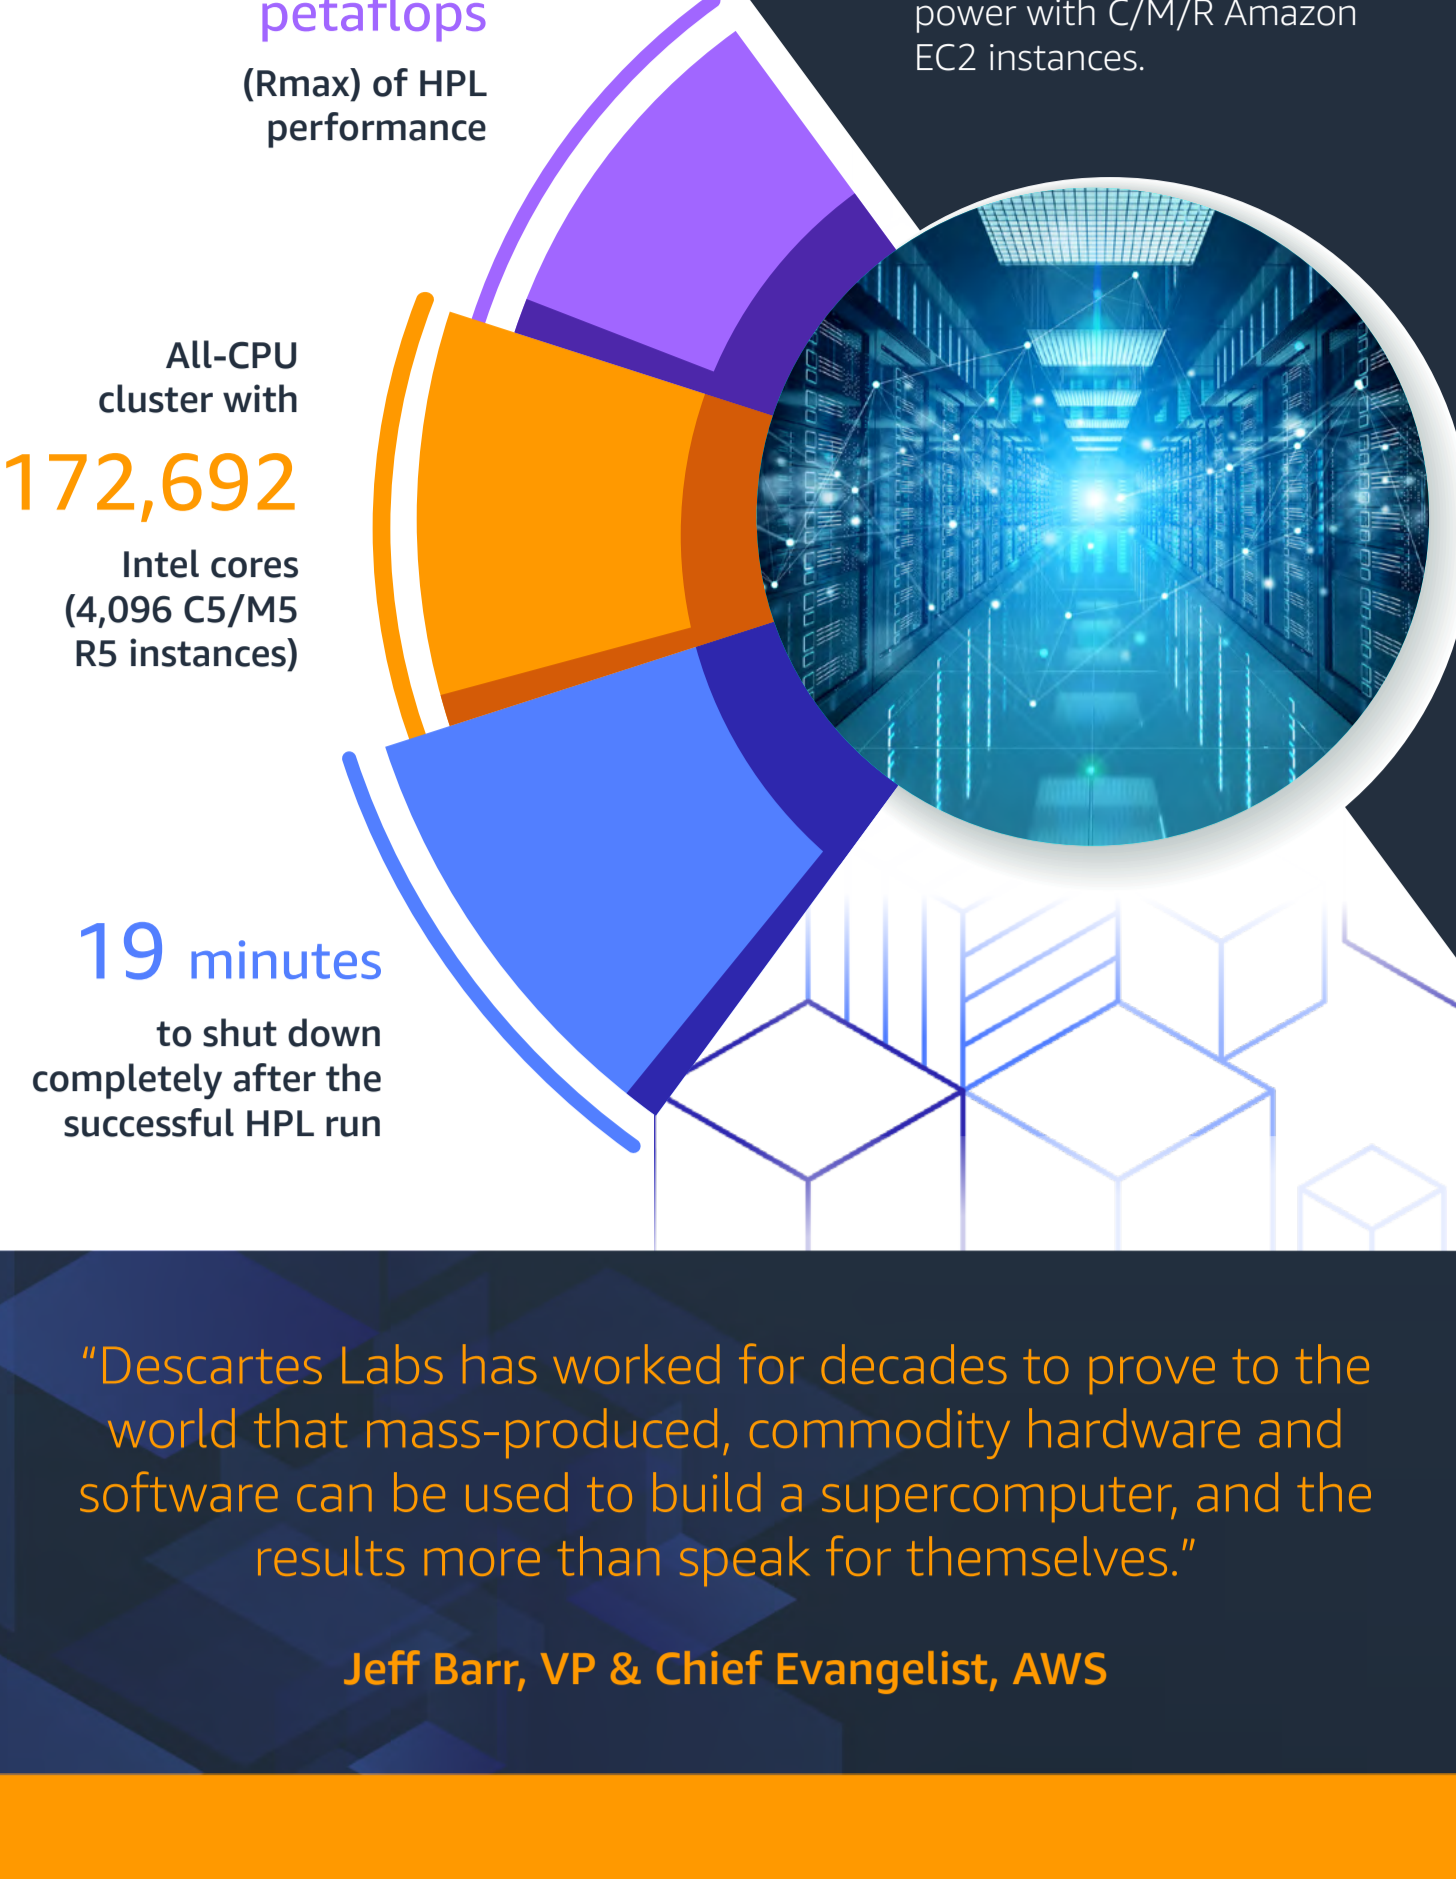




The future of HPC is here: Virtually limitless HPC+ML. On-demand.

High performance computing (HPC) continues to develop at incredible pace, and the convergence of HPC and machine learning—and even quantum computing—is opening up new possibilities. Gain insights faster, and scale up your compute with AWS' virtually unlimited compute capacity.

What would you do with 172,692 cores?



