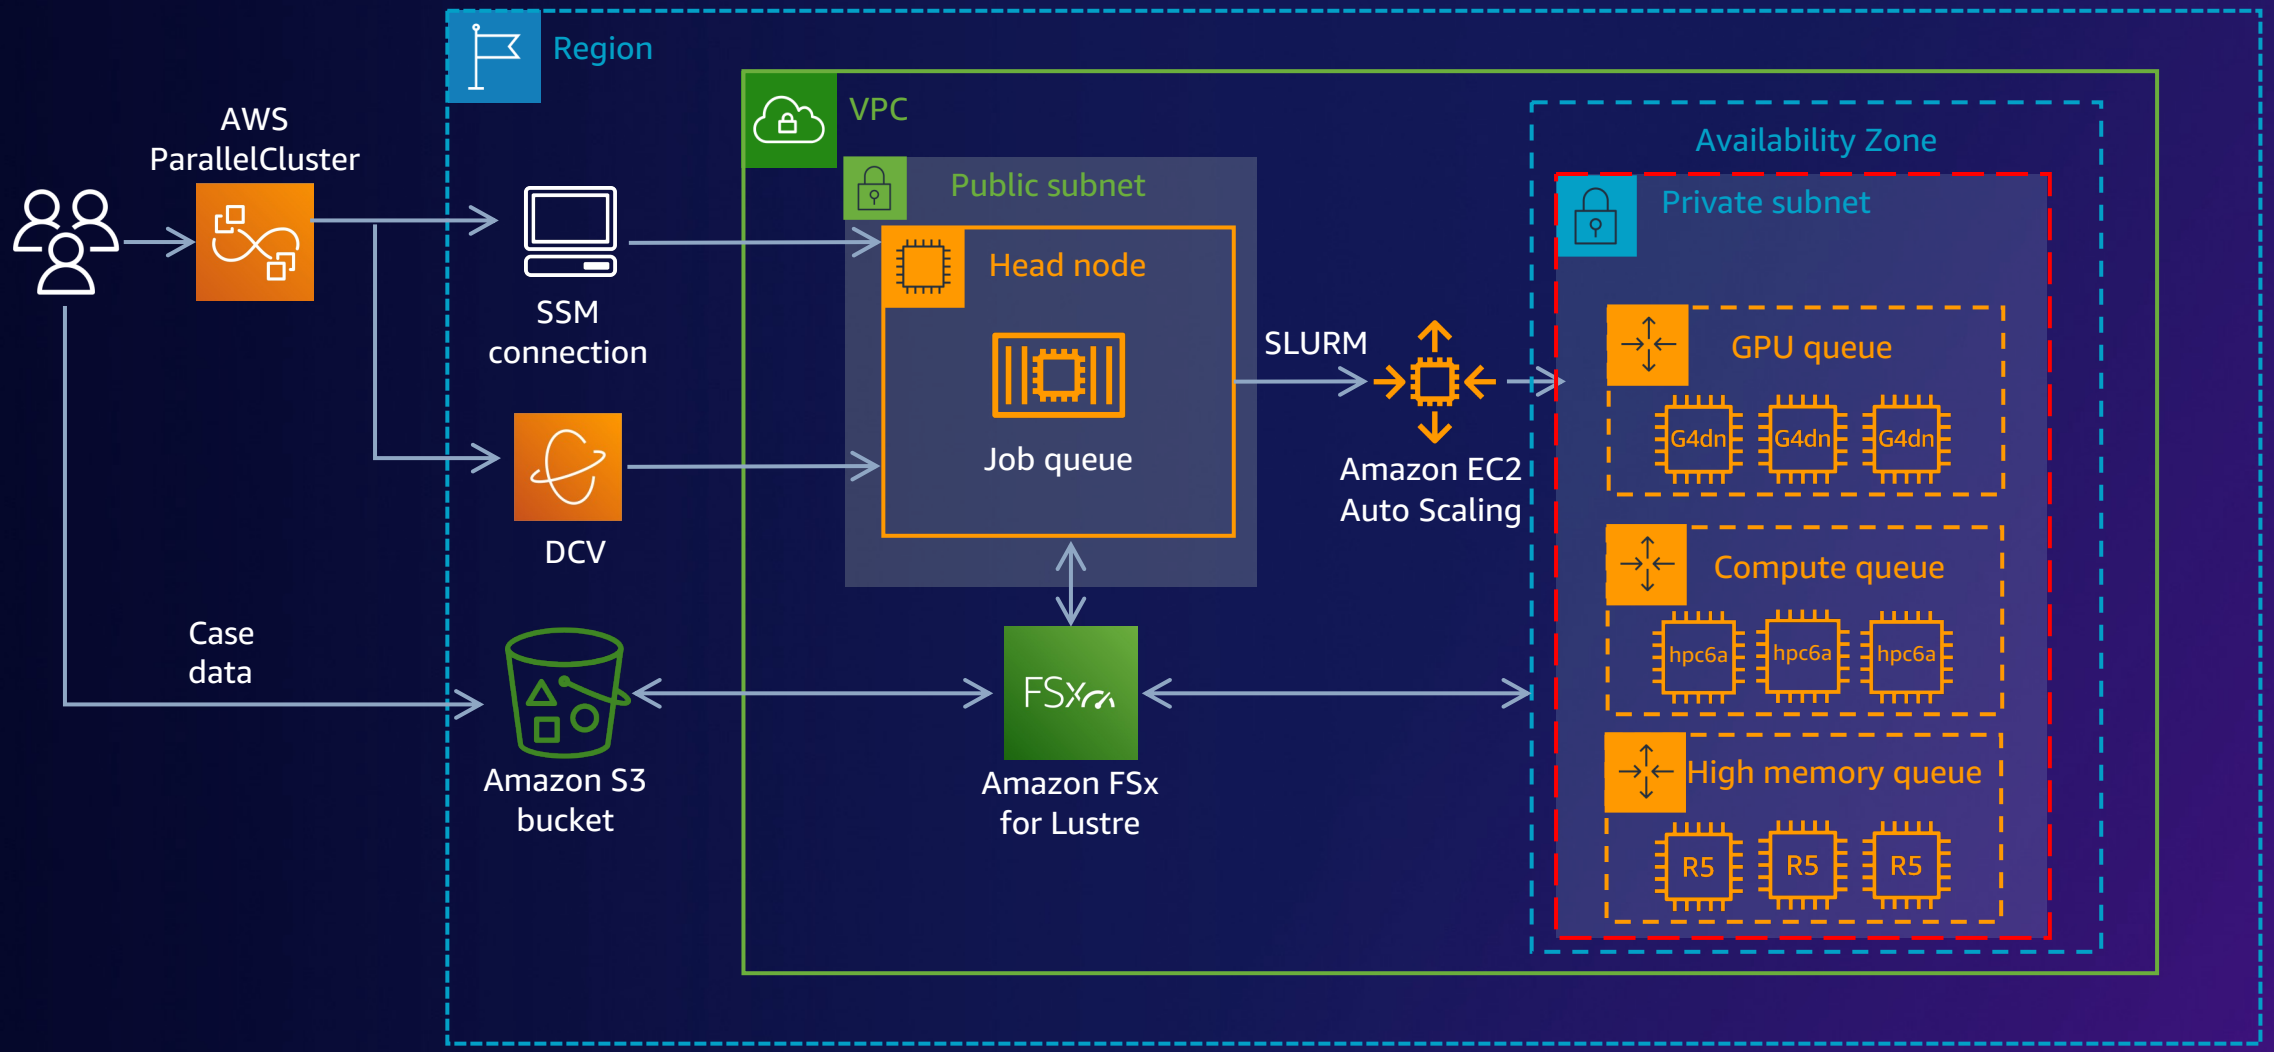

Instance type selection

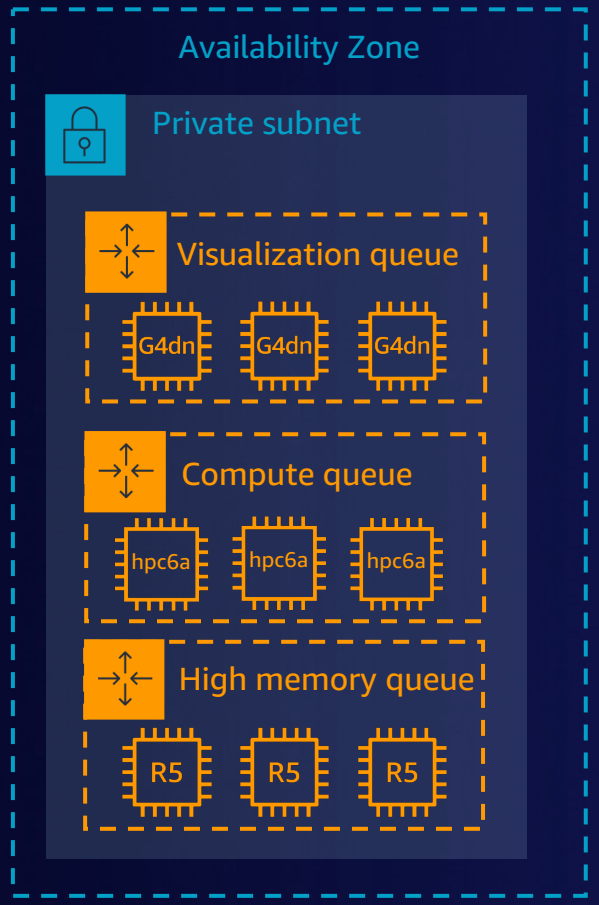

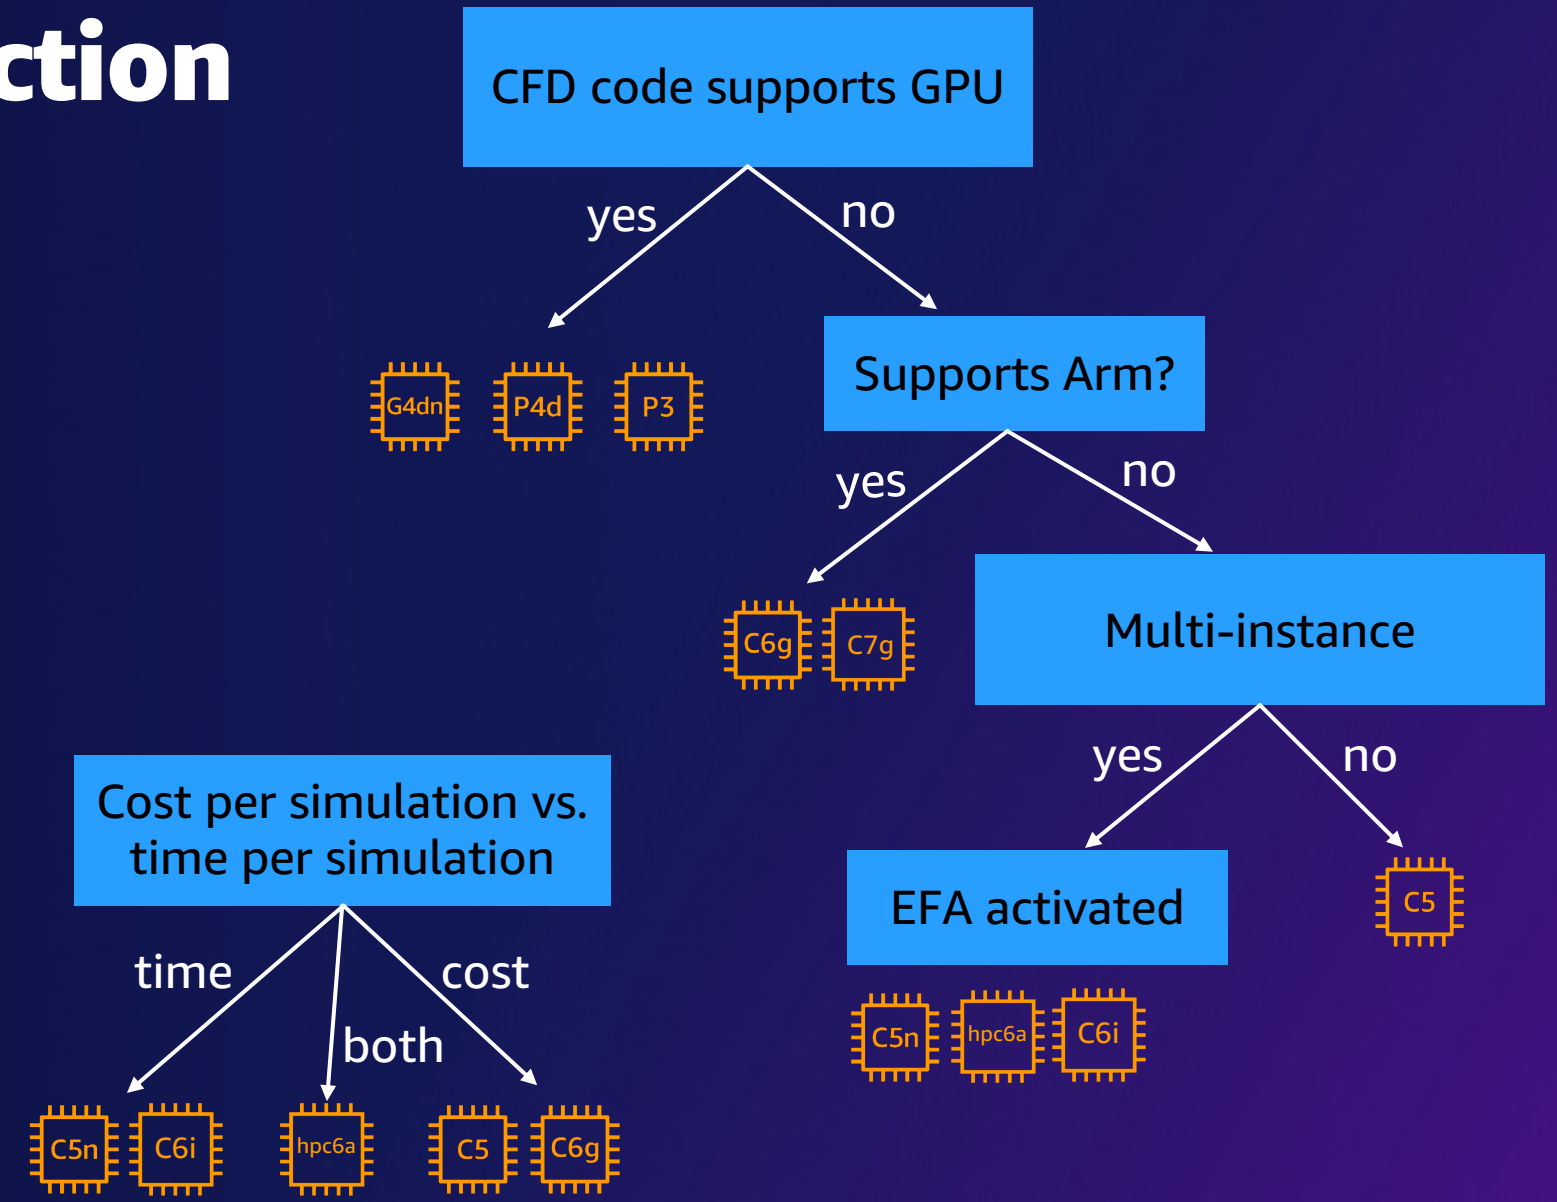

Amazon EC2 Hpc6a instances

DESIGNED TO DELIVER THE BEST PRICE PERFORMANCE FOR COMPUTE-INTENSIVE, HIGH PERFORMANCE COMPUTING WORKLOADS IN AMAZON EC2

AMD third-generation EPYC Milan processors, 96 cores, up to 3.6 GHz frequency, and 384 GB of RAM

Up to 65% better price performance over comparable Amazon EC2 x86-based compute-intensive instances

Elastic Fabric Adapter (EFA) is activated by default for 100 Gbps networking for inter-instance communications

AWS Regions at launch: US East (Ohio) and AWS GovCloud (US-West)

