Scaling research in the cloud

With AWS high performance compute and data sharing

Steve Gillard, Senior Solutions Architect, AWS

November 2020

What is the problem?

On-premises

Fixed & limited compute allocation for research

Strict corporate IT control mechanisms

Limited access to / ability to share data & results

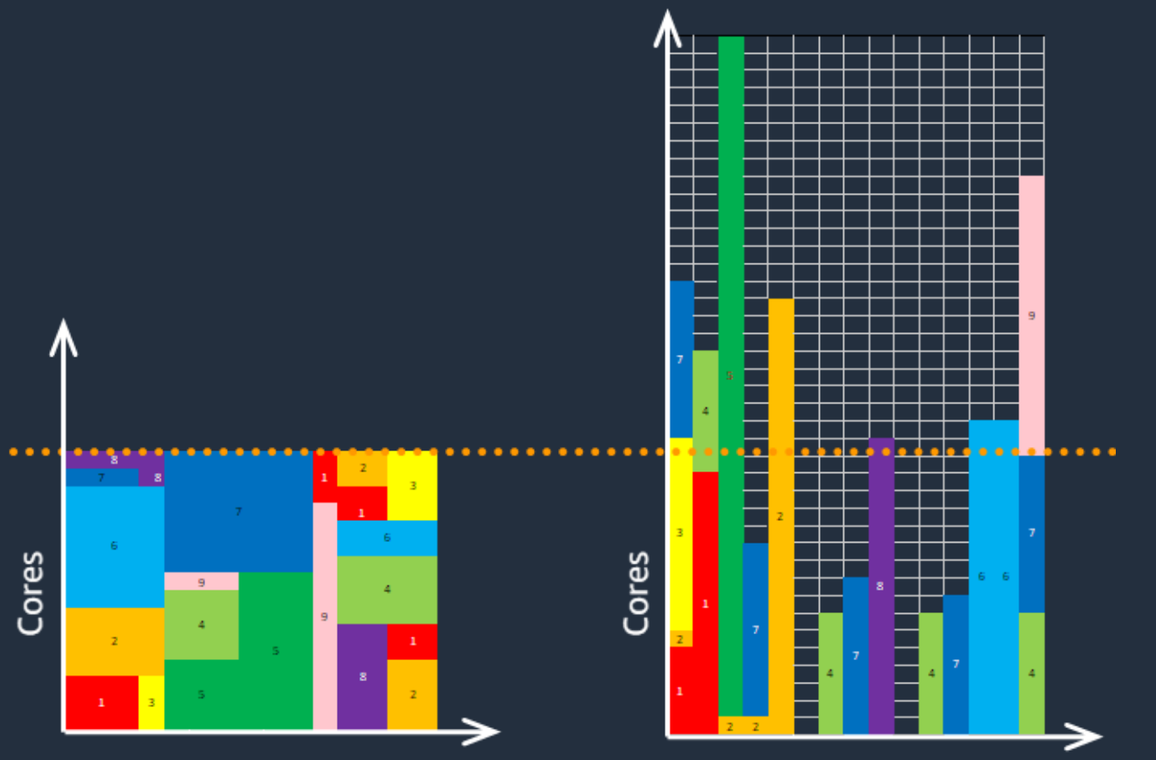

Cloud

Capacity on demand, elastic scalability

Guardrails, superior visibility and control

Securely share data, bring compute to the data

aws

Pace of Innovation

2020
- “P4d” instances with NVIDIA A100 Tensor Core GPUs and 400Gbps network bandwidth
- AWS Graviton2 ARM instances
- 2nd Gen AMD EPYC instances

2019
- AWS Inferentia instances for machine learning inference
- “P3” and “G4” instances with NVIDIA V100 Tensor Core and T4 GPUs
- Expanded range of bare metal instances (compute and memory optimized)

2018
- “A1” first generation AWS ARM instances
- 1st Gen AMD EPYC instances
- “C5n” instances with up to 100Gbps network bandwidth
- Elastic Fabric Adapter – high throughput, low latency networking
- “F1” FPGA instances

2017
- “C5” compute optimized instances and “I3” bare metal instances
- AWS Nitro system introduced

Elastic Fabric Adapter (EFA)

High throughput, low latency networking

