AWS re:Invent

NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

CMP402

A day in the life of a cloud-empowered aerodynamics engineer

Dr. Neil Ashton

Principal Computational Engineering Solution Architect

AWS

Sean Smith

Sr. Solutions Architect HPC

AWS

Agenda

- Introduction to engineering challenges
- What is CFD and the goal of digital certification
- CFD on AWS Reference Architecture
- Live Demo using Siemens Simcenter STAR-CCM+
- Optimizing your cloud CFD setup (Compute, storage, visualization, cost, ML)

How do you design the next generation of X?


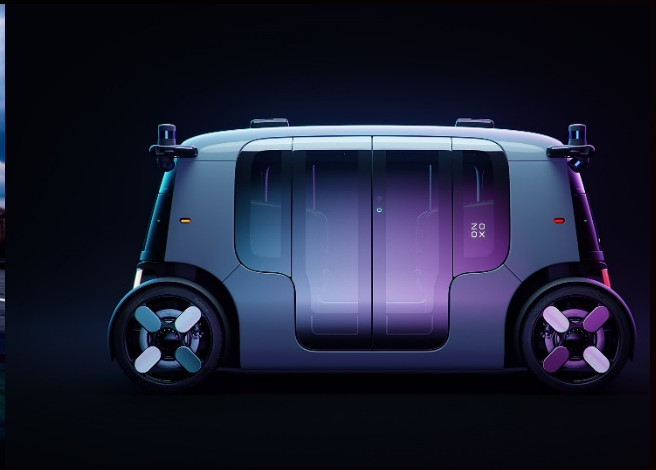

Companies need a way to assess the performance and efficiency of new designs

Aerodynamics is key, especially in the age of electric vehicles

They want to go from conceptual design to full design in the shortest possible time with the least expense

Digital certification is the dream for many industries to reduce cost and time compared to physical tests; let's explore that further

Computational fluid dynamics – CFD

No manufacture

Cheap & fast*

Accurate?

Navier-Stokes Equations (1822) – *Claude-Louis Navier & George Gabriel Stokes*

$$\frac{\partial u_i}{\partial x_i} = 0$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j}$$

Key tool for digital certification

Why high-fidelity CFD? – Automotive

- Improved **accuracy** over low-fidelity methods (e.g., RANS)
- Transient nature means **5x to 20x** more computationally expensive i.e 50k core-hours to 500k core-hours per sim.
- Growing desire from motorsport, automotive, or aerospace to move

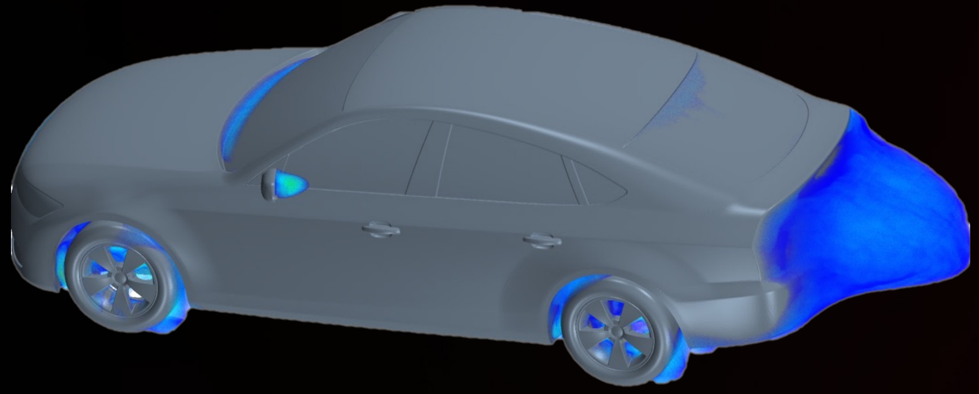

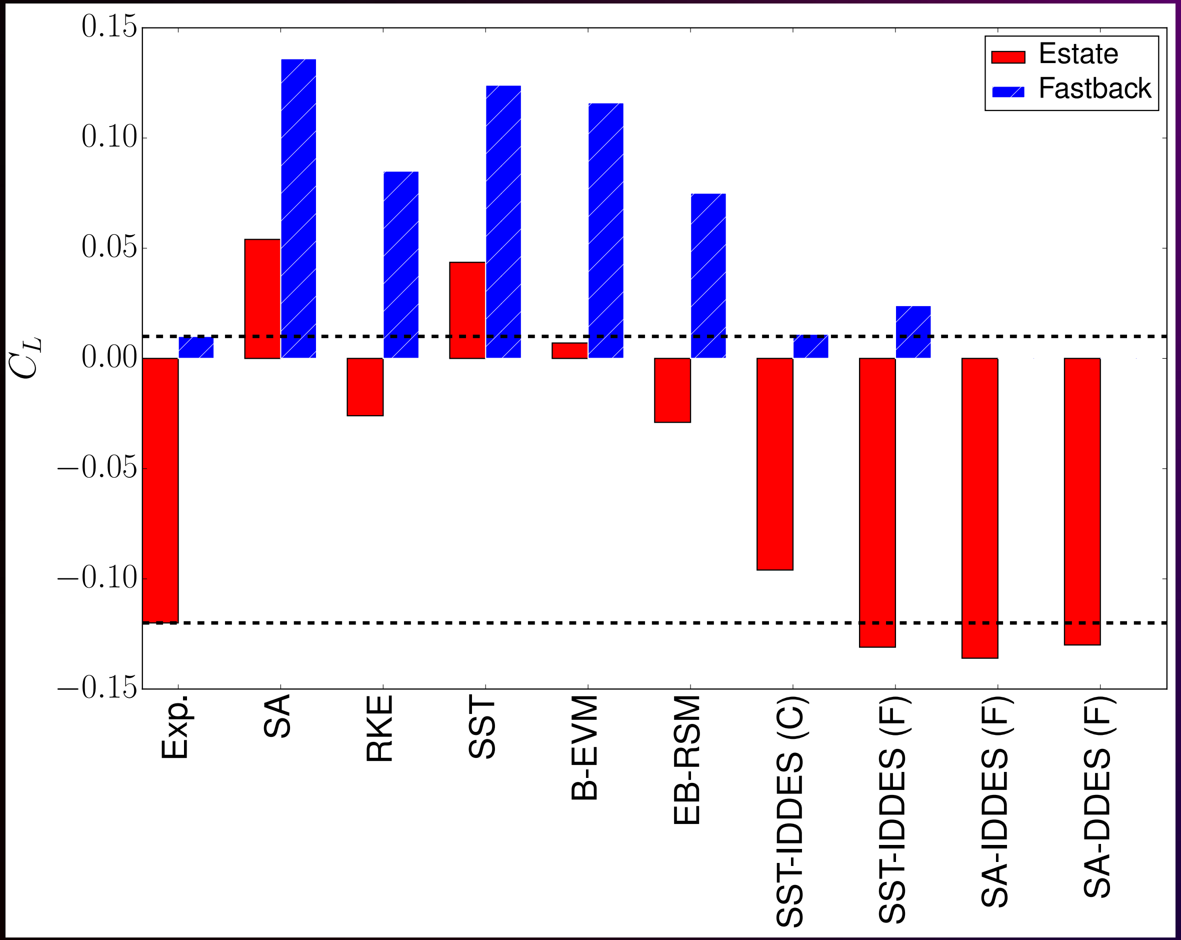

N. Ashton, A. West, S. Lardeau, A. Revell, *Assessment of RANS and DES methods for realistic automotive models*, Computers. Fluids. 128 (2016) 1–15. doi:10.1016/j.compfluid.2016.01.008.

https://autocfd.eng.ox.ac.uk to learn more

Why high-fidelity CFD? – Aerospace

4th High Lift Prediction Workshop

Advance the state of the art for CFD with Boeing/NASA and many others

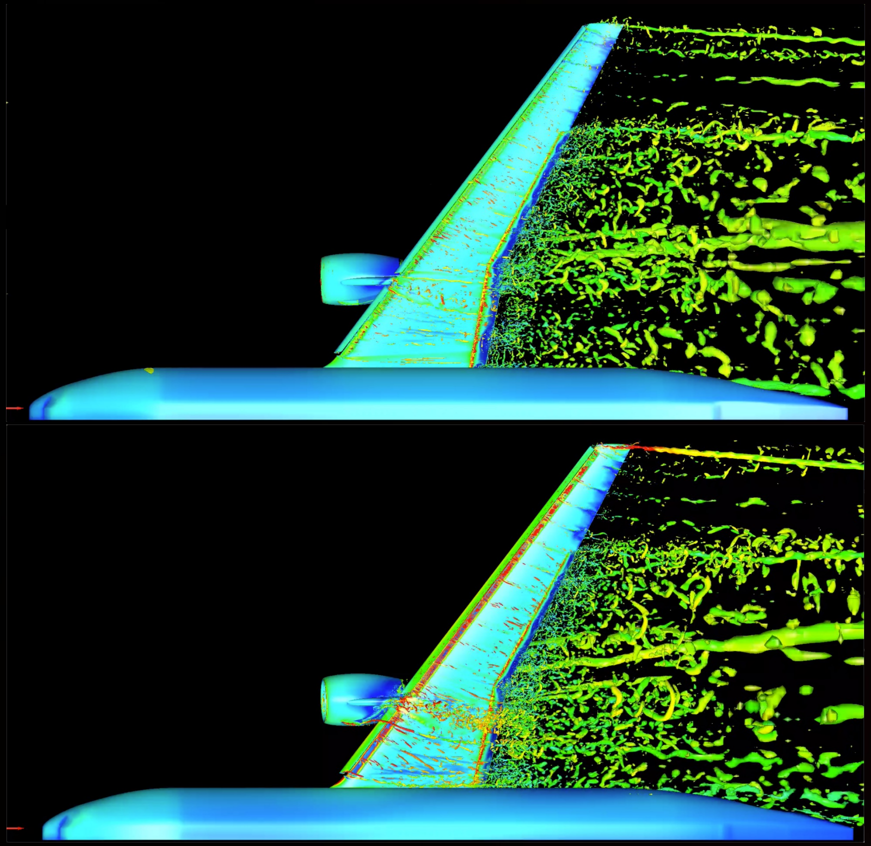

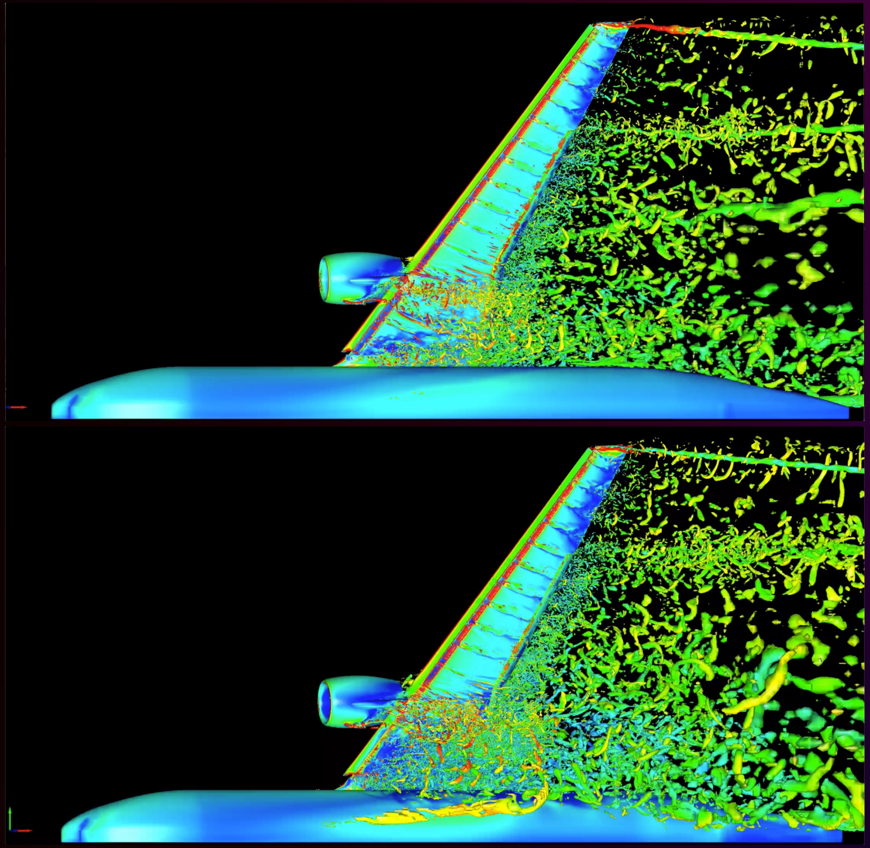

https://hiliftpw.larc.nasa.gov to learn more

Why high-fidelity CFD?