Multi-node

- For multi-node, the network plays a key role
- CFD splits the solution into lots of partitions and then communicates the values around neighboring cells
- Key metric is cells per core

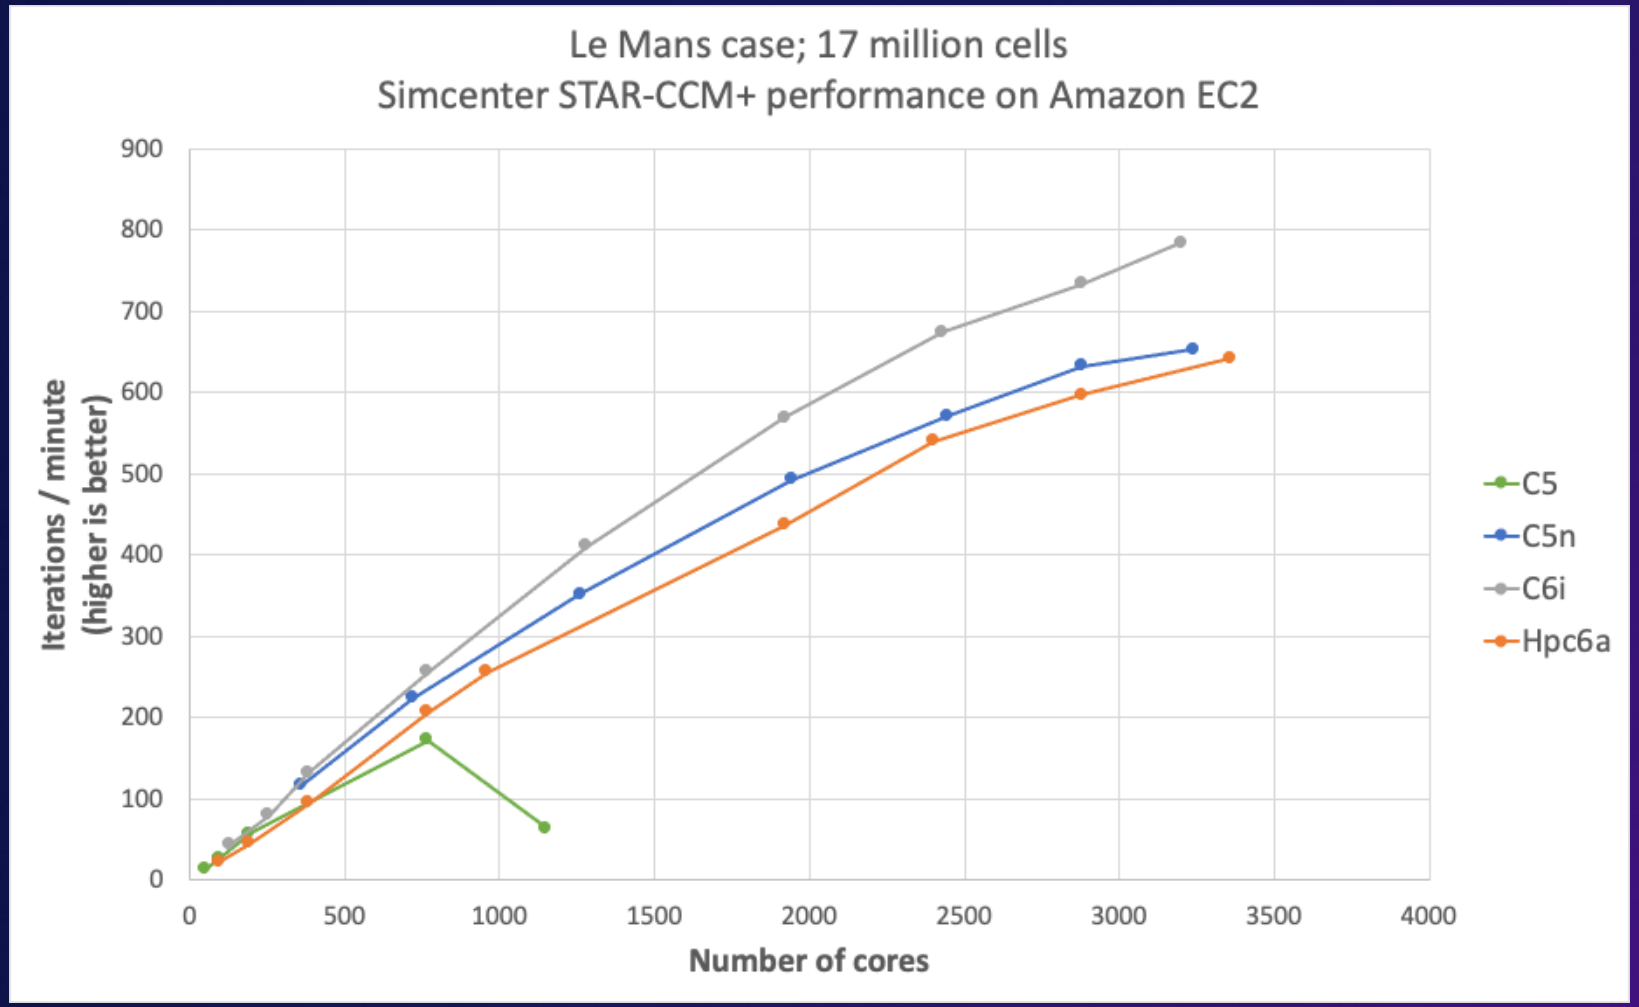

Cells per core

To achieve best parallel efficiency, we recommend greater than **50K cells/core**

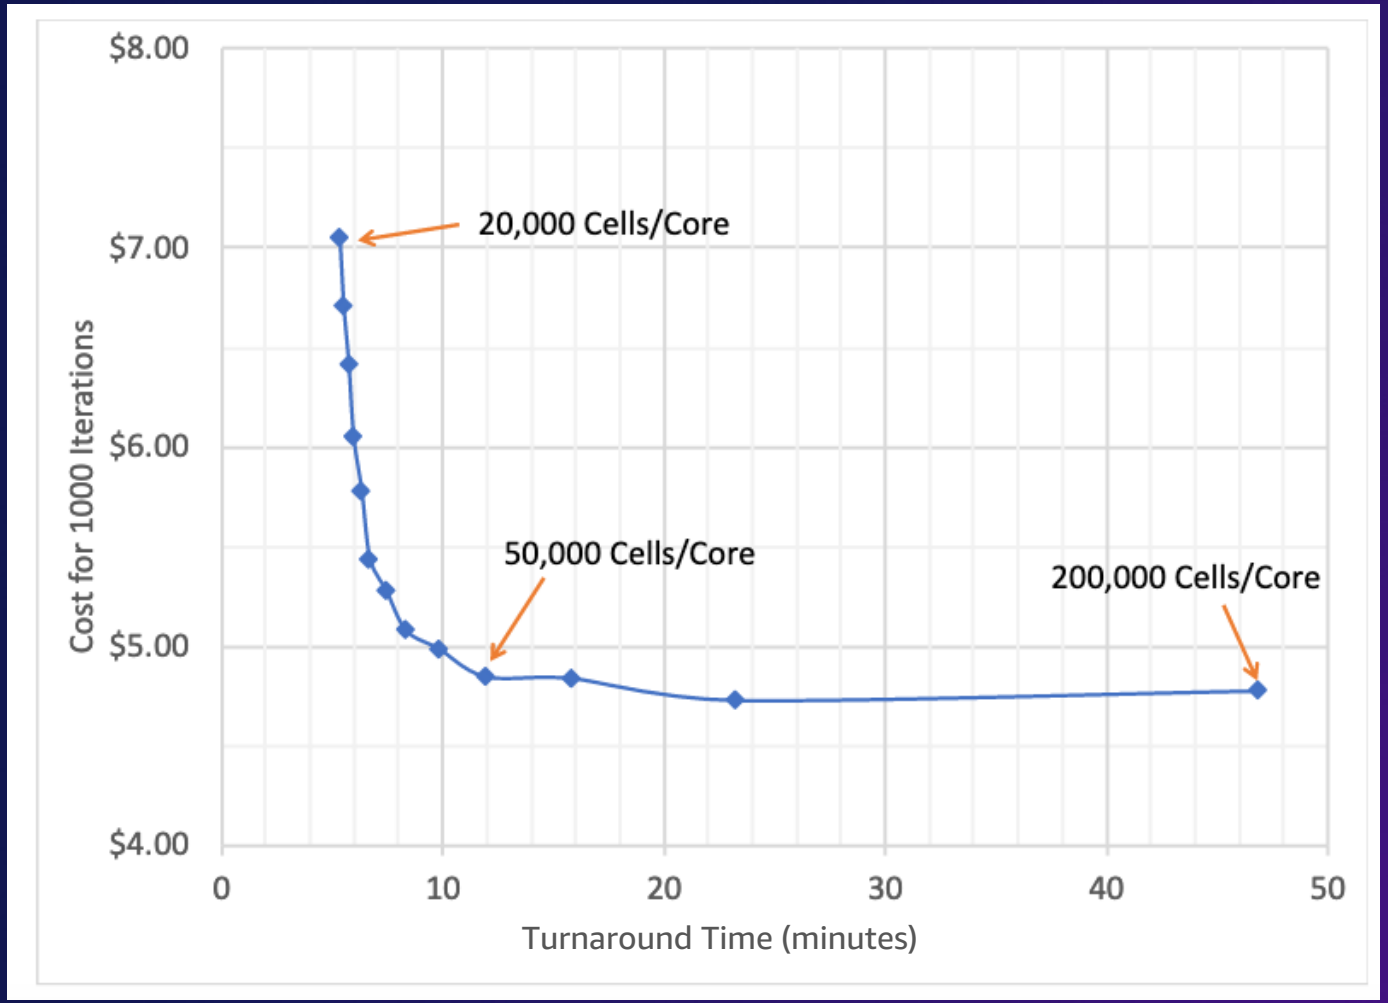

Reference architecture

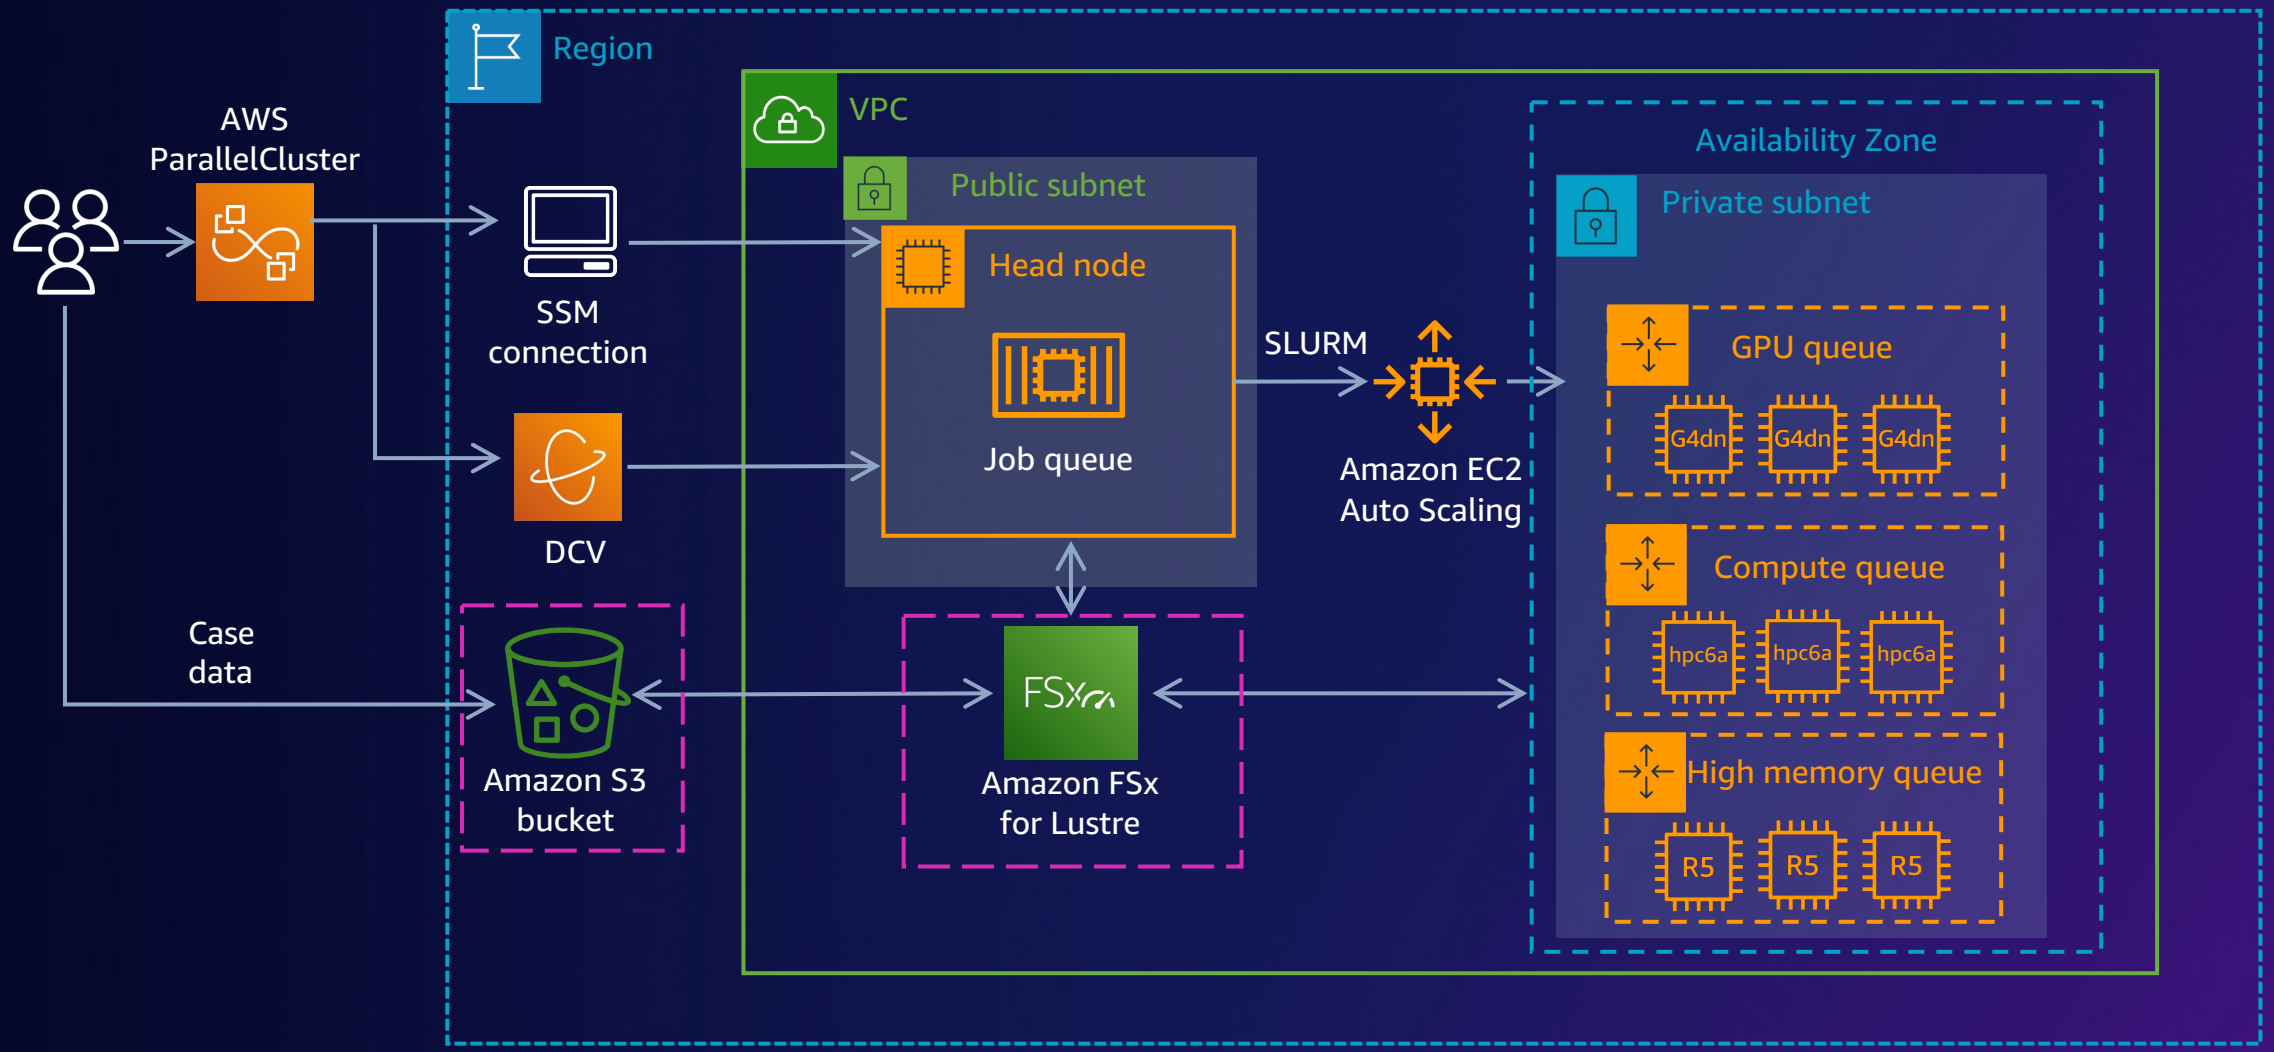

Storage options optimized for price performance

Storage type	Baseline throughput per TiB	Price per GB/s-hour
HDD persistent	12 MB/s	$2.85
	40 MB/s	$2.85
NEW SSD persistent	125 MB/s	$1.59
	250 MB/s	$1.15
	500 MB/s	$0.93
	1,000 MB/s	$0.82

