
AWS
re:Invent

CMP206

HPC on AWS: Innovating without infrastructure constraints

Ian Colle

General Manager, AWS Batch, HPC, and Visualization

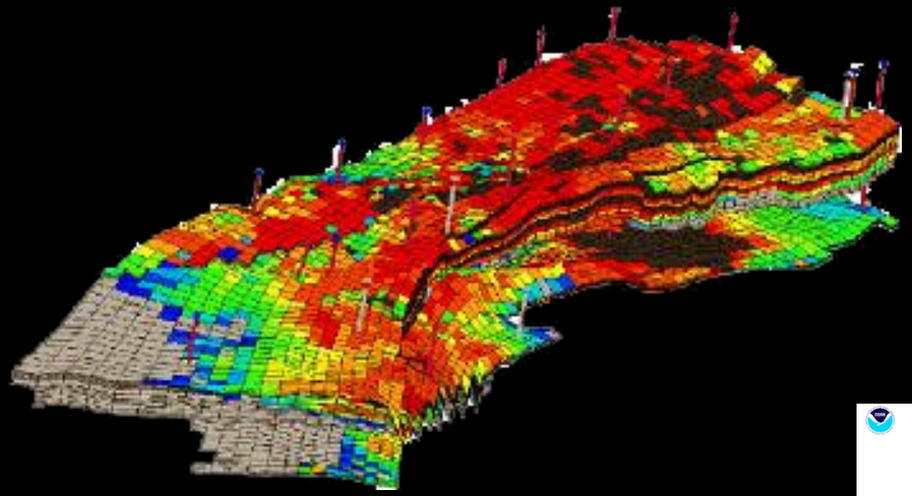
AWS

AWS re:Invent

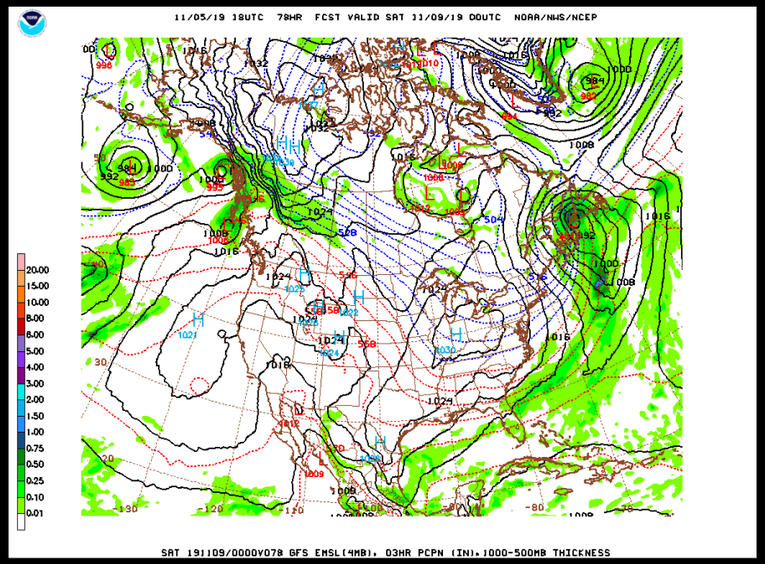
© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Any type of HPC workload is possible on AWS