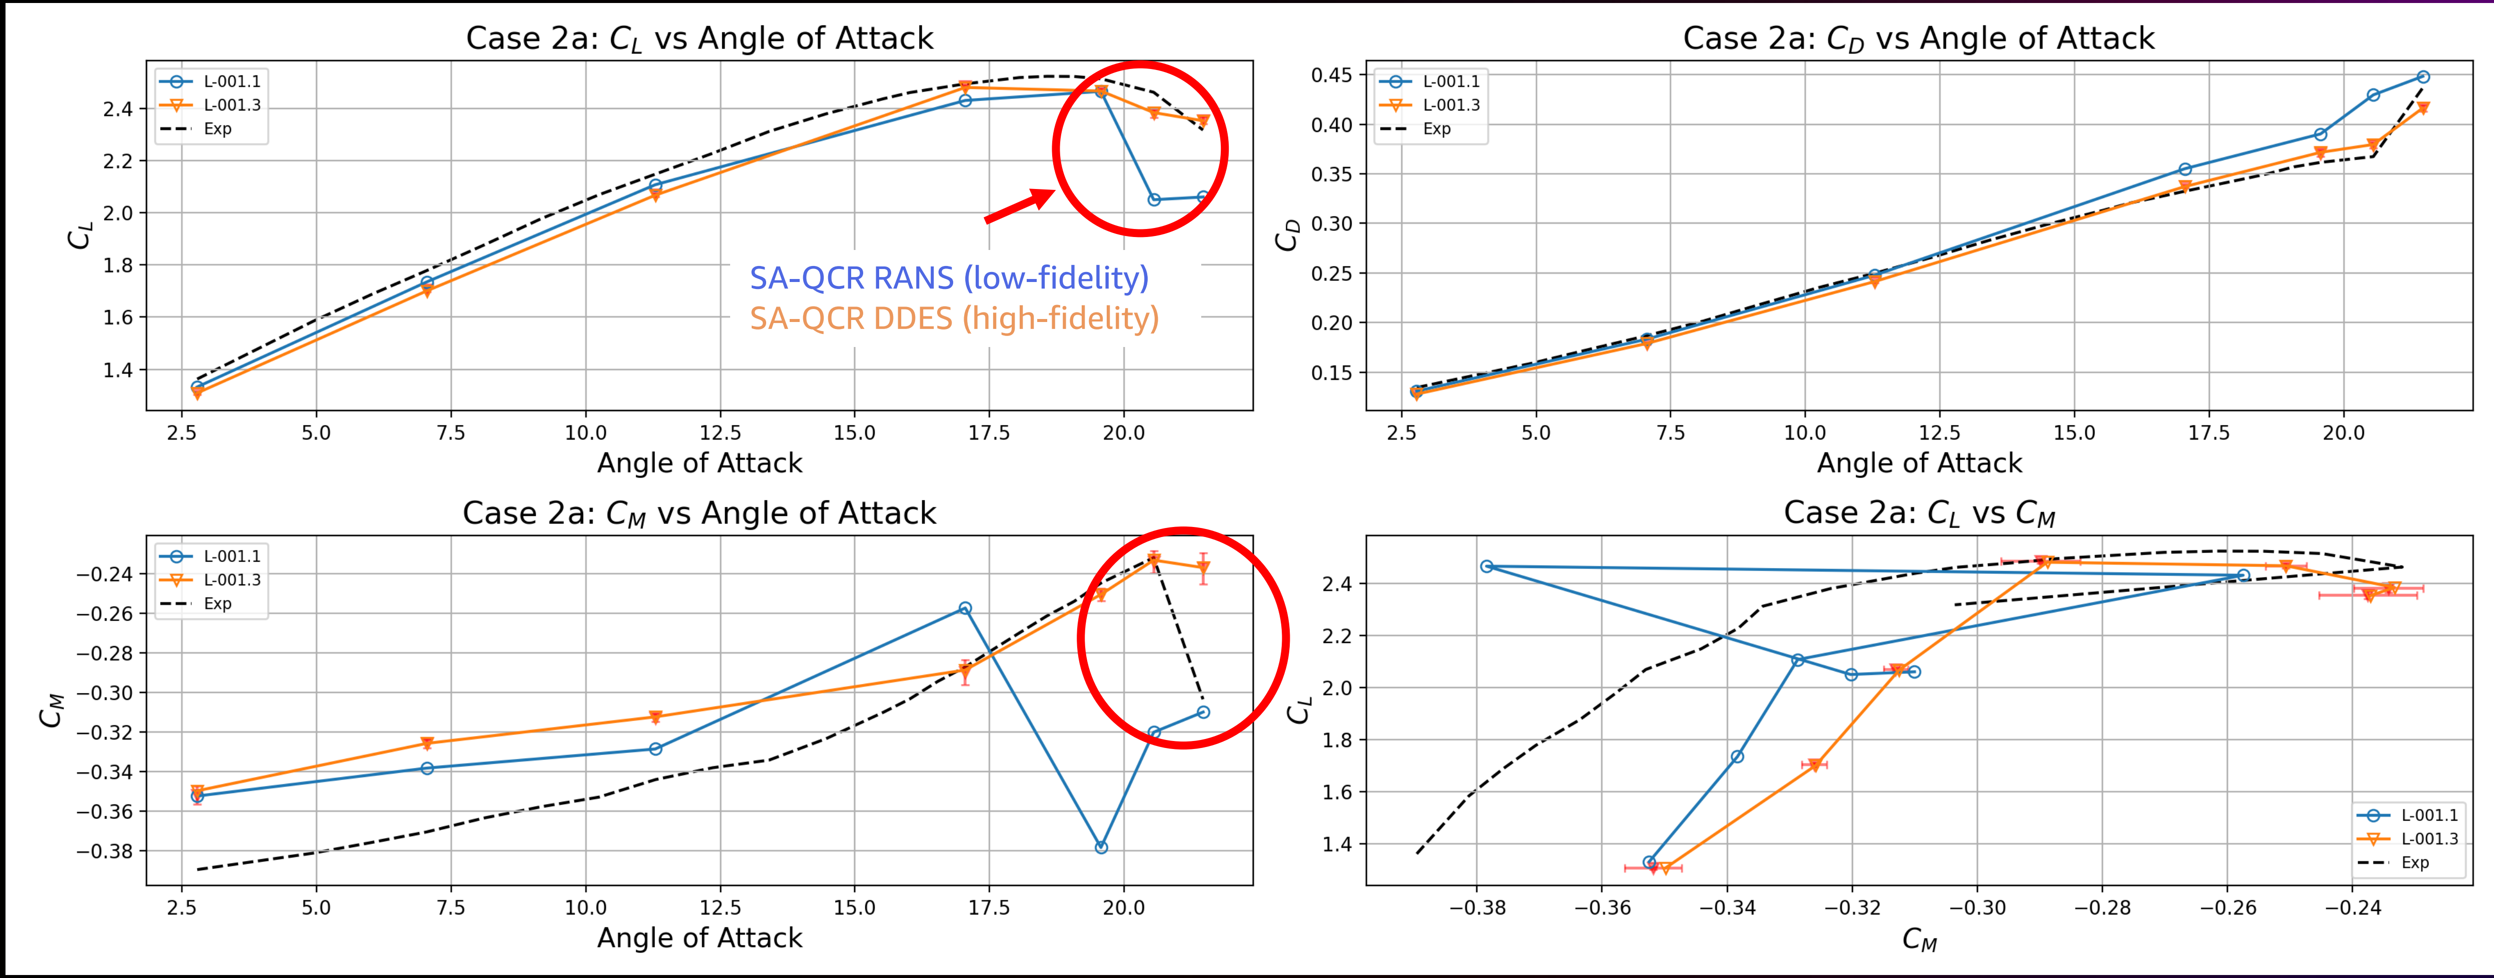

Supporting evidence

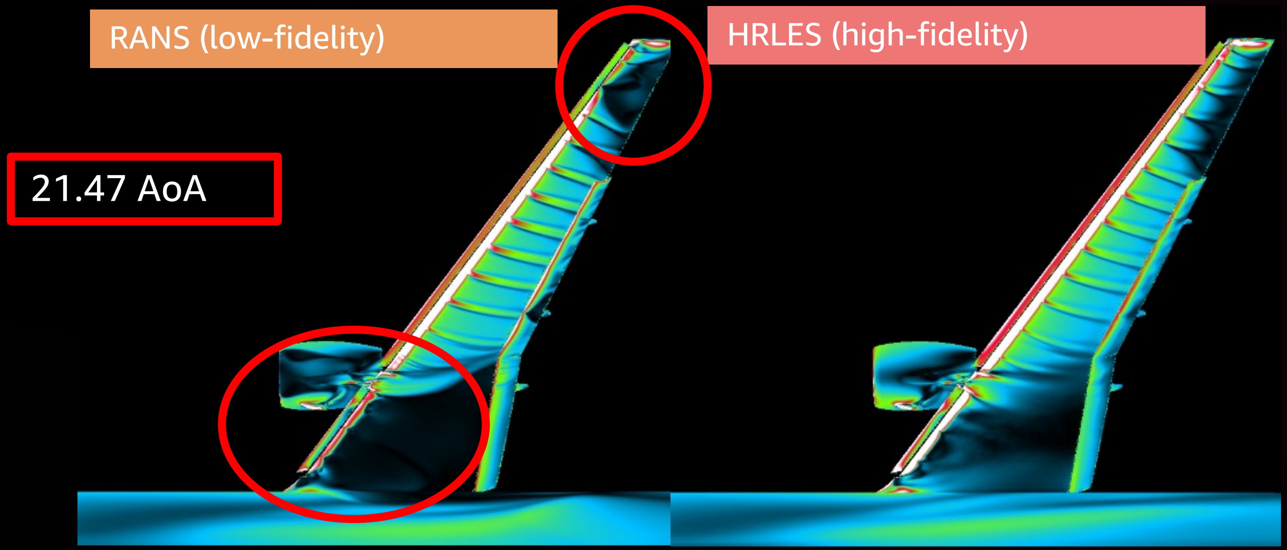

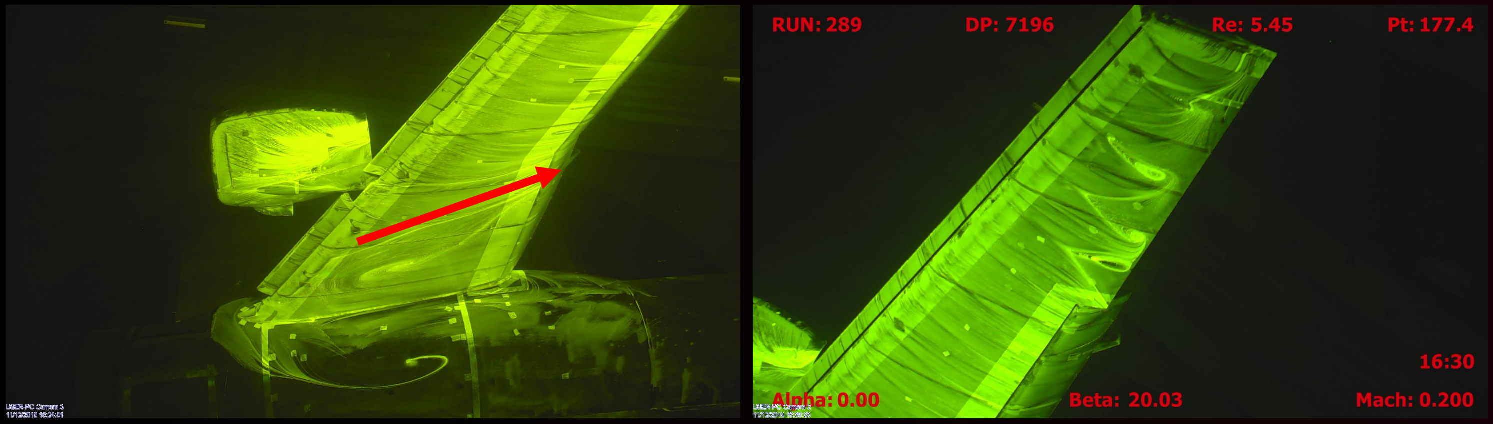

Improved accuracy over low-fidelity methods (for example, RANS)

RANS predicts too much outboard and inboard separation

Transient nature means 5x to 20x more computationally expensive i.e 25k to 500k core-hours. 4000 cores for >12hrs

HPC can be the bottleneck

- As we move to higher fidelity methods to move towards purely digital design and digital certification, jobs are going to need thousands of cores per job or tens of GPUs
- Cloud provides the agility and capacity as you need ever greater HPC
- But how does this work in practice? What would the day look like for a cloud-enabled aerodynamics engineer? How do you make sure you have the optimum setup?
- Lets look at the overall architecture and then a live interactive demo!