Fast file interface for data on S3: How it works

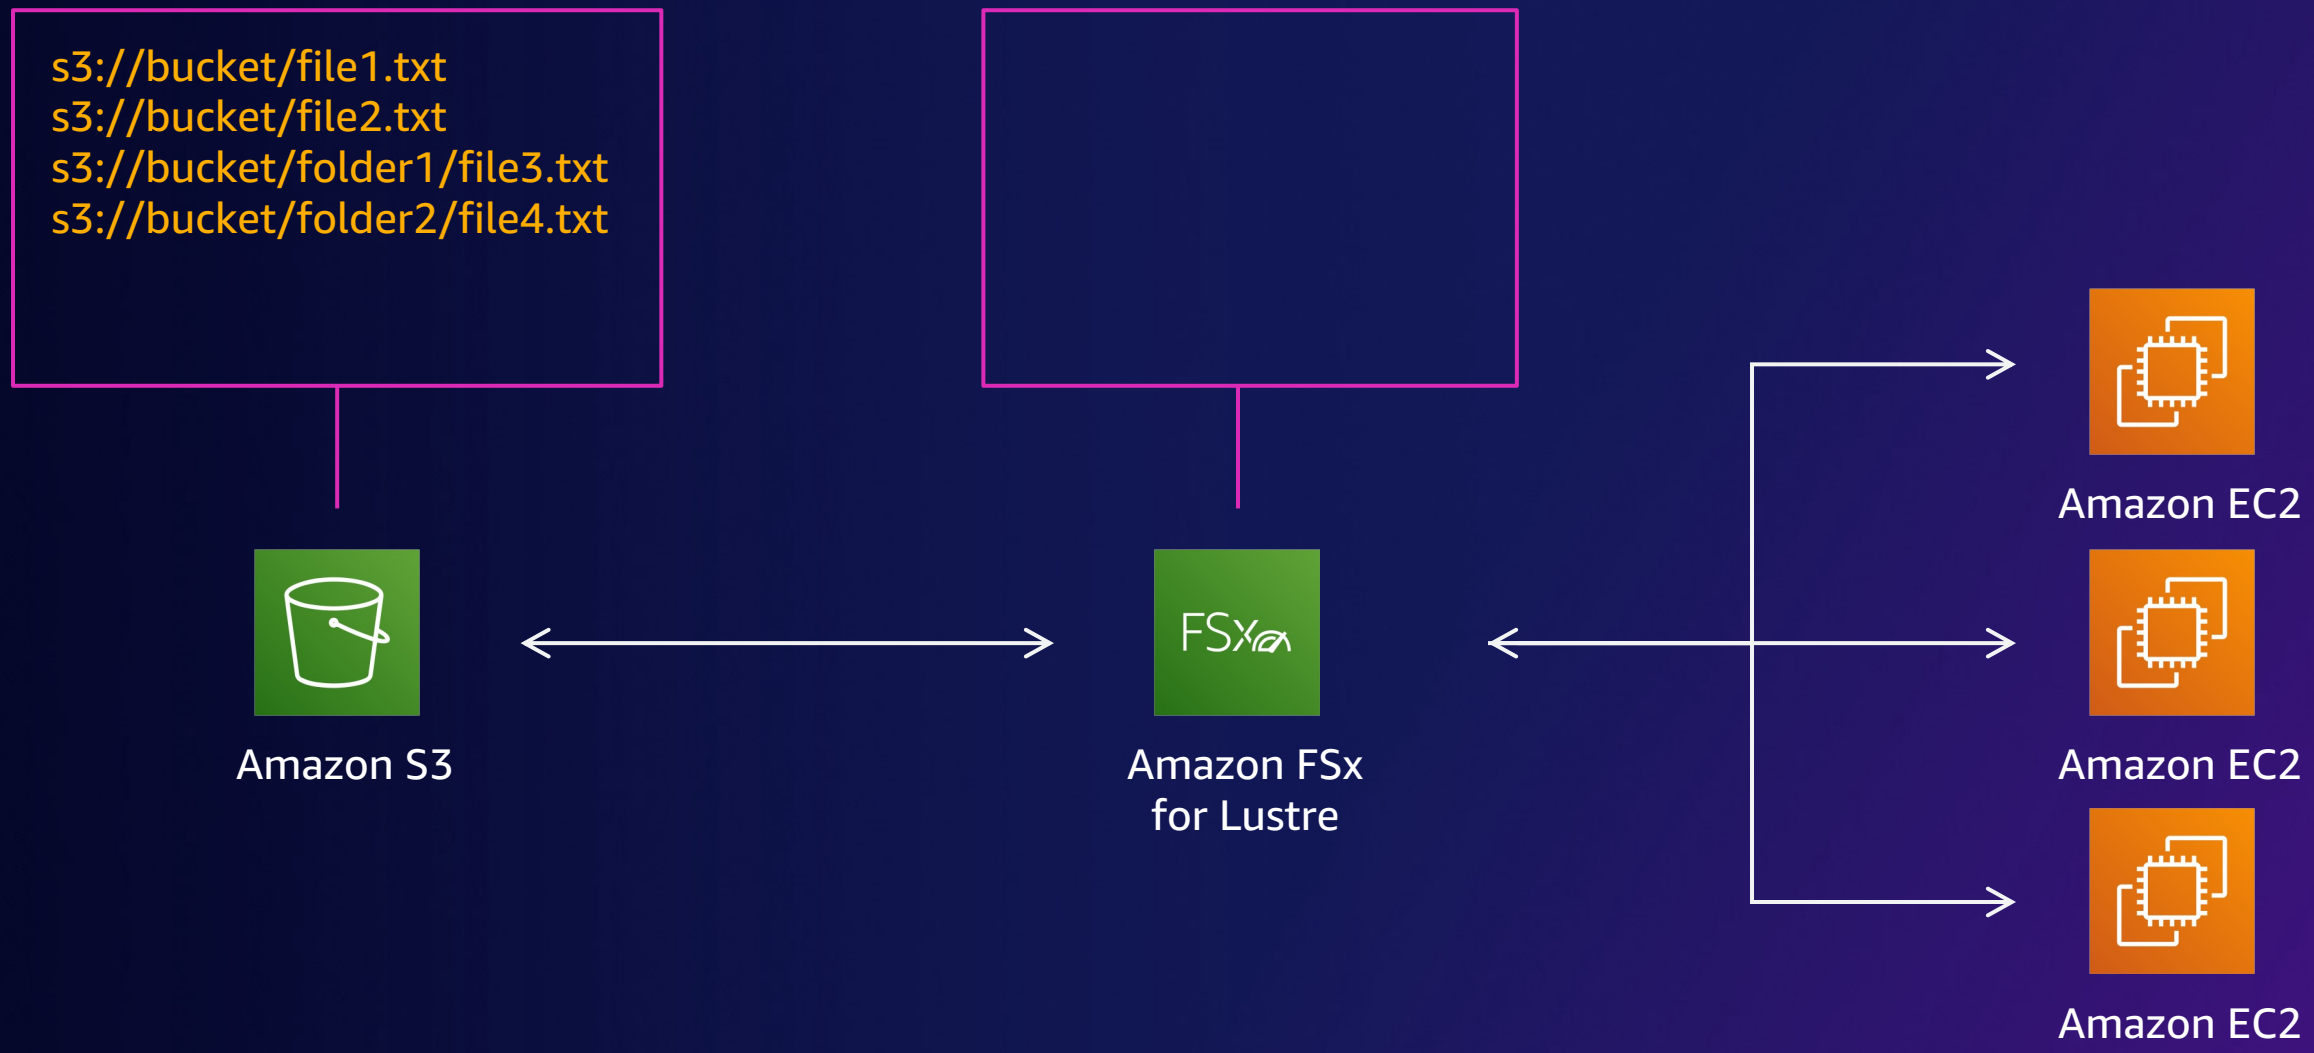

S3 objects appear on Amazon FSx file systems

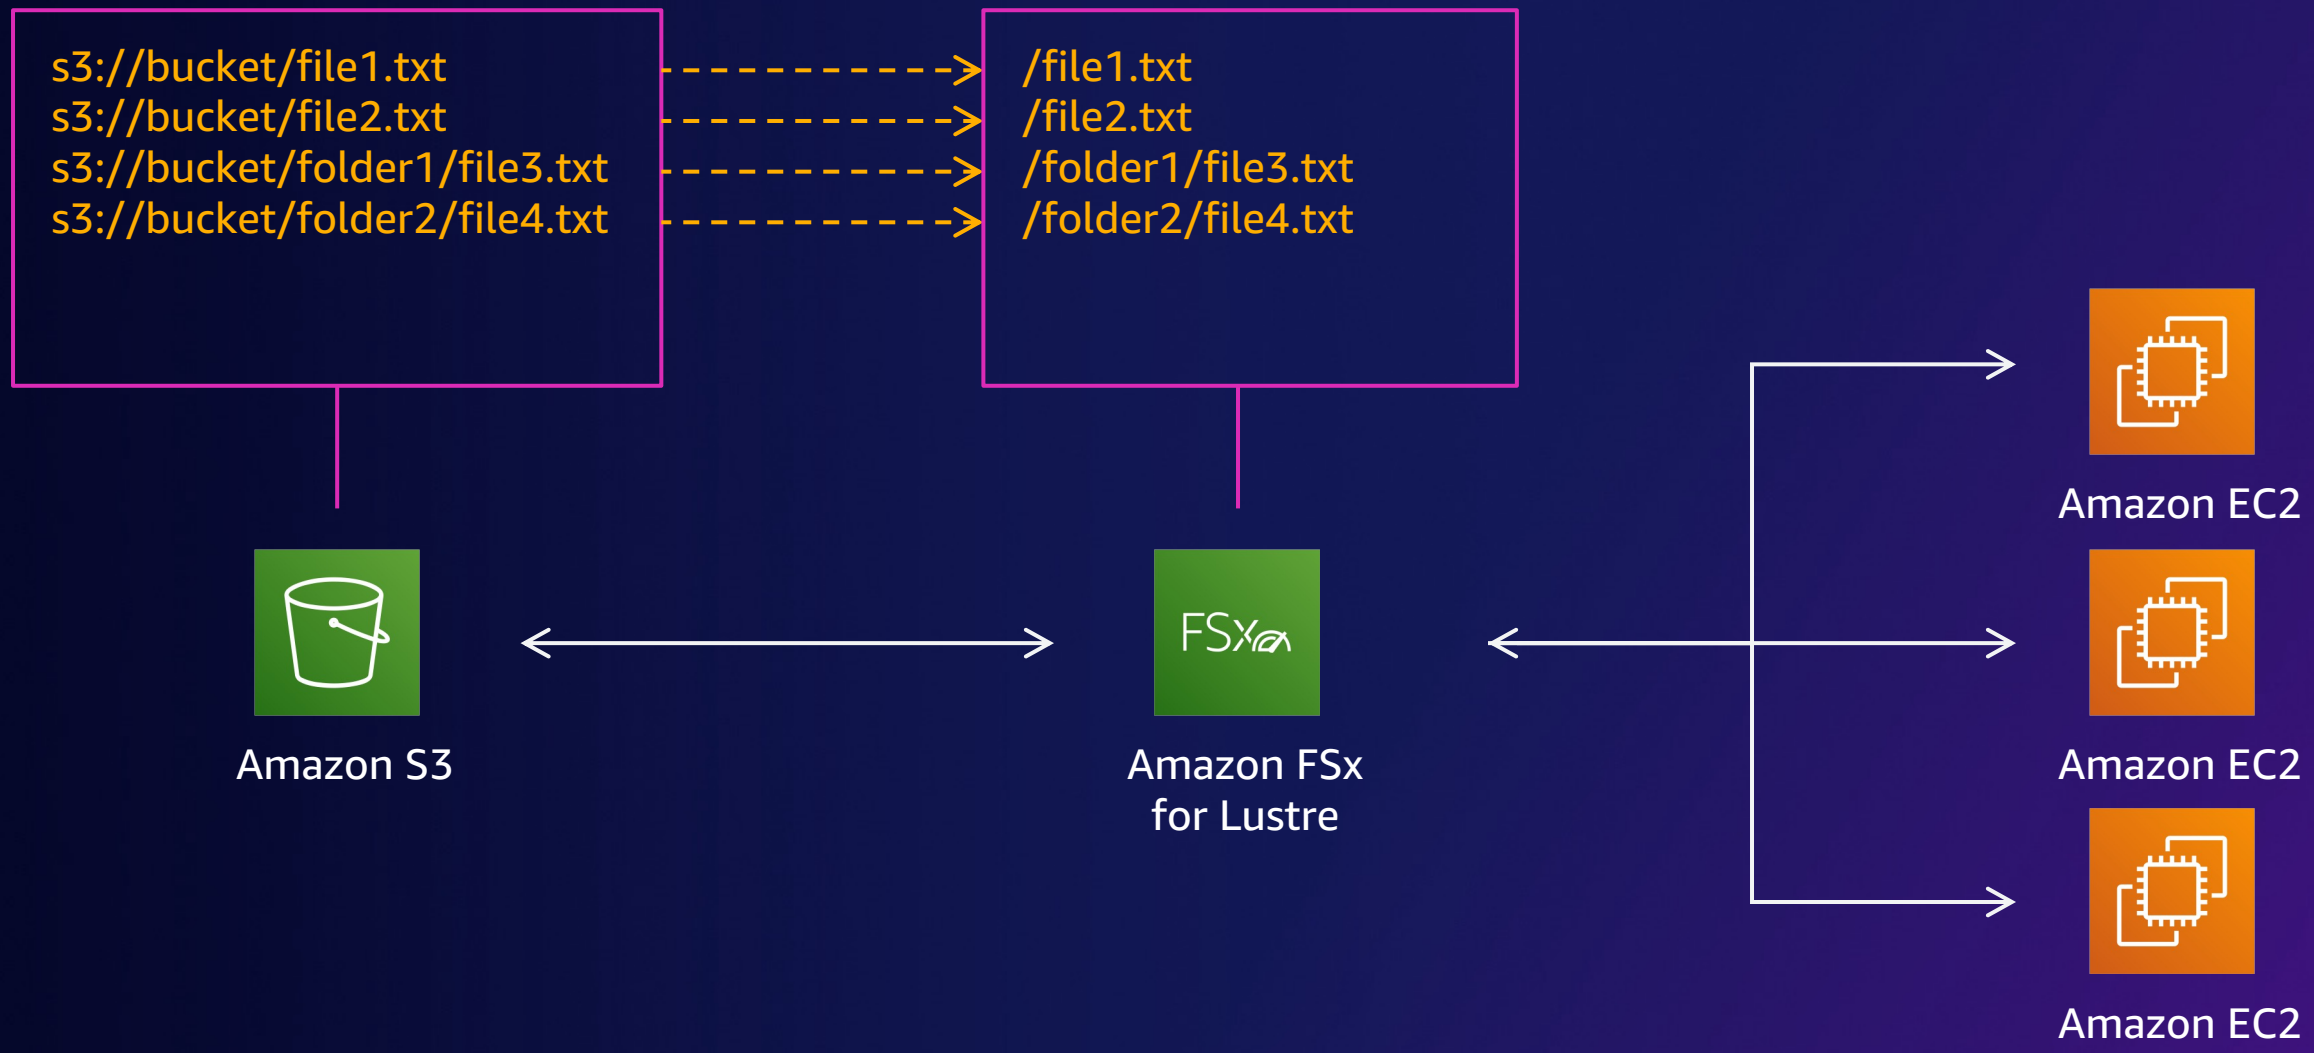

Updates on S3 are imported to FSx for Lustre

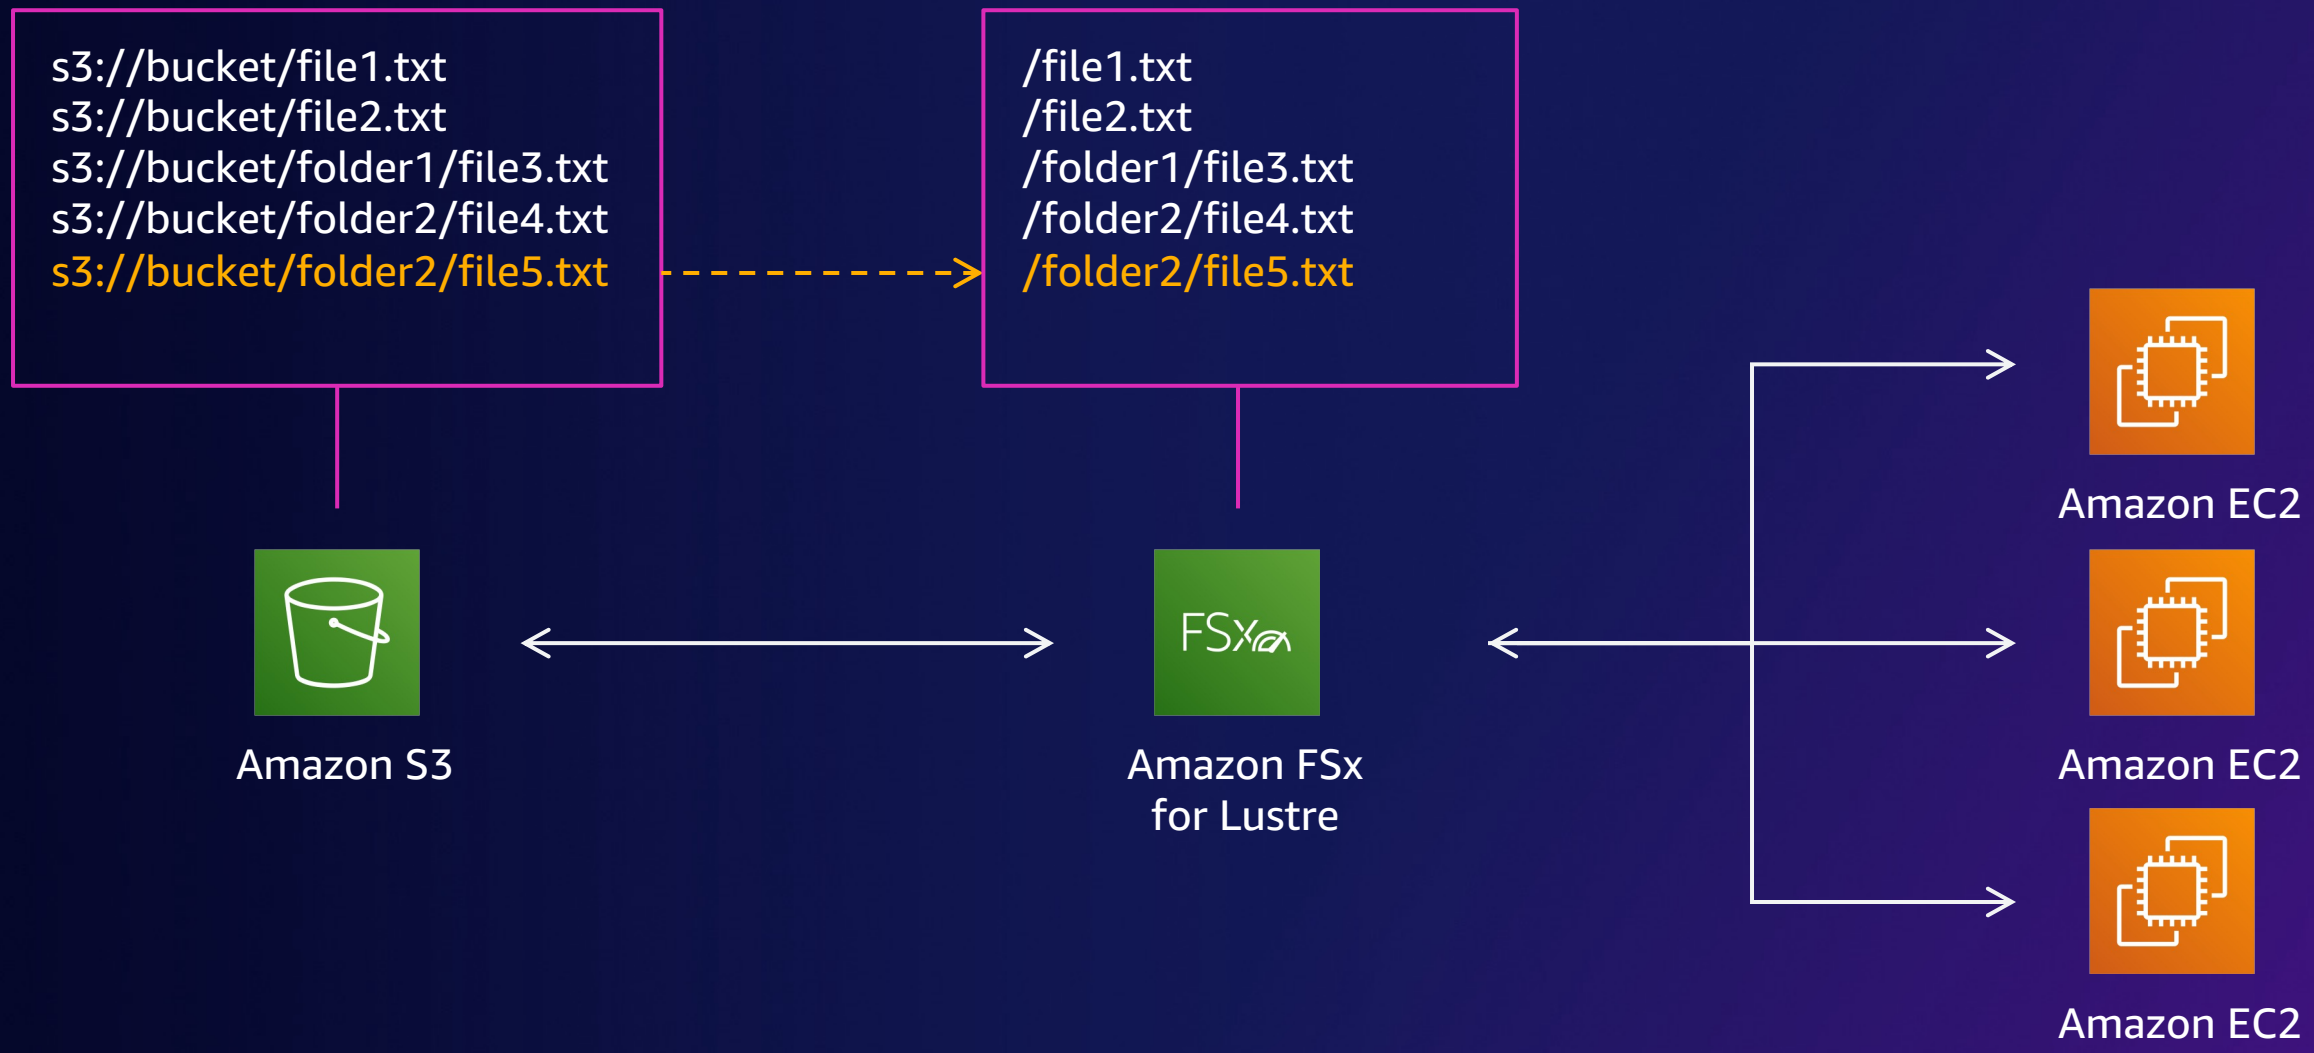

Updates on FSx for Lustre are exported to S3

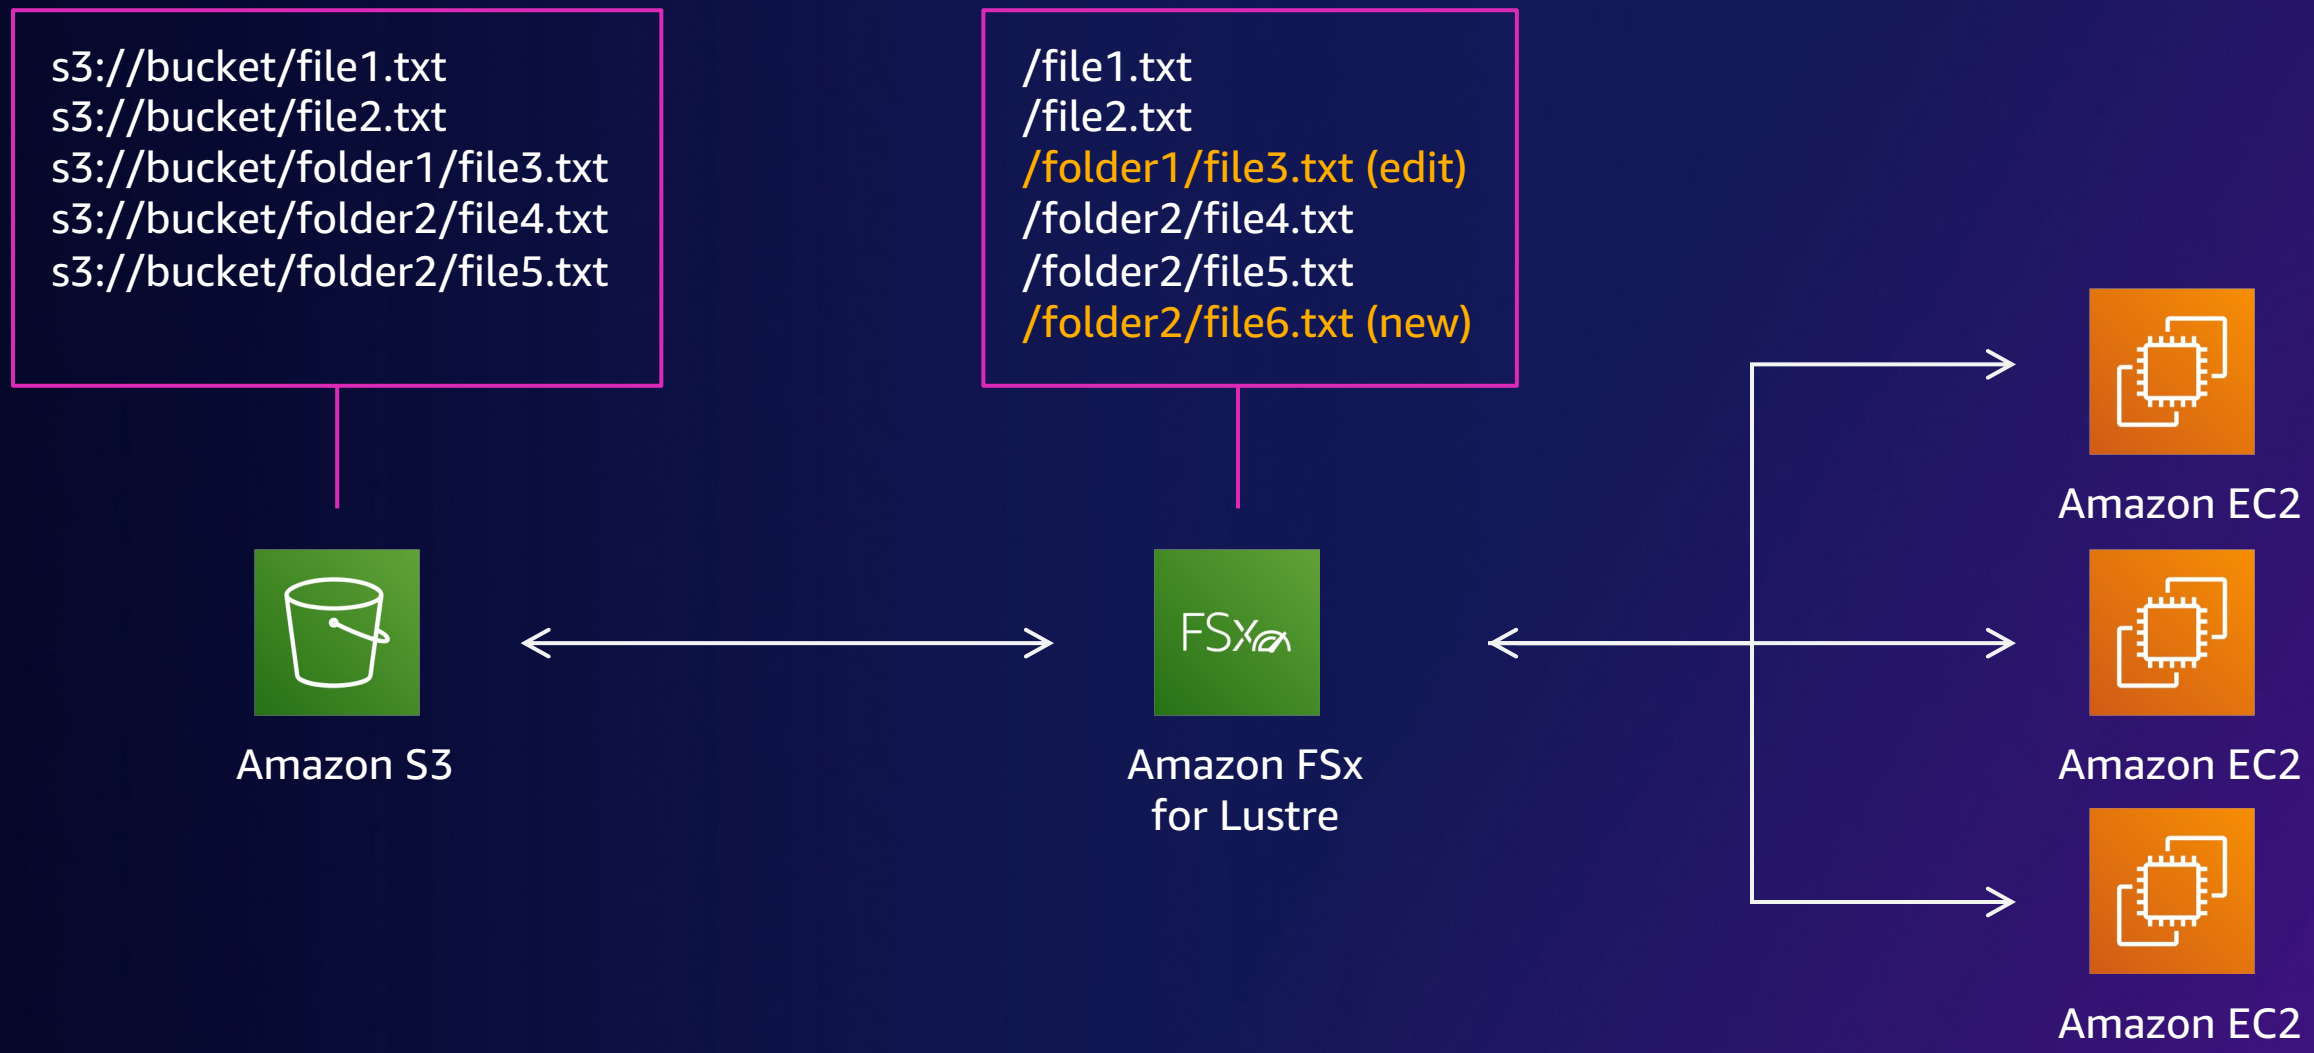

Updates on FSx for Lustre are exported to S3

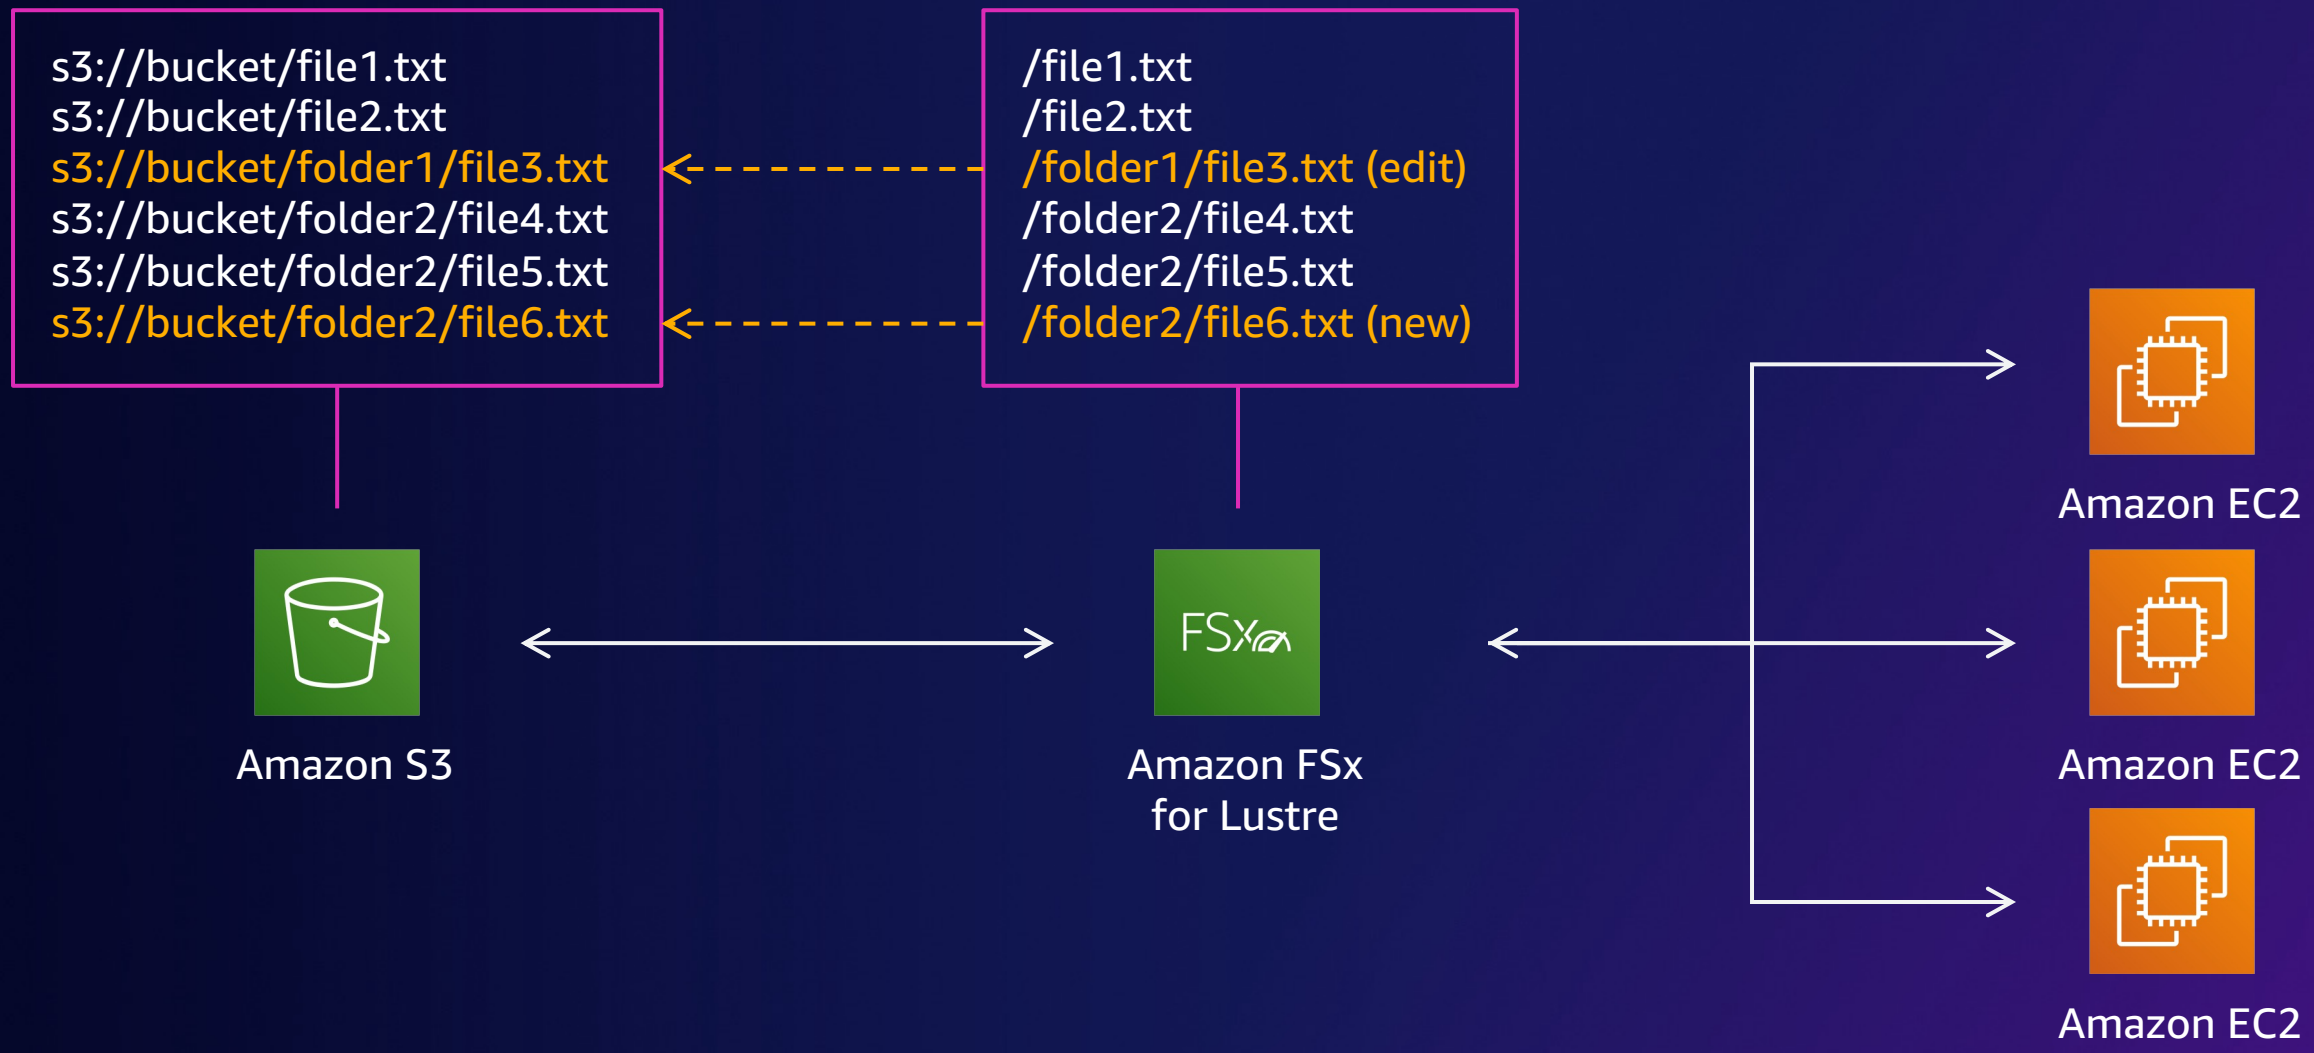