"Descartes Labs has worked for decades to prove to the world that mass-produced, commodity hardware and software can be used to build a supercomputer, and the results more than speak for themselves."

Jeff Barr, VP & Chief Evangelist, AWS

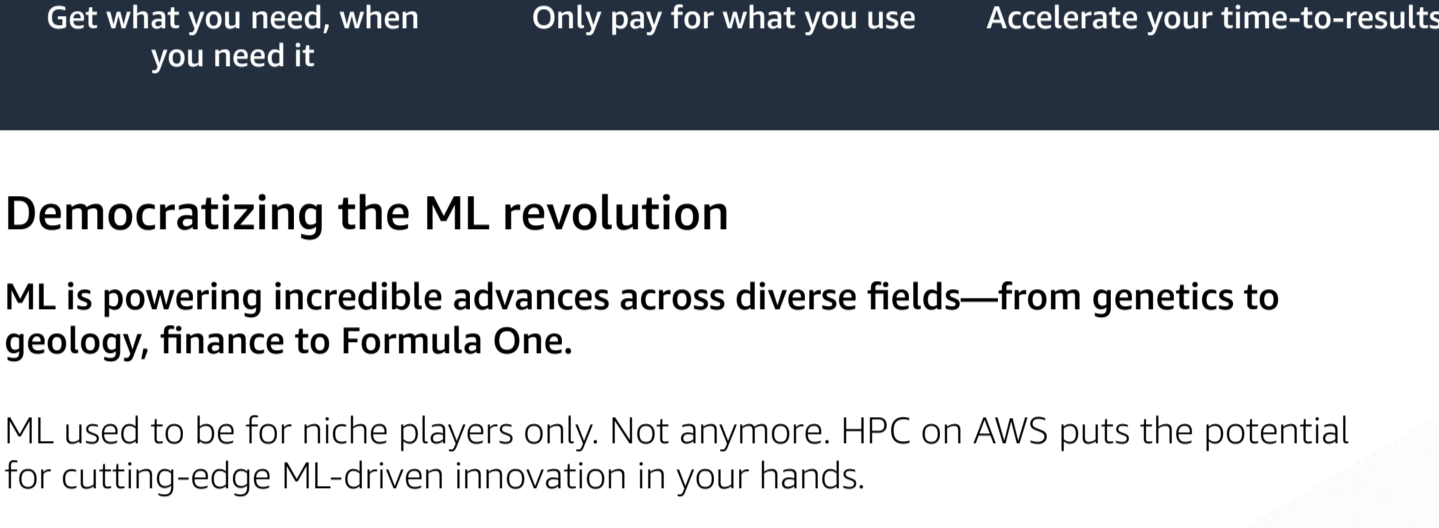
Get all the HPC capacity you need, when you need it

AWS HPC is enabling customers to create multiple "world firsts" on their own.



Any HPC workload is possible on AWS

Whether it's a Formula One team designing winning race cars with CFD, AstraZeneca running billions of genome sequences to create life saving treatment, or Maxar changing the landscape of weather forecasting... Whatever the HPC workload you can run it on AWS.



Computational fluid dynamics

Reservoir simulations

Weather forecasting

Elastic compute power on-demand

With HPC on AWS, you get rapid and secure access to thousands of cores when you need them. And only when you need them.