How do we run CFD on AWS? Reference architecture

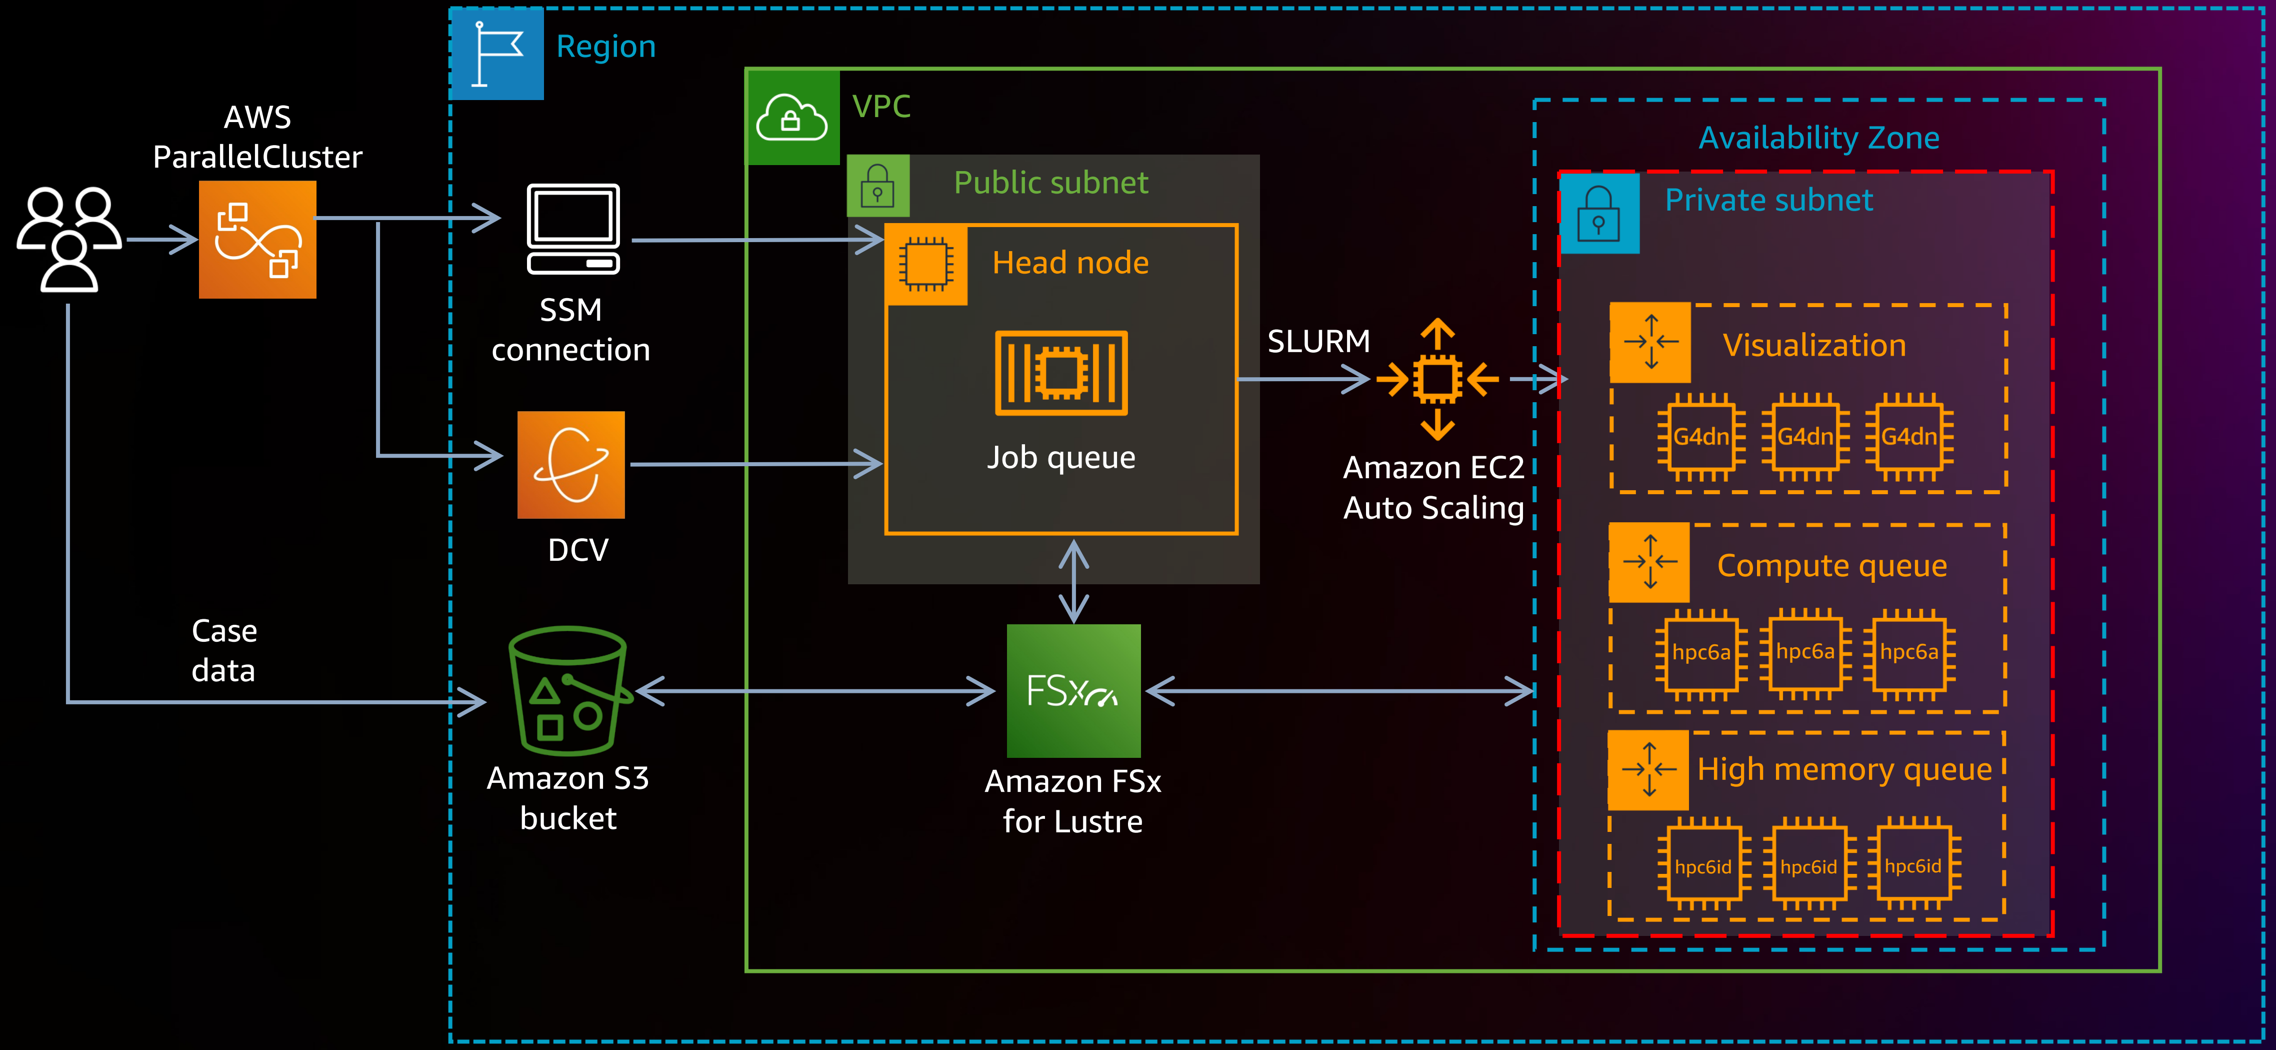

Demo

Instance type selection

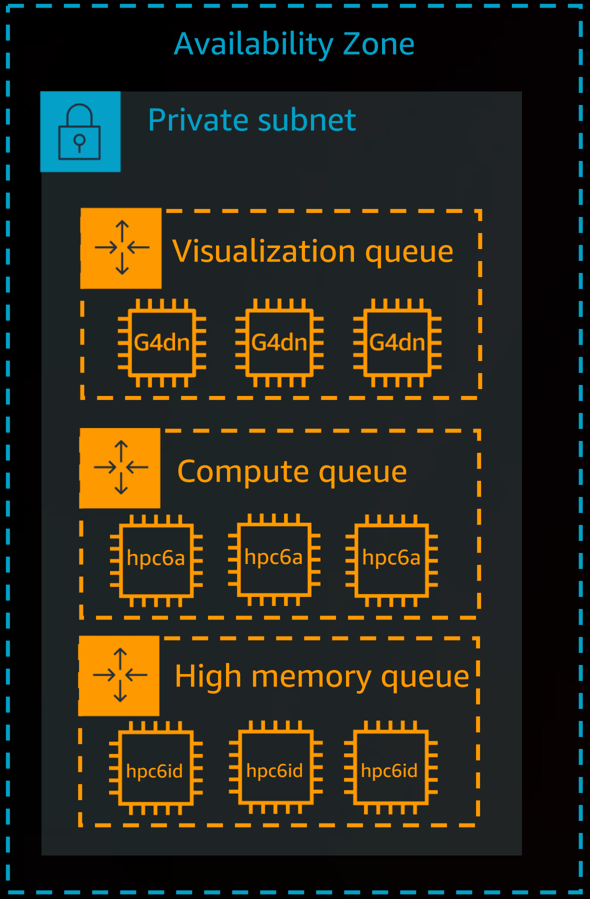

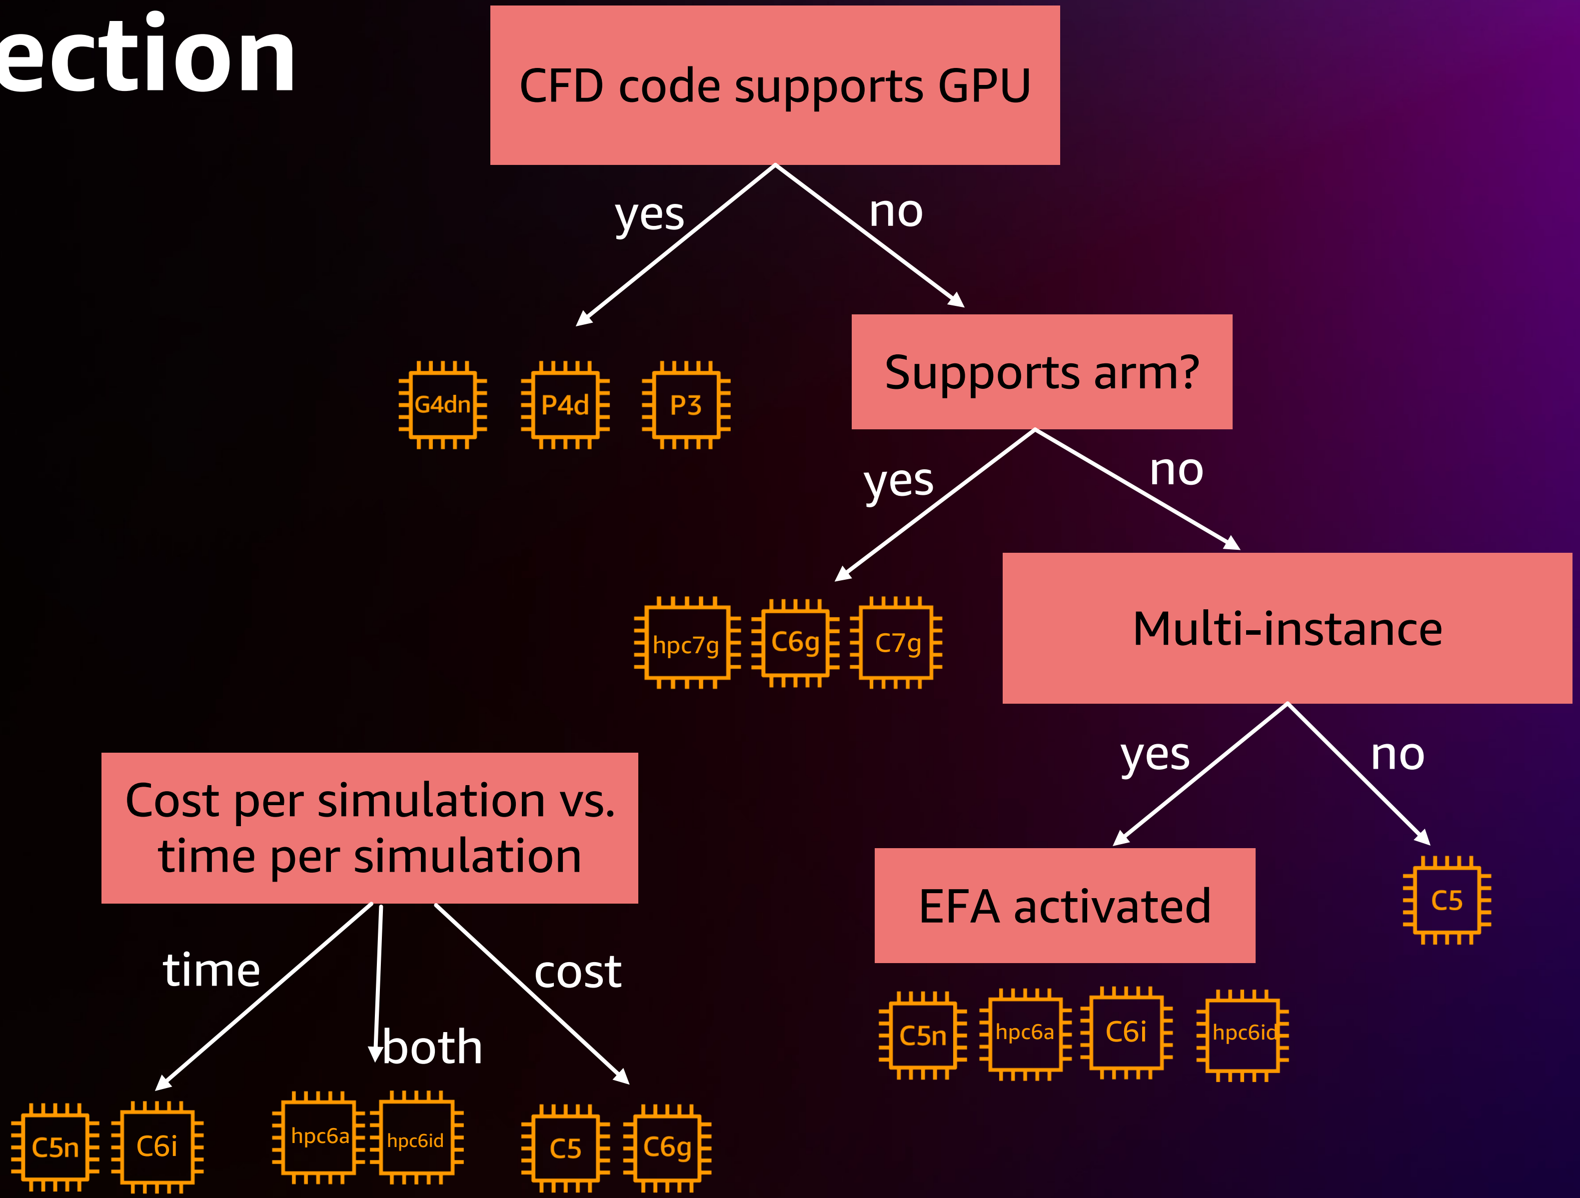

Multi-node

- 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

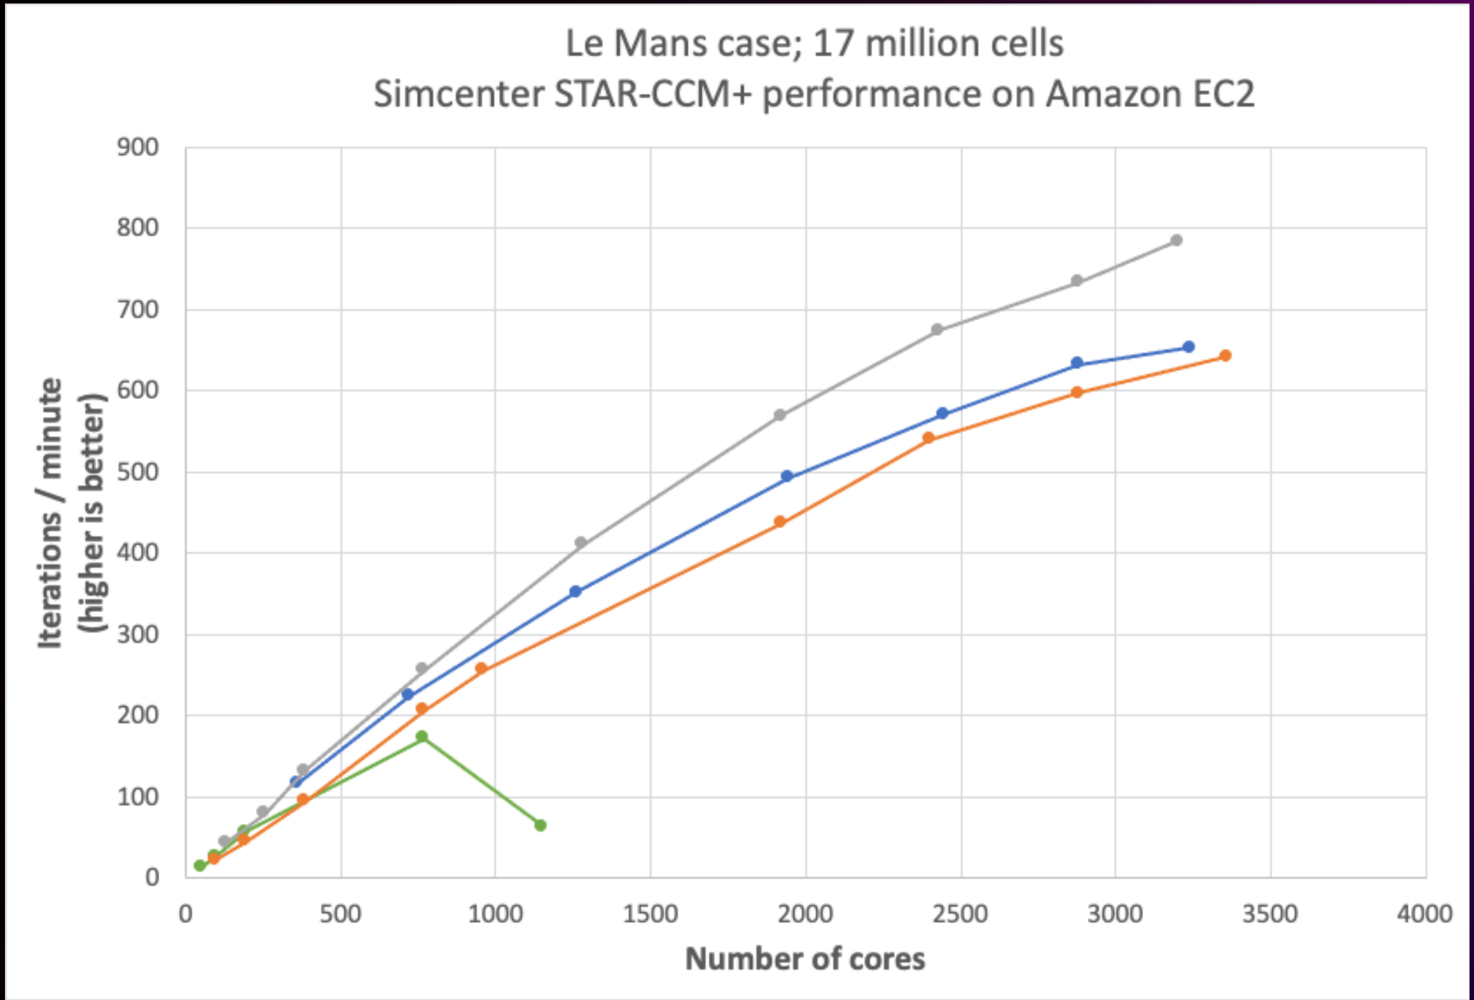

Cells per core

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

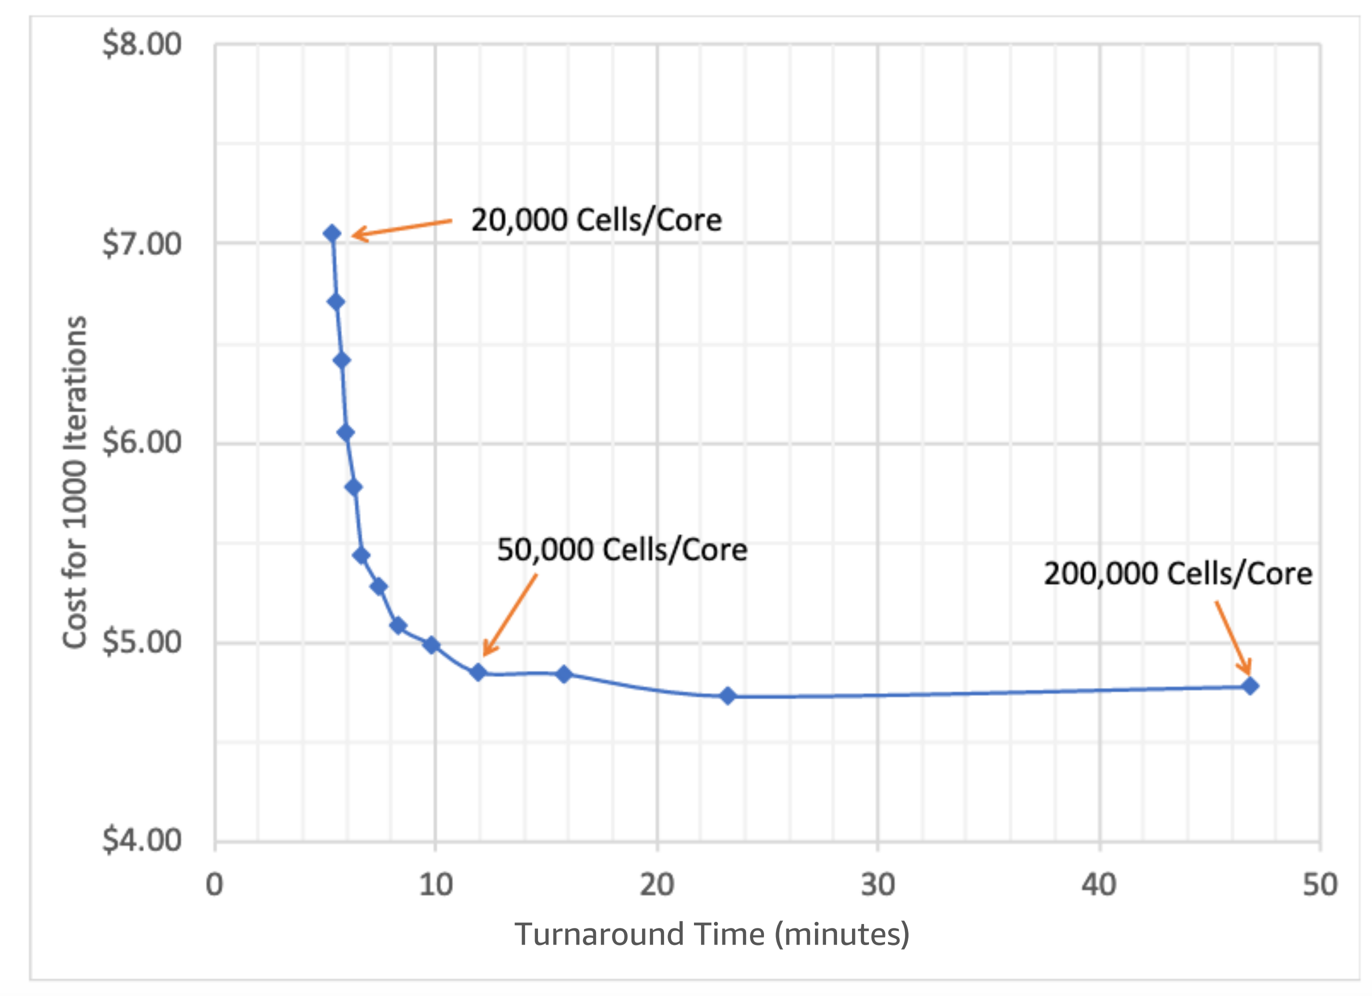

Time to solution vs. cost per simulation