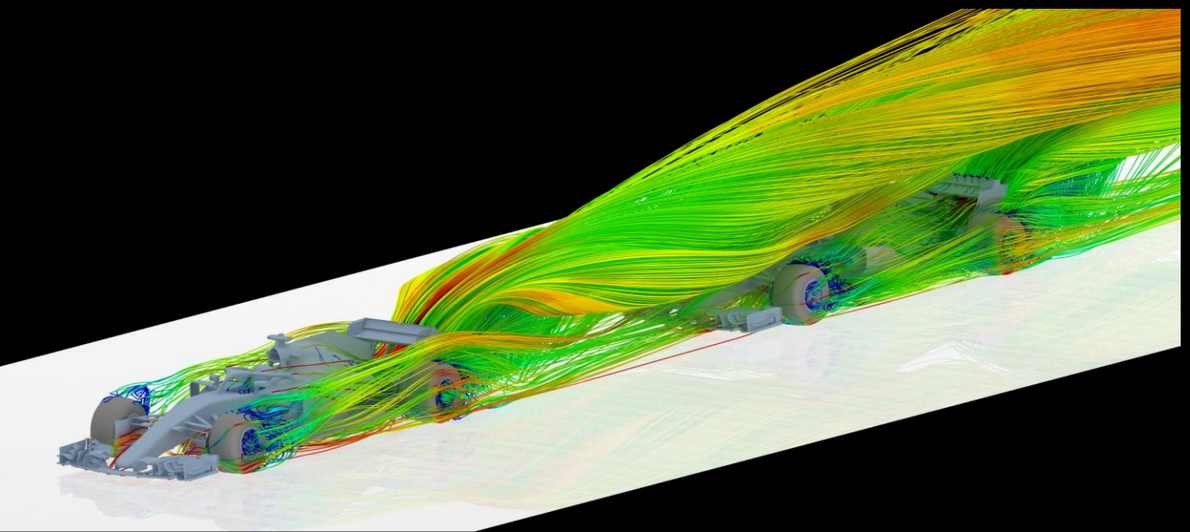
Reservoir simulations



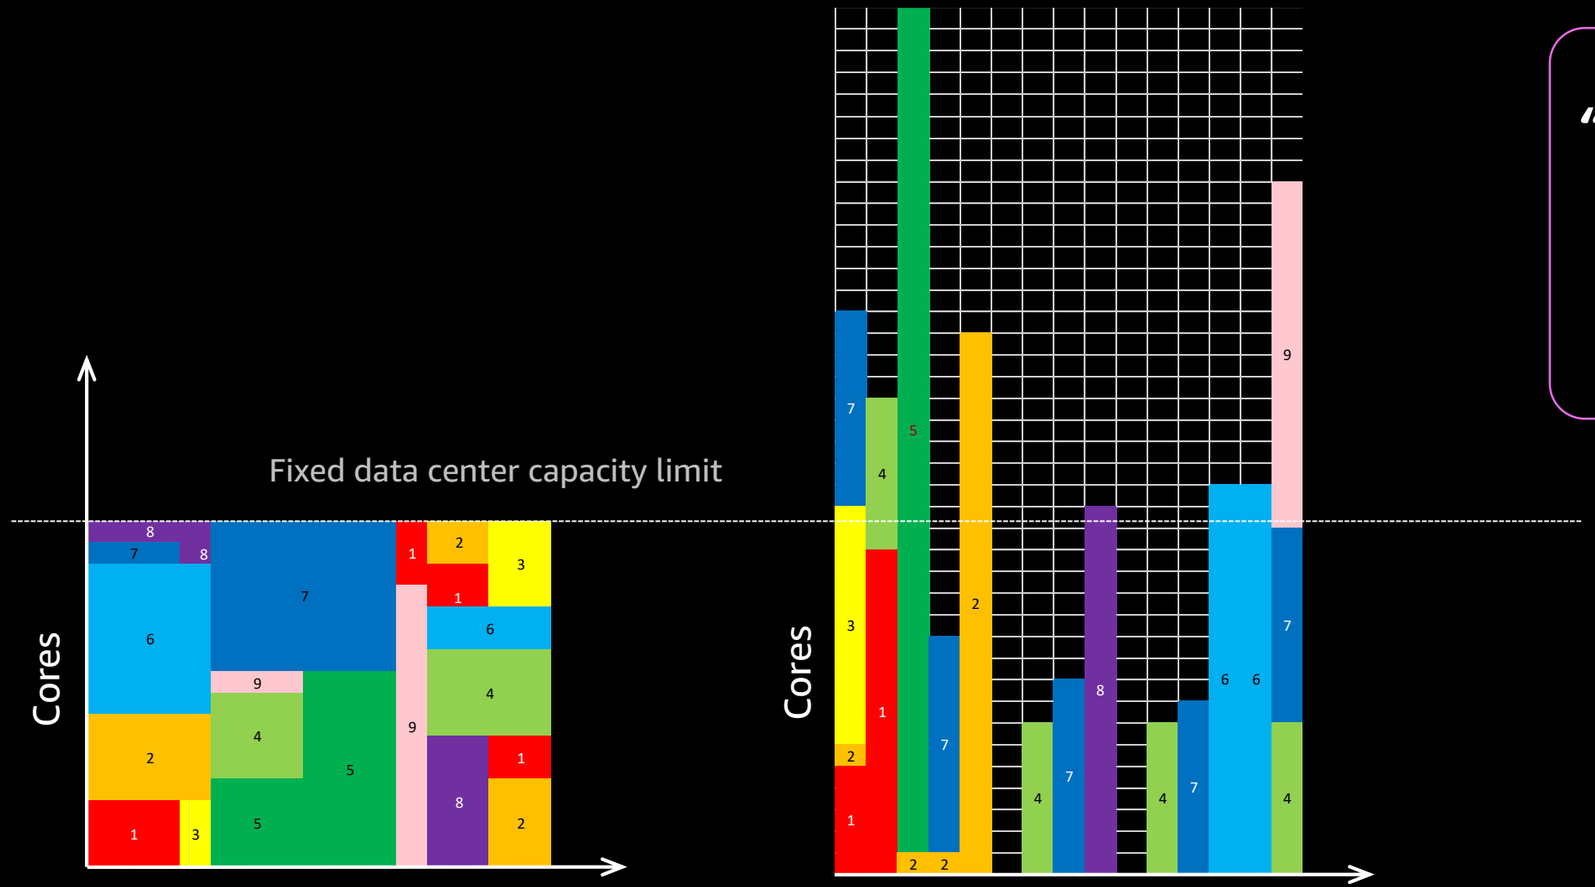
Weather forecasting



Computational Fluid Dynamics



The metric for success should be time-to-results



Finite capacity, usually with long queues to wait in

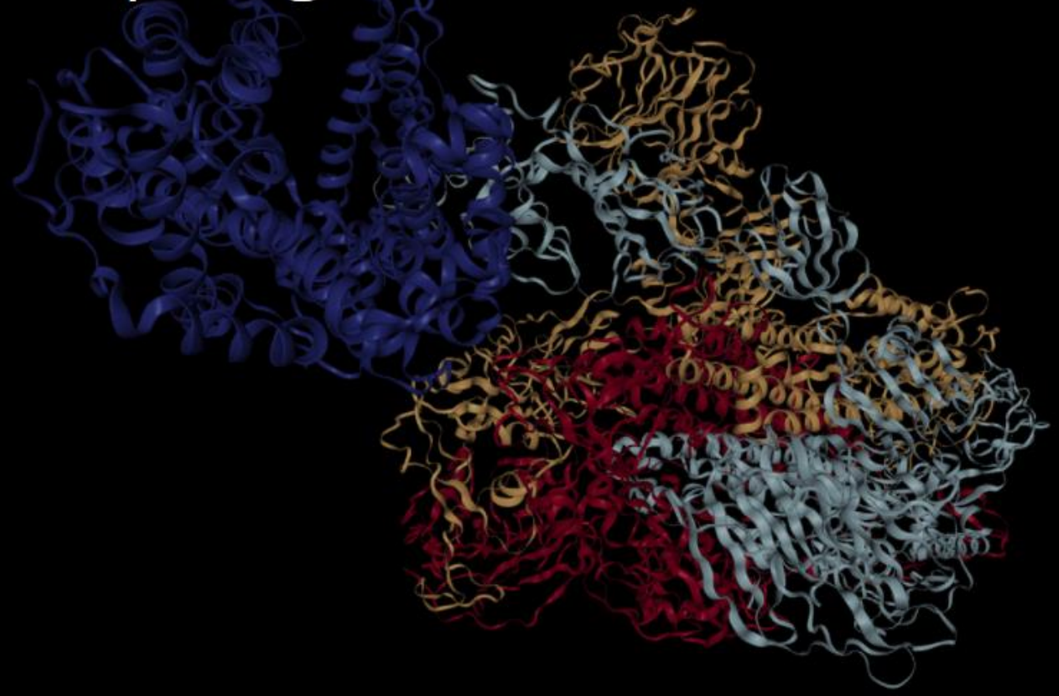
Massive capacity when needed to speed up time-to-results, and agile environment when additional hardware and software experimentation is needed

> *"For every $1 spent on HPC, businesses see $507 in incremental revenues and $47 in incremental profit."**

*Source: Hyperion ROI Study (http://www.hyperionresearch.com/roi-with-hpc/)

How has HPC impacted the worldwide response to COVID-19?

The COVID-19 High Performance Computing Consortium



AWS provides no-cost access to HPC capacity to select research projects

- Over **$1.5M in AWS credits** across **17 projects** since March 2020
- Offering free technical support

