

The background features a dark blue gradient with abstract geometric shapes. On the left, a large triangle is formed by a vertical orange line and a diagonal orange line, with a blue-to-orange gradient fill. On the right, a curved orange-to-blue gradient shape sweeps across the frame. A thin blue line forms a large rectangle in the lower half of the image.

# AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

CMP405

# The future of computational fluid dynamics in the cloud

Sean Smith

HPC Solution Architect  
AWS

Karthik Raman

Senior HPC Solution 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
- **60 hours** to **10 hours** thanks to AWS



# Agility = the right tool for the job

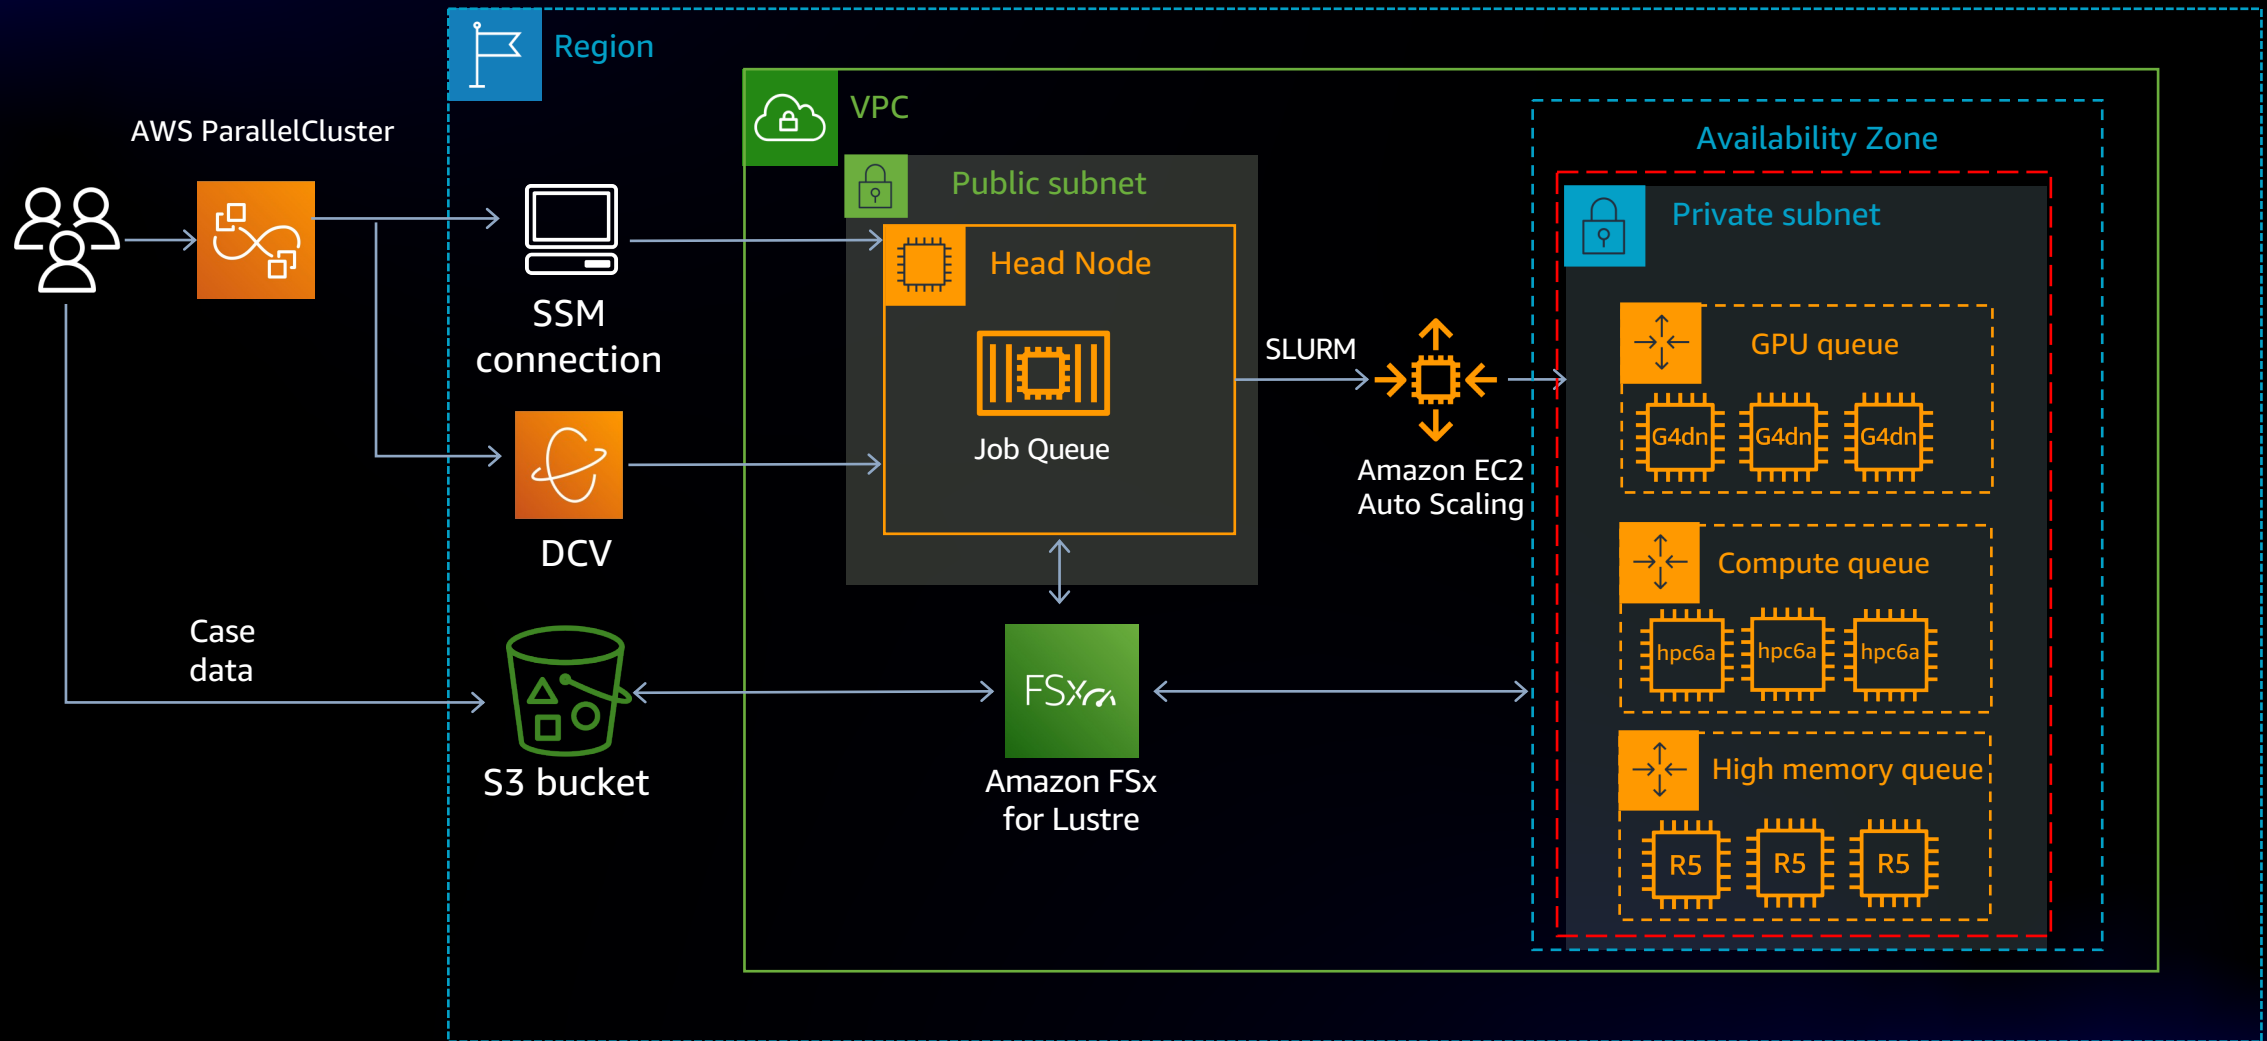


Company	Simulator(s)	Hardware
Ansys	Fluent	CPU, GPU
Dassault	Abaqus	CPU, GPU
Siemens	Star-CCM+	CPU
LSTC/Ansys	LS-DYNA	CPU, GPU
OpenFOAM Foundation	OpenFOAM	CPU, GPU
Comsol	Multiphysics	CPU
Altair	AcuSolve	CPU