- 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 Amazon 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

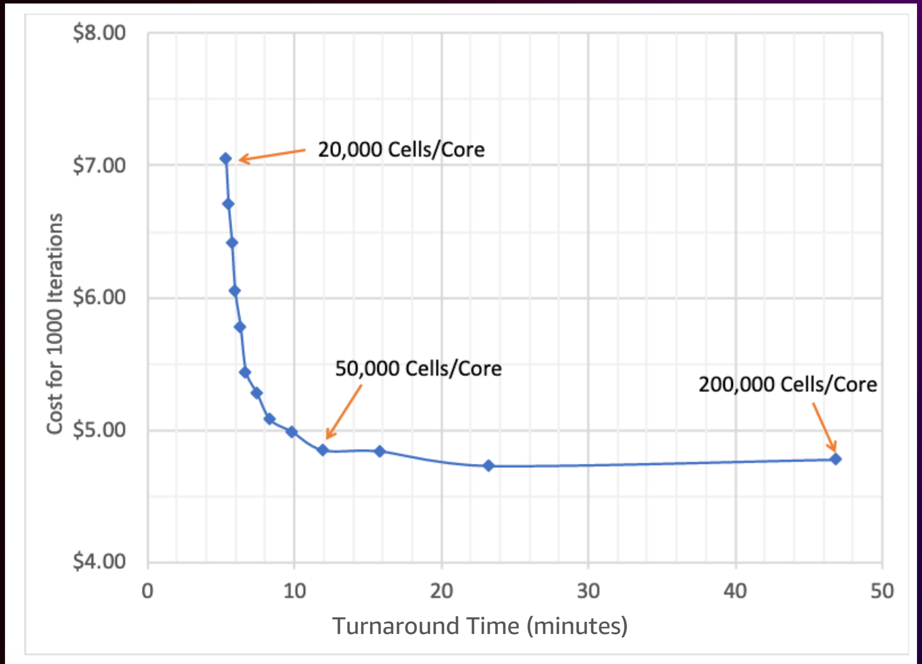

Cost models

Scale using Spot, On-Demand, or both

Use Reserved Instances/Saving Plans for known/steady-state workloads

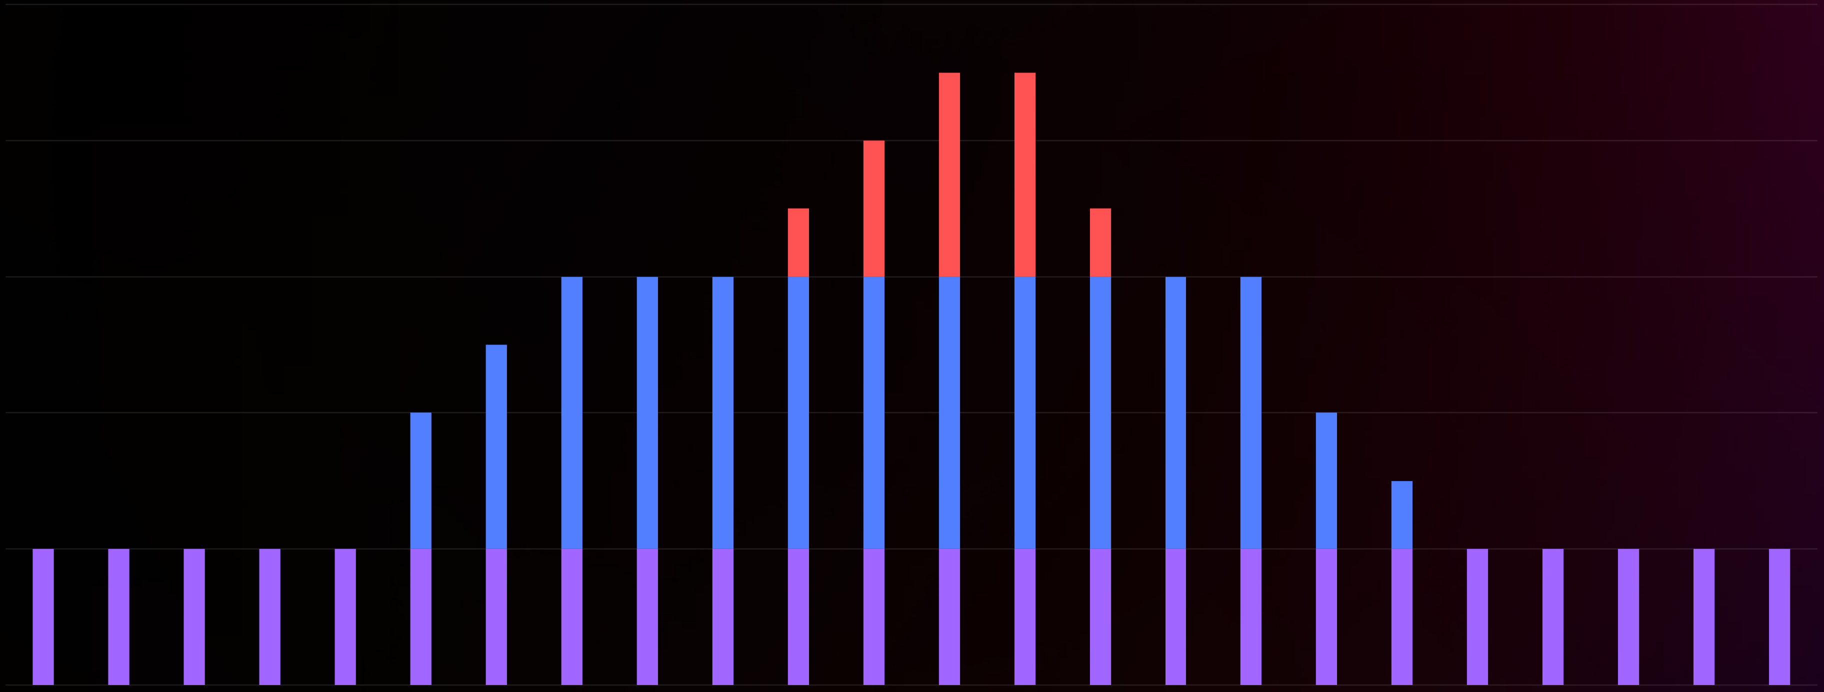

Key resources

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

DCV queues
https://swsmith.cc/posts/dcv-visualization-queue.html

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/

Q&A

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!