novel techsciences

Covid19Tree

Moderna uses AWS to accelerate its COVID-19 vaccine development

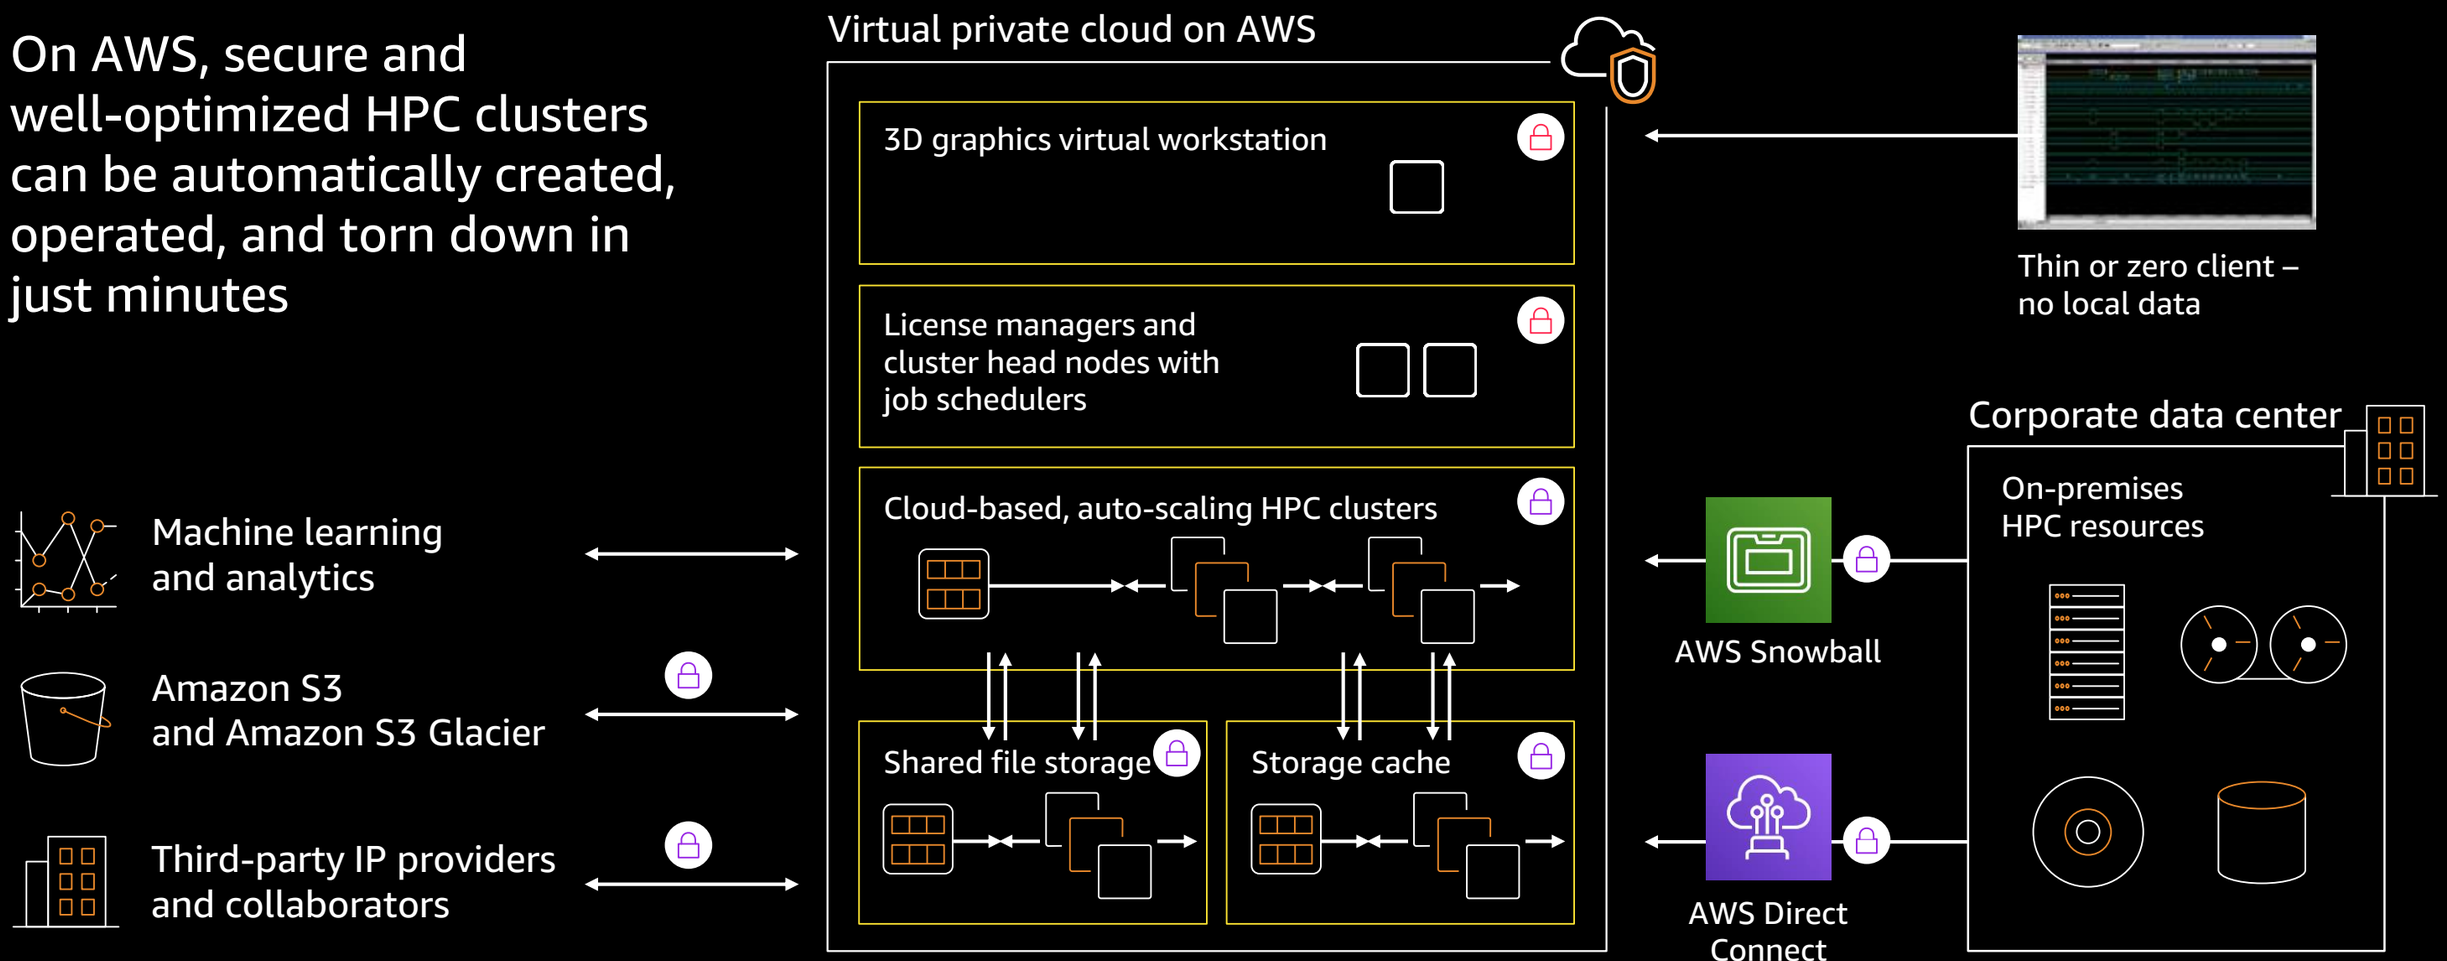


"AWS's breadth and depth of services are supporting our mission to create a new generation of medicines for patients and are instrumental in our quest to develop a vaccine for COVID-19 and other life-threatening diseases."

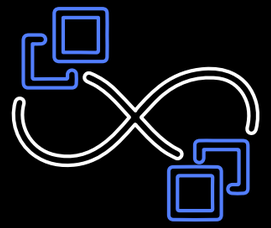
Stéphane Bancel
Moderna CEO

HPC on AWS

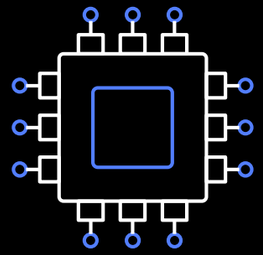
On AWS, secure and well-optimized HPC clusters can be automatically created, operated, and torn down in just minutes



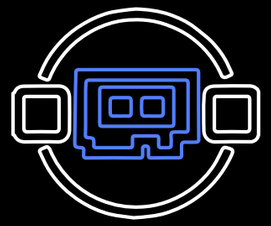
Key services that enable HPC on AWS



AWS ParallelCluster



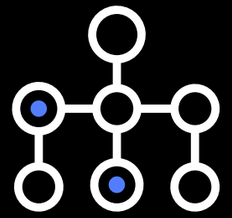
Amazon EC2



Elastic Fabric Adapter (EFA) + 100 Gbps networking



Amazon FSx for Lustre

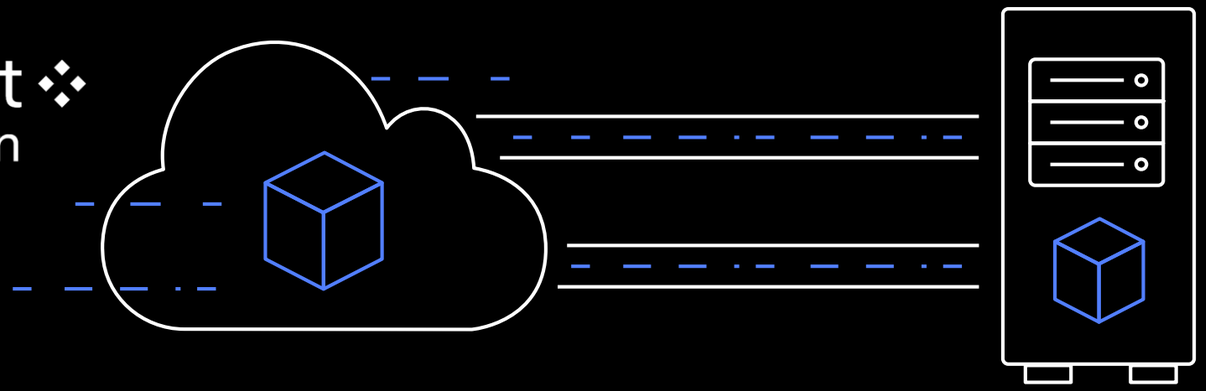


AWS Batch

Run your HPC workloads with the price performance you expect and the security you demand

AWS ParallelCluster

intel select solution



One-stop shop to set up your HPC cluster

Integrated with AWS services you need

Amazon FSx for Lustre

Amazon EC2 instances

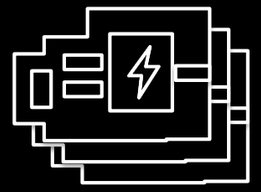
EFA

NICE DCV

AWS Batch

Bare-metal-like performance: Nitro Gen 4

Nitro Cards



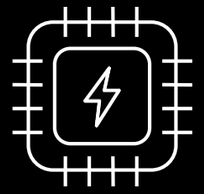
Amazon VPC networking

Amazon Elastic Block Store (Amazon EBS)

Instance storage

System controller

Nitro Security Chip

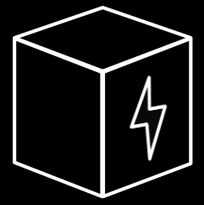


Integrated into motherboard

Protects hardware resources

Hardware root of trust

Nitro Hypervisor



Lightweight hypervisor

Memory and CPU allocation

Bare-metal-like performance

Broadest and deepest platform choice

General purpose

Burstable

Compute intensive

Memory intensive

Storage (high I/O)

Dense storage

GPU compute

Graphics intensive

+

Choice of processor
(AWS, Intel, AMD)

Fast processors
(up to 4.5 GHz)

High memory footprint
(up to 24 TiB)

Instance storage
(HDD and SSD)

Accelerated computing
(ASIC, GPUs, and FPGA)

Networking NEW
(up to 400 Gbps)

Bare metal

Size
(Nano to 48xlarge)