Cloud optimized recipe

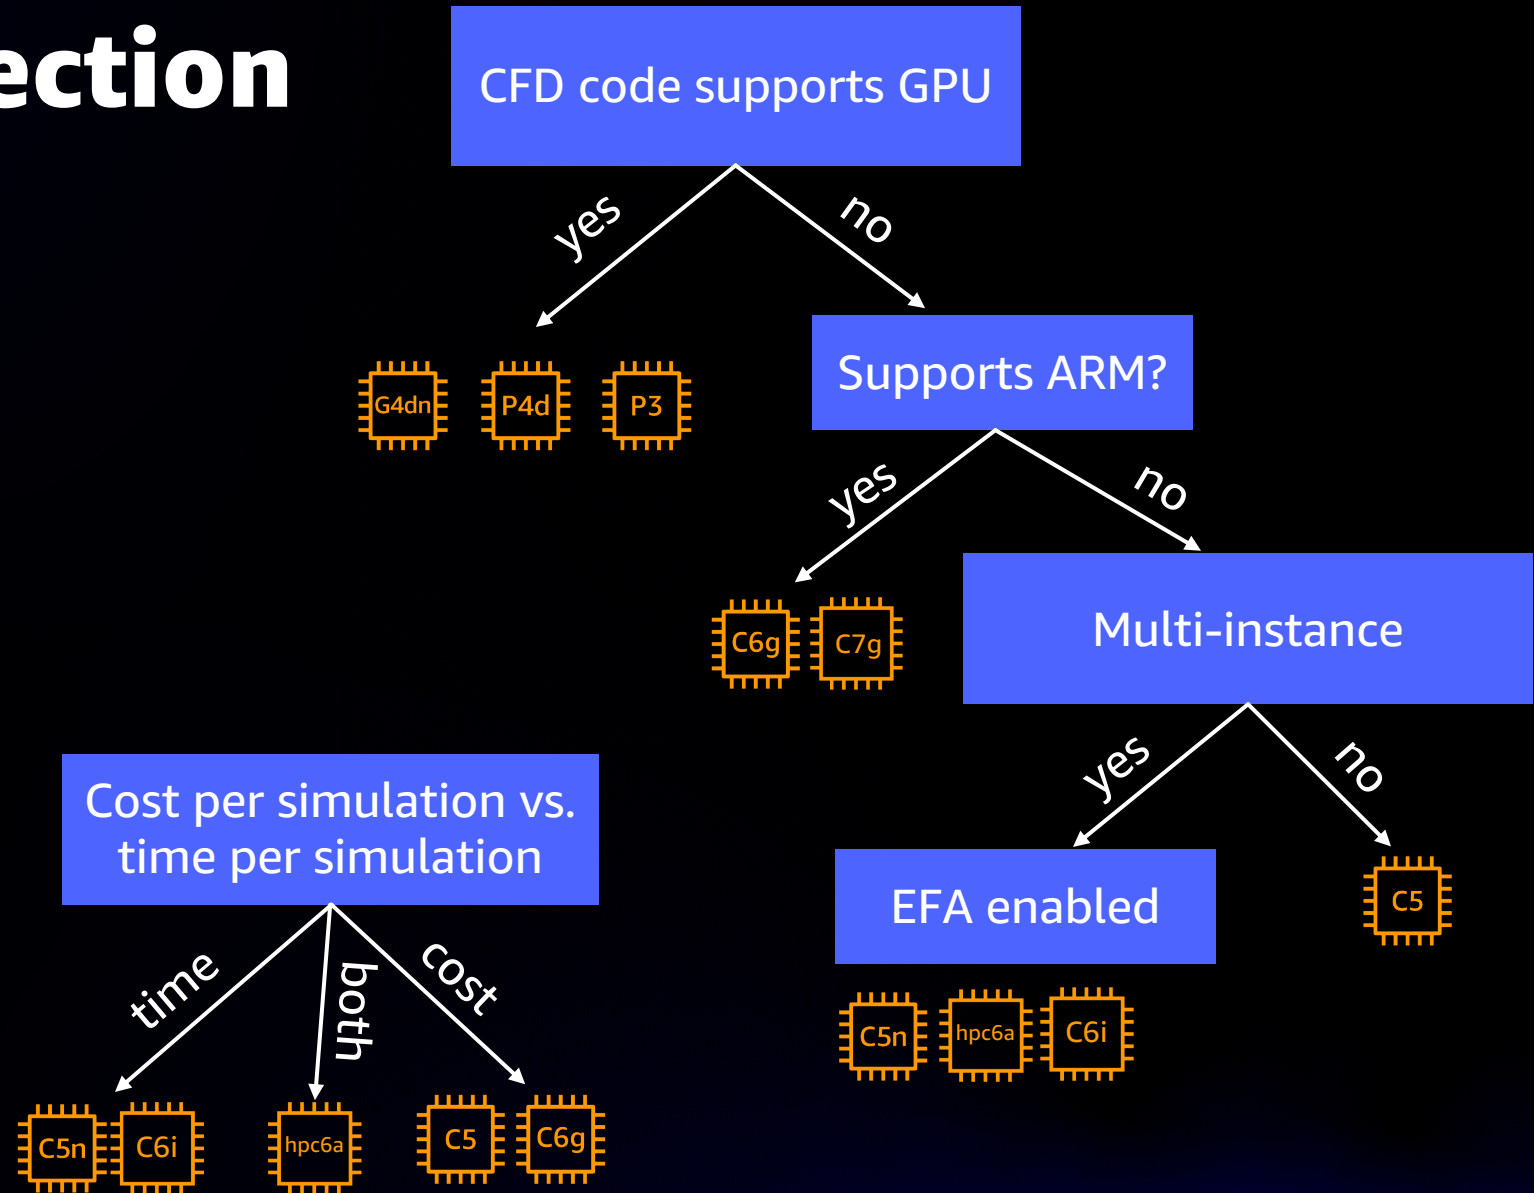
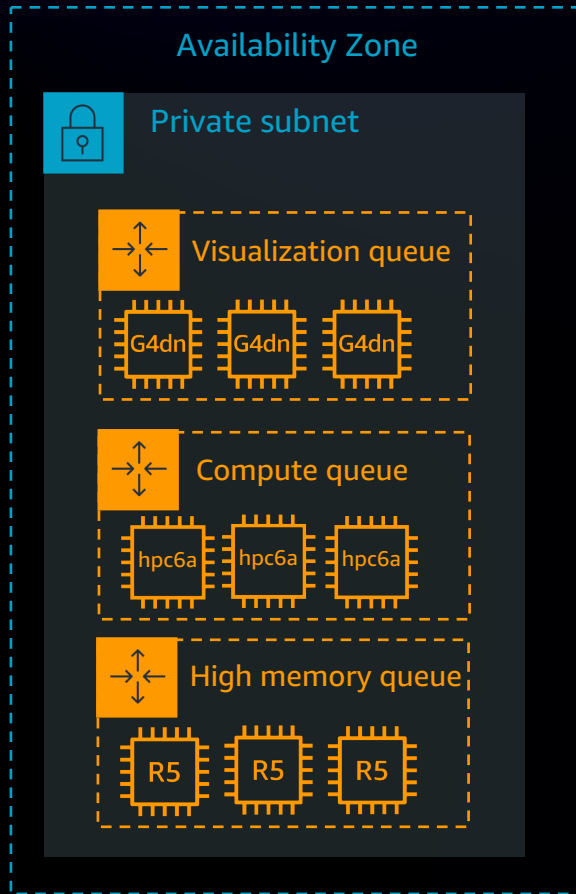


# Reference architecture





# 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 3<sup>rd</sup> Gen EPYC Milan processors, 96 cores, up to 3.6GHz frequency, and 384GB of RAM

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

Elastic Fabric Adapter enabled by default for 100 Gbps networking for inter-instance communications

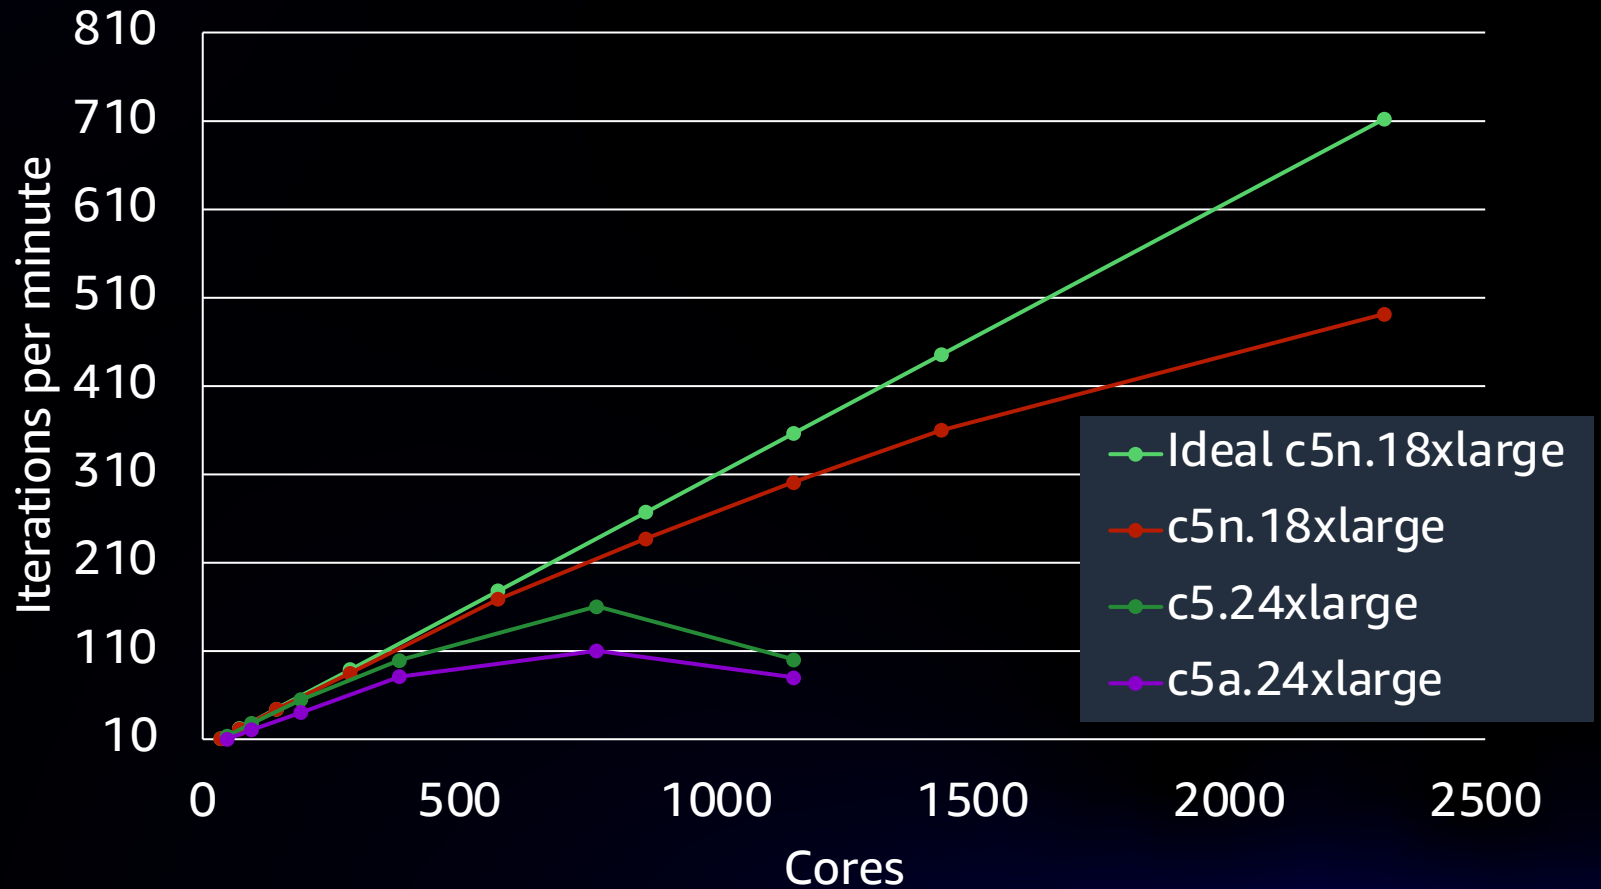
AWS Regions at launch: US East (Ohio) and GovCloud US West



# Multi-node


- For multi-node, the **network bandwidth** 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**