Get what you need, when you need it



Only pay for what you use

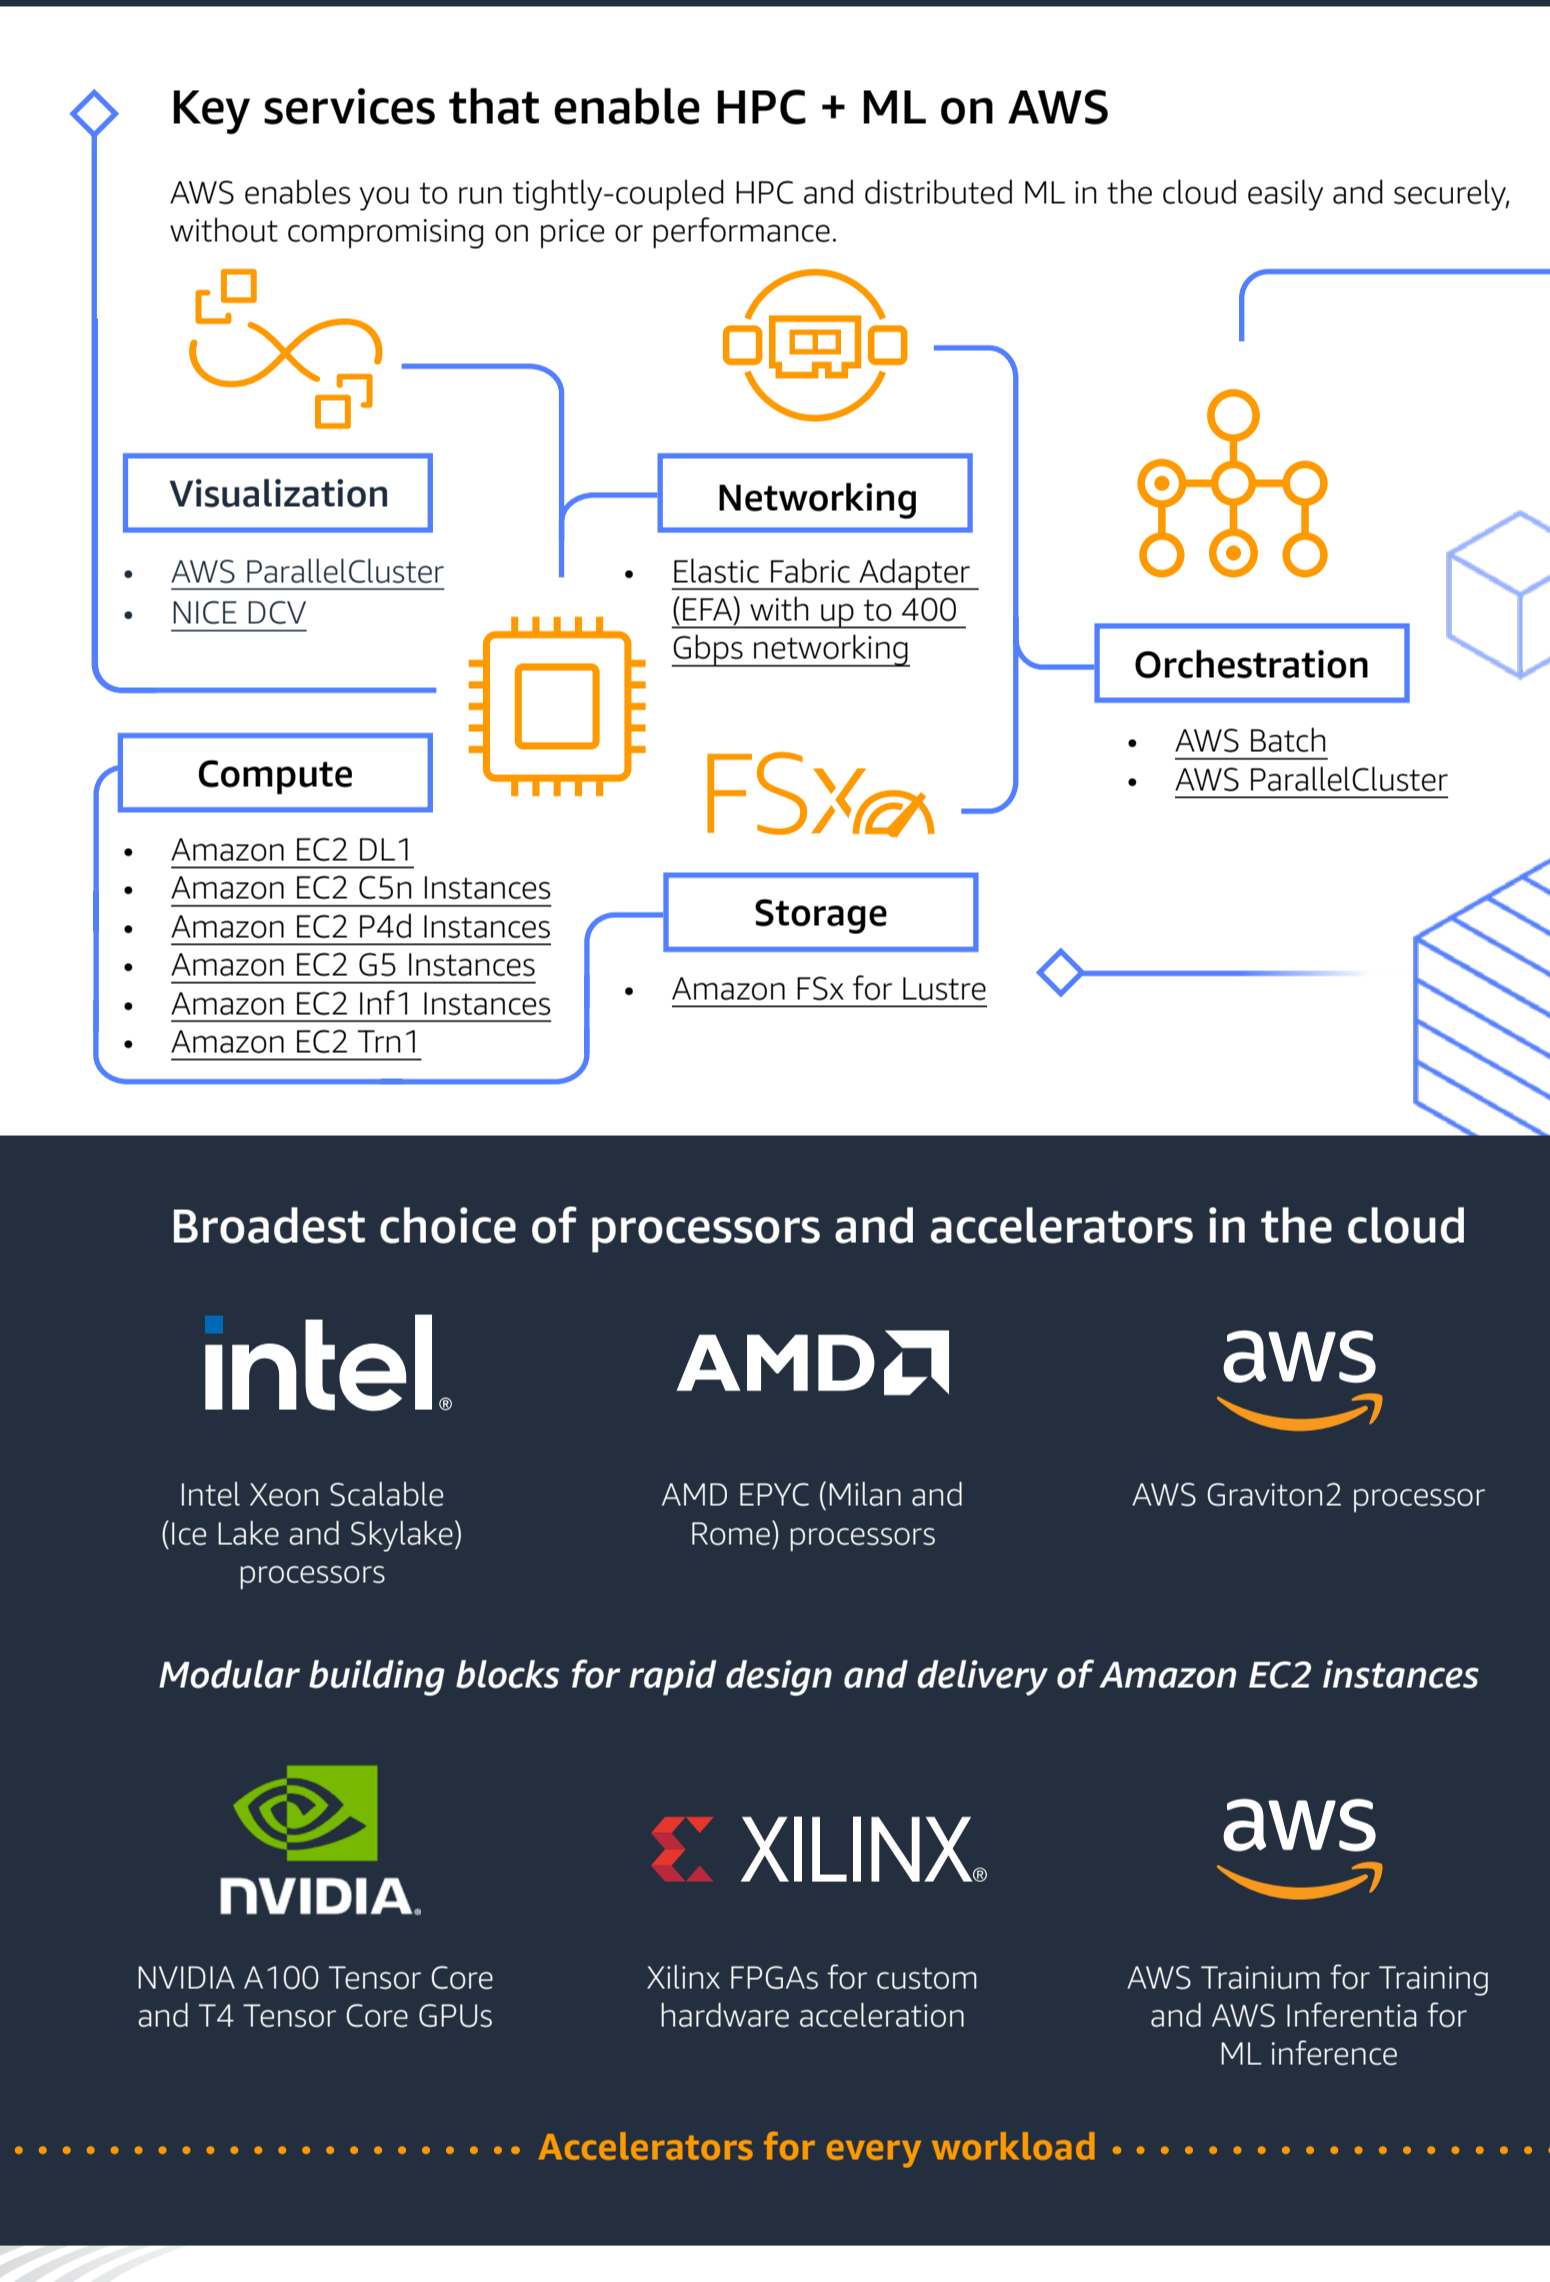


Accelerate your time-to-results

Democratizing the ML revolution

ML is powering incredible advances across diverse fields—from genetics to geology, finance to Formula One.

ML used to be for niche players only. Not anymore. HPC on AWS puts the potential for cutting-edge ML-driven innovation in your hands.




More choice. More flexibility. More innovation.