+

Amazon EBS

Amazon Elastic Graphics

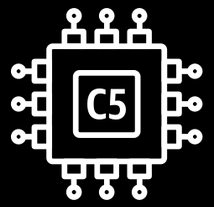
Amazon Elastic Inference

=

350+

instance types

for virtually every workload and business need



Amazon EC2 C5n instances

amazon Prime Air

Two key HPC-related features

More memory bandwidth

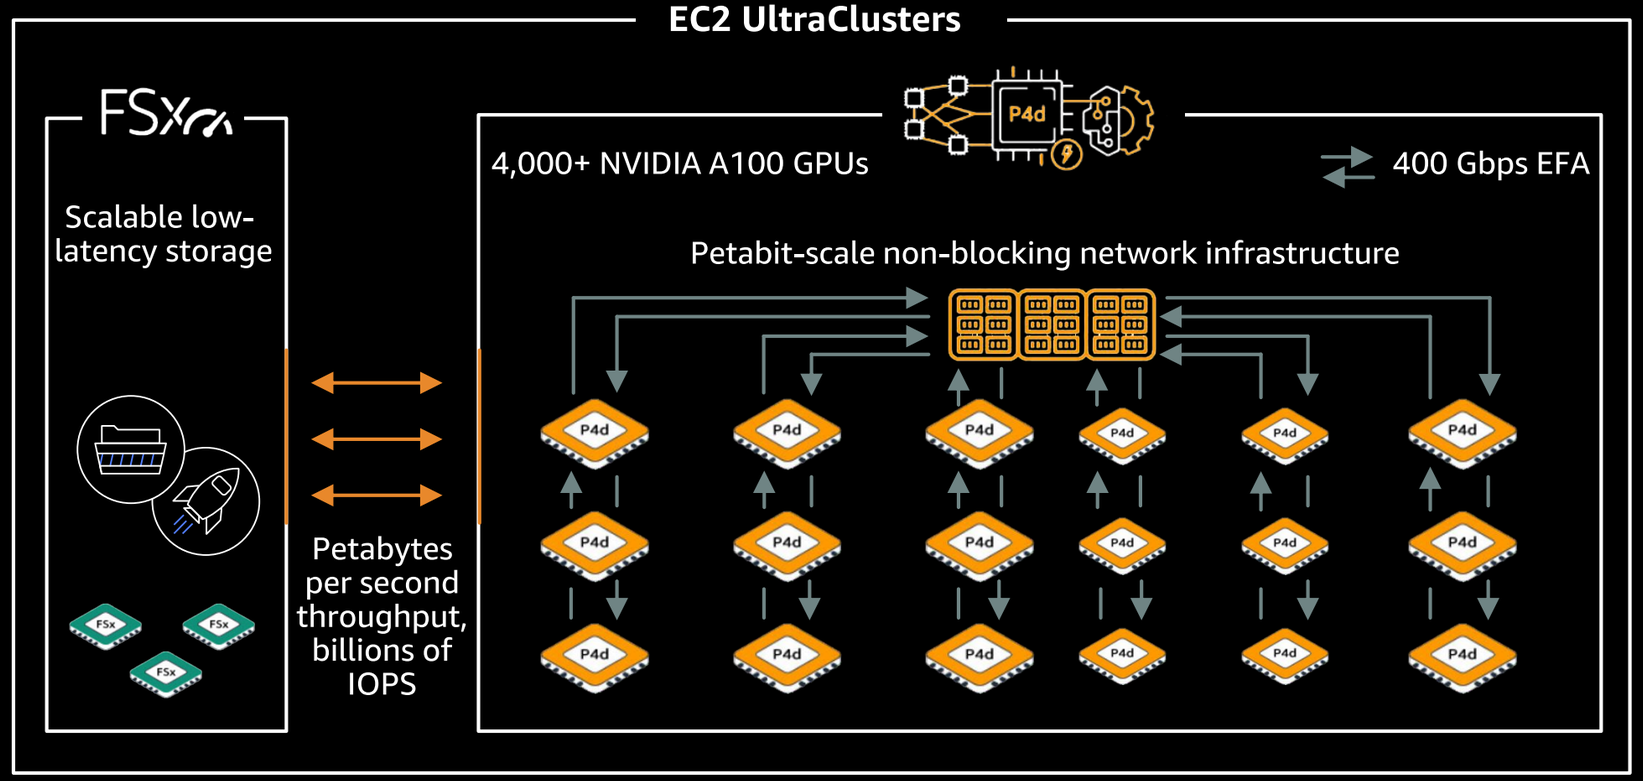
100 Gbps network throughput

Other 100 Gbps instances powered by AWS Nitro System

P3dn, I3en, M5n, R5n, G4dn

EC2 UltraClusters of P4d: Supercomputing-class performance for deep learning workflows

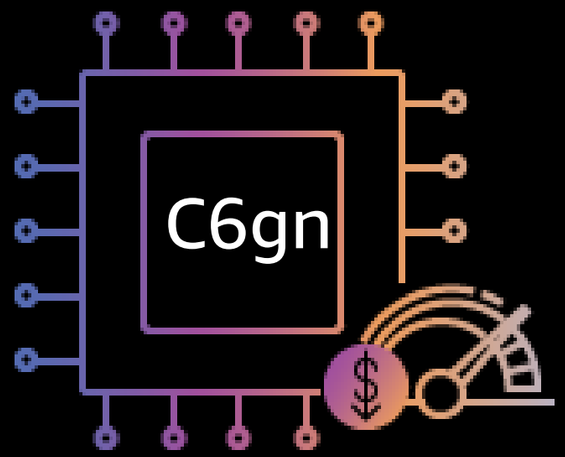
NEW!



- Based on NVIDIA A100 GPUs
- Availability to scale-out to large clusters for distributed training
- 2.5x better deep learning performance and 60% lower cost to train
- EC2 UltraClusters with EFA enables 400 Gbps and allows you to scale to over 4,000 GPUs

Introducing Amazon EC2 C6gn instances powered by AWS Graviton2 processors

NEW!



Up to 40% better price performance over C5n instances

100 Gbps networking

Ideal for compute-optimized workloads that need high networking throughput and can use Arm

“Based on the cost savings we’ve observed in our benchmarking, we plan to include Graviton2-based workflows for a range of our CFD simulations.”

Elastic Fabric Adapter

SRD protocol

NEW!