STAR-CCM+ v15.02.009 -LeMans 17M cells -  
OpenMPI

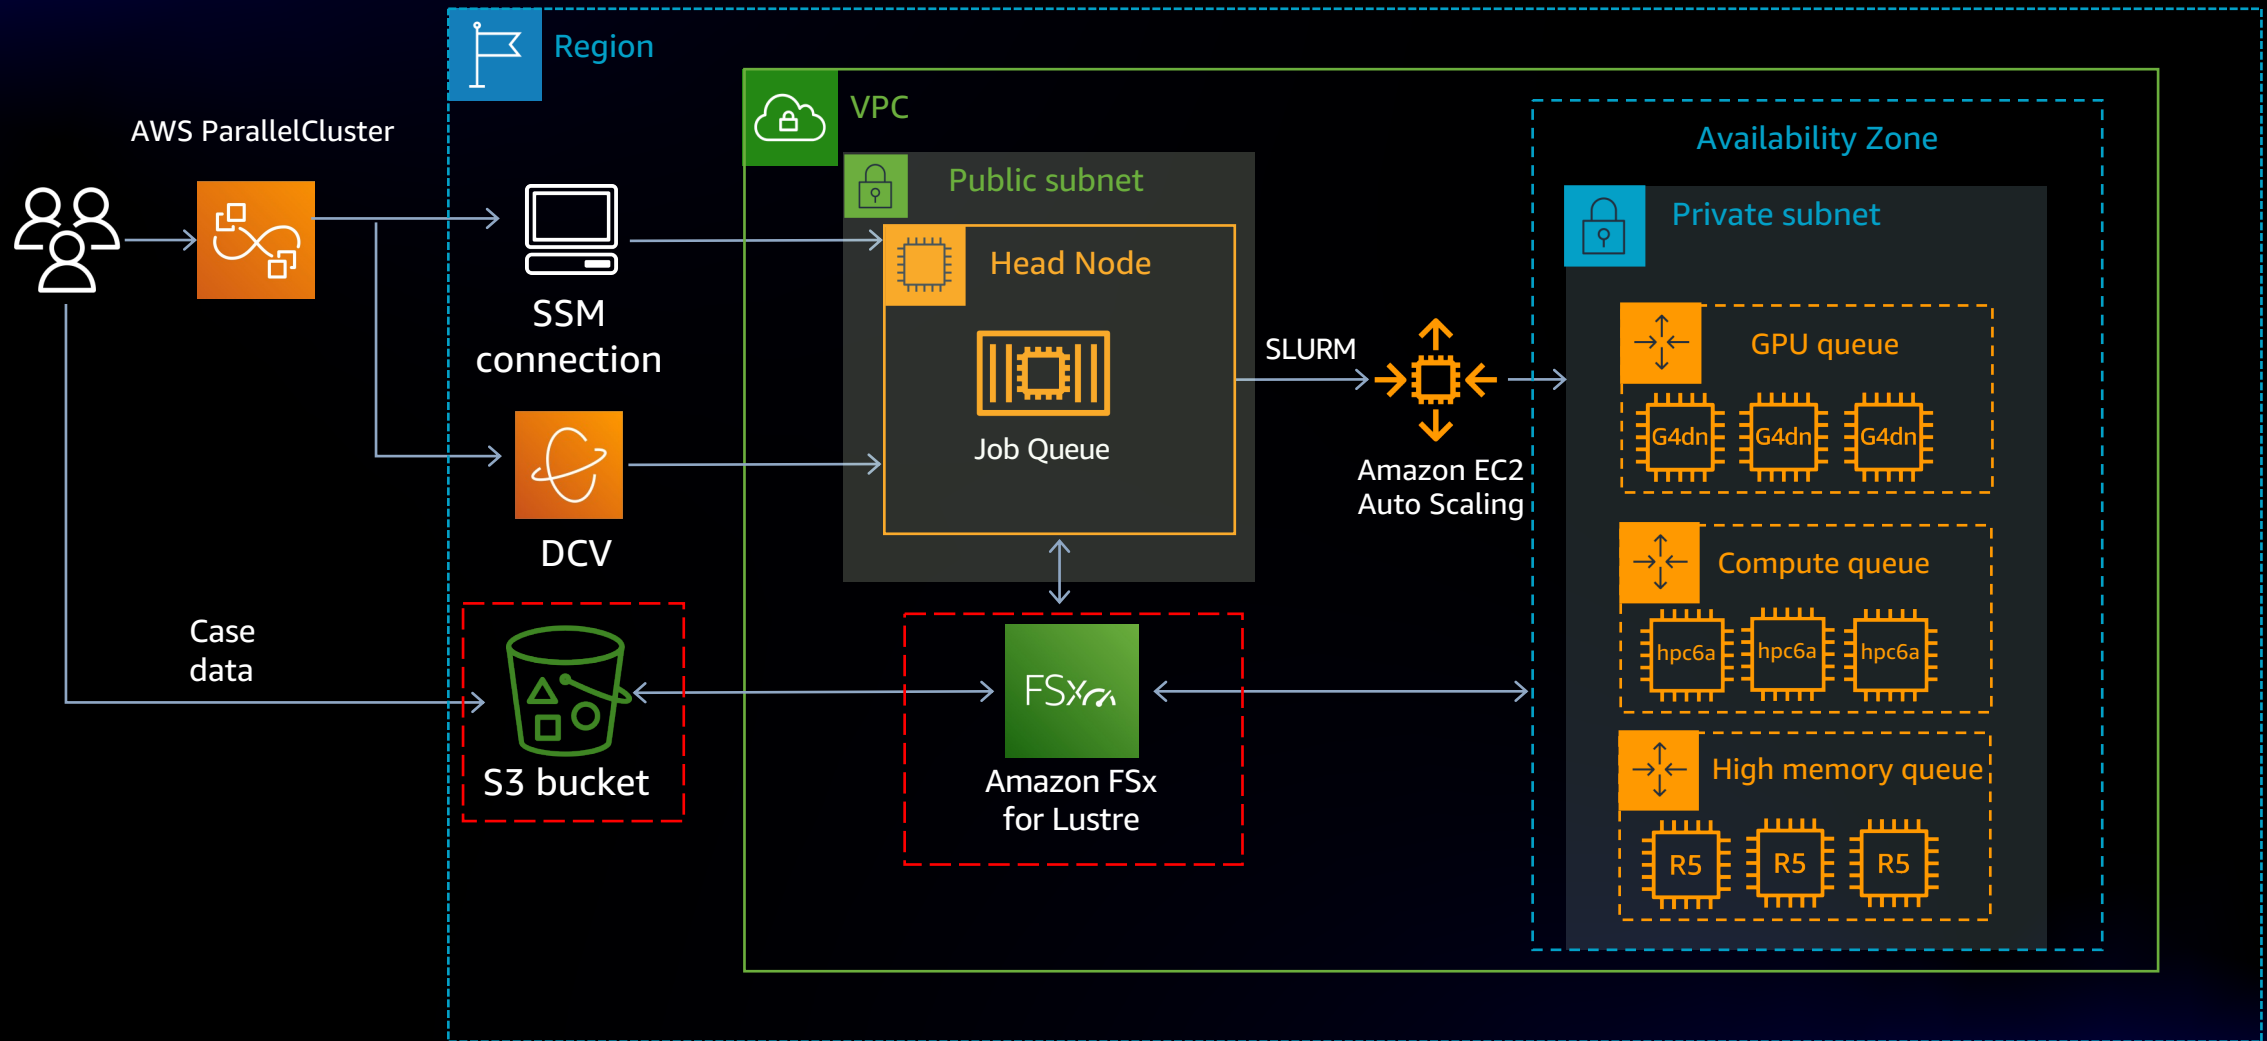


# Cells per core













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