AWS ML solutions are enabling organizations to:




Reduce training time by **50%**



Deliver **3x** faster network throughput



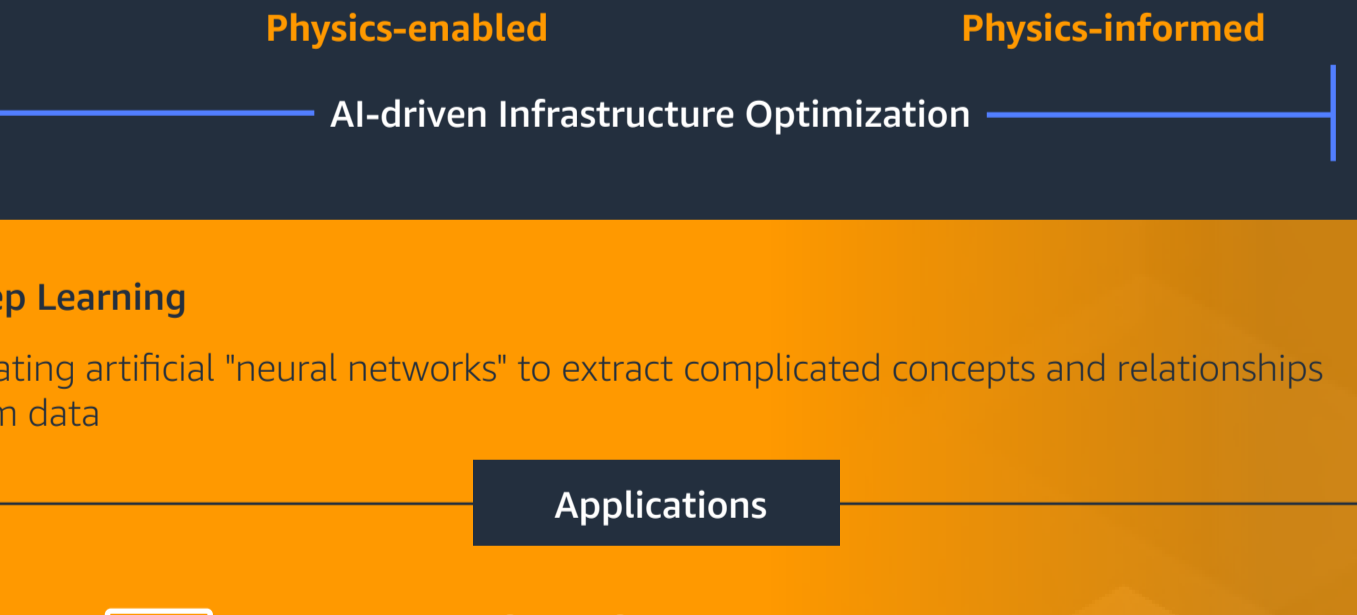
Provide **90%** scaling efficiency



Improve price and performance by **25%**

Key services that enable HPC + ML on AWS


AWS enables you to run tightly-coupled HPC and distributed ML in the cloud easily and securely, without compromising on price or performance.



Broadest choice of processors and accelerators in the cloud



Intel Xeon Scalable (Ice Lake and Skylake) processors



AMD EPYC (Milan and Rome) processors

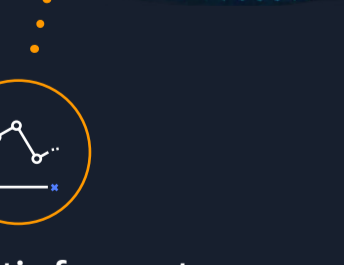


AWS Graviton2 processor

Modular building blocks for rapid design and delivery of Amazon EC2 instances



NVIDIA A100 Tensor Core and T4 Tensor Core GPUs




Xilinx FPGAs for custom hardware acceleration



AWS Trainium for Training and AWS Inferentia for ML inference

Accelerators for every workload

HPC accelerates...



Drug discovery | Cryo-EM process

- Increasing structural resolution with sub-cellular detail
- Optimizing price performance with nearly 500 instances types to choose from
- Turbo-charging drug discovery—from years to weeks



Autonomous vehicle navigation | Zoox

- Developing localization, perception, path planning, & control
- Petabytes of data processed per vehicle
- Millions of miles of simulations running continuously-massive HPC+ML at scale
- Spiky workloads requiring best-in-class infrastructure on-demand



Race car R&D | Formula One team

- Modeling turbulence with computational fluid dynamics
- Running 1,000s of designs to create the optimal car—beyond the capabilities of on-premises equipment
- AWS HPC+ML speeds up the process from months to hours

HPC & ML powers a broad set of new applications

Deep Learning

Creating artificial "neural networks" to extract complicated concepts and relationships from data