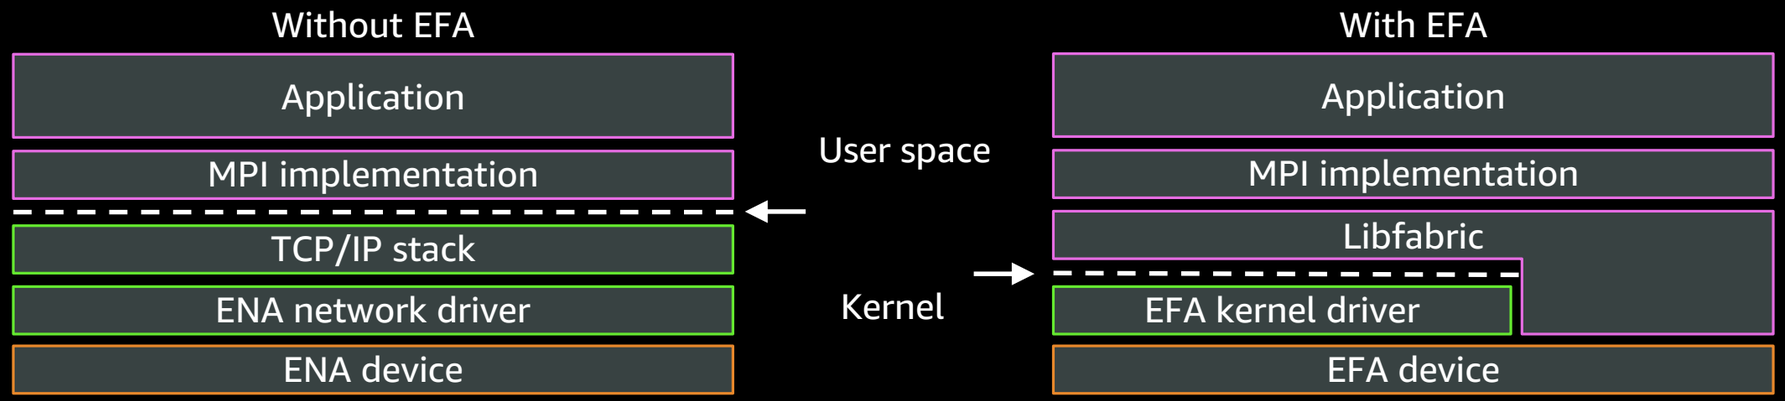
Up to 400 Gbps networking bandwidth

Proving myths about latency constraints wrong

CFD

Seismic

Weather modeling



INEOS | TEAM UK

INEOS uses EFA to run CAE / CFD workloads for boat designs

Purpose-built: EFA

Elastic Fabric Adapter



OS bypass

GPUdirect and RDMA

Libfabric core supports wide array of MPIs and NCCL

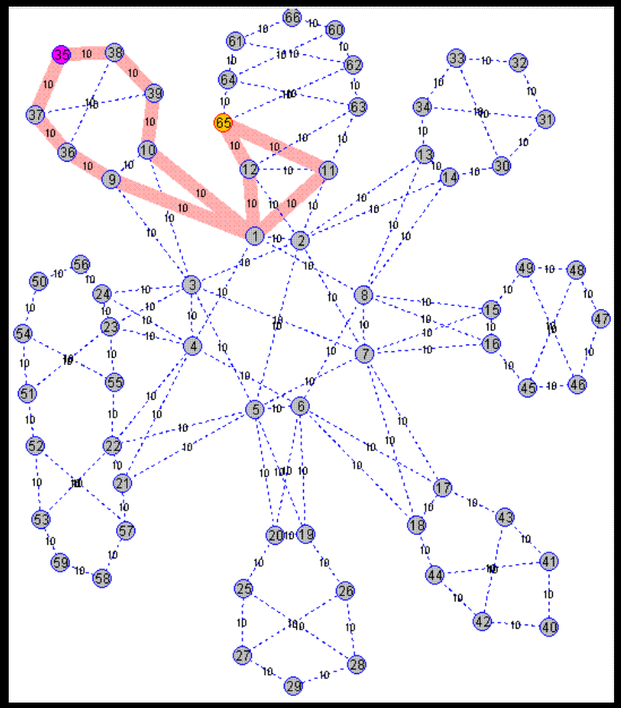
Scalable Reliable Datagram



OCMP-enabled packet spraying

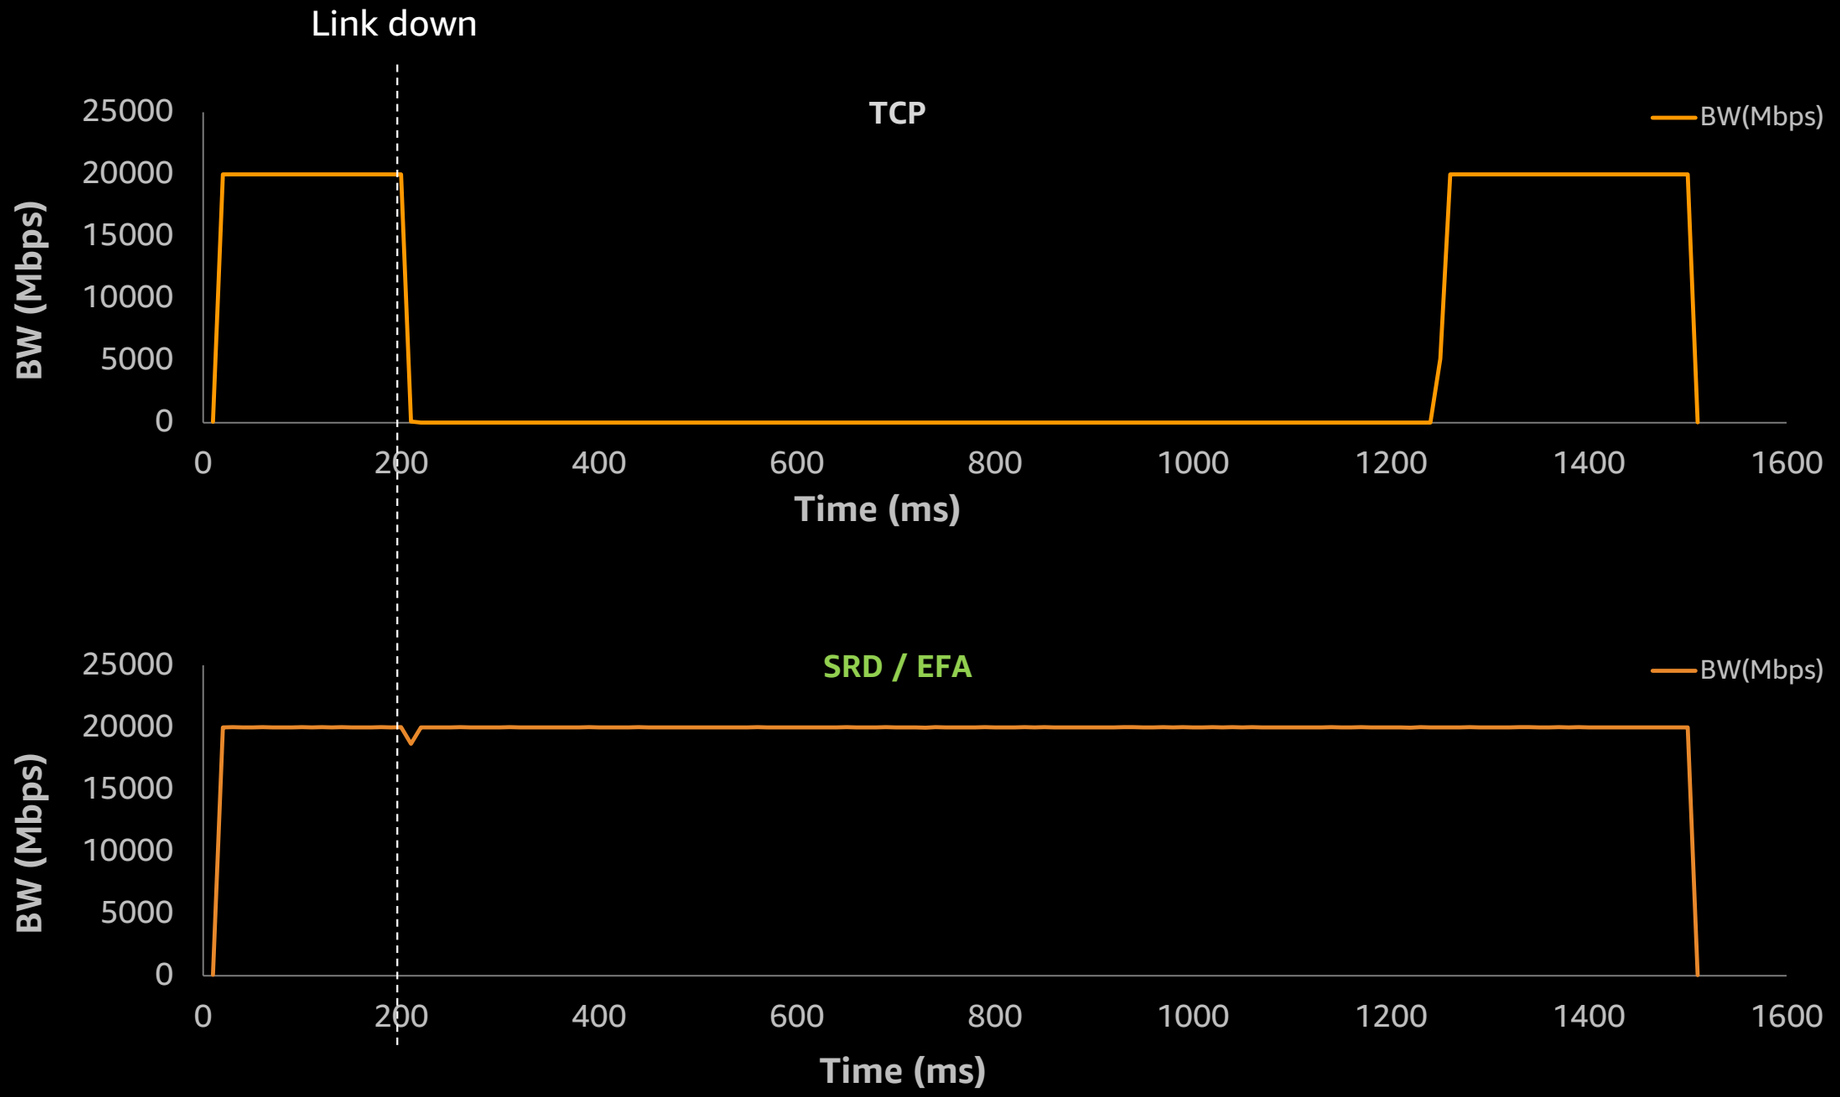
Cloud-scale congestion control

Fast recovery from packet loss or link failure



Thanks to Wikipedia and Peter Ashwood-Smith for the snappy animated GIF explaining OCMP