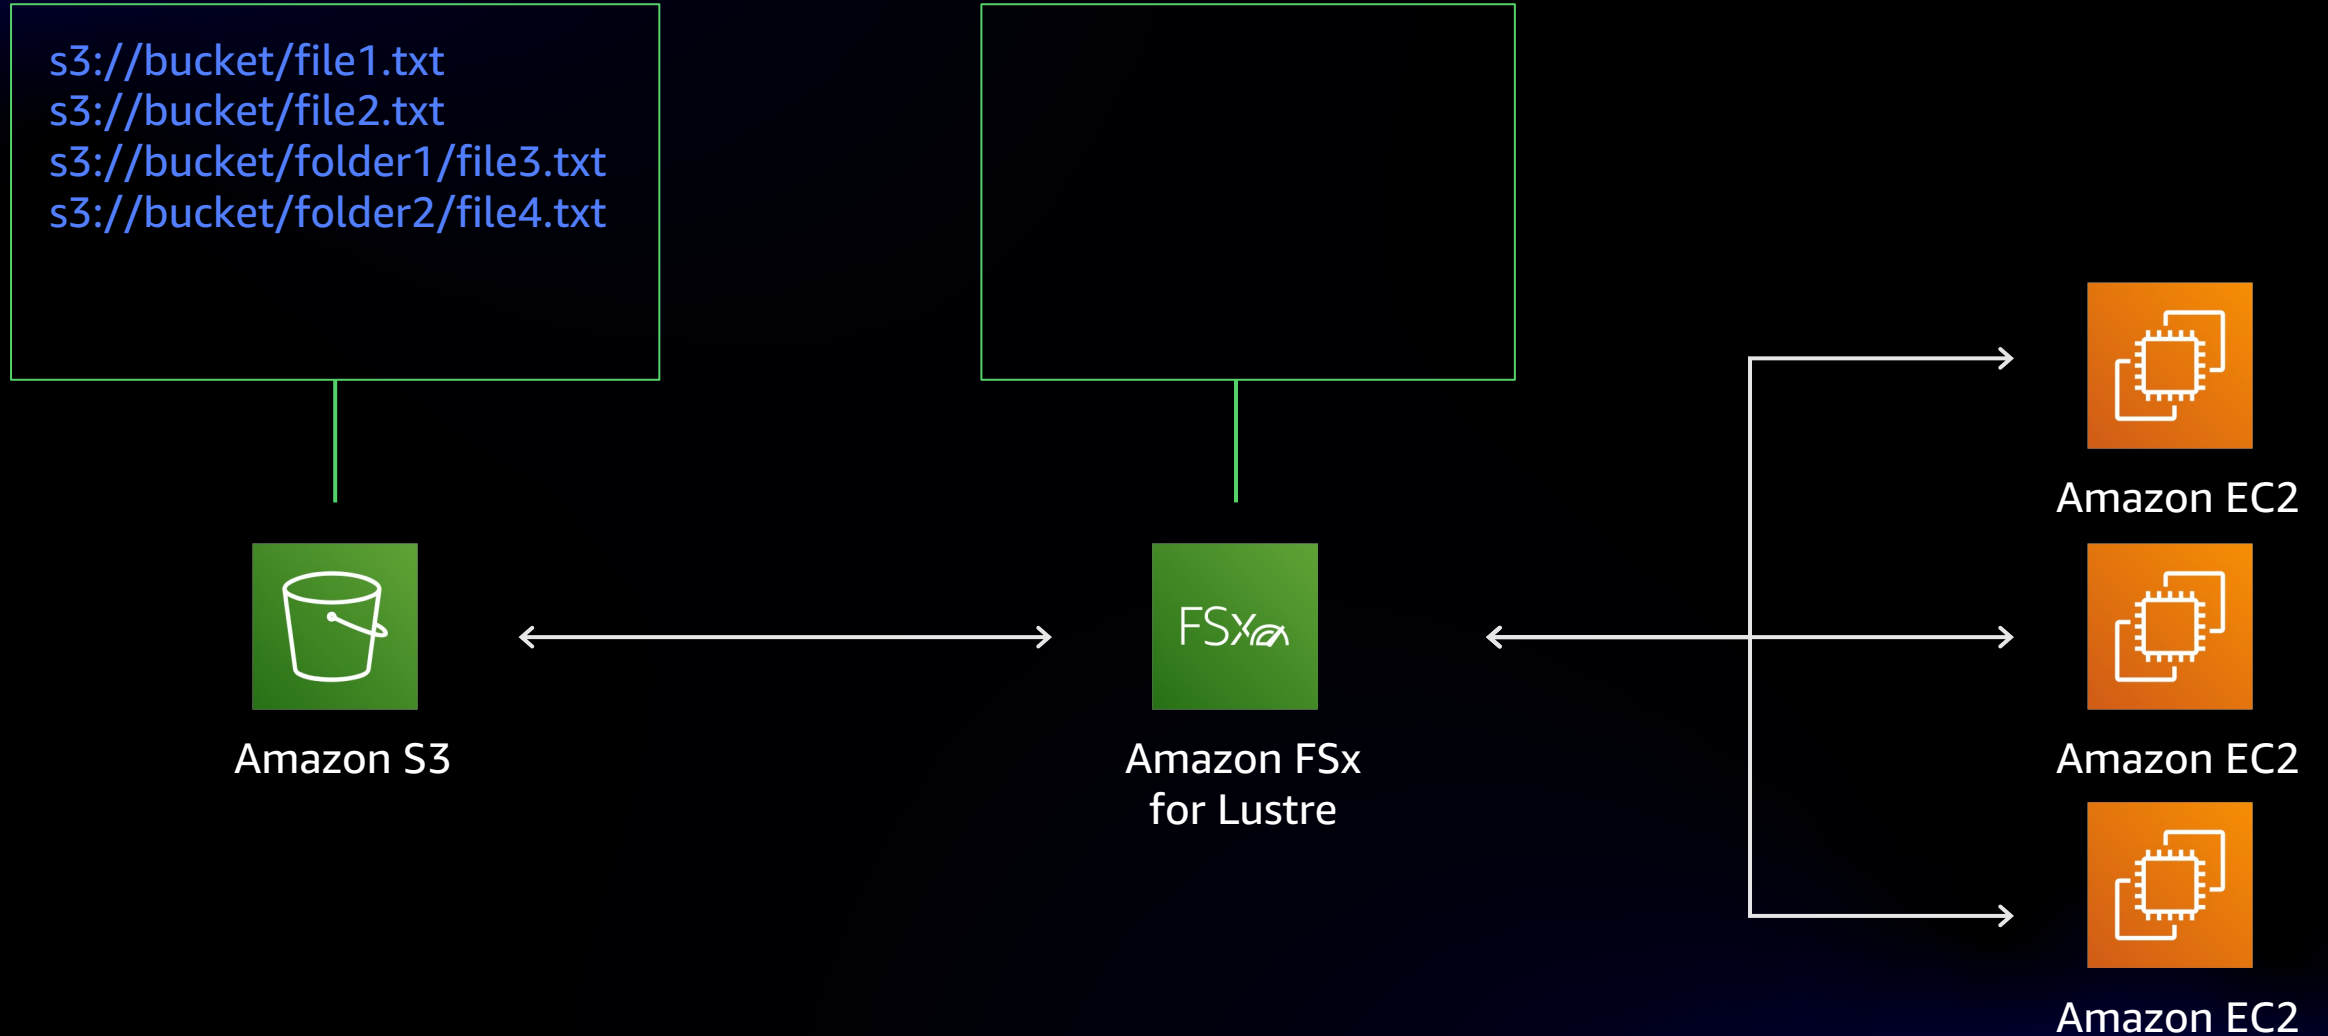
# Reference architecture



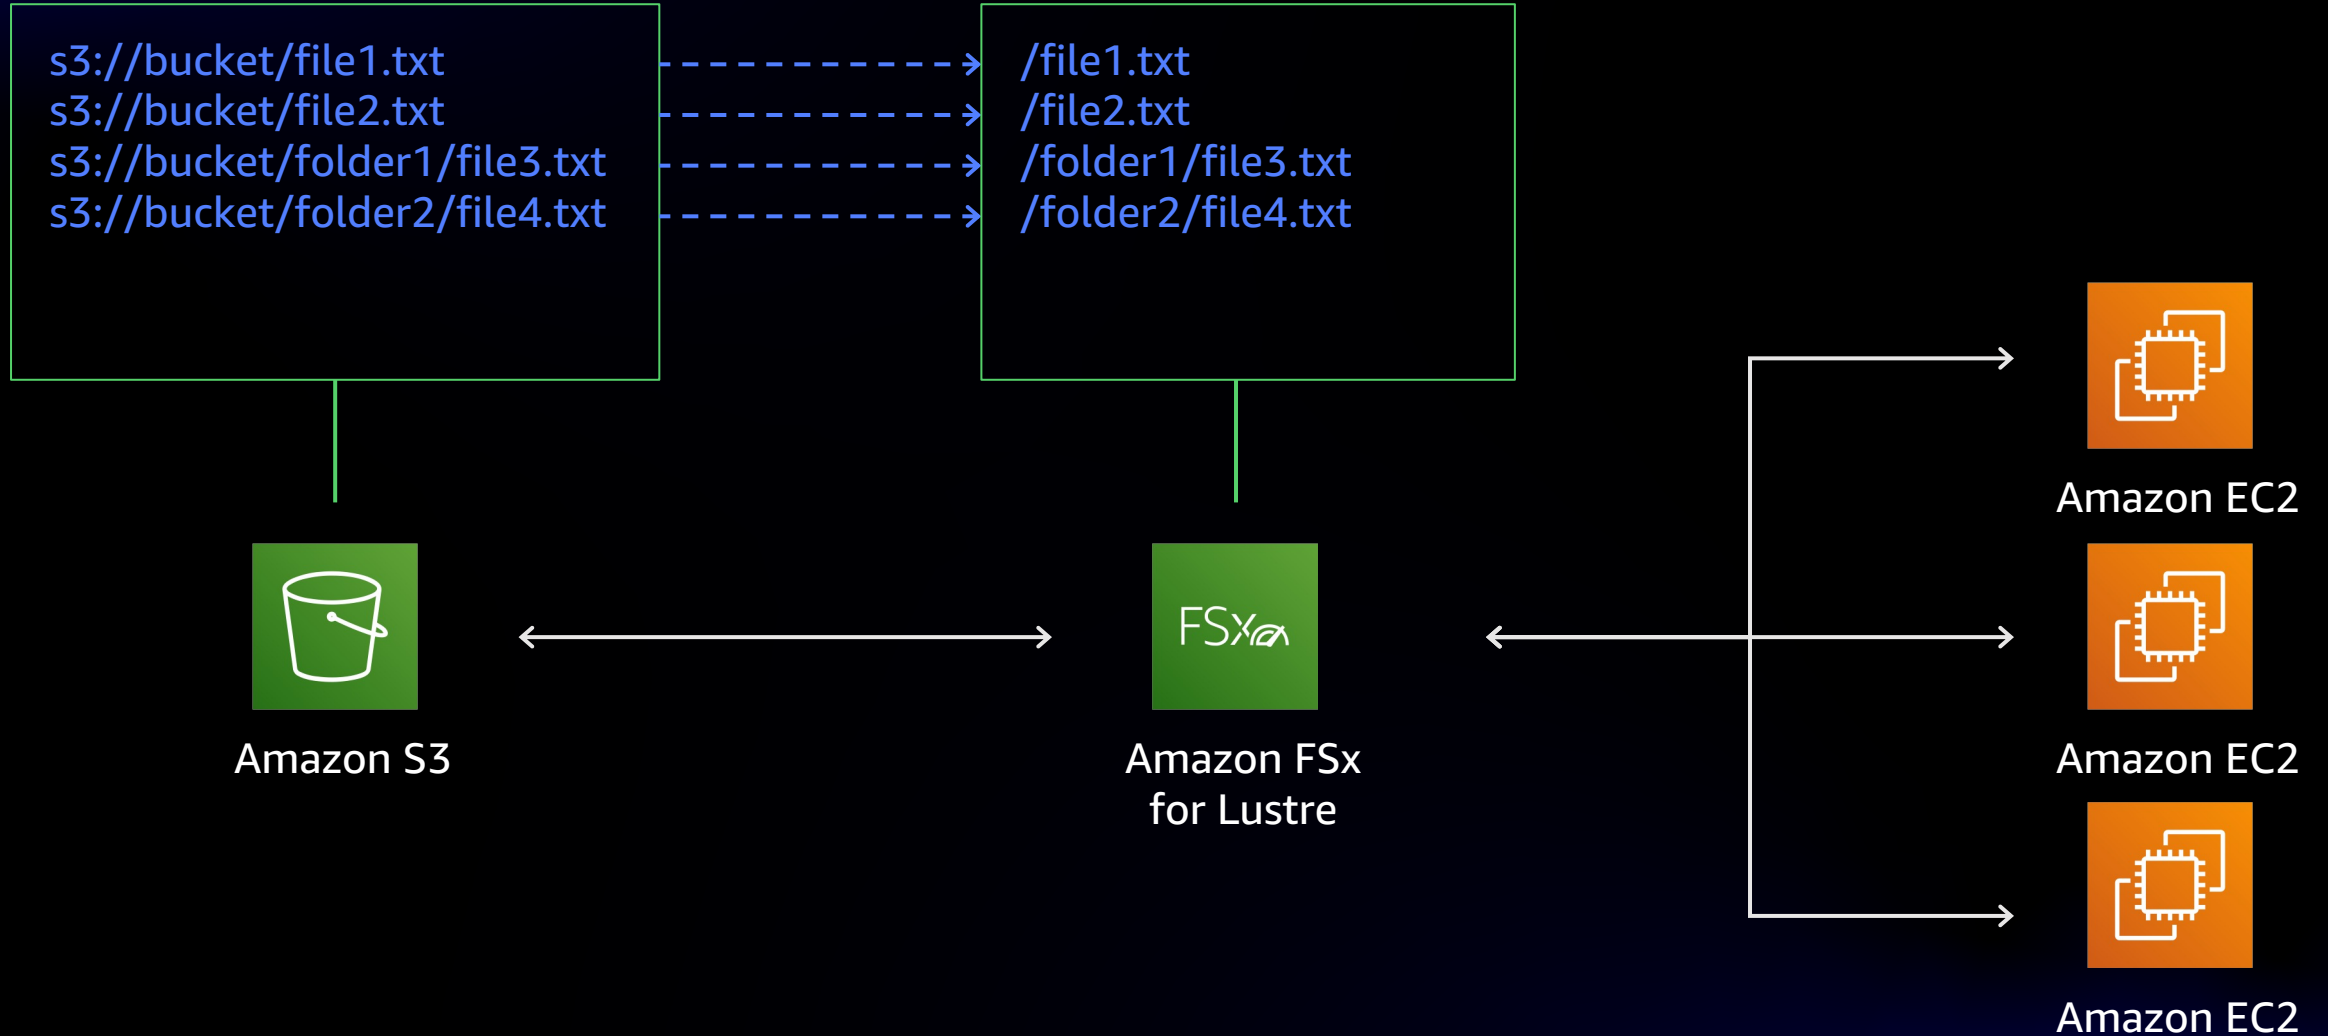
# Storage options optimized for price performance