Spin up or down with compute resources

s3://bucket/file1.txt
s3://bucket/file2.txt
s3://bucket/folder1/file3.txt
s3://bucket/folder2/file4.txt
s3://bucket/folder2/file5.txt
s3://bucket/folder2/file6.txt

Amazon S3

/file1.txt
/file2.txt
/folder1/file3.txt (edit)
/folder2/file4.txt
/folder2/file5.txt
/folder2/file6.txt (new)

Spin down resources between workloads

Amazon FSx
for Lustre

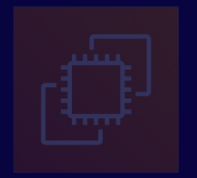

Amazon EC2

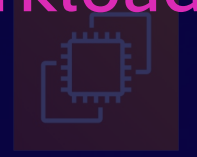

Amazon EC2

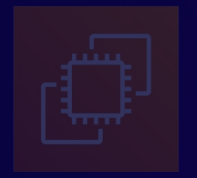

Amazon EC2

Reference architecture

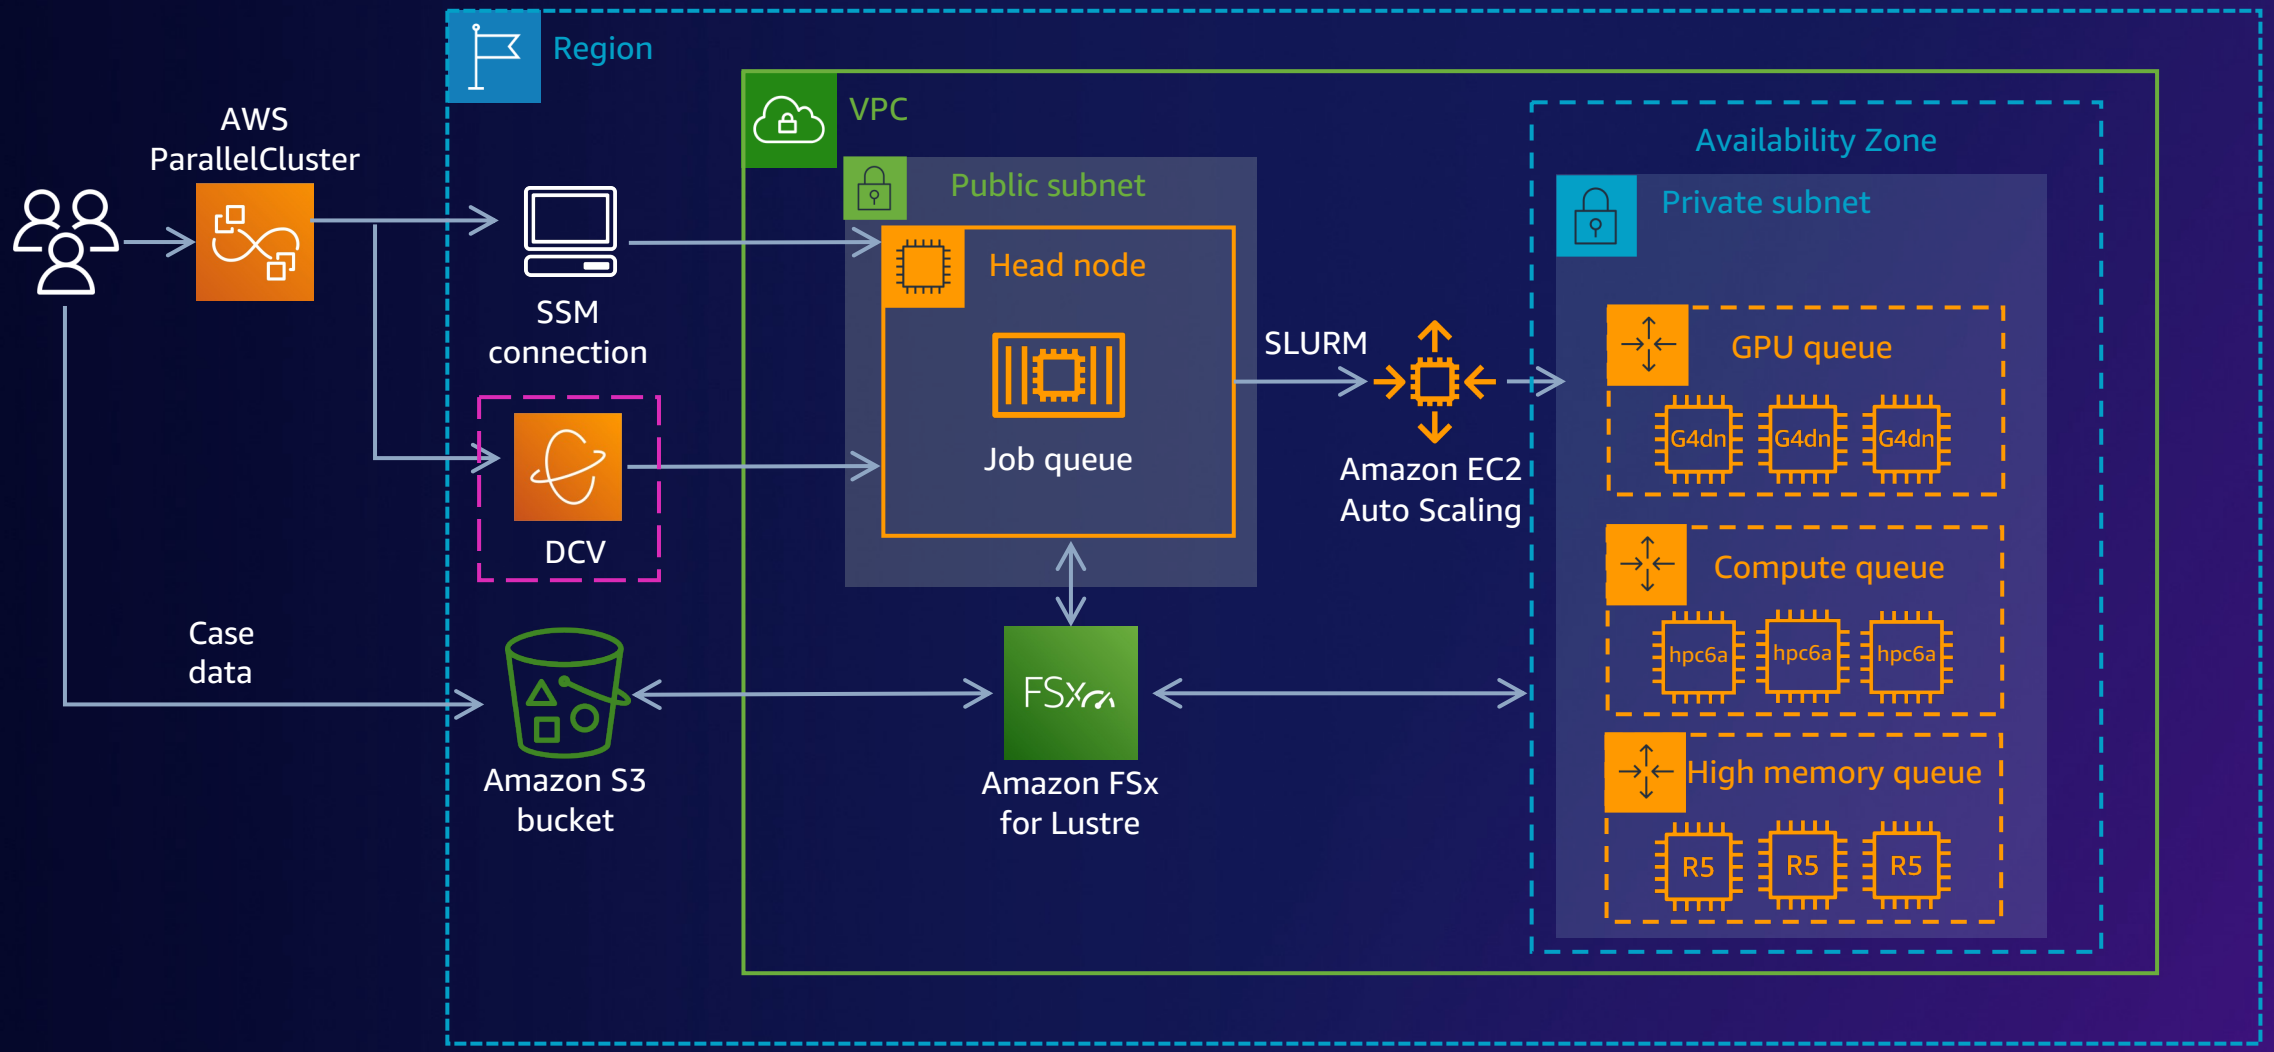

NICE DCV encrypts and transports pixels to devices

Local clients

Transfer pixels, not data

Remote servers

Users can access, manipulate, and share business-critical information, regardless of their location, over LAN or WAN networks