Dr. Neil Ashton
neashton@amazon.com

Sean Smith
seaam@amazon.com

Please complete the session survey in the **mobile app**

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

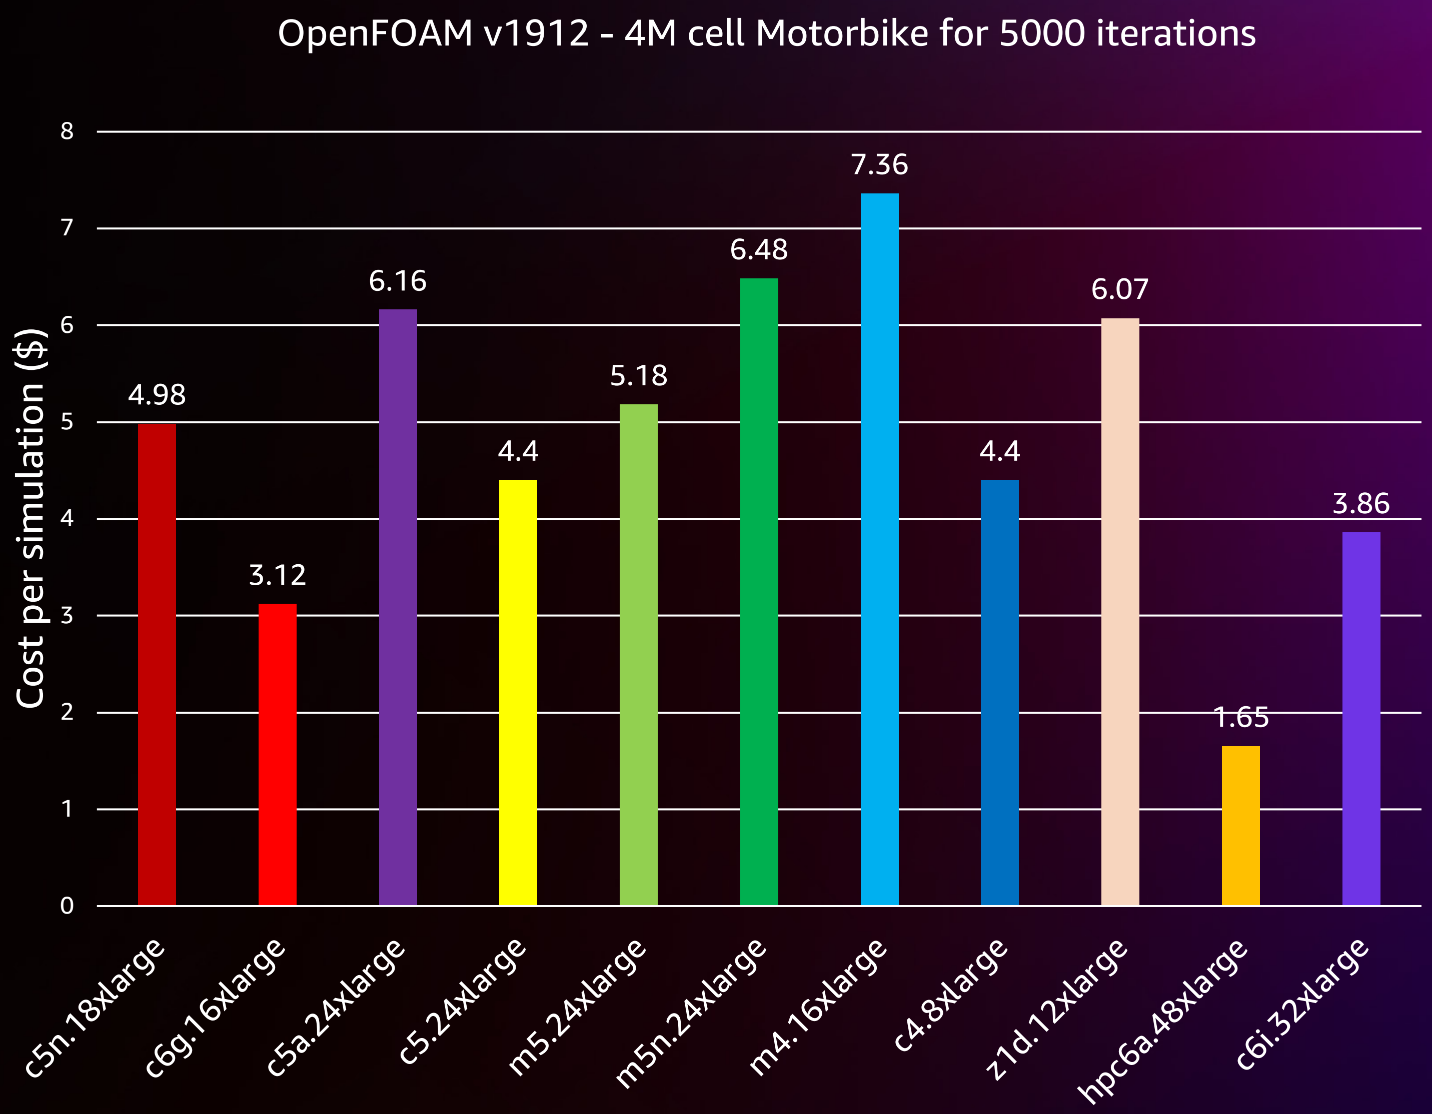

Scaling on AWS – OpenFOAM

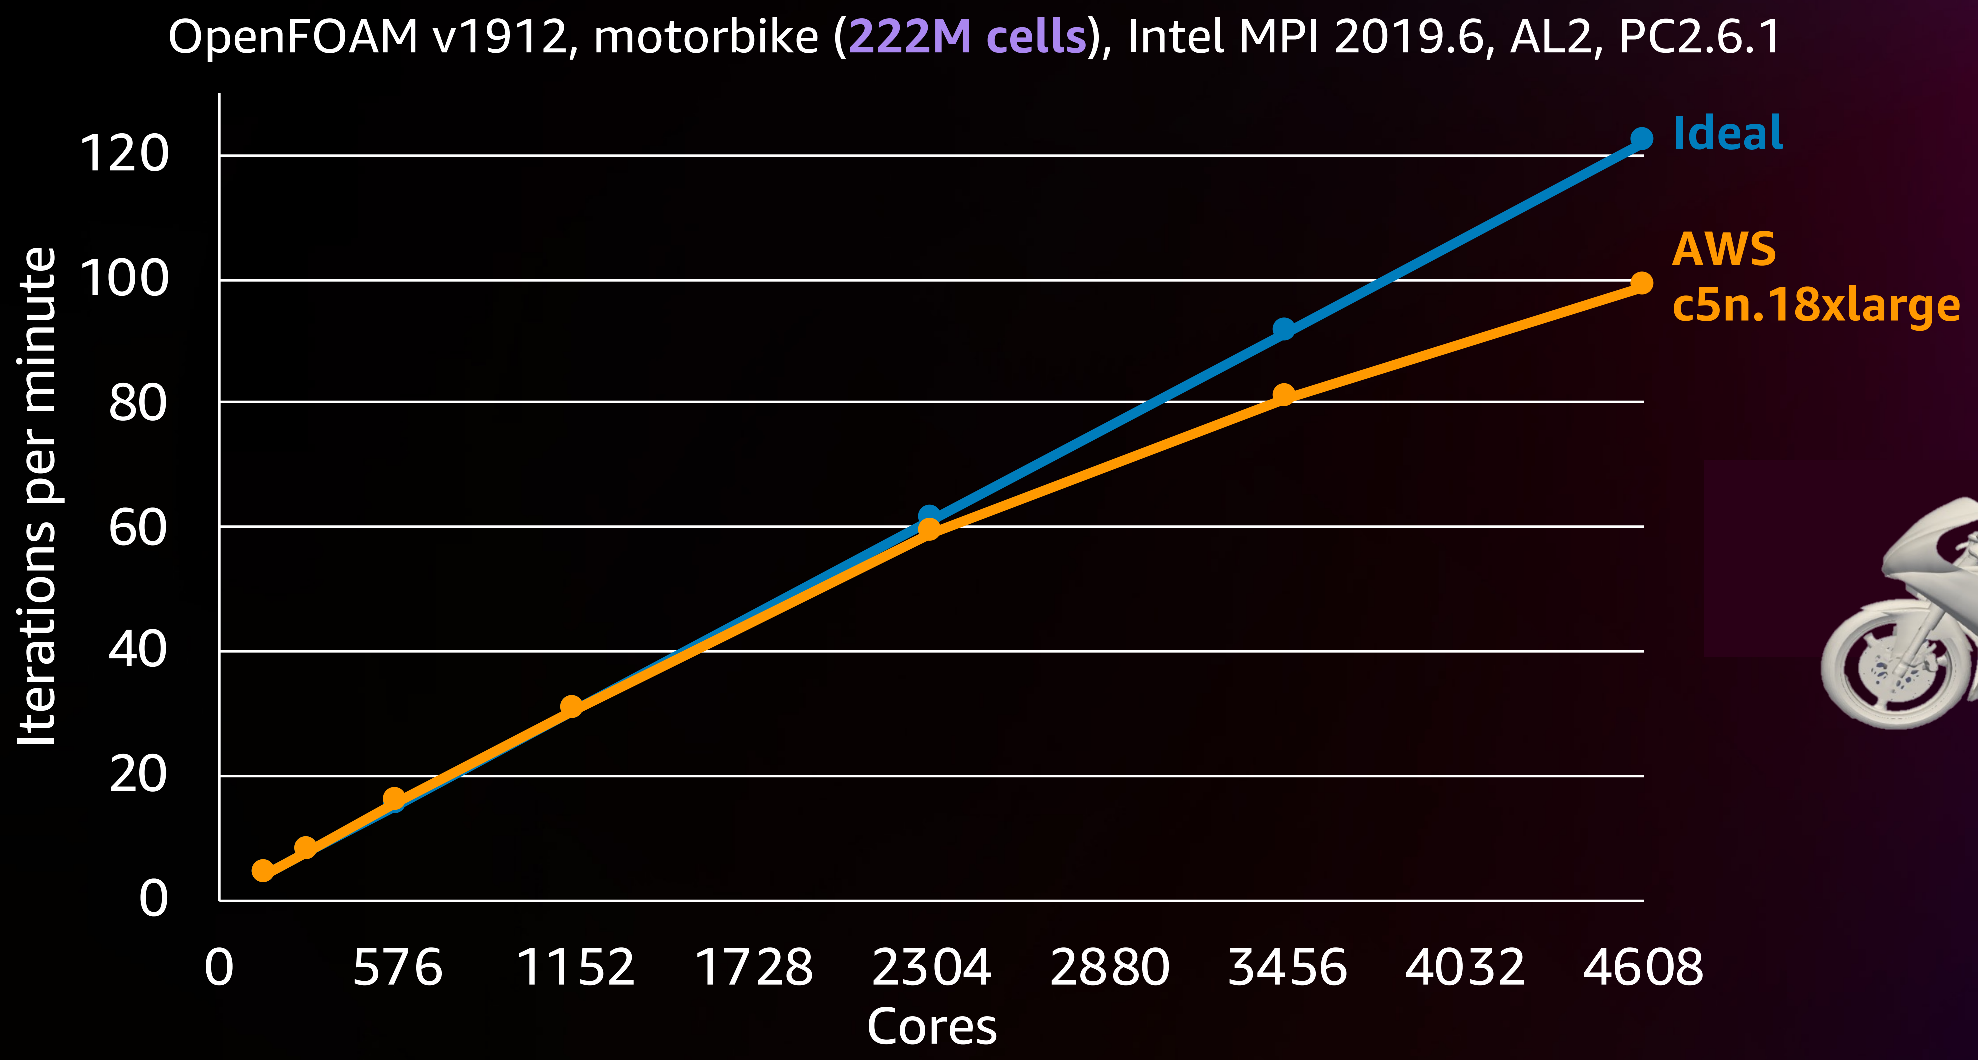

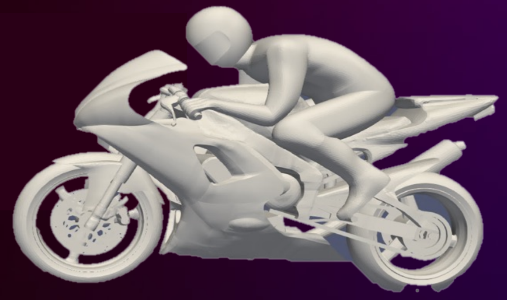

Scaling on AWS – STAR-CCM+