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	
	1000 MB/s		\$0.82	

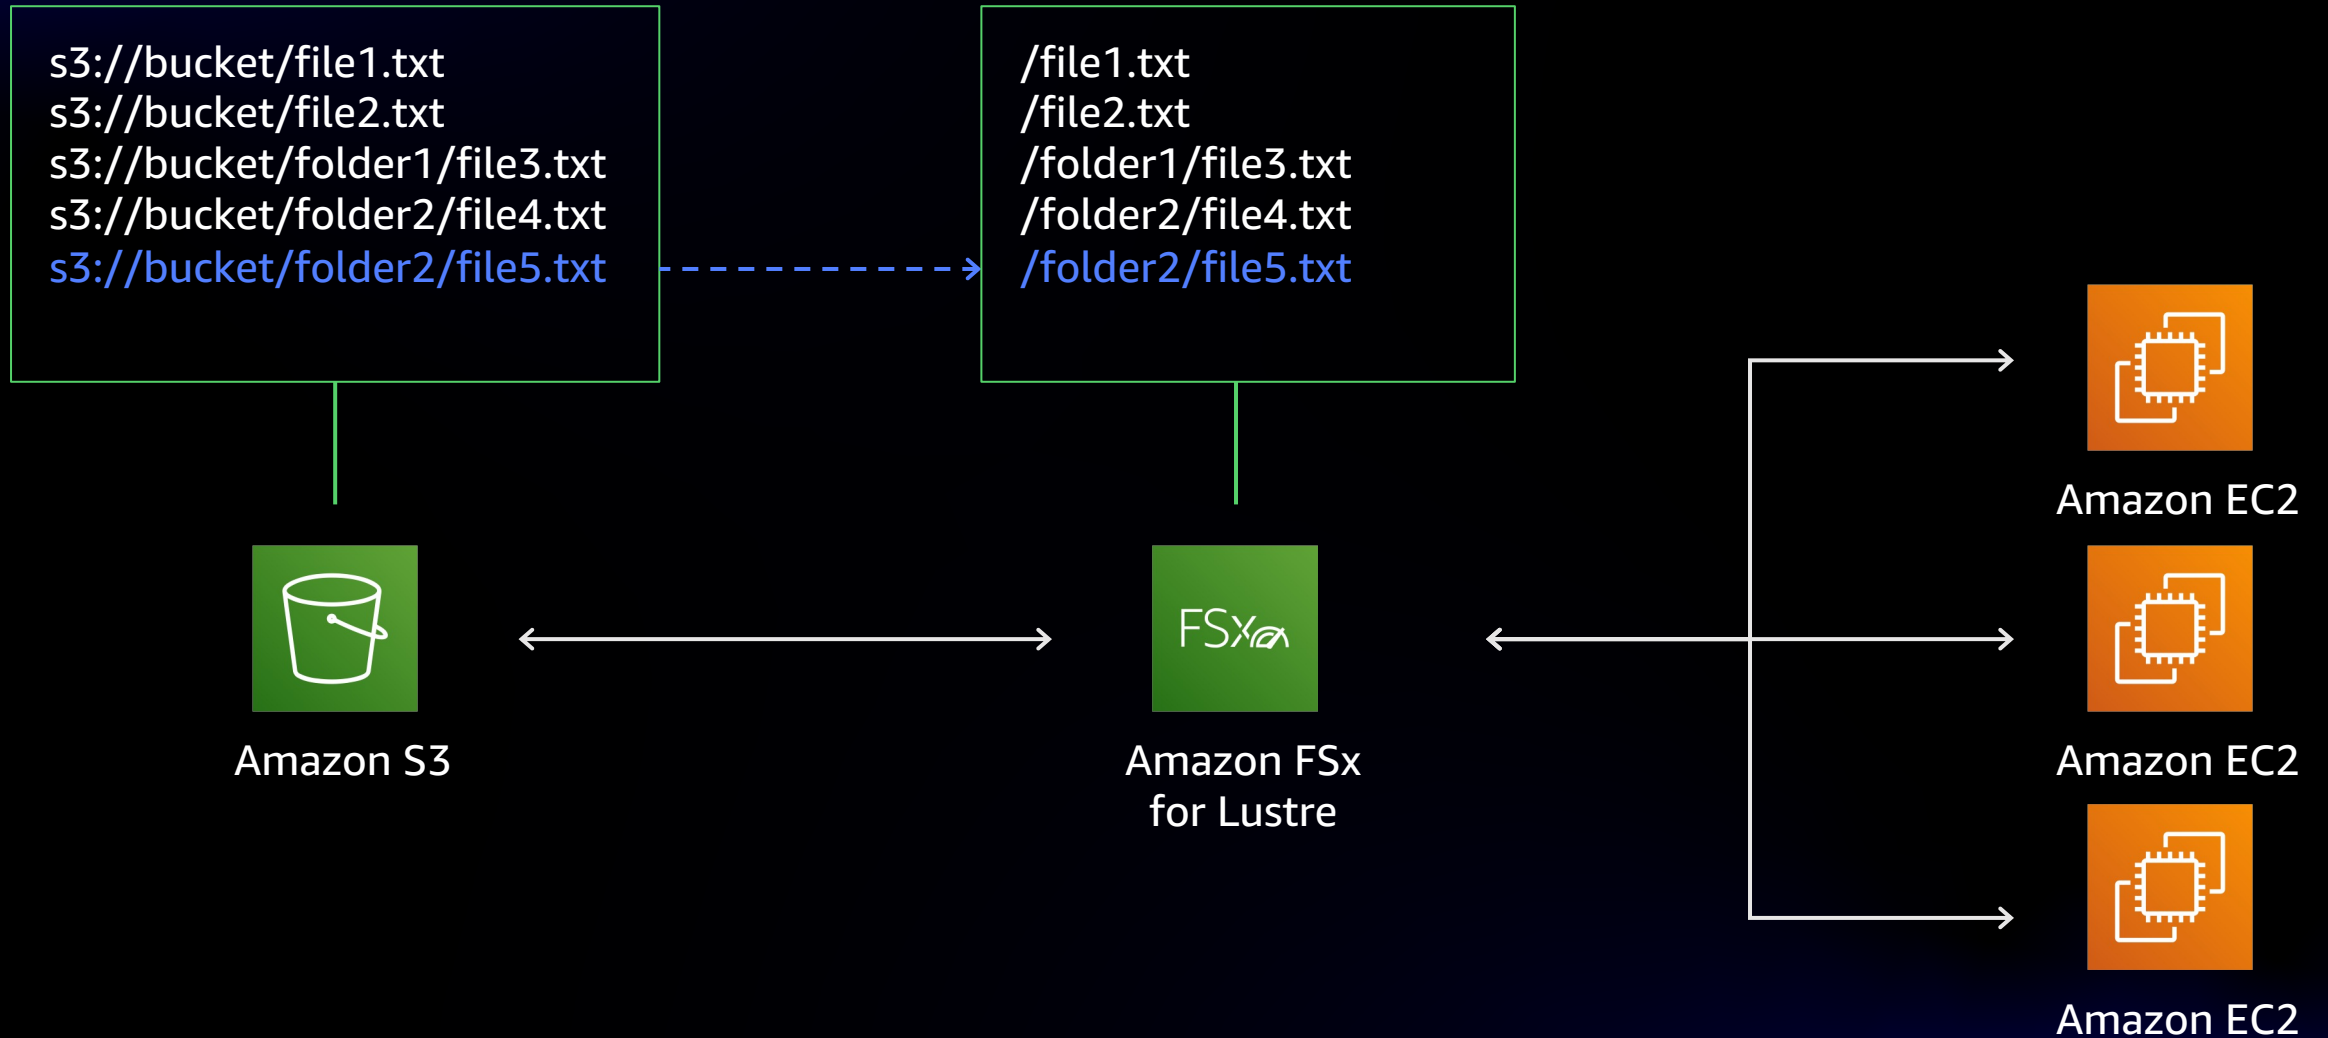
# Fast file interface for data on S3: **How it works**