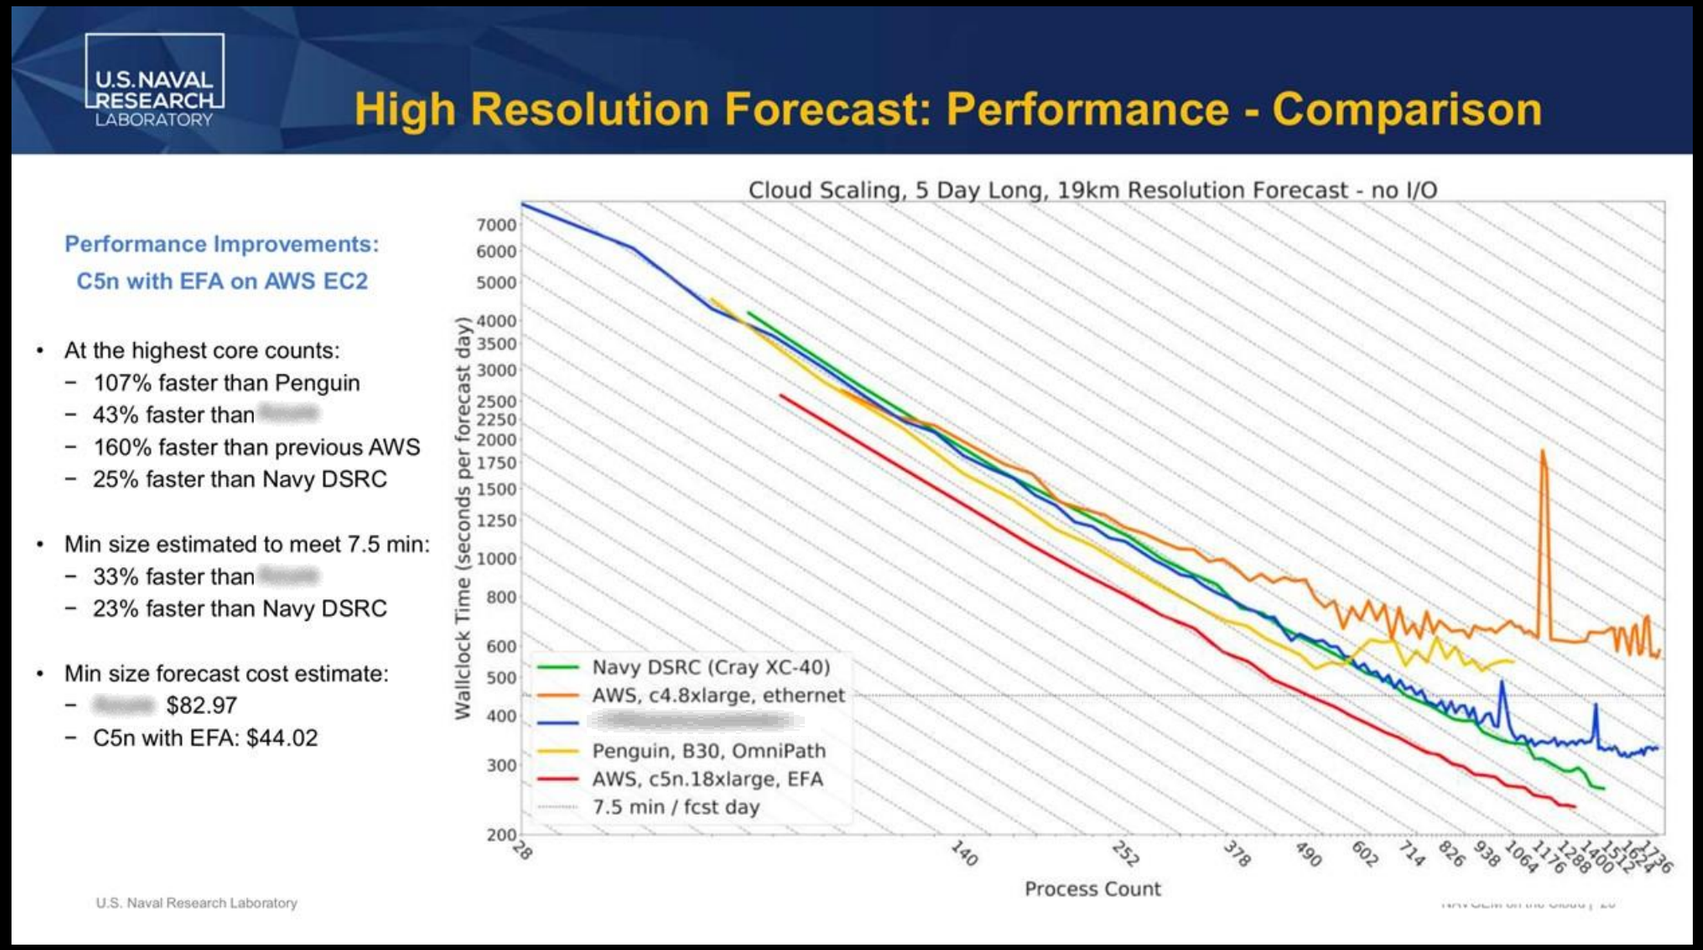
Graceful error recovery: Throughput



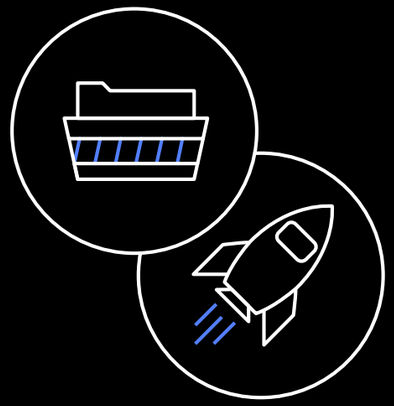
Delivering performance where it matters

THE ONLY BENCHMARK THAT MATTERS IS YOUR APPLICATION

U.S. Naval Research Lab scaled its tightly-coupled weather codes on EFA-enabled Intel instances, beating all its other platforms

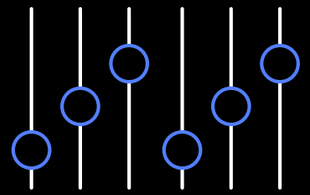


Amazon FSx for Lustre



High and scalable performance

Parallel file system



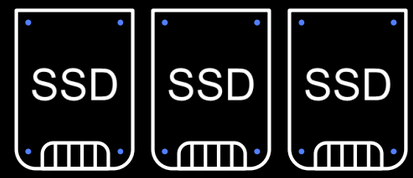


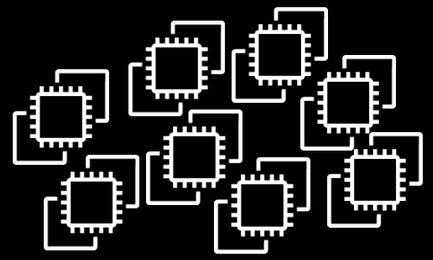
+100 GiB/s throughput

Millions of IOPS

Consistent submillisecond latencies

SSD-based



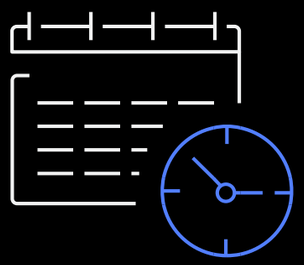


Supports concurrent access from hundreds of thousands of cores

TOYOTA
RESEARCH INSTITUTE

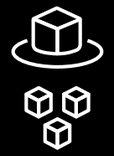
Toyota Research Institute uses FSx for Lustre for their large-scale machine learning workloads

AWS Batch



NEW!