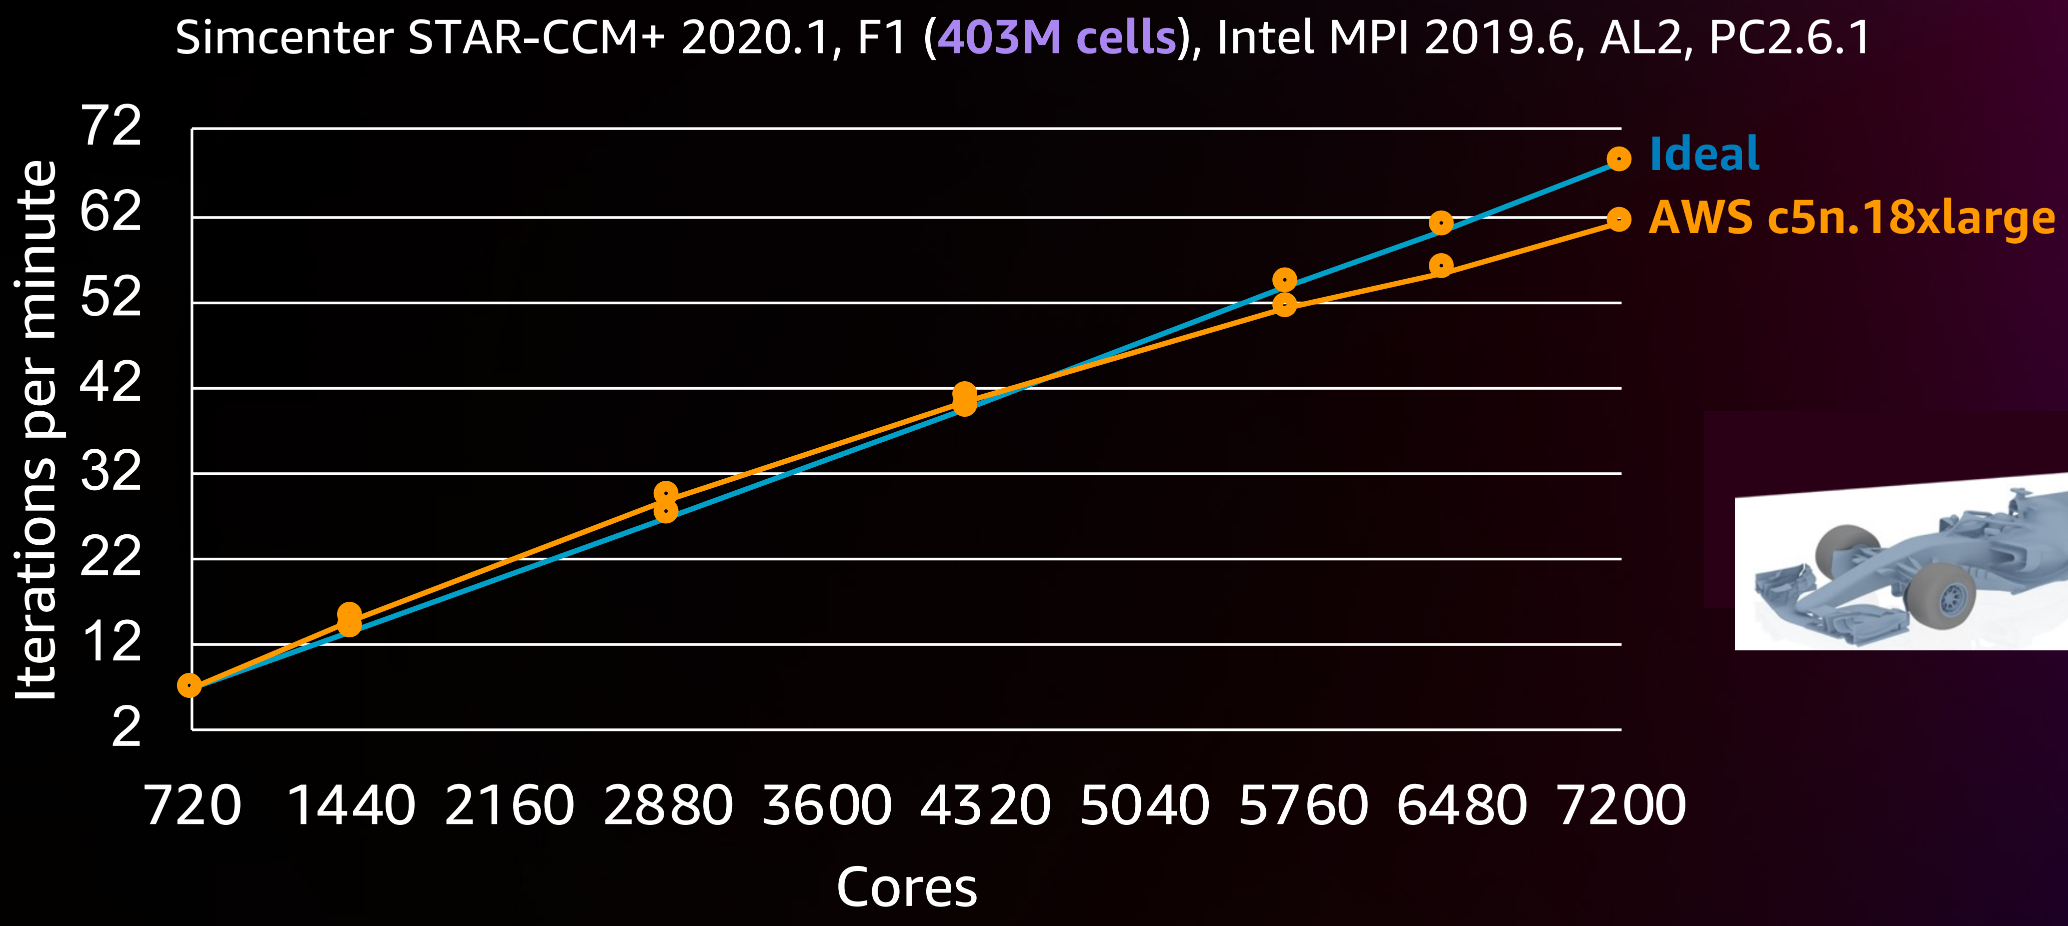

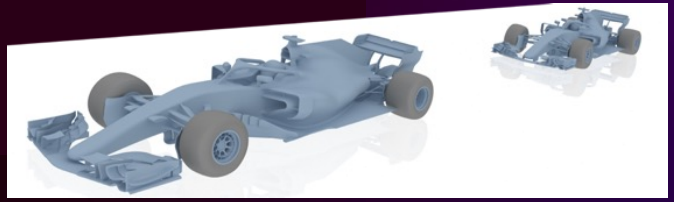

Scaling on AWS – Fluent

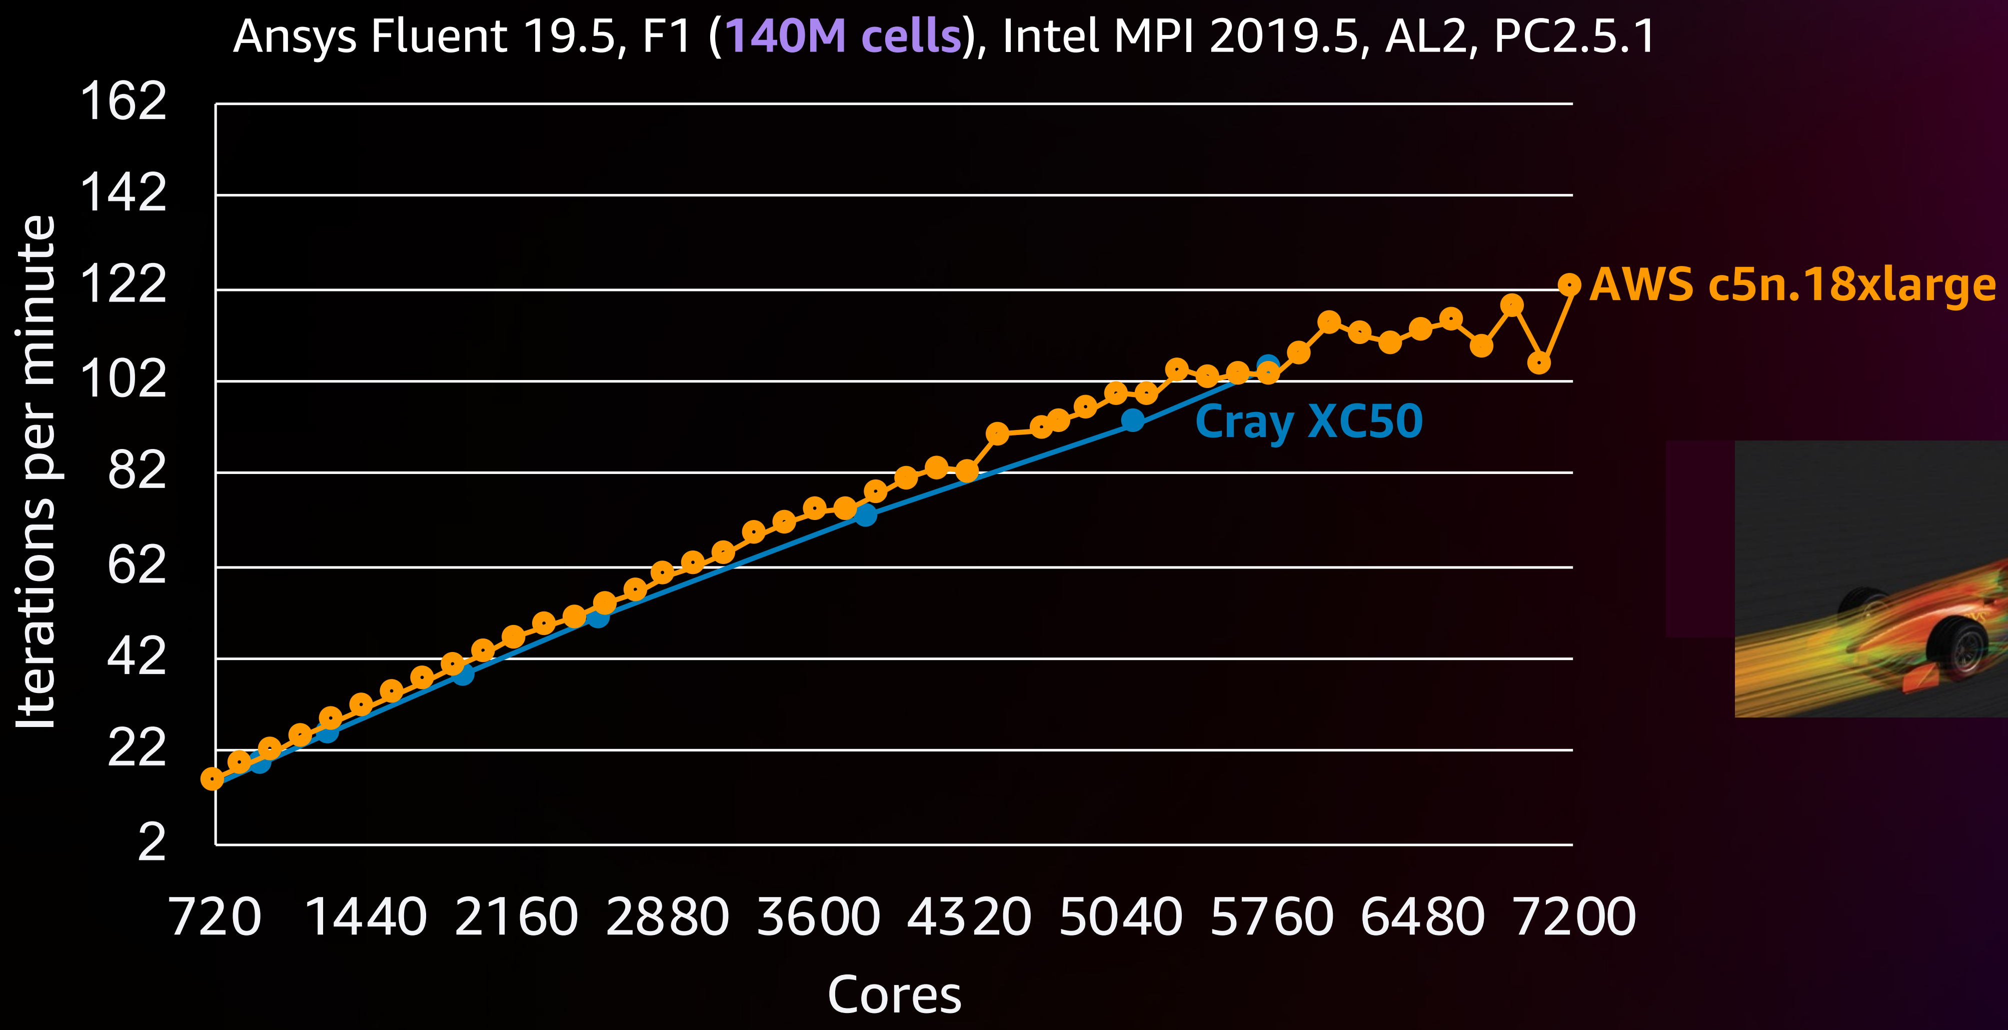

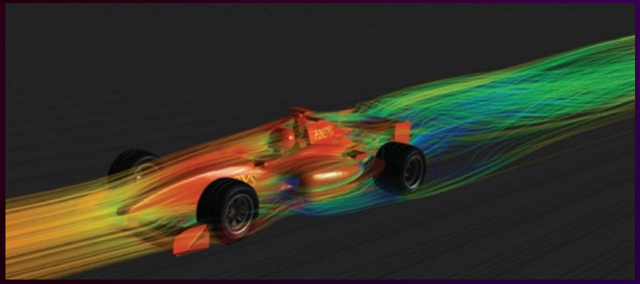

Reference architecture

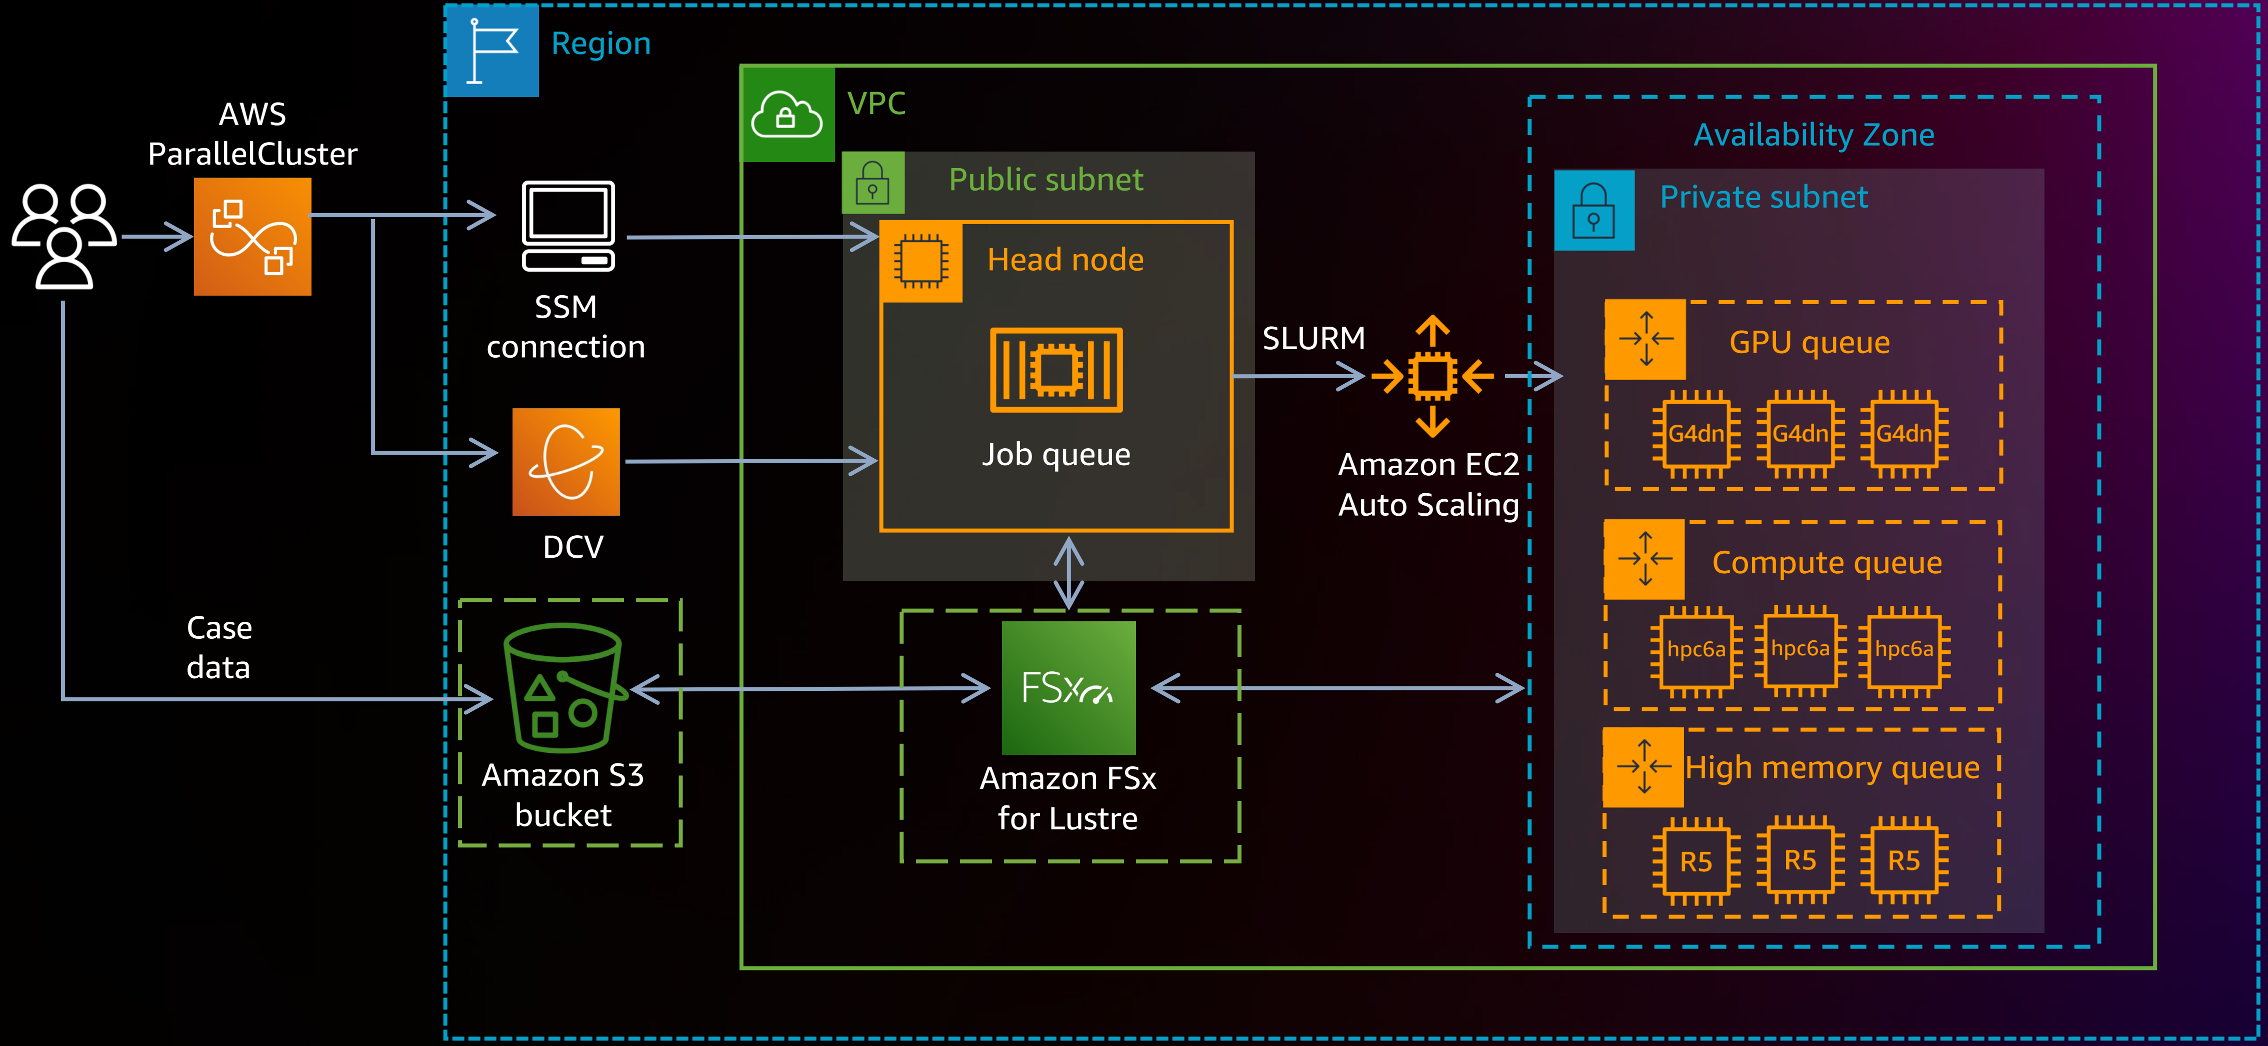

Storage options optimized for price performance

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

Updates on FSx for Lustre are exported to S3