# S3 objects appear on FSx file systems

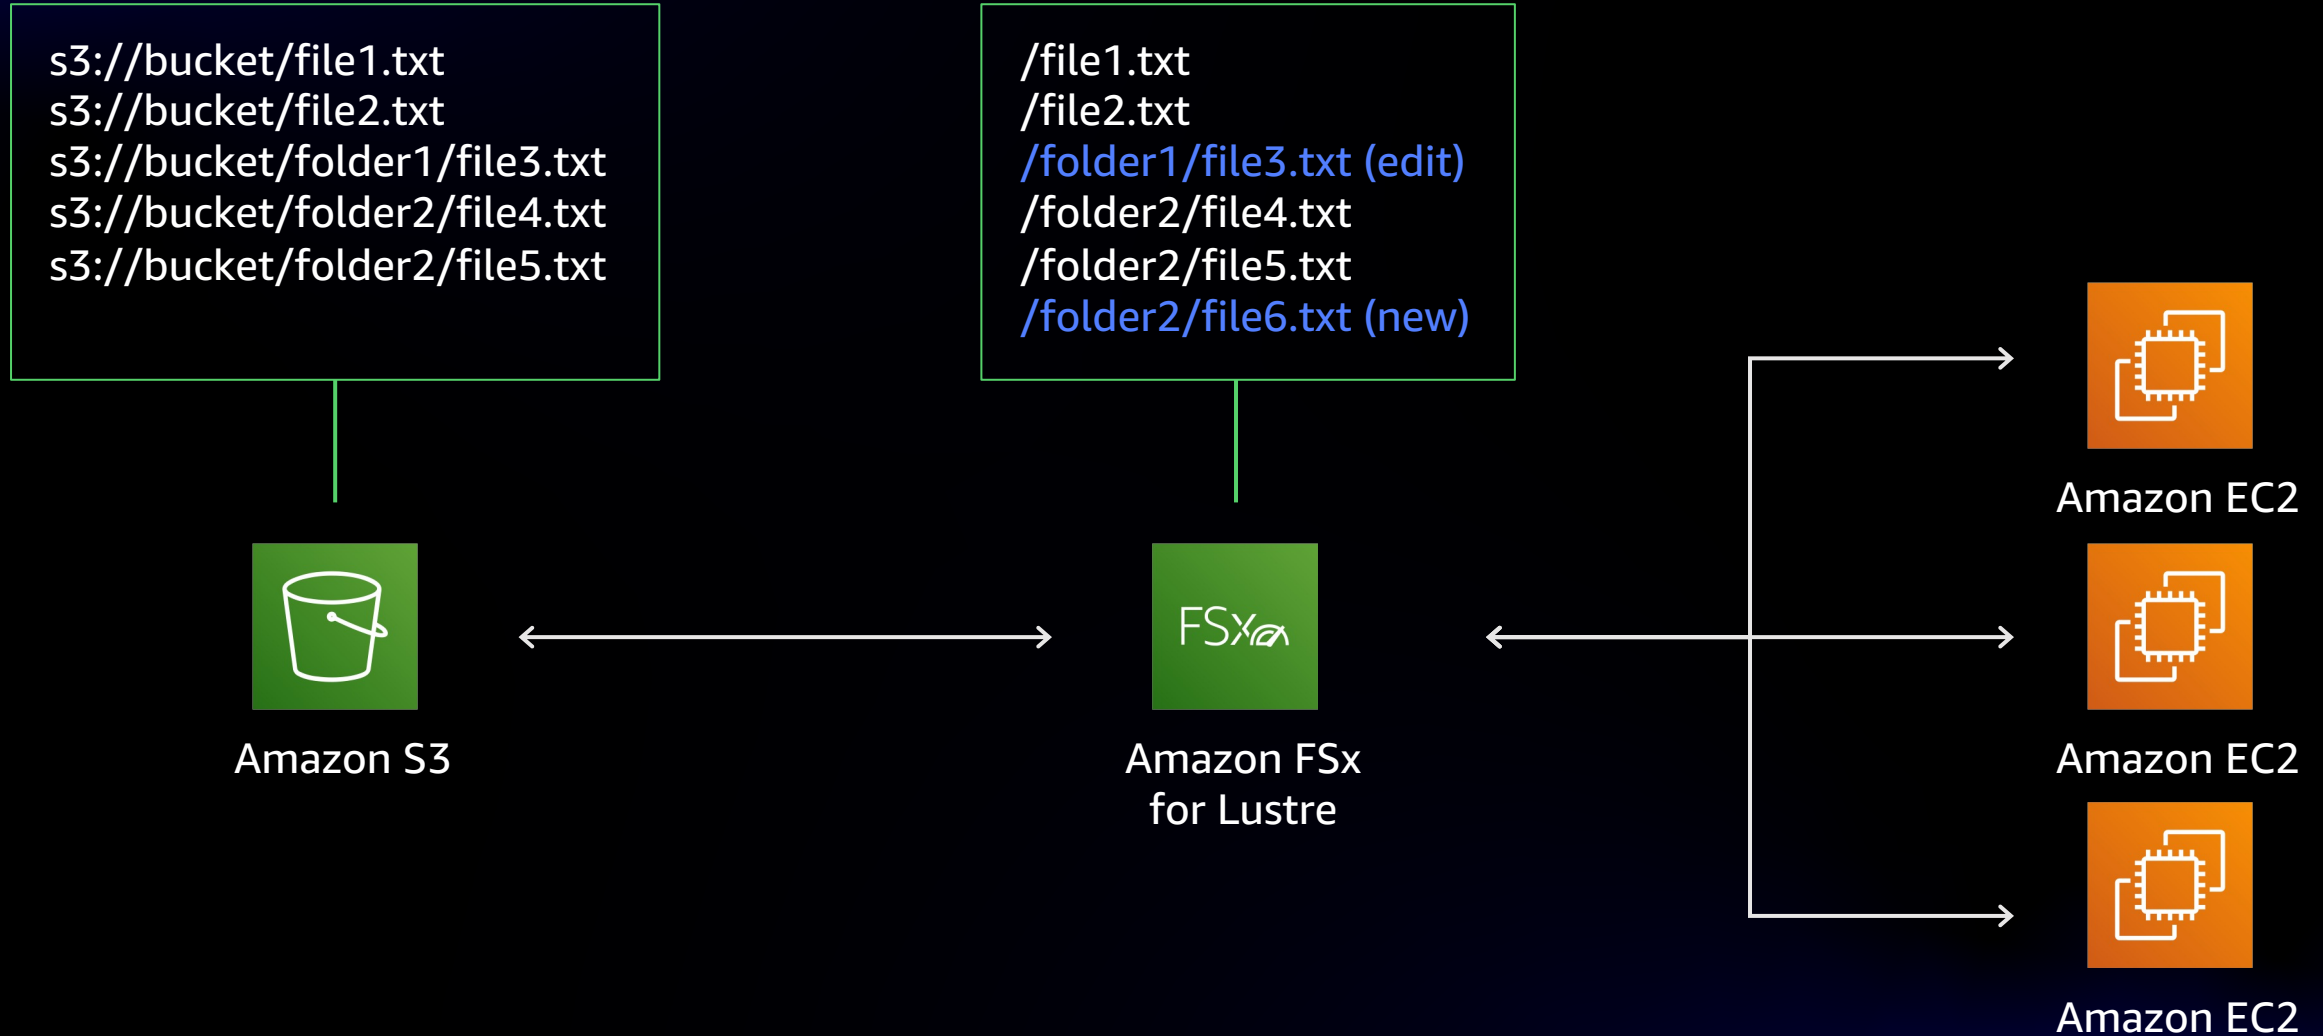


# Updates on S3 are imported to FSx for Lustre

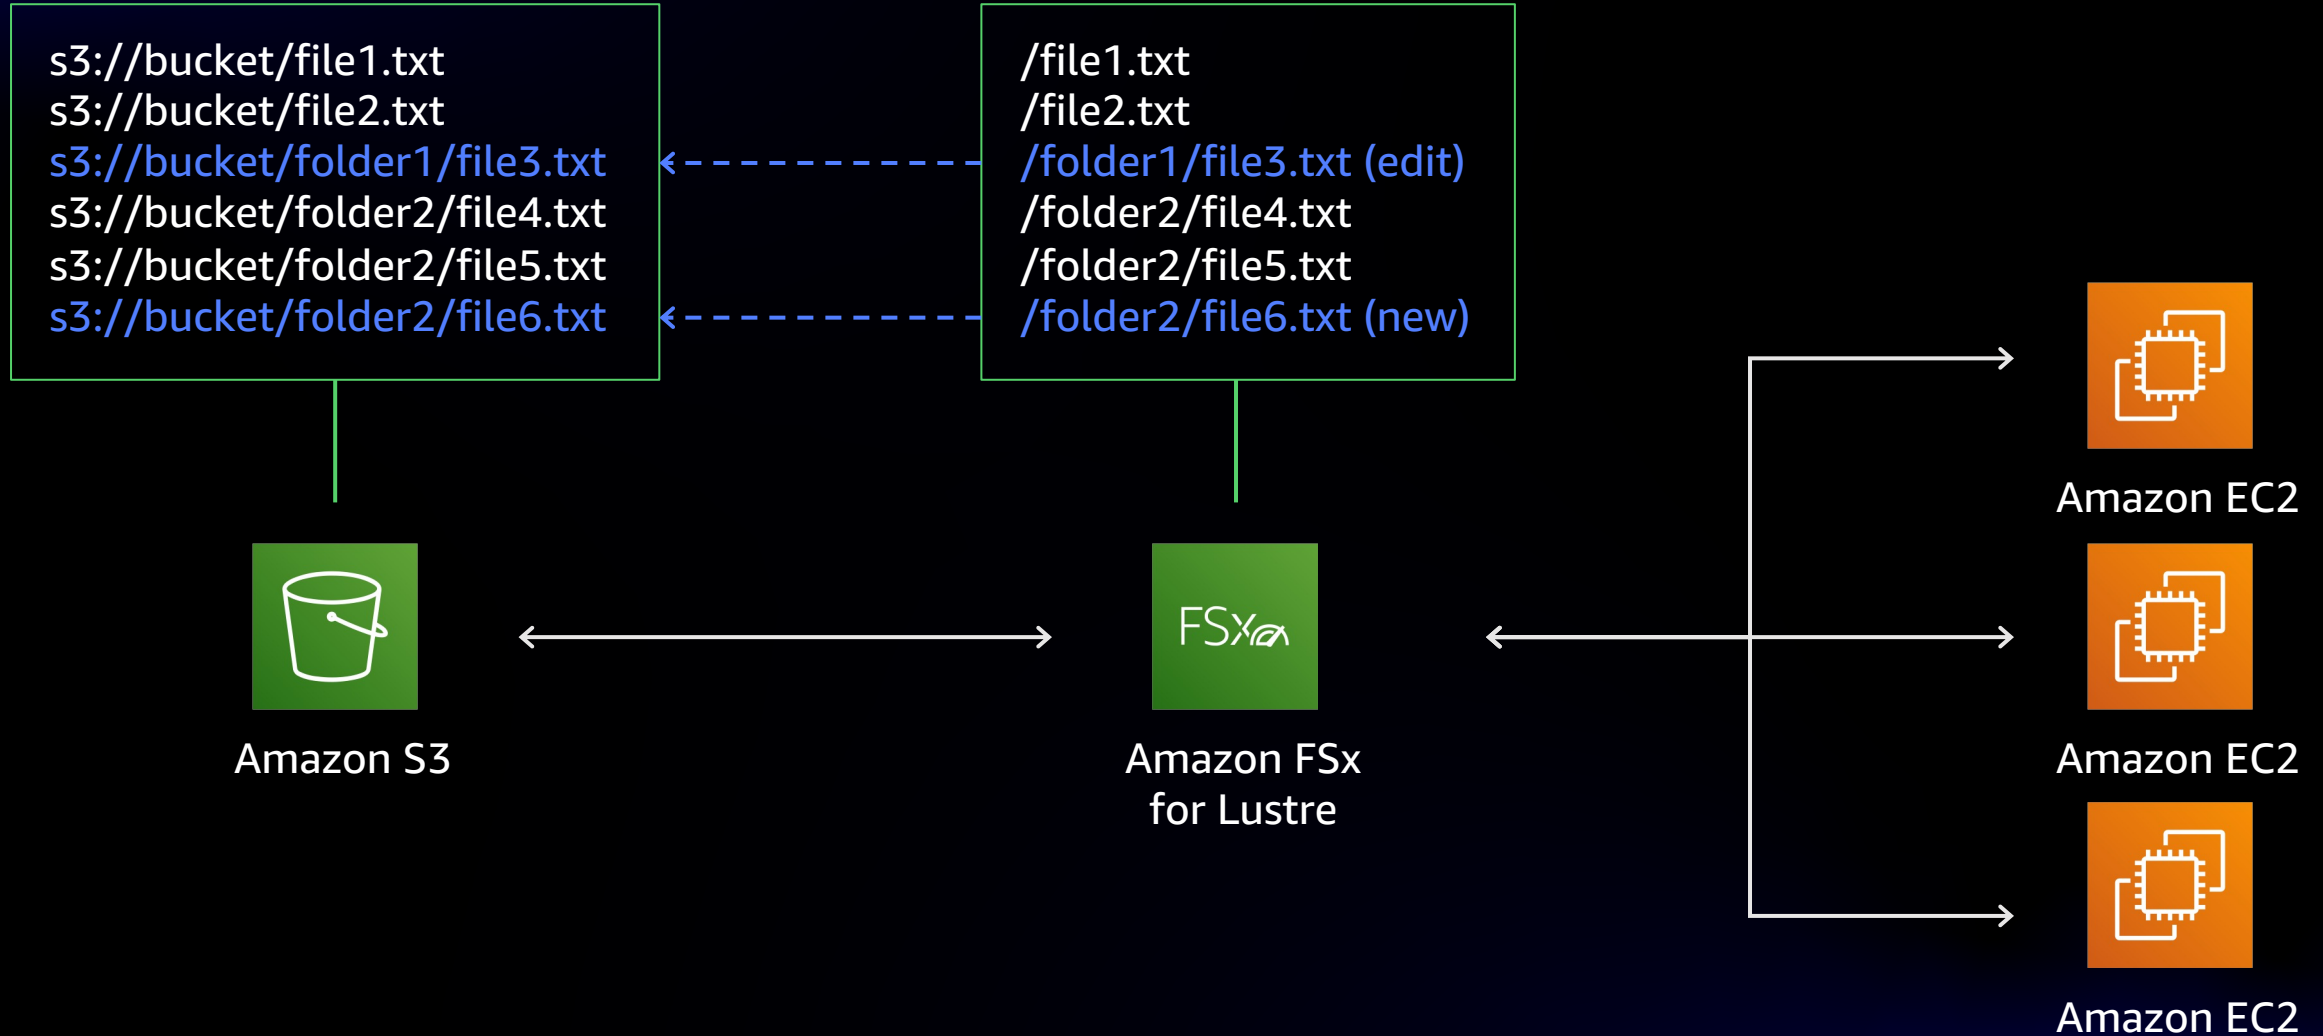




# Updates on FSx for Lustre are exported to S3



# Updates on FSx for Lustre are exported to S3



# Spin up/spin 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)
```




Amazon FSx  
for Lustre


Spin down resources between workloads



Amazon EC2

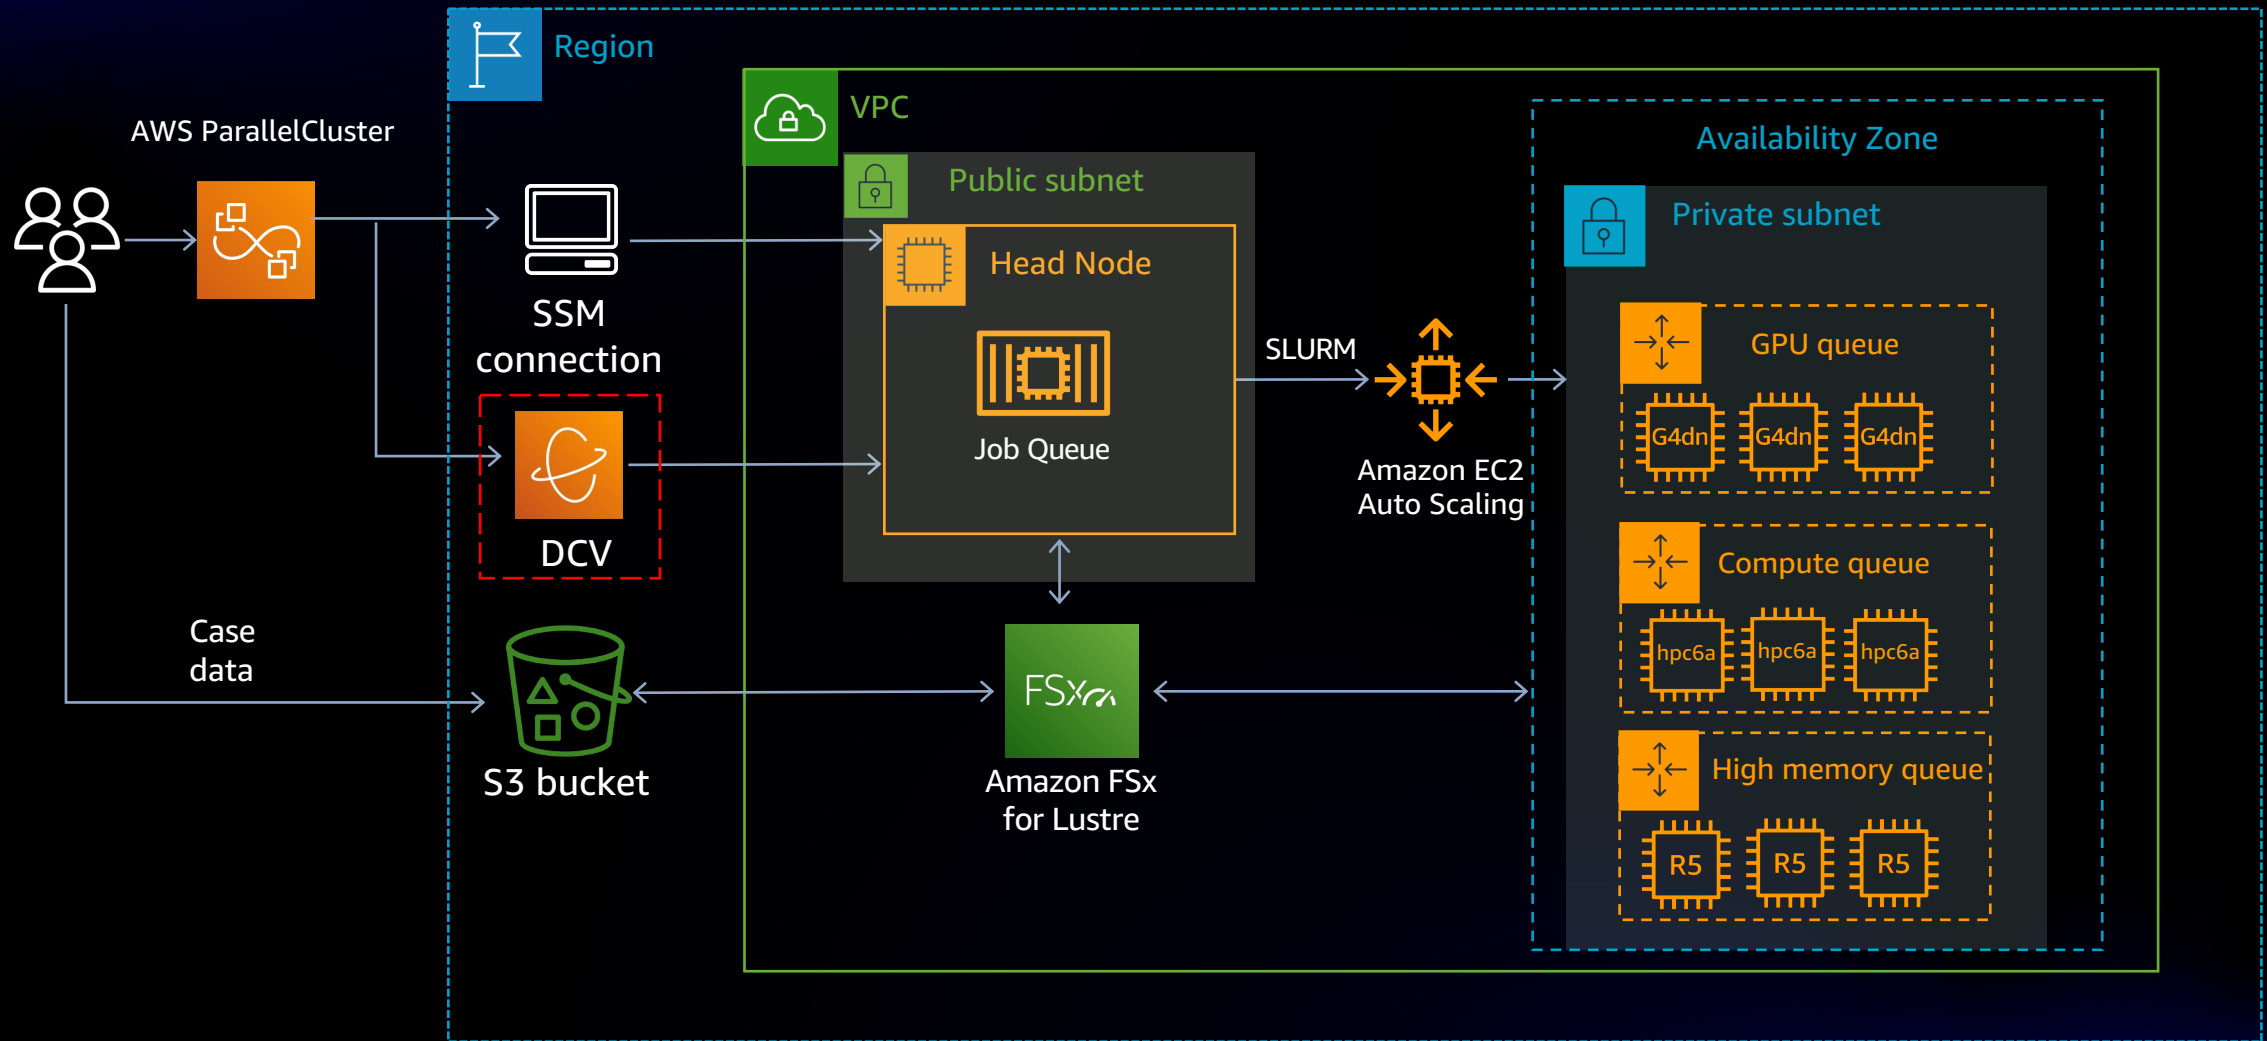


Amazon EC2




Amazon EC2

# Reference architecture

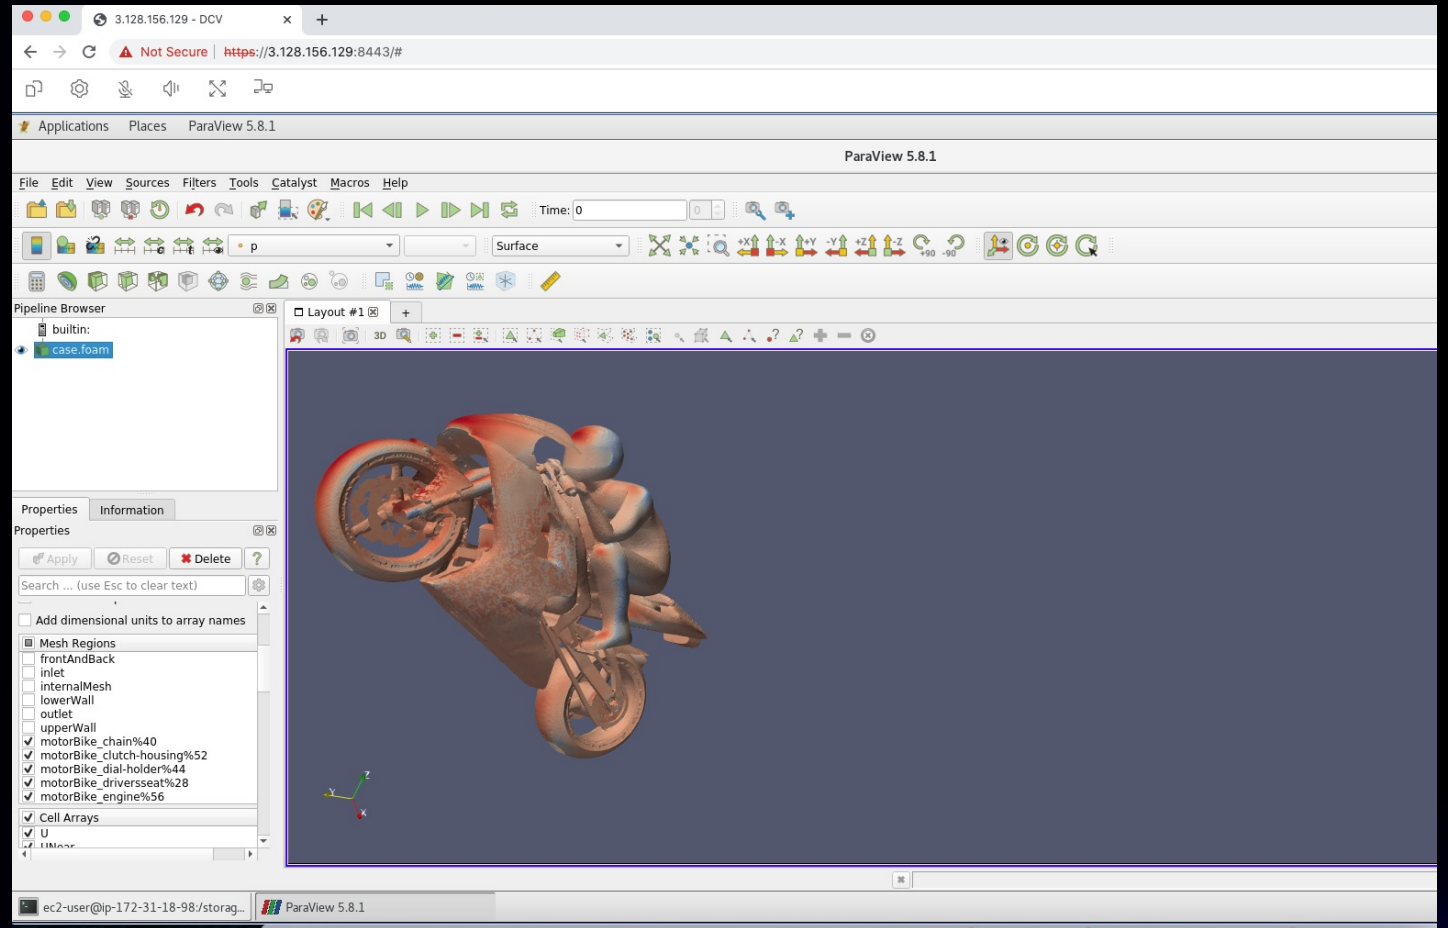


# NICE DCV encrypts and transports pixels to devices



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


# DCV enables pre & post processing in the cloud



# Instances

- Many users (especially commercial) want the best price/performance
- **C6g.16xlarge** is the clear winner who can use Arm – not all codes can support
- C5.24xlarge is still the best with c4.8xlarge an option but for x2 slower turnaround