AWS Fargate



Event



Changes in data state



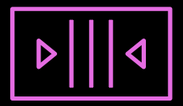
Scheduled triggers



Requests to endpoints



Job queue





Auto Scaling

Compute



Your code



Execution

Services (anything)



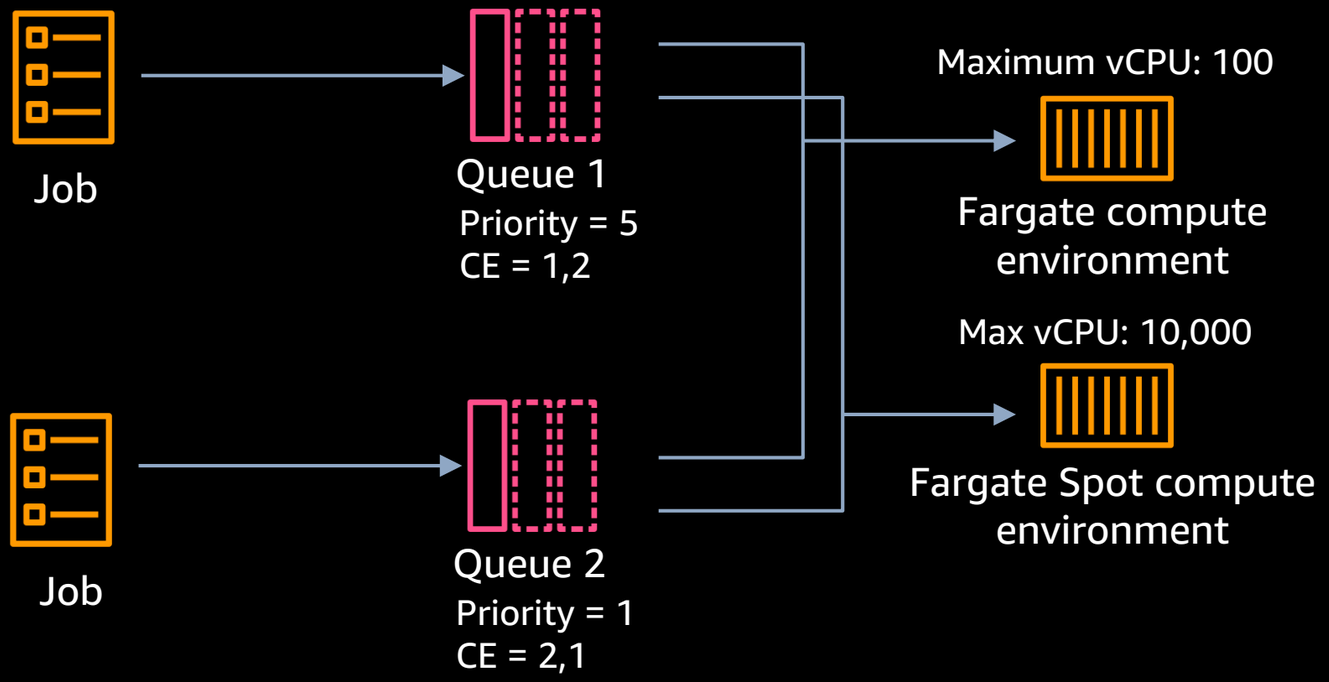


An Intel Company

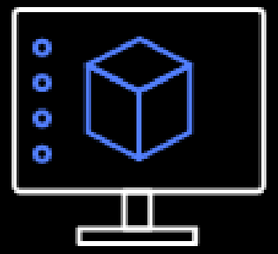
Mobileye runs 200–300 thousand cores on Spot Instances daily with AWS Batch for their AV simulations

How AWS Batch with Fargate works

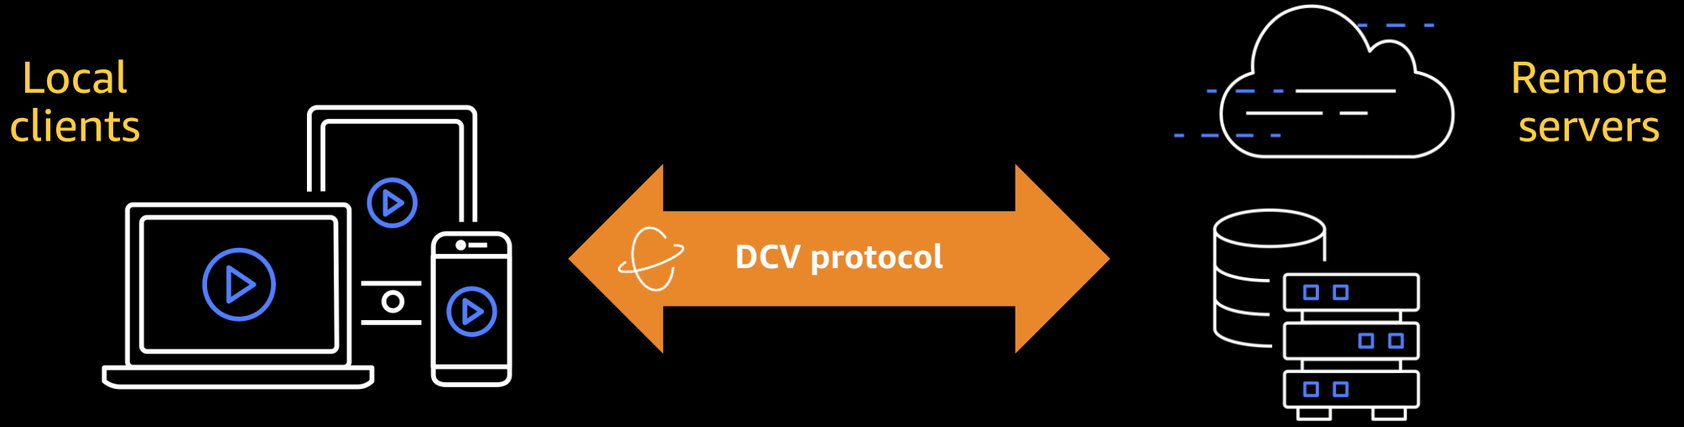
- Set maximum vCPUs to run on Fargate
- AWS Batch starts submitting to Fargate Spot
- If a job fails, AWS Batch will automatically retry (if desired)
- Can also specify dependencies, managed by AWS Batch



NICE DCV



Stream pixels, not data



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

Volkswagen automotive engineers use NICE DCV to access CAE applications running in the company’s on-premises HPC cluster