DCV activates pre- & post-processing in the cloud

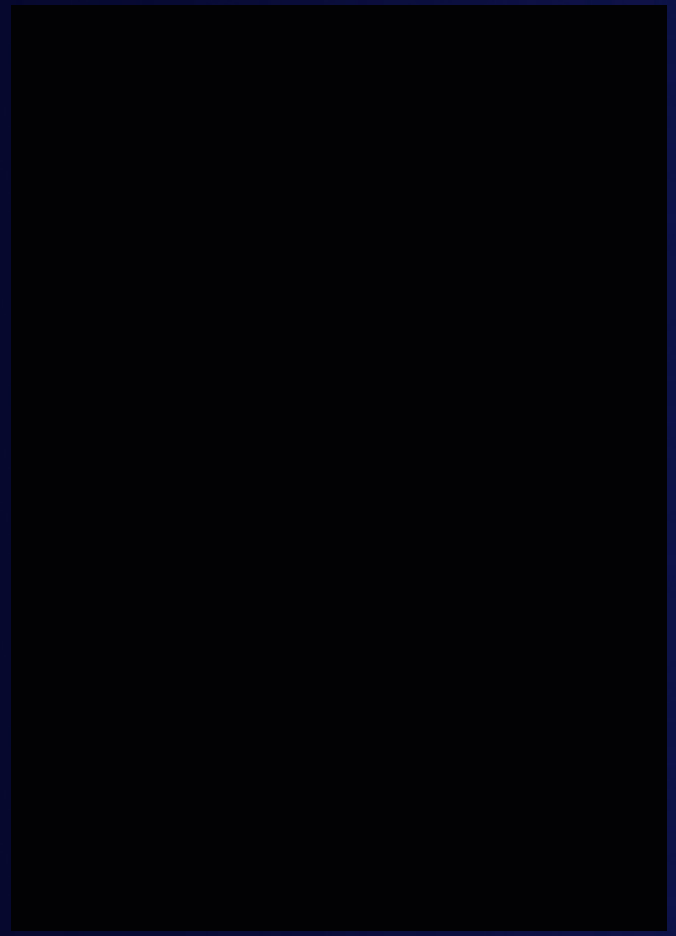

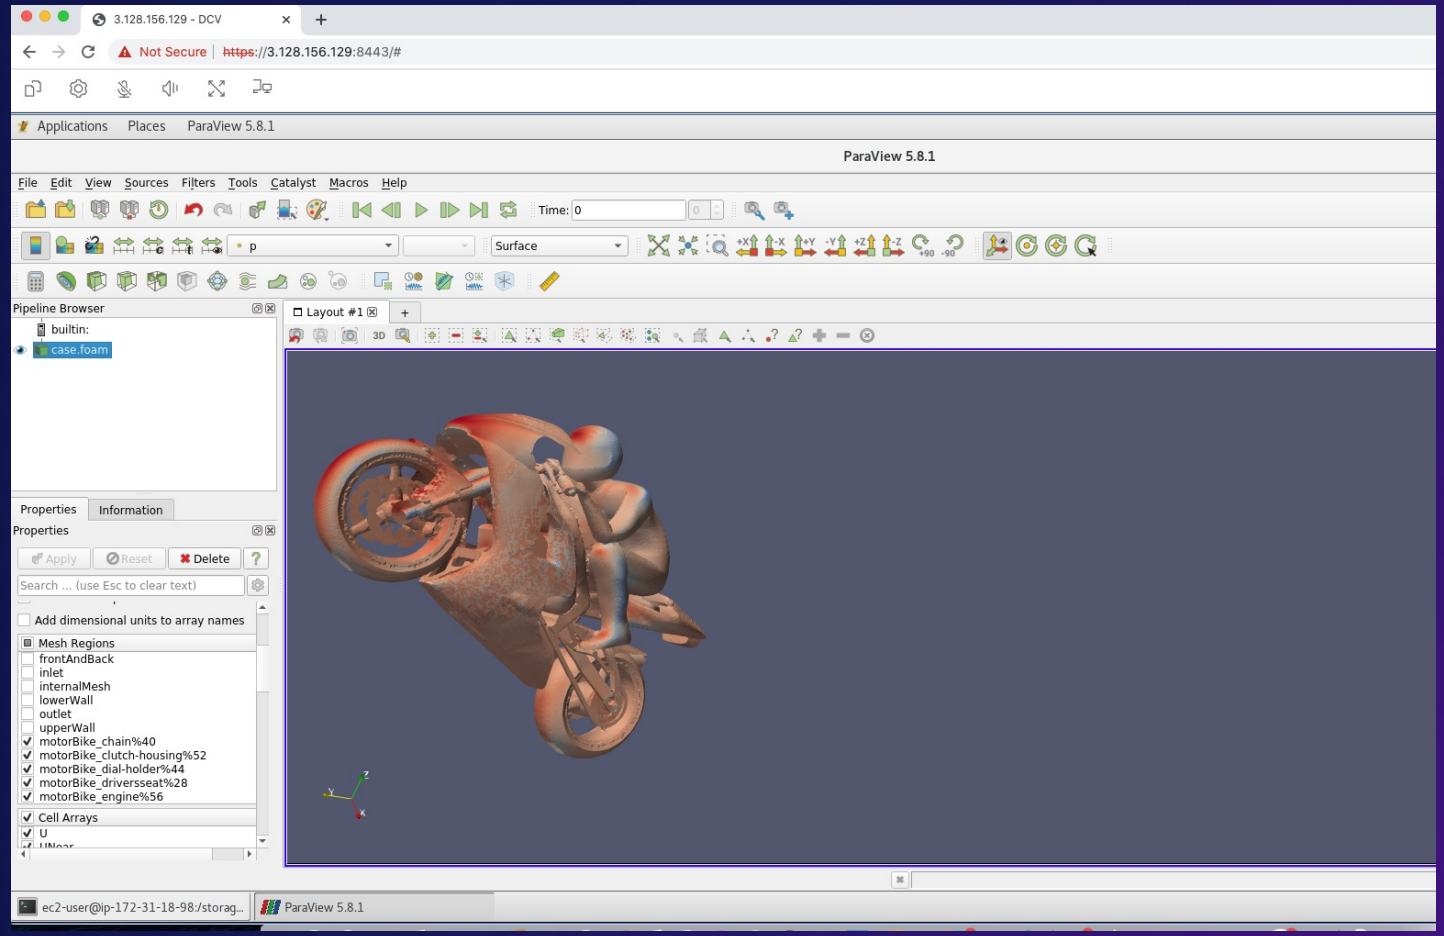

Instances

- Many users (especially commercial) want the best price/performance
- hpc6a.48xlarge is the clear winner
- c6i.32xlarge is the second best option among the x86-based instances
- c6g.16xlarge is a good second option for codes which support Arm

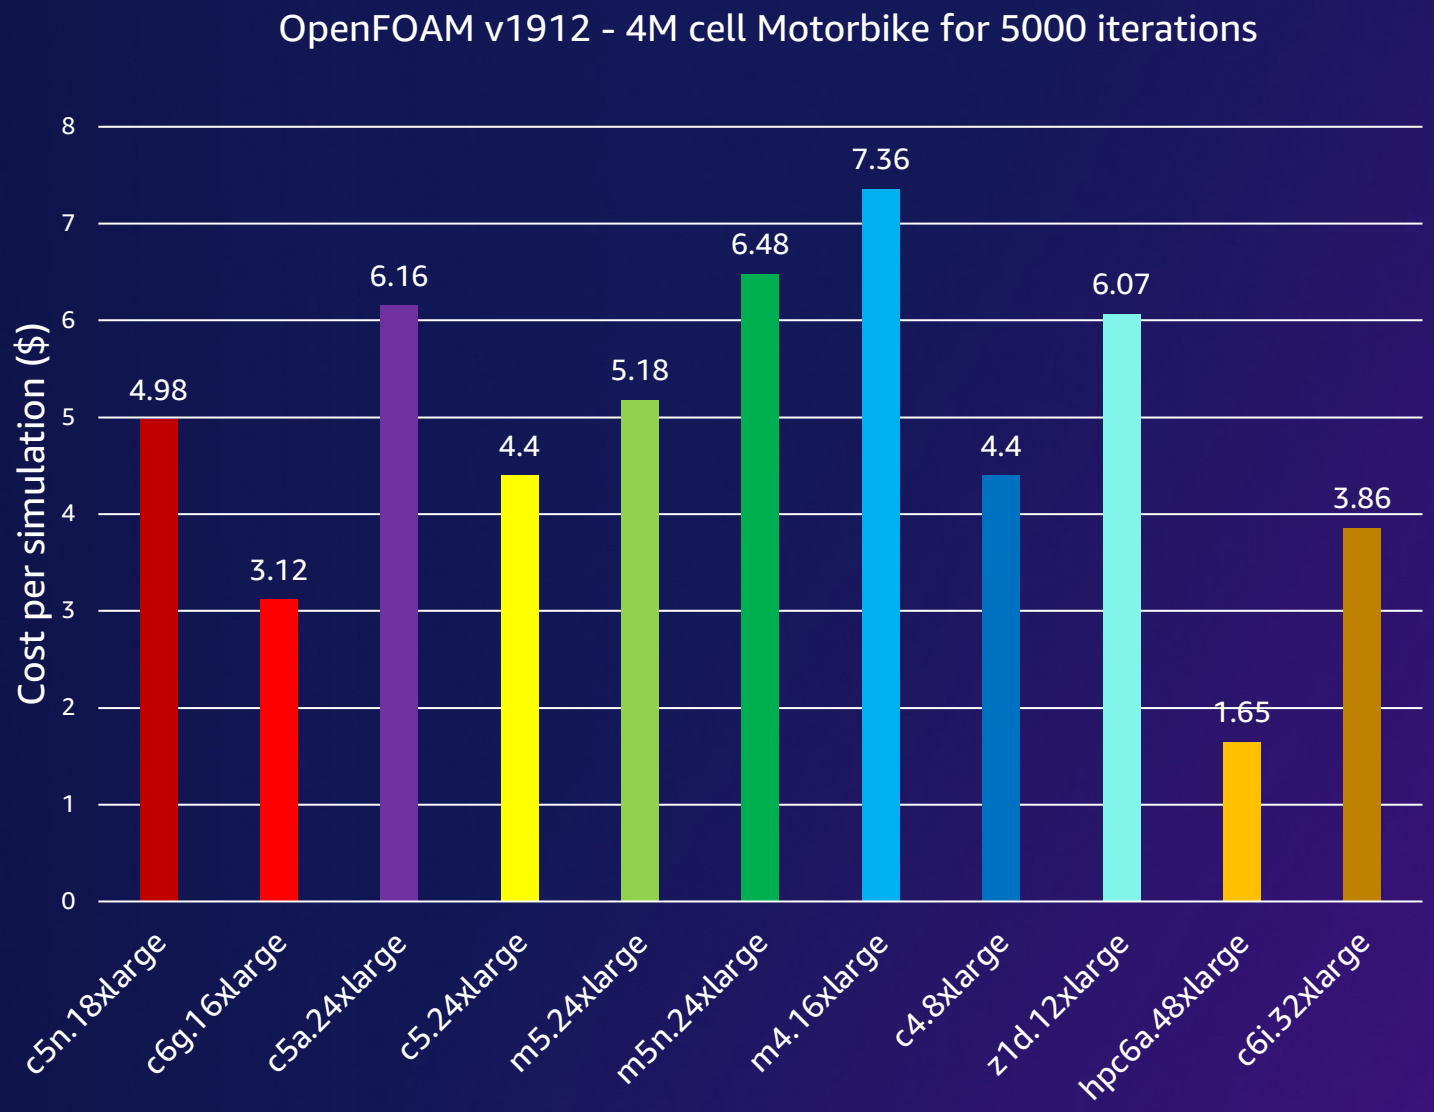

Time to solution vs. cost per simulation

- Optimize compute instance selection and number of instances
- Choose the optimal cells/core for your model
- Optimize storage by using FSx for Lustre as scratch space and S3 for long-term storage
- Next, look at pricing models such as Spot or Reserved Instances