OpenFOAM v1912 - 4M cell Motorbike for 5000 iterations







# 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 Lustre as scratch space and S3 for long term storage
- Next look at pricing models such as Spot or Reserved Instances

Evaluate the tradeoff of time to solution vs. cost for scaling




# Innovating with ML



**Model accuracy continues to increase with additional data**

## Applications


Natural  
language  
processing



Image/  
video analysis




Autonomous  
vehicle systems



Recommendation  
systems



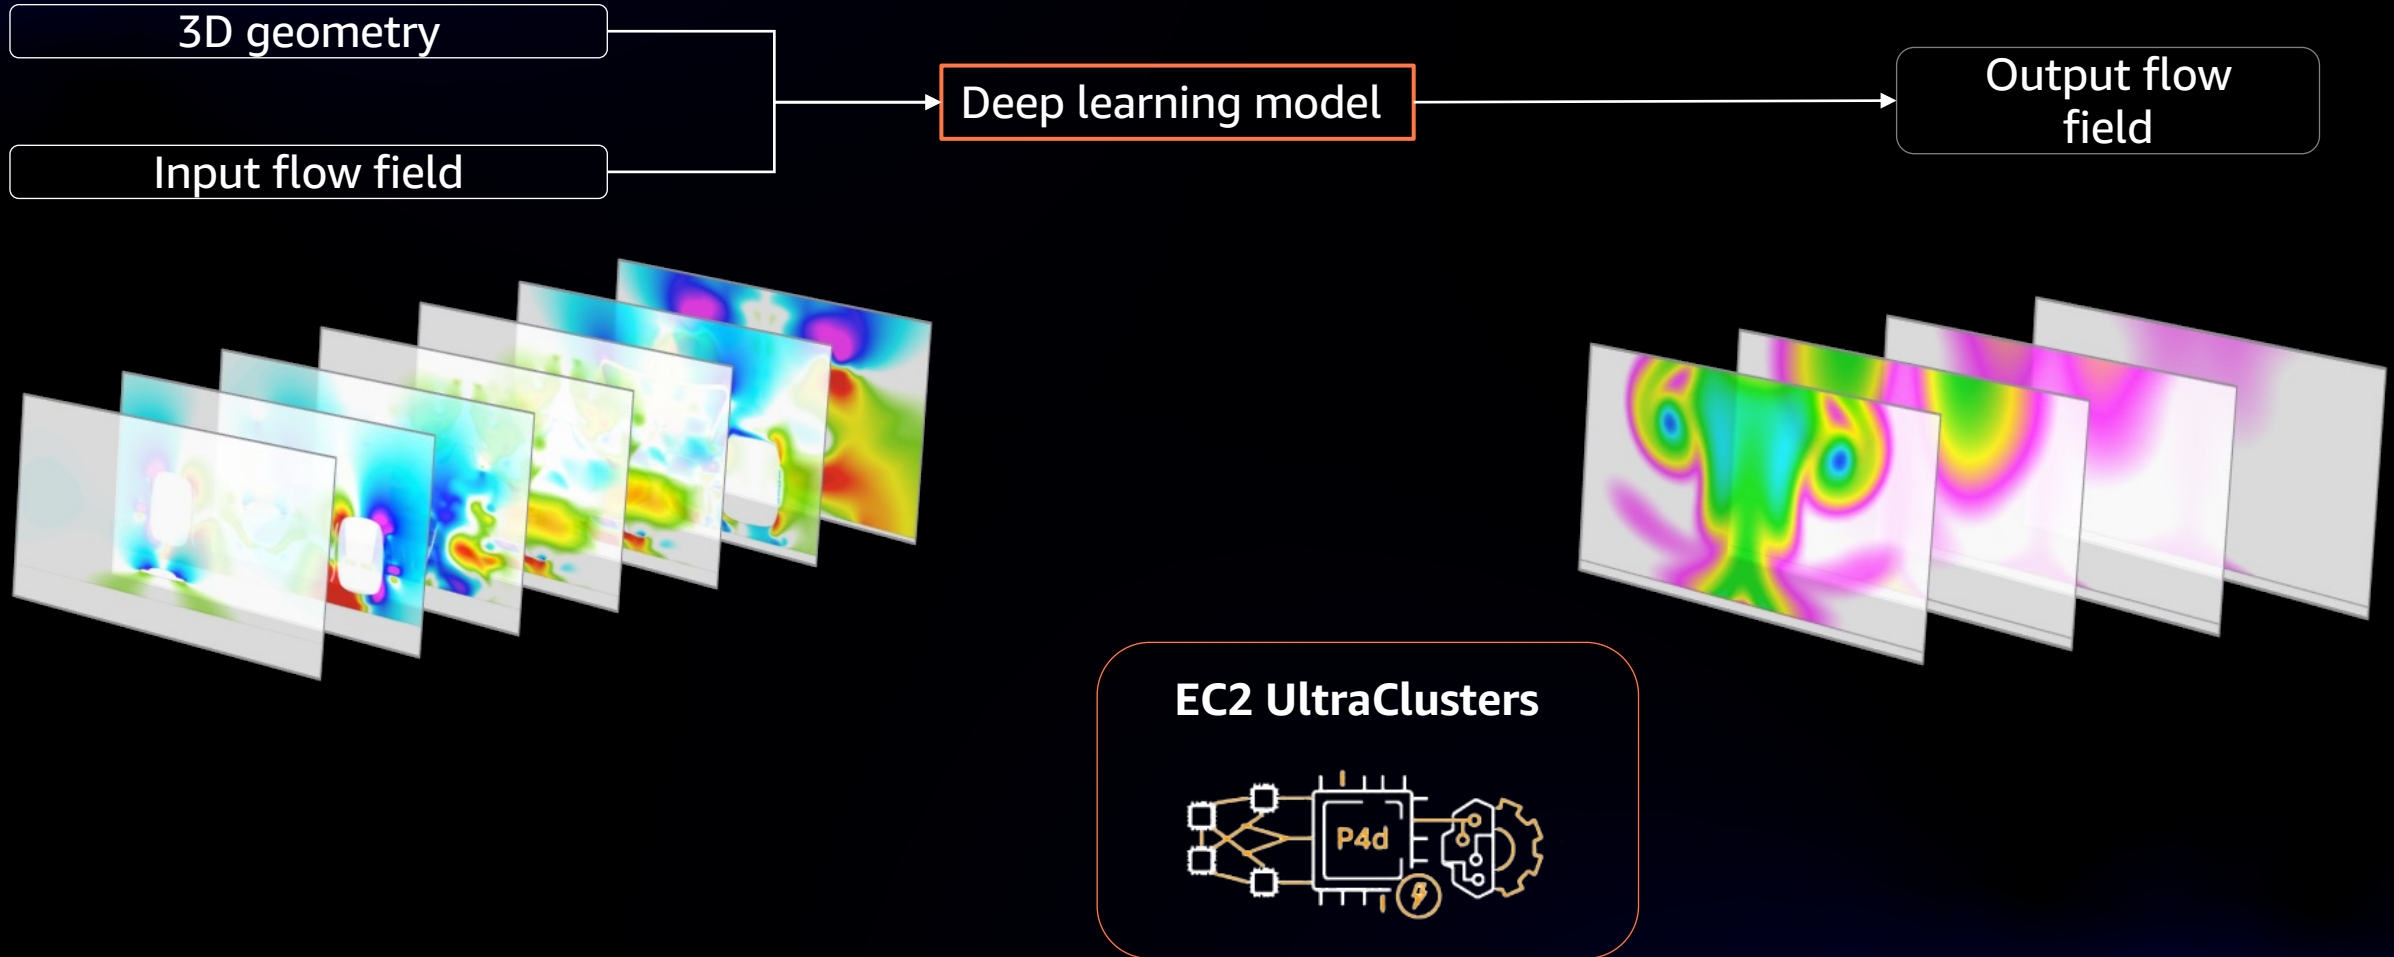
# Demonstrating ML + HPC: Extract relationships from HPC simulation results



## Problem statement

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

# ML + HPC solution



# Call to action

**CMP324-R2:** Scaling machine learning workloads on Amazon EC2

Date: Thursday, December 2, 1:00–2:00 PM

Location: Level 1, Alliance 311, Caesars Forum

**CMP312:** The future of HPC is looking a lot like ML

Date: Thursday, December 2, 3:15–4:15 PM