HPC workloads across industries



FRED HUTCH
CURES START HERE™

Life sciences



Fannie Mae®

Financial services

HALLIBURTON

Oil & gas



AUTODESK®

Design & engineering

MAXAR
TECHNOLOGIES

Climate & earth sciences



Level 5

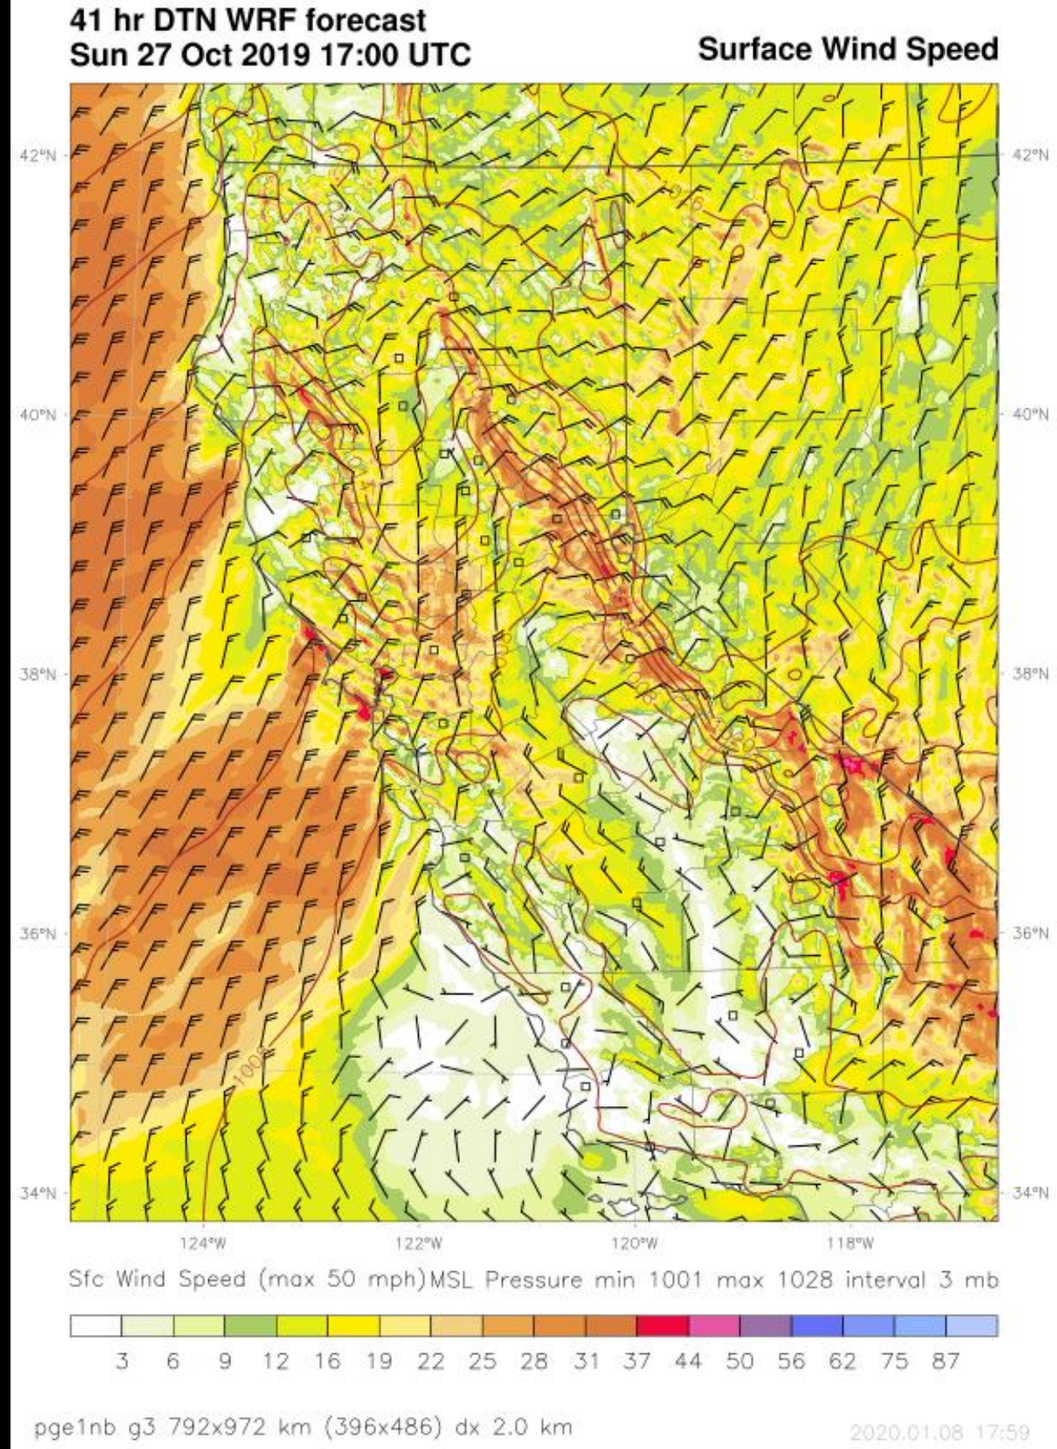
Autonomous vehicles

Time-sensitive weather forecasting using AWS HPC

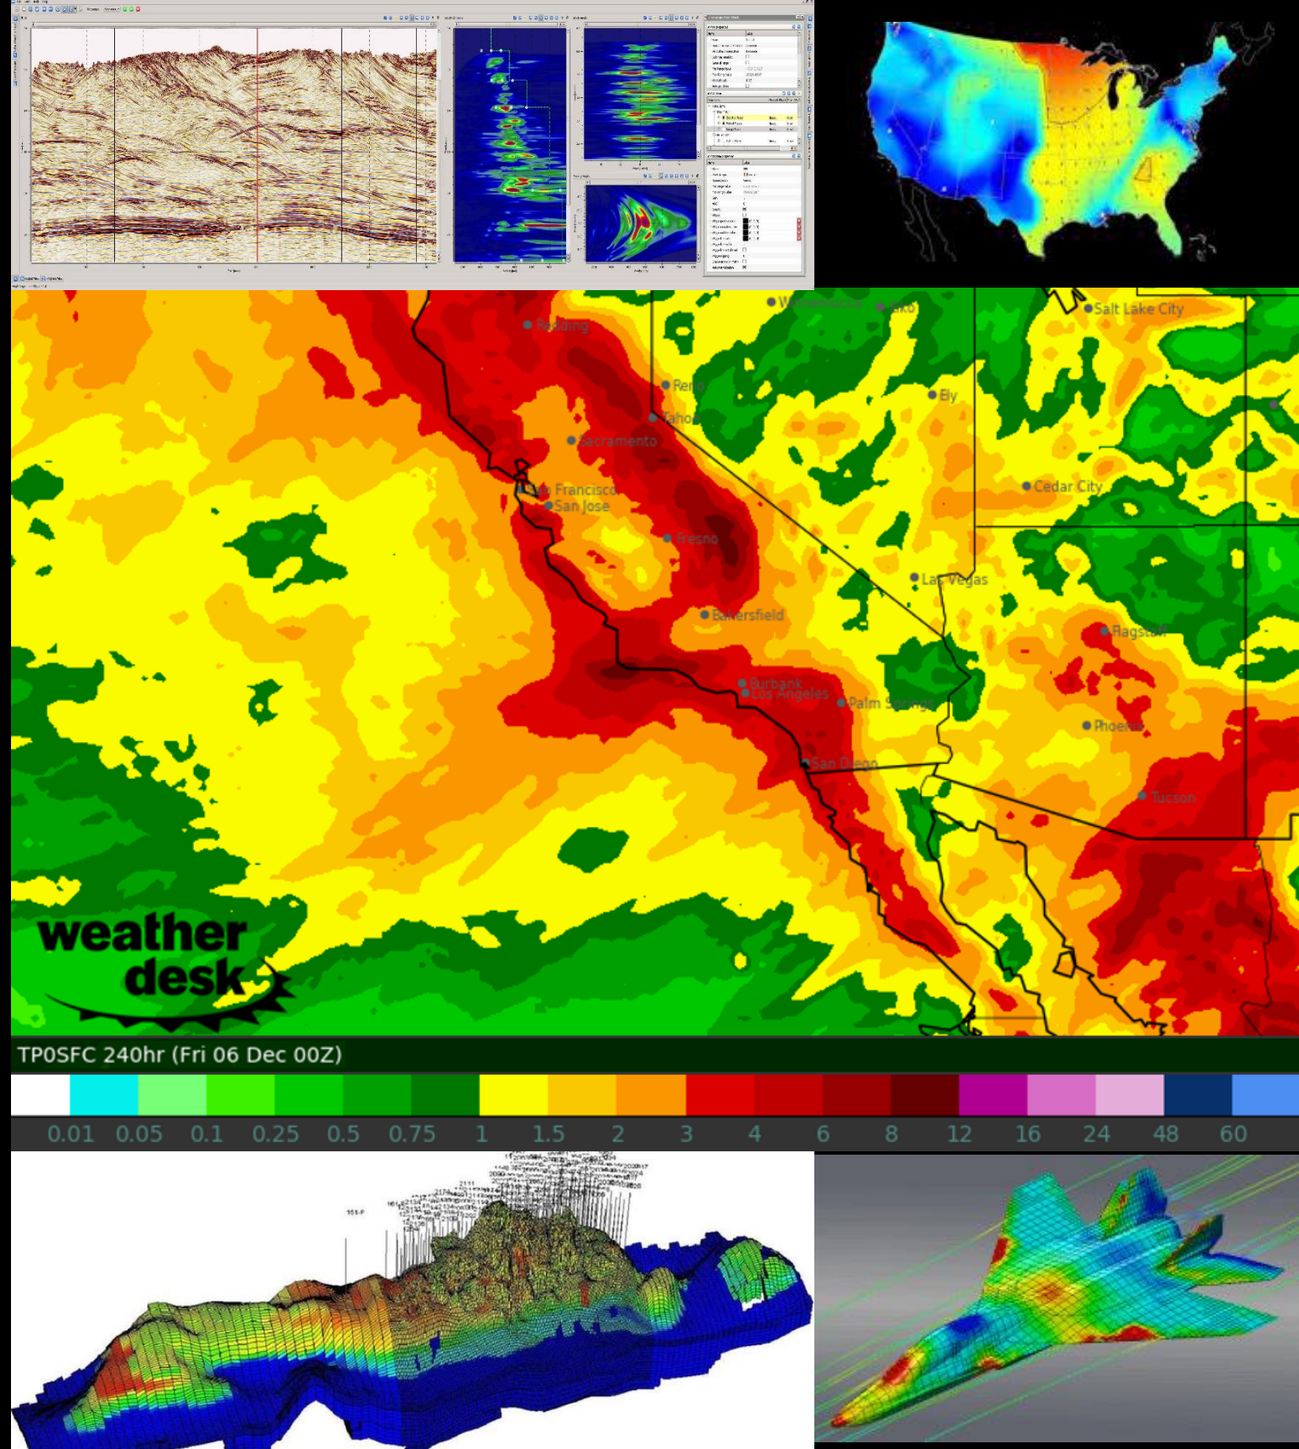
DTN°

- 2x price performance improvement across instance types
- 33% faster results on C5n instances (vs. C5 instances) even at low node counts using EFA and AWS ParallelCluster
- End users have more time to make critical business and safety decisions

Lowered compute costs by 57%



What business problems will you solve with this level of price performance?



Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.