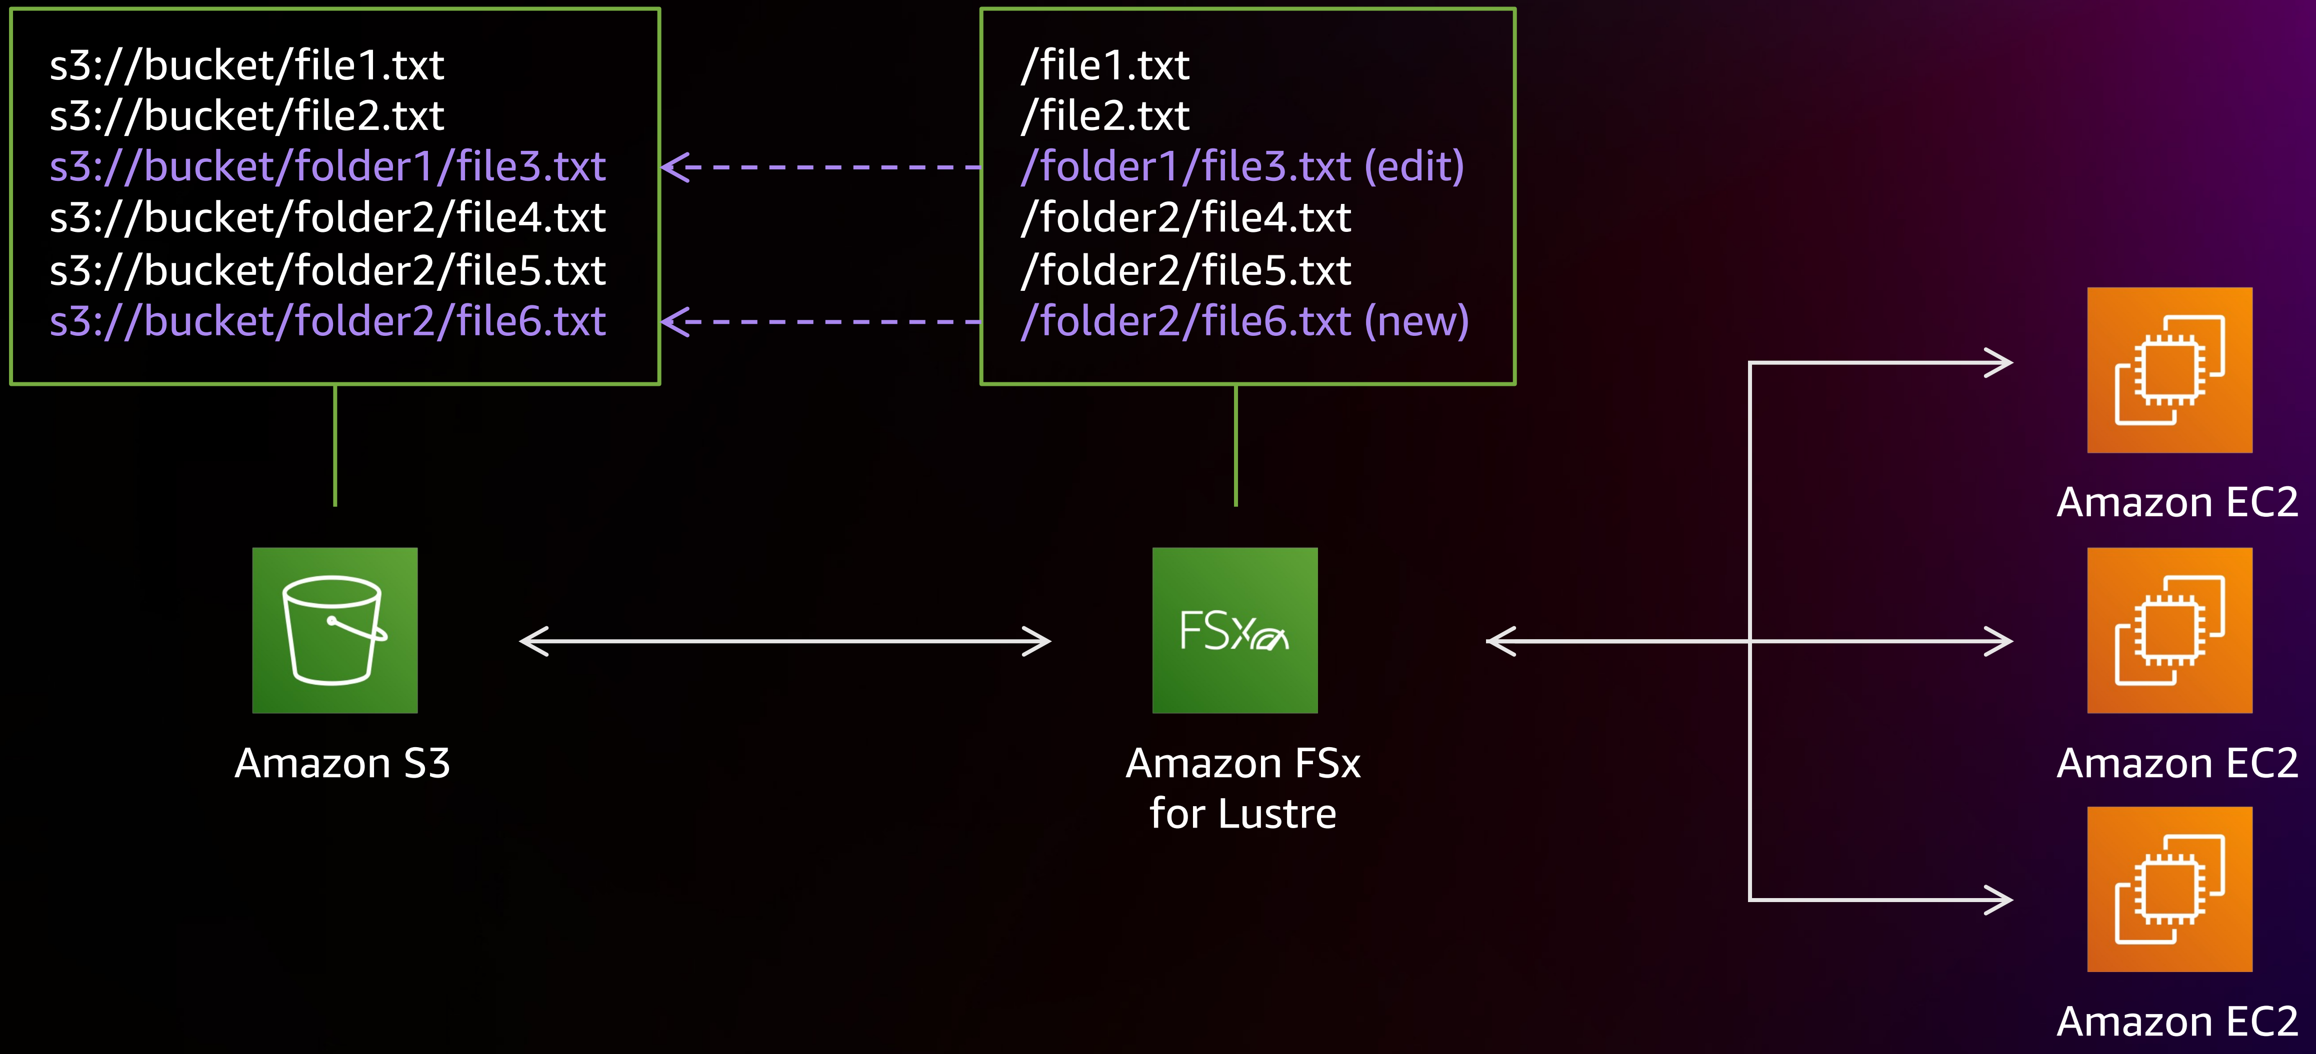

Spin up or down with compute resources

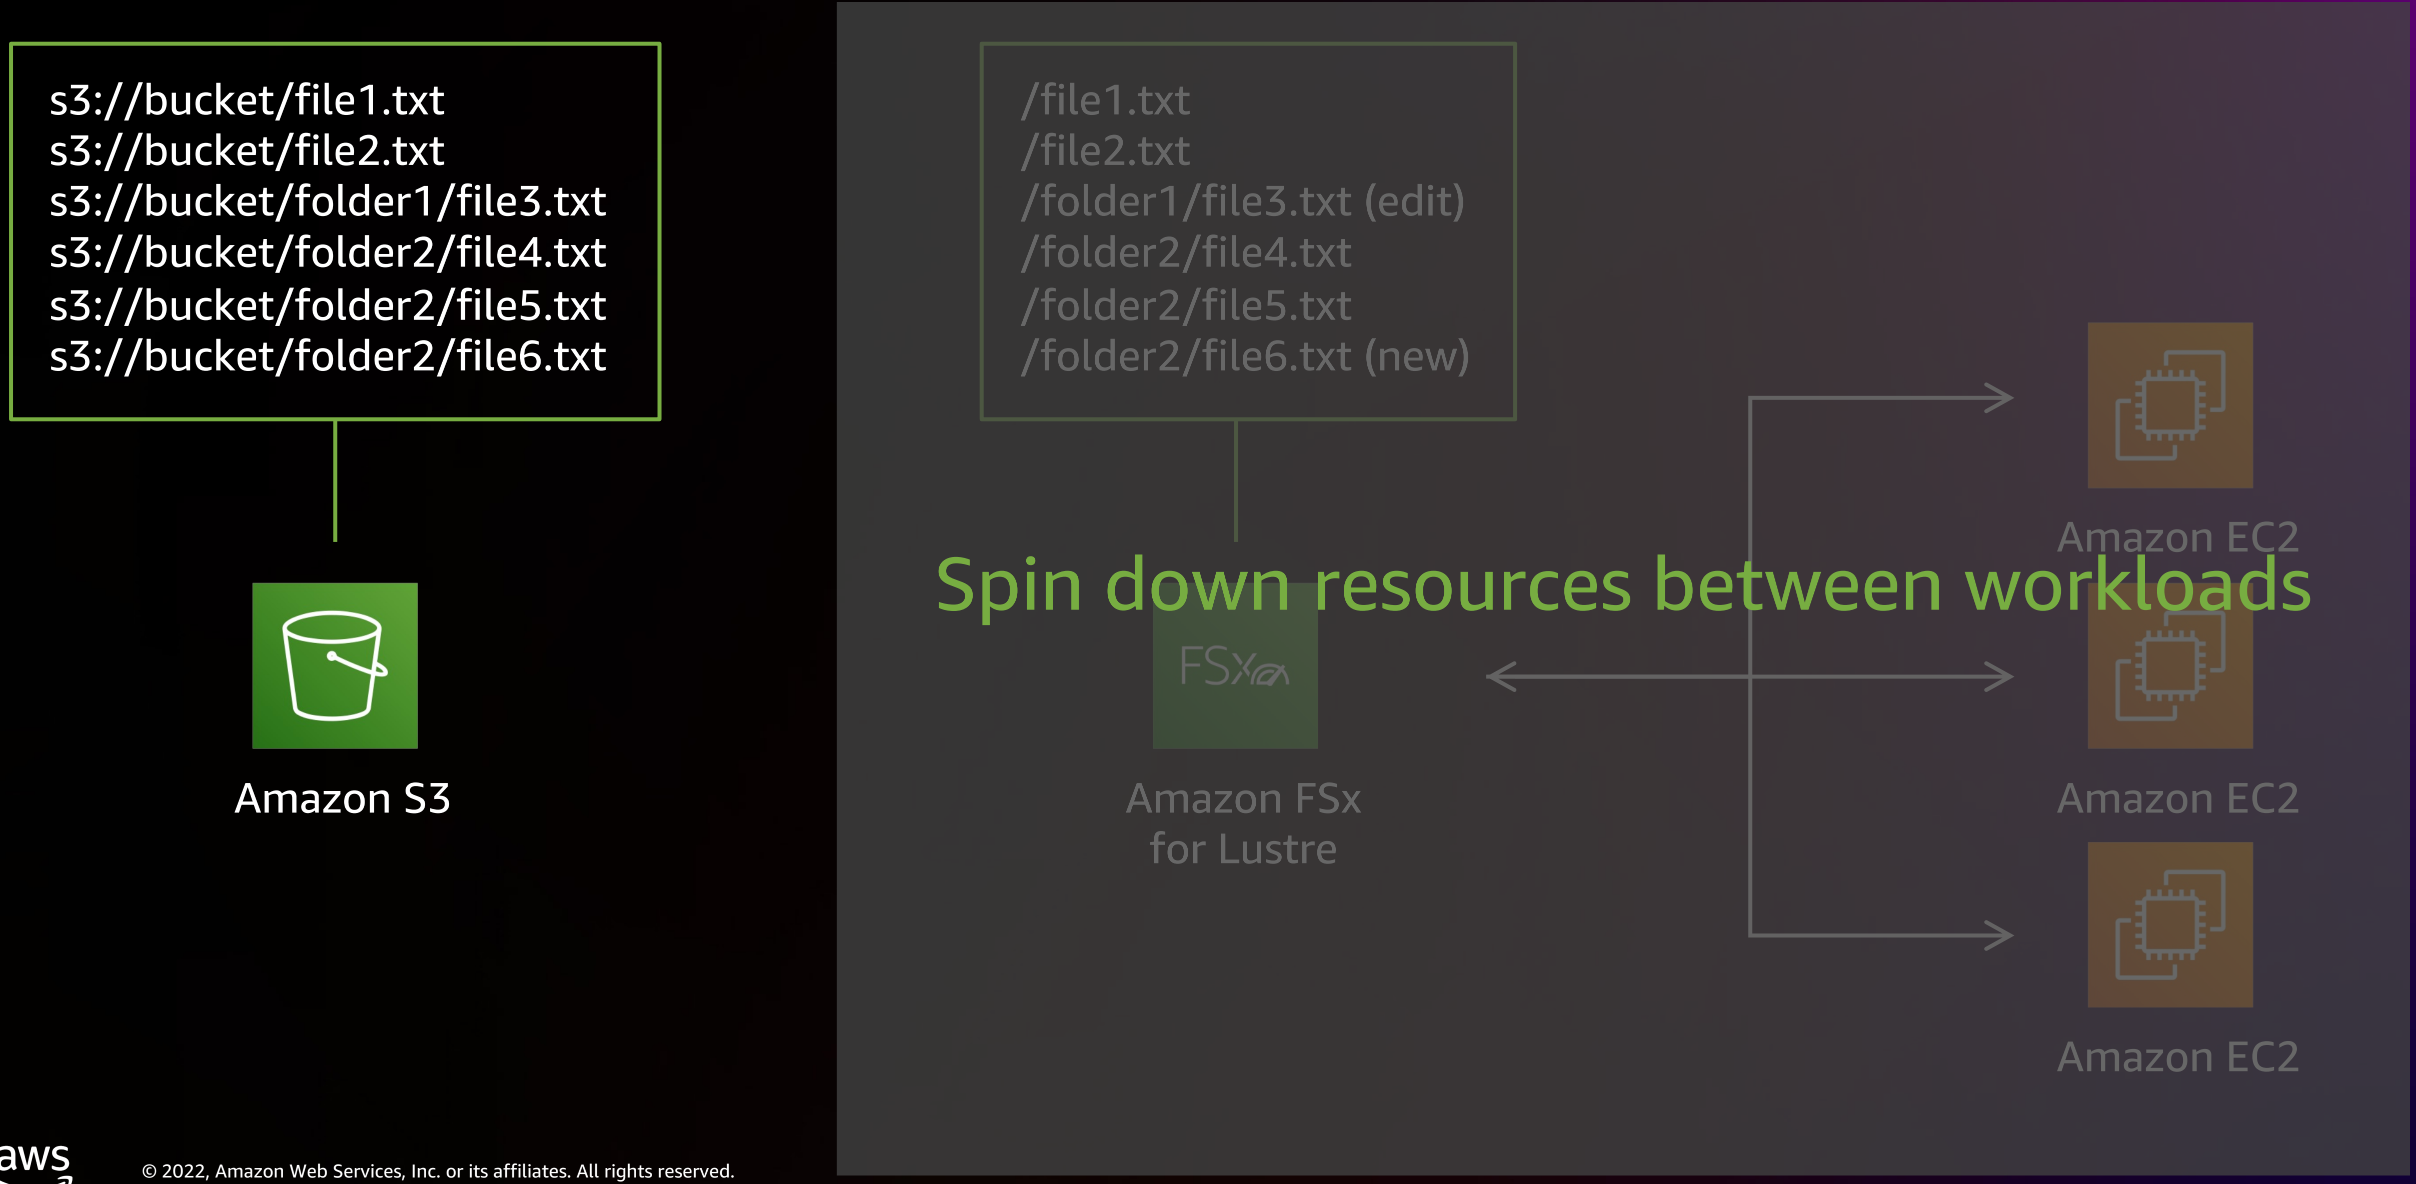

Innovating with ML

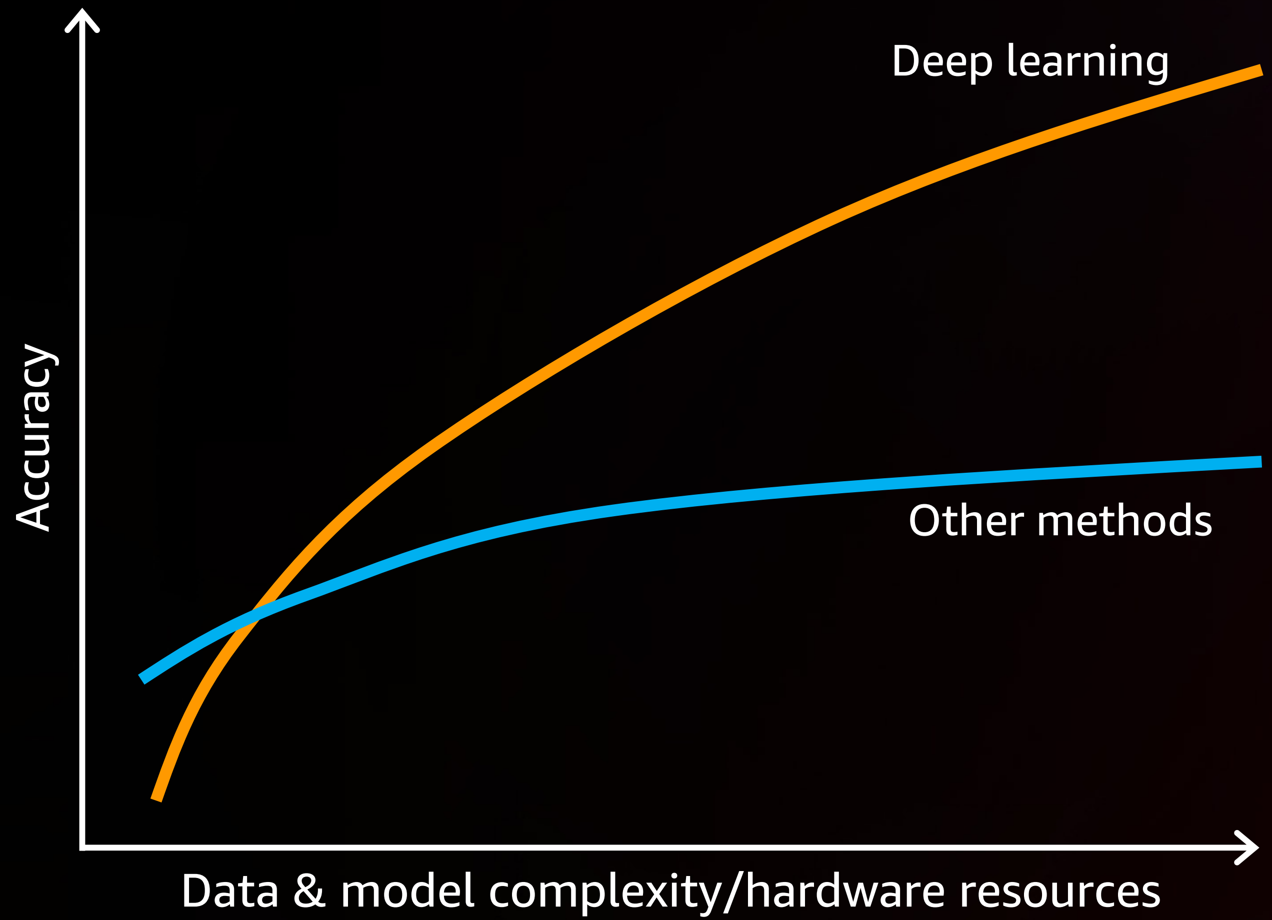

Applications

Natural language processing

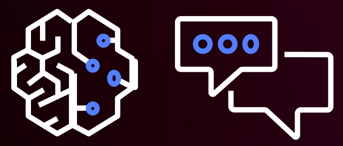

Image/ video analysis

Autonomous vehicle systems

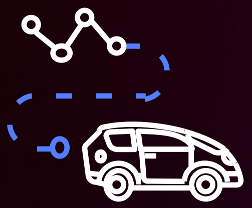