Location: Level 1, Summit 231, Content Hub, Blue Screen, Caesars Forum

# Key resources

CFD on AWS

[aws.amazon.com/hpc/cfd](https://aws.amazon.com/hpc/cfd)

AWS CFD workshops

[cfd-on-pcluster.workshop.aws](https://cfd-on-pcluster.workshop.aws)

CFD on AWS whitepaper



[d1.awsstatic.com/whitepapers/computational-fluid-dynamics-on-aws.pdf](https://d1.awsstatic.com/whitepapers/computational-fluid-dynamics-on-aws.pdf)

Content is being actively developed – please check back frequently



# Scaling on AWS – OpenFOAM



OpenFOAM v1912 - MotorBike (222M cells) - IntelMPI 2019.6 - AL2 - PC2.6.1





# Scaling on AWS – STAR-CCM+



Simcenter STAR-CCM+ 2020.1 - F1 (403M cells) - IntelMPI 2019.6 - AL2 - PC2.6.1



- Ideal
- AWS c5n.18xlarge

# Scaling on AWS – Fluent

ANSYS Fluent 19.5 - F1 (140M cells) - IntelMPI 2019.5 - AL2 - PC2.5.1



# Thank you!

Karthik Raman

[akraman@amazon.com](mailto:akraman@amazon.com)

Sean Smith

[seam@amazon.com](mailto:seam@amazon.com)





Please complete  
the session survey