Evaluate the trade-off of time to solution vs. cost for scaling

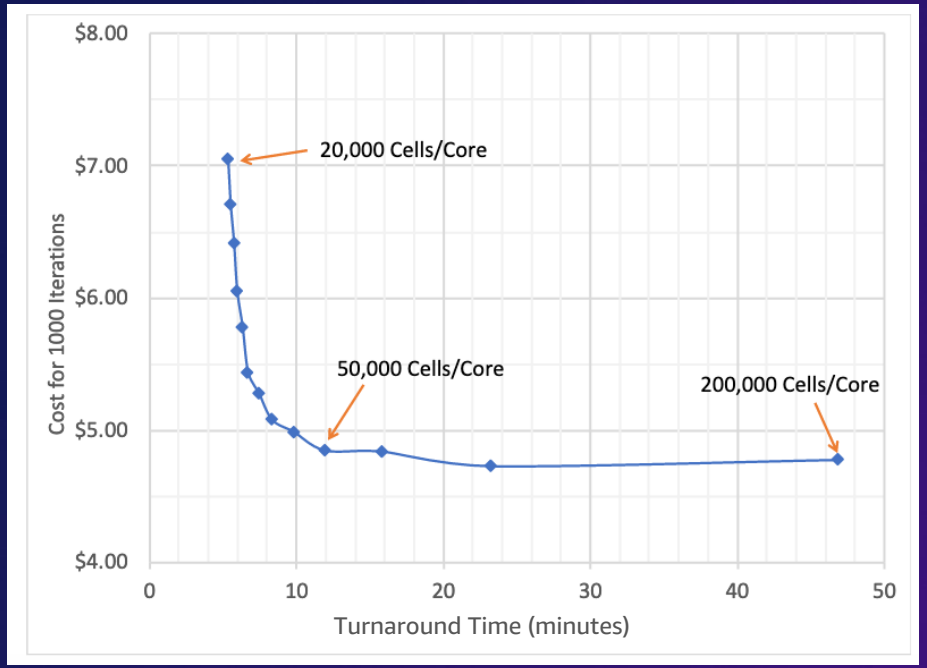

Innovating with ML

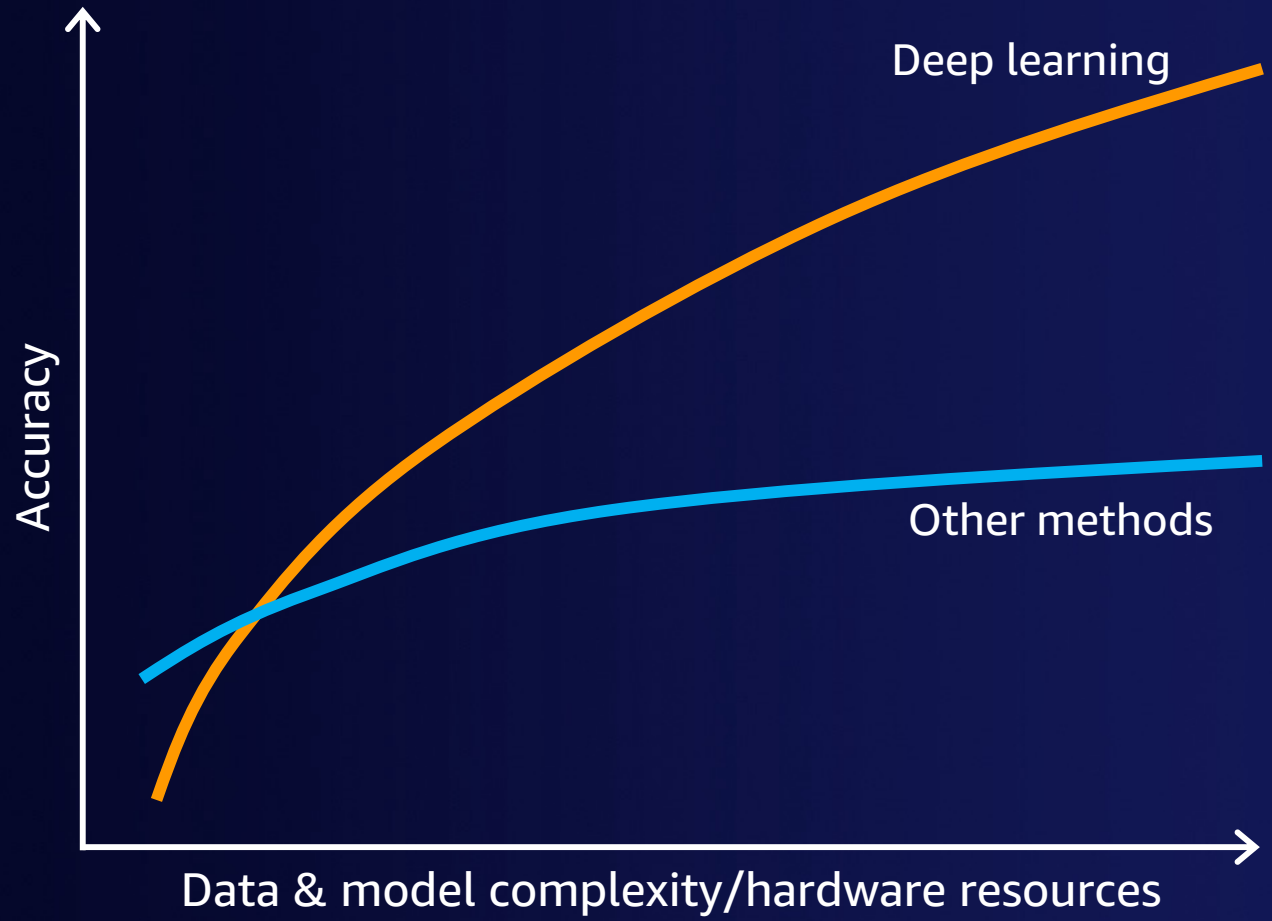

Applications

Natural language processing

Image/ video analysis

Autonomous vehicle systems

Recommendation systems

Model accuracy continues to increase with additional data

Demonstrating ML and HPC: Extract relationships from HPC simulation results

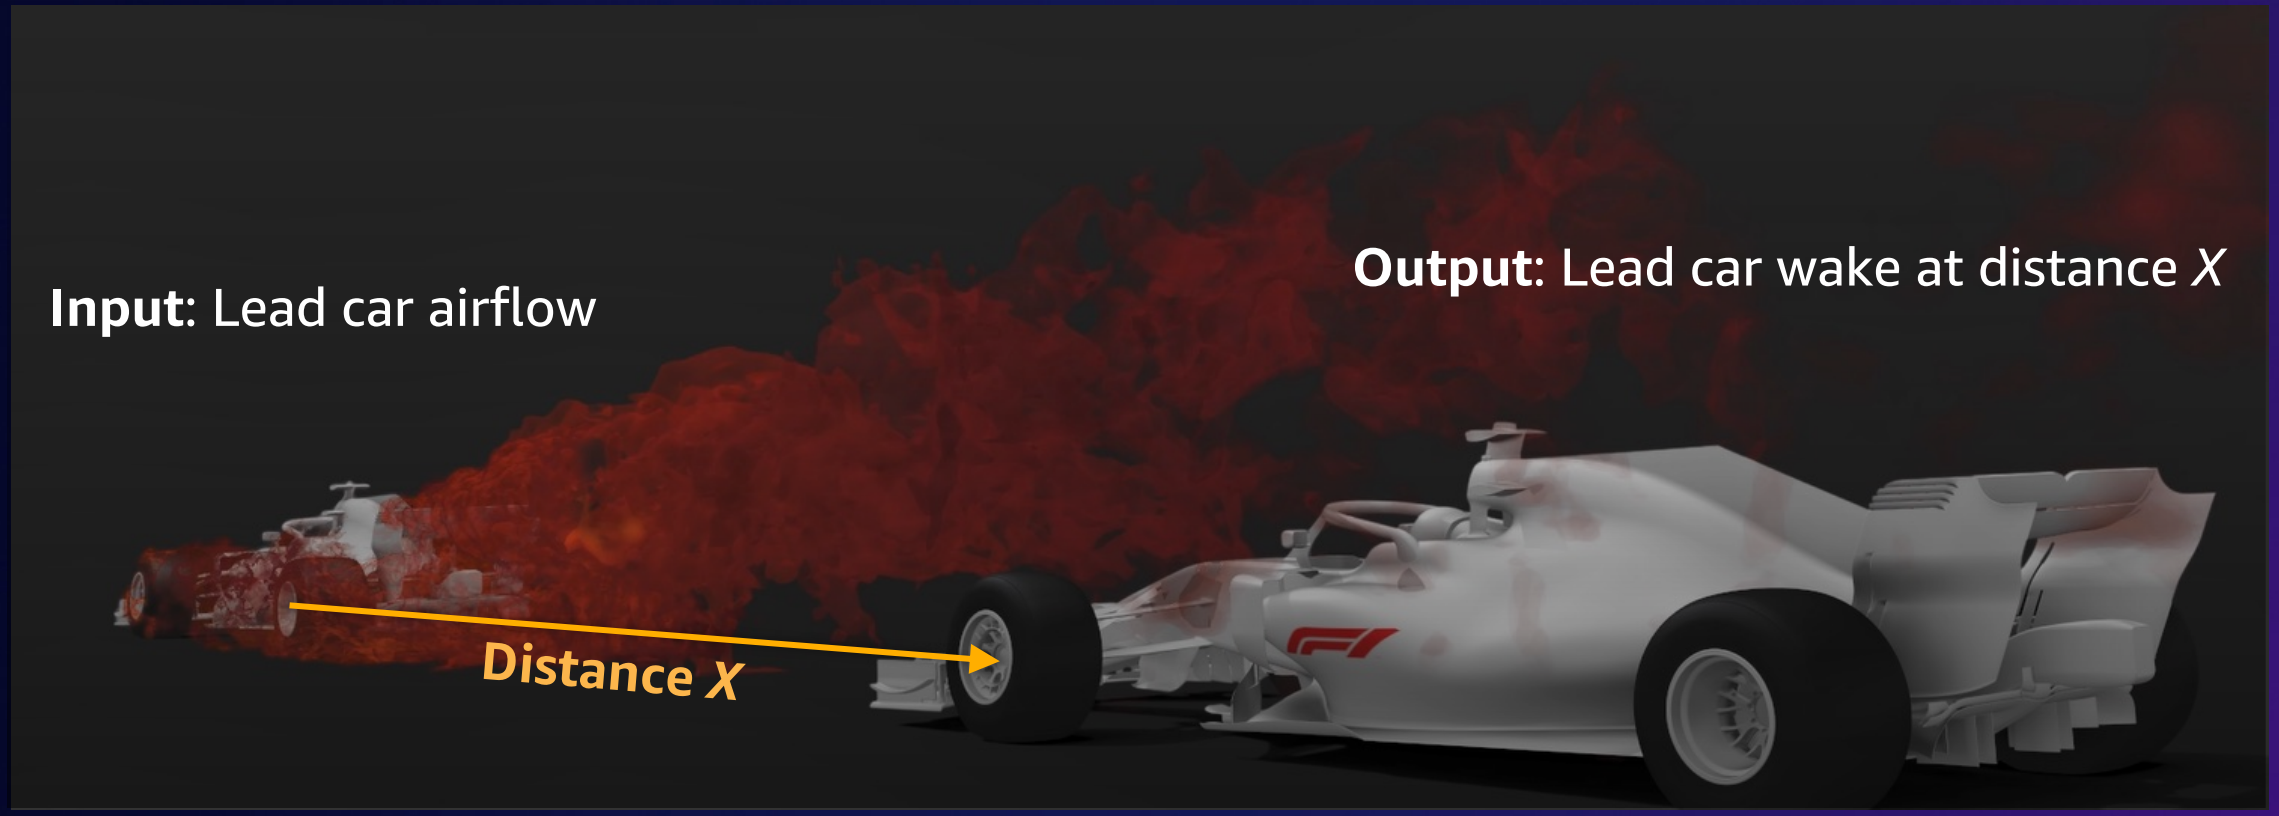

Problem statement:

Determine optimal lead car features to obtain target flow profile at distance *X*

ML and HPC solution

3D geometry

Input flow field

Deep learning model

Output flow field

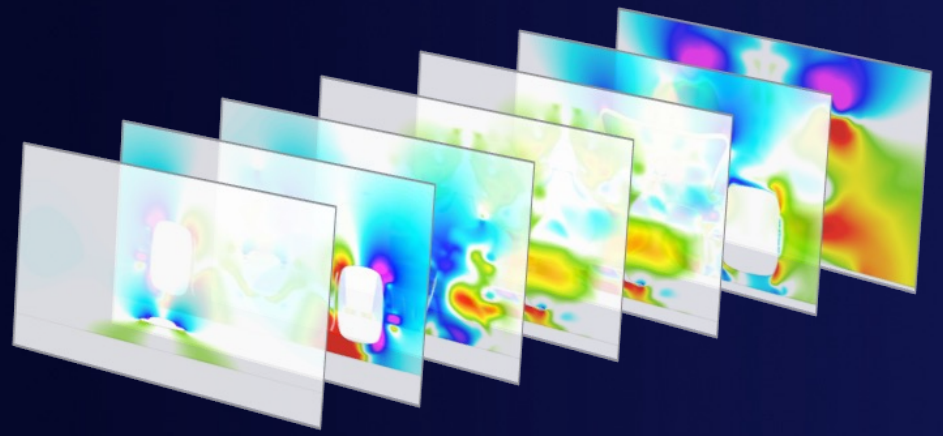

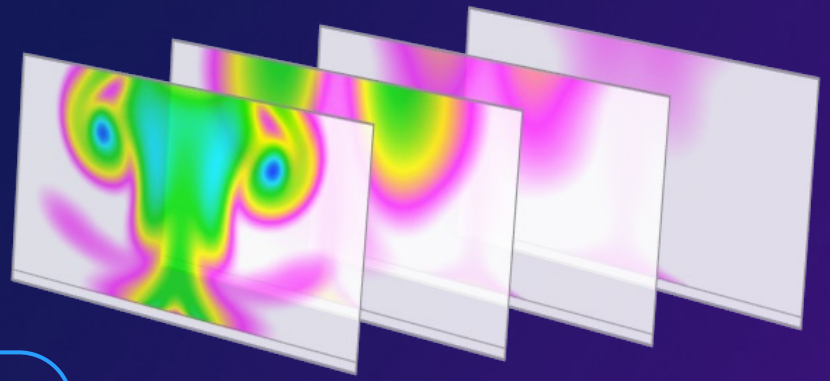

EC2 UltraClusters

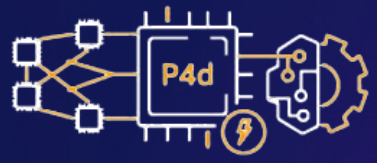

Key resources

CFD on AWS
aws.amazon.com/hpc/cfd

AWS CFD workshops
cfd-on-pcluster.workshop.aws

CFD on AWS whitepaper
d1.awsstatic.com/whitepapers/computational-fluid-dynamics-on-aws.pdf