Recommendation systems

Model accuracy continues to increase with additional data

Demonstrating ML and HPC: Extract relationships from HPC simulation results

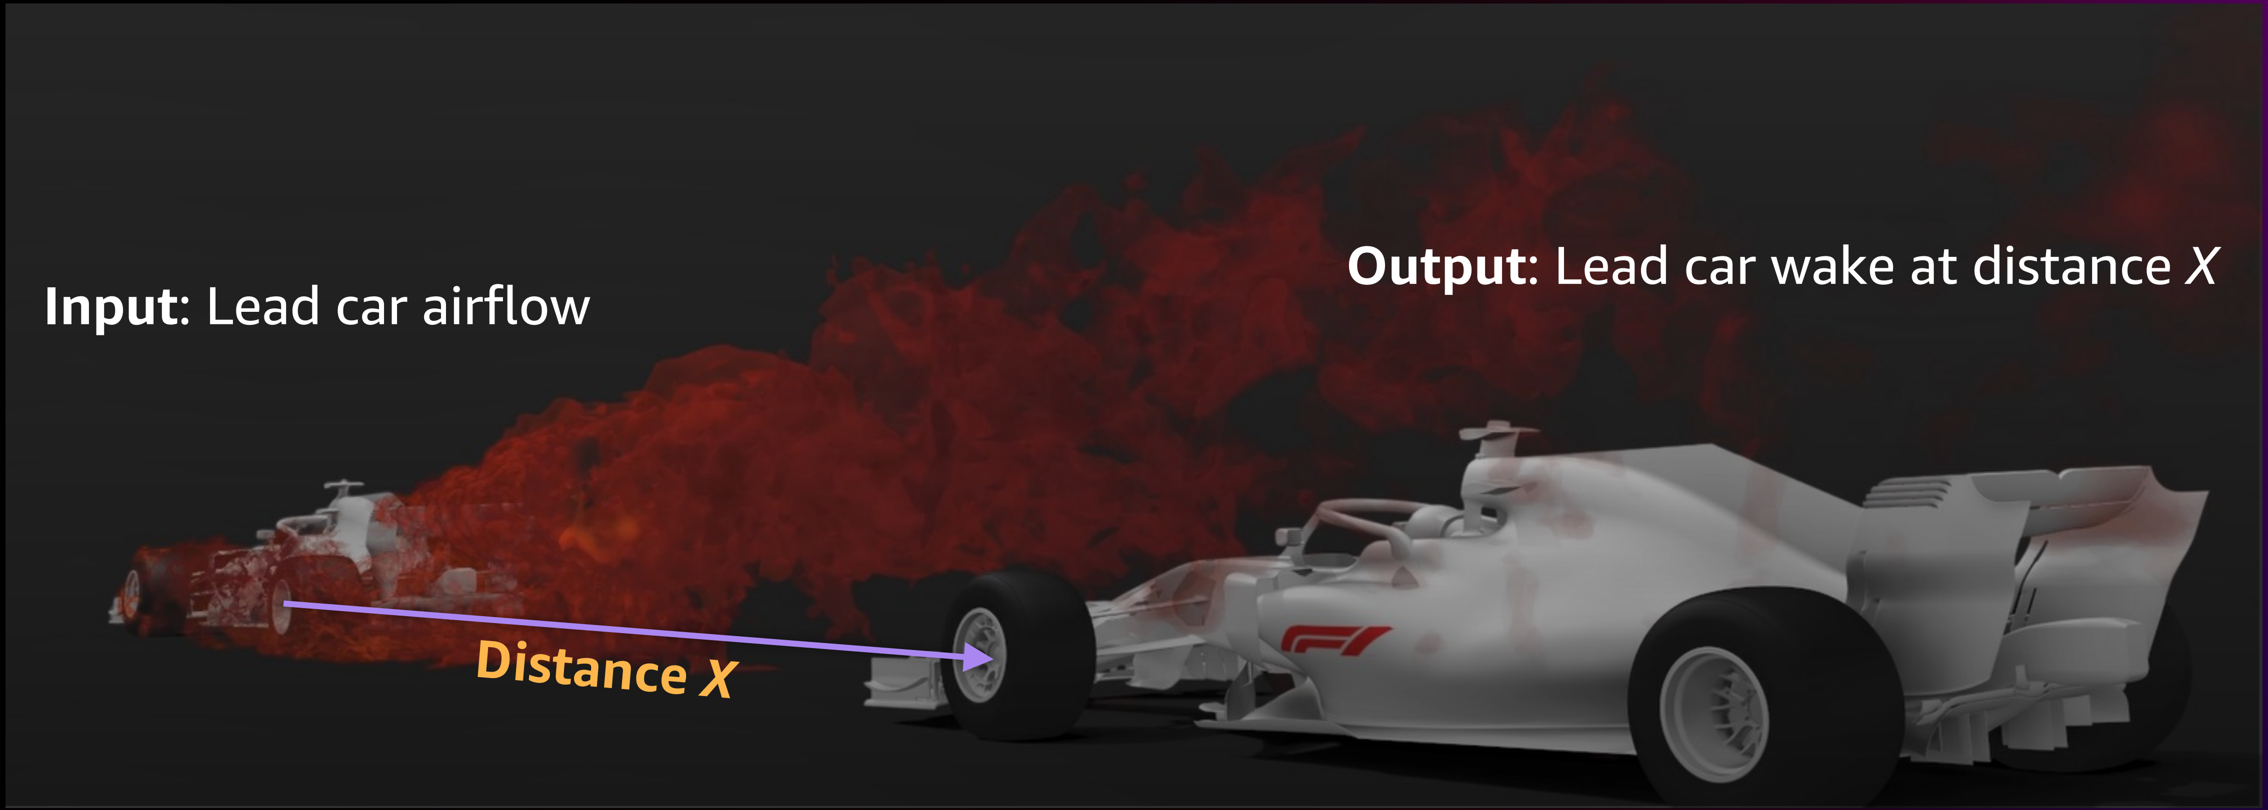

Problem statement:

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

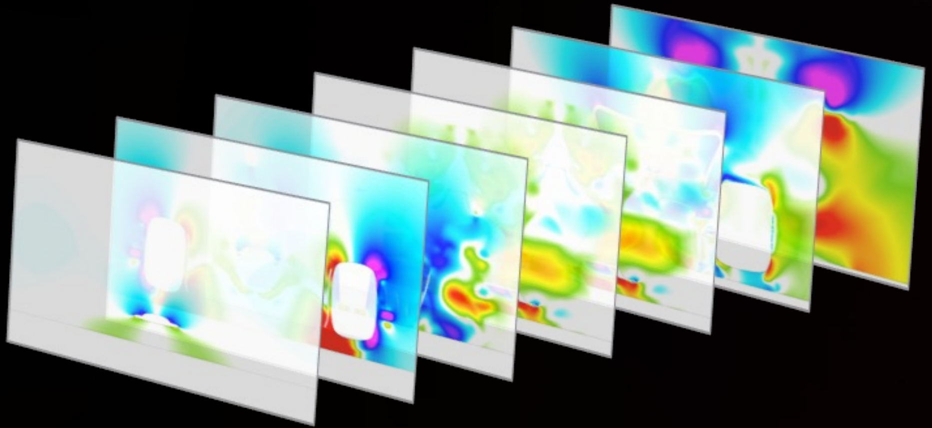

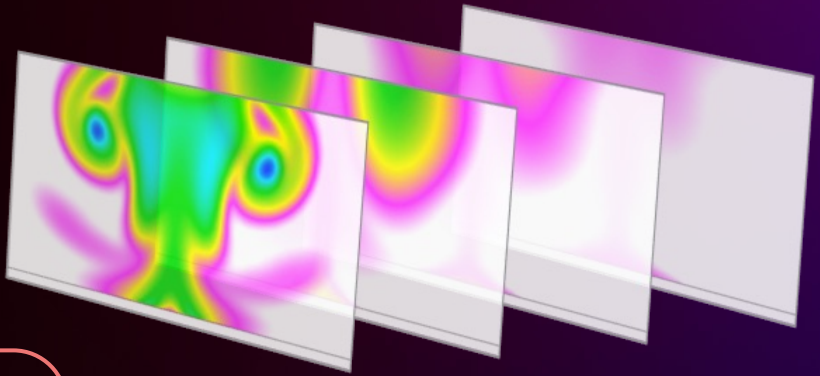

EC2 UltraClusters