Applications

Natural Language processing

Image/video analysis

Autonomous Vehicle systems

Recommendation systems


Digital Twins

A "living" digital representation of a physical system, process (or entity) that enables many business outcomes:


HPC + ML drives digital twins at scale



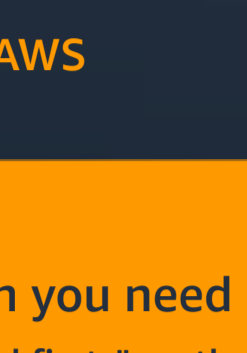
What-if analysis




Uncertainty quantification



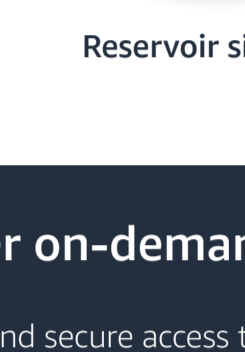
Probabilistic forecasts



Tuned model

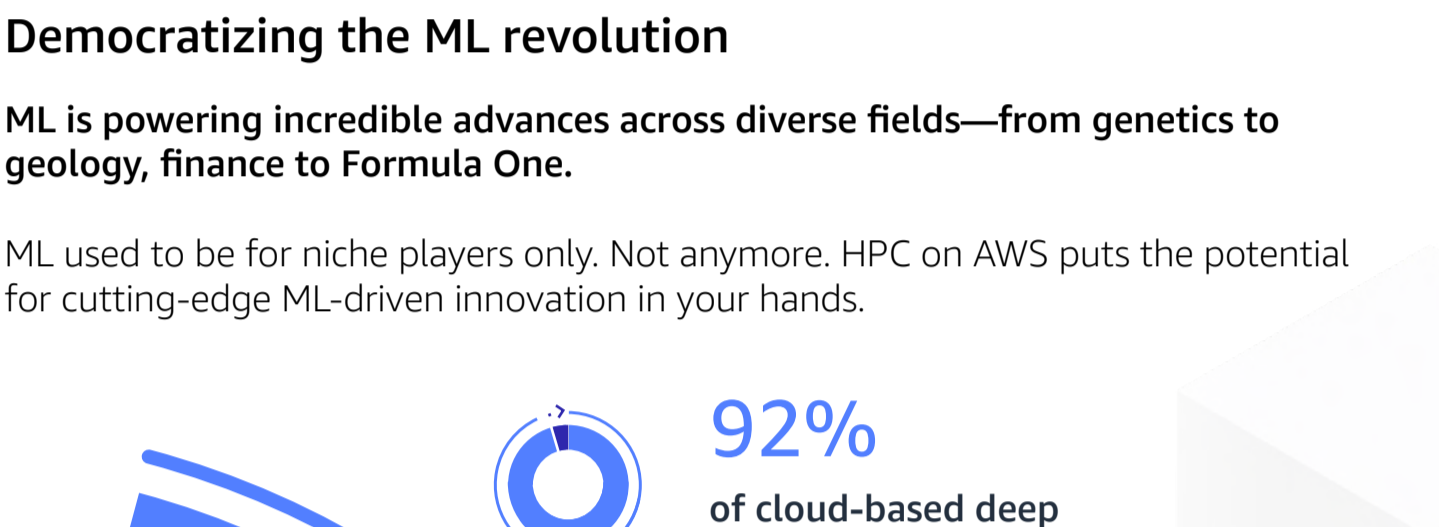


Sensitivity



Global optimization

Quantum Computing holds the promise to exponentially speed up solutions to hard problems



Computational chemistry

Machine learning

Optimization

AWS customers can run quantum classical simulations on AWS today.

Modern HPC + ML on AWS enables multiple paradigms simultaneously




Tightly-coupled workloads




Loosely-coupled workloads



Accelerated computing



Visualization



AI/ML



High volume data analytics

All the solutions you need. Limitless possibilities.

This is just a snapshot of what's possible now.

Don't wait for the future, make it happen today!

Learn more <https://aws.amazon.com/hpc>

¹ <https://nucleonresearch.com/research/angle/guidebook-deep-learning-on-aws-2/>