AWS HPC Blog
https://aws.amazon.com/blogs/hpc/

Content is being actively developed – please check back frequently

Scaling on AWS – OpenFOAM

OpenFOAM v1912, motorbike (**222M cells**), Intel MPI 2019.6, AL2, PC2.6.1

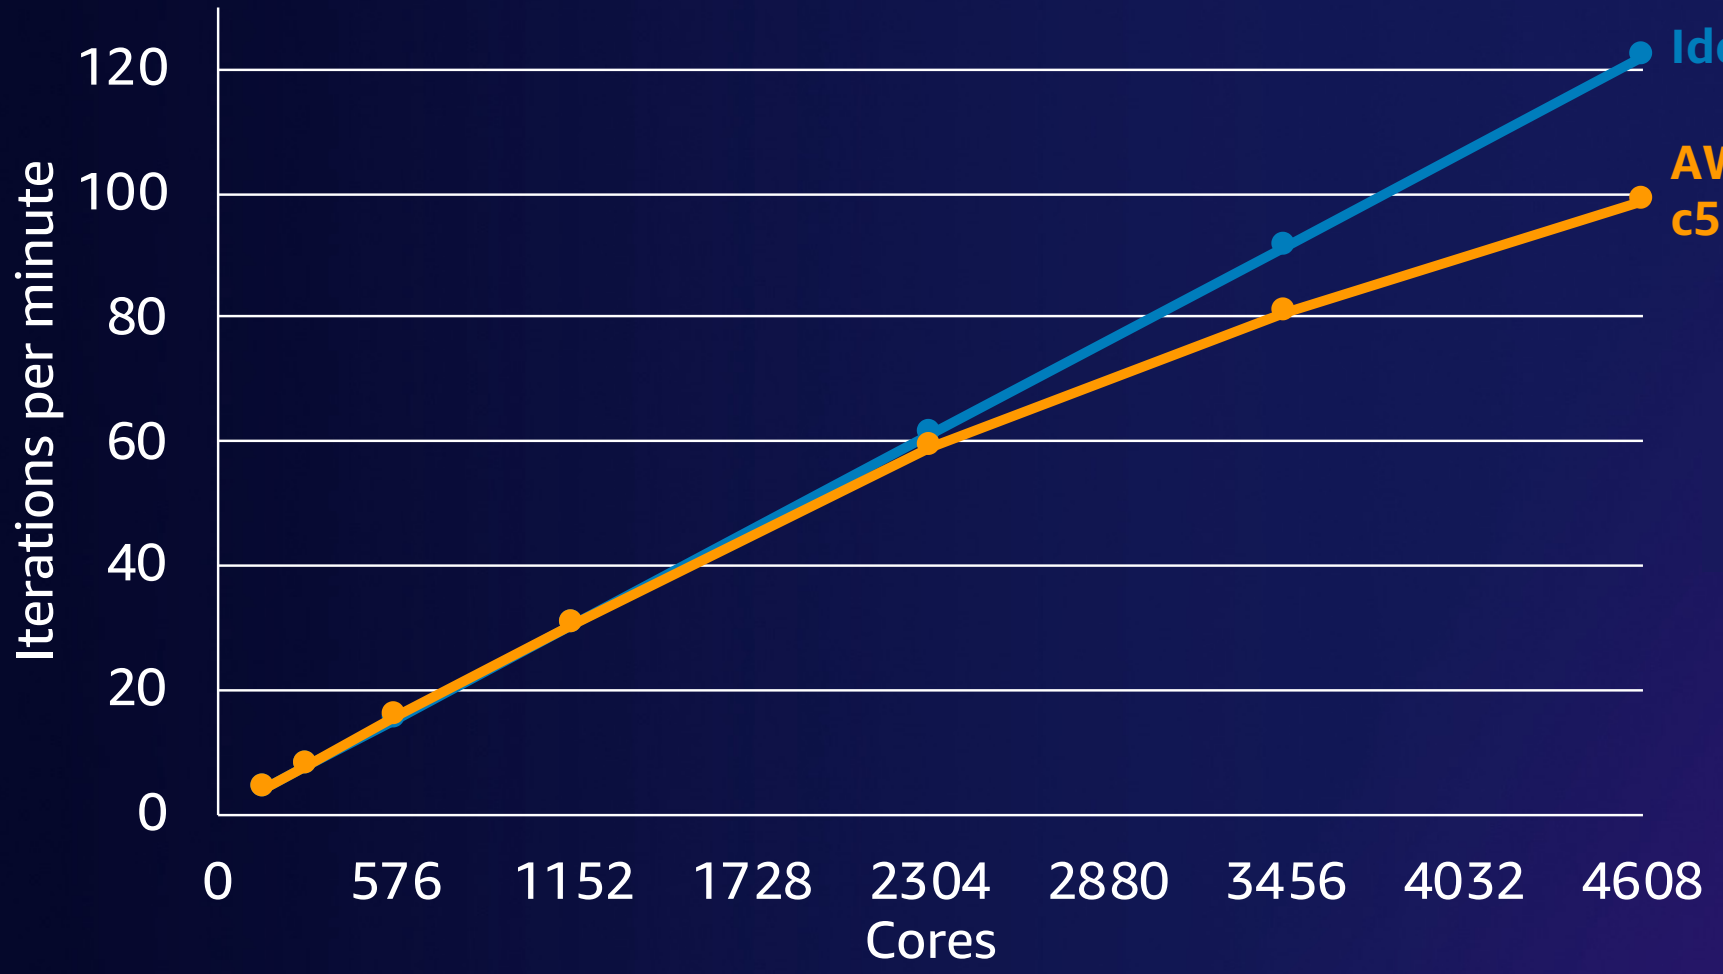

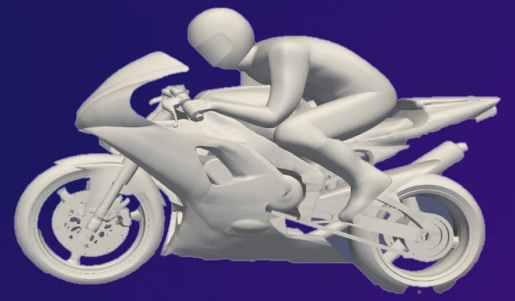

Scaling on AWS – STAR-CCM+

Simcenter STAR-CCM+ 2020.1, F1 (**403M cells**), Intel MPI 2019.6, AL2, PC2.6.1

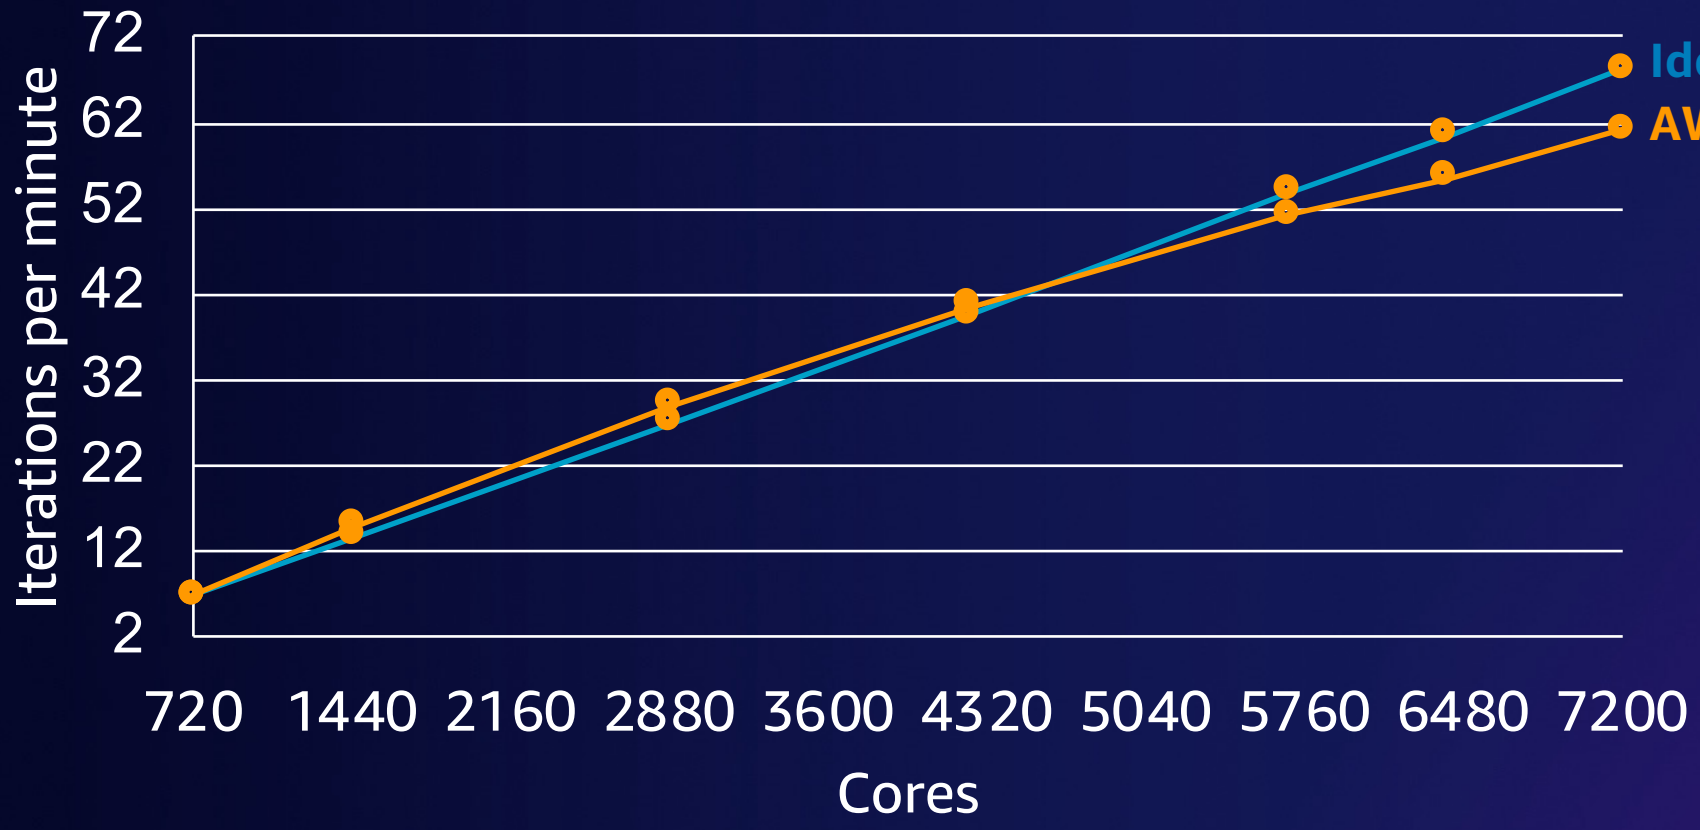

Scaling on AWS – Fluent

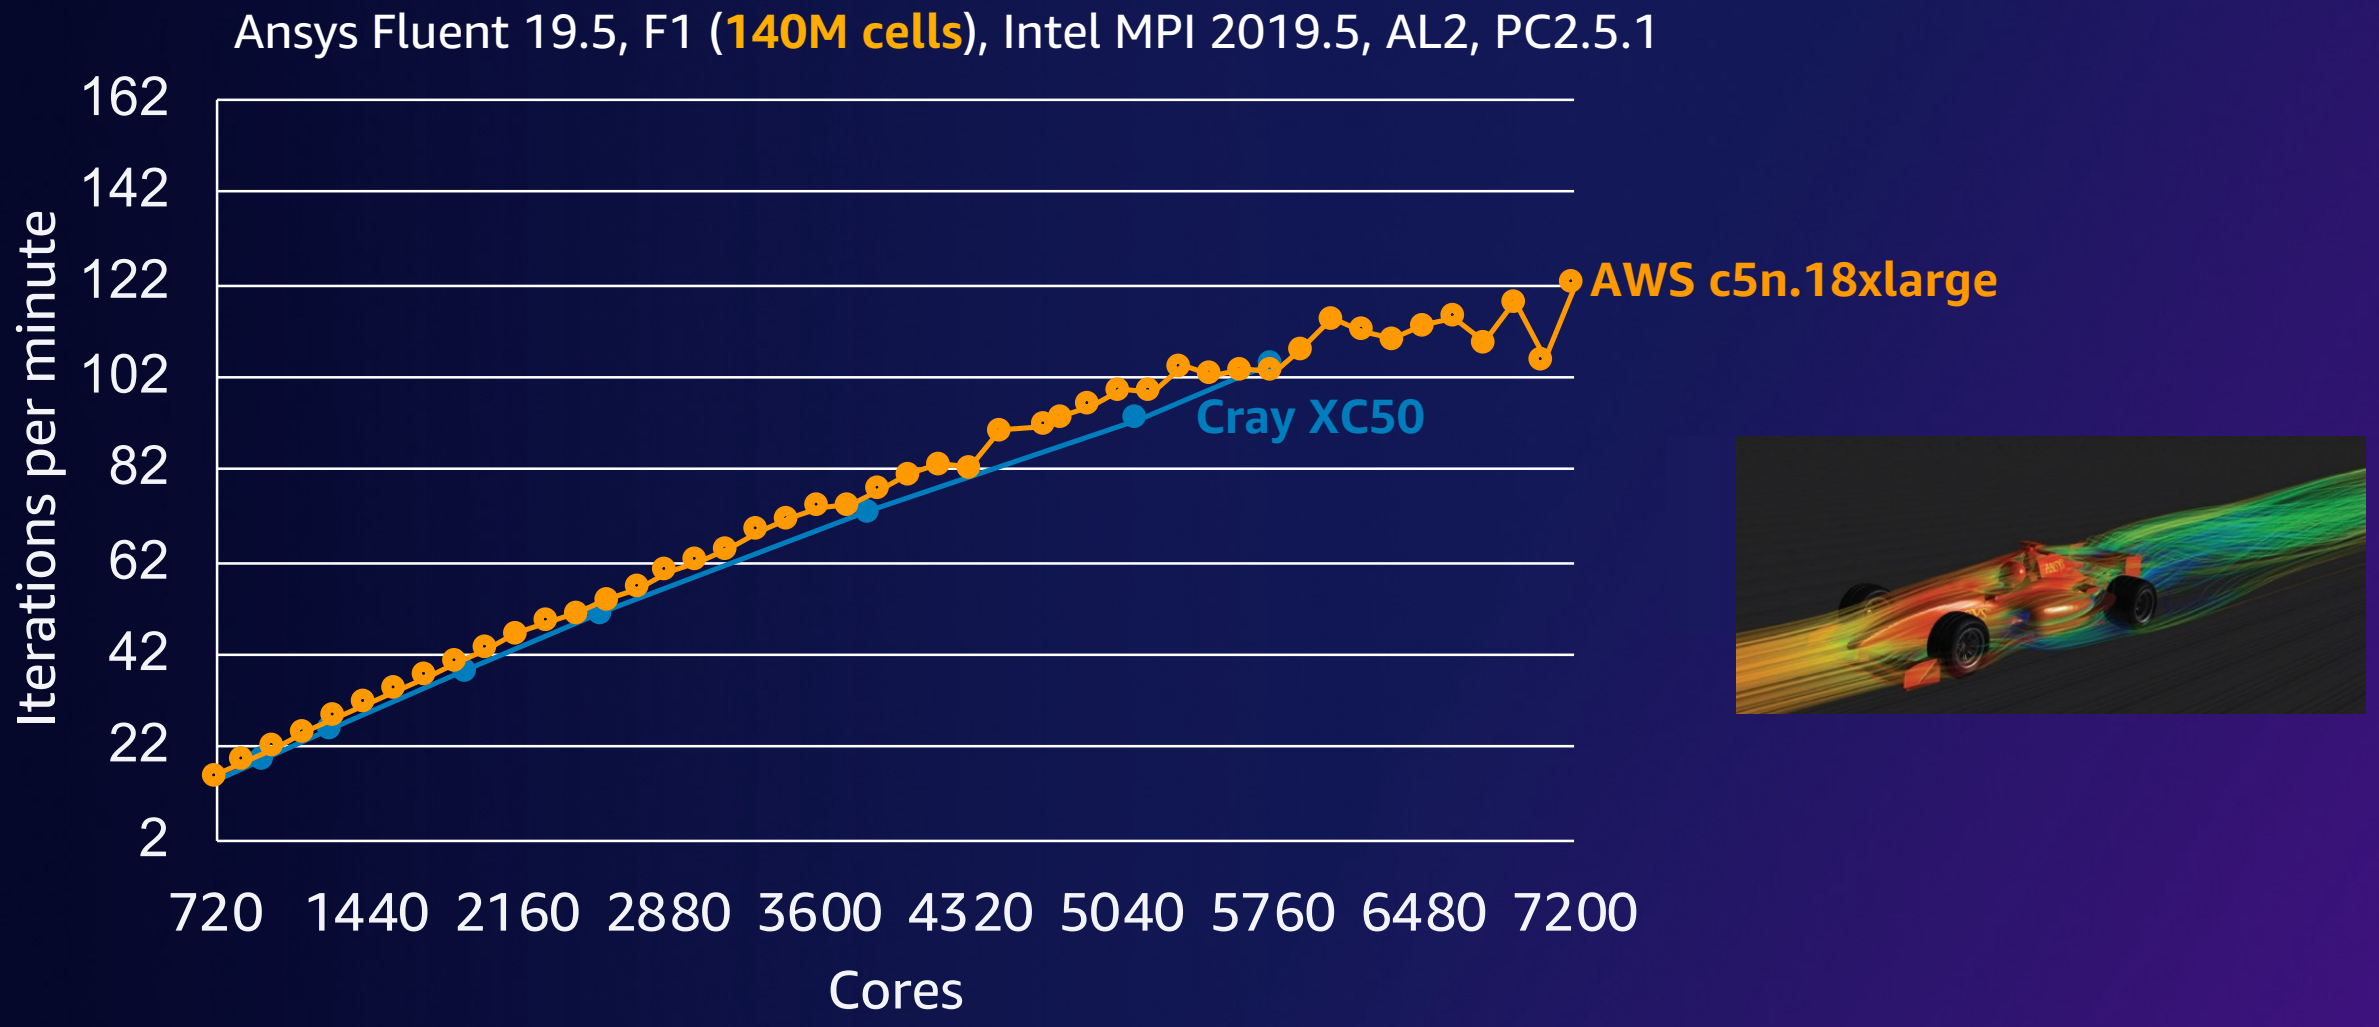

Learn in-demand AWS Cloud skills

AWS Skill Builder

Access **500+ free** digital courses and Learning Plans

Explore resources with a variety of skill levels and **16+** languages to meet your learning needs

Deepen your skills with digital learning on demand

Train now

AWS Certifications

Earn an industry-recognized credential

Receive Foundational, Associate, Professional, and Specialty certifications

Join the **AWS Certified community** and get exclusive benefits

Access **new** exam guides

Thank you!

Sean Smith
seaam@amazon.com

Karthik Raman
akraman@amazon.com

Please complete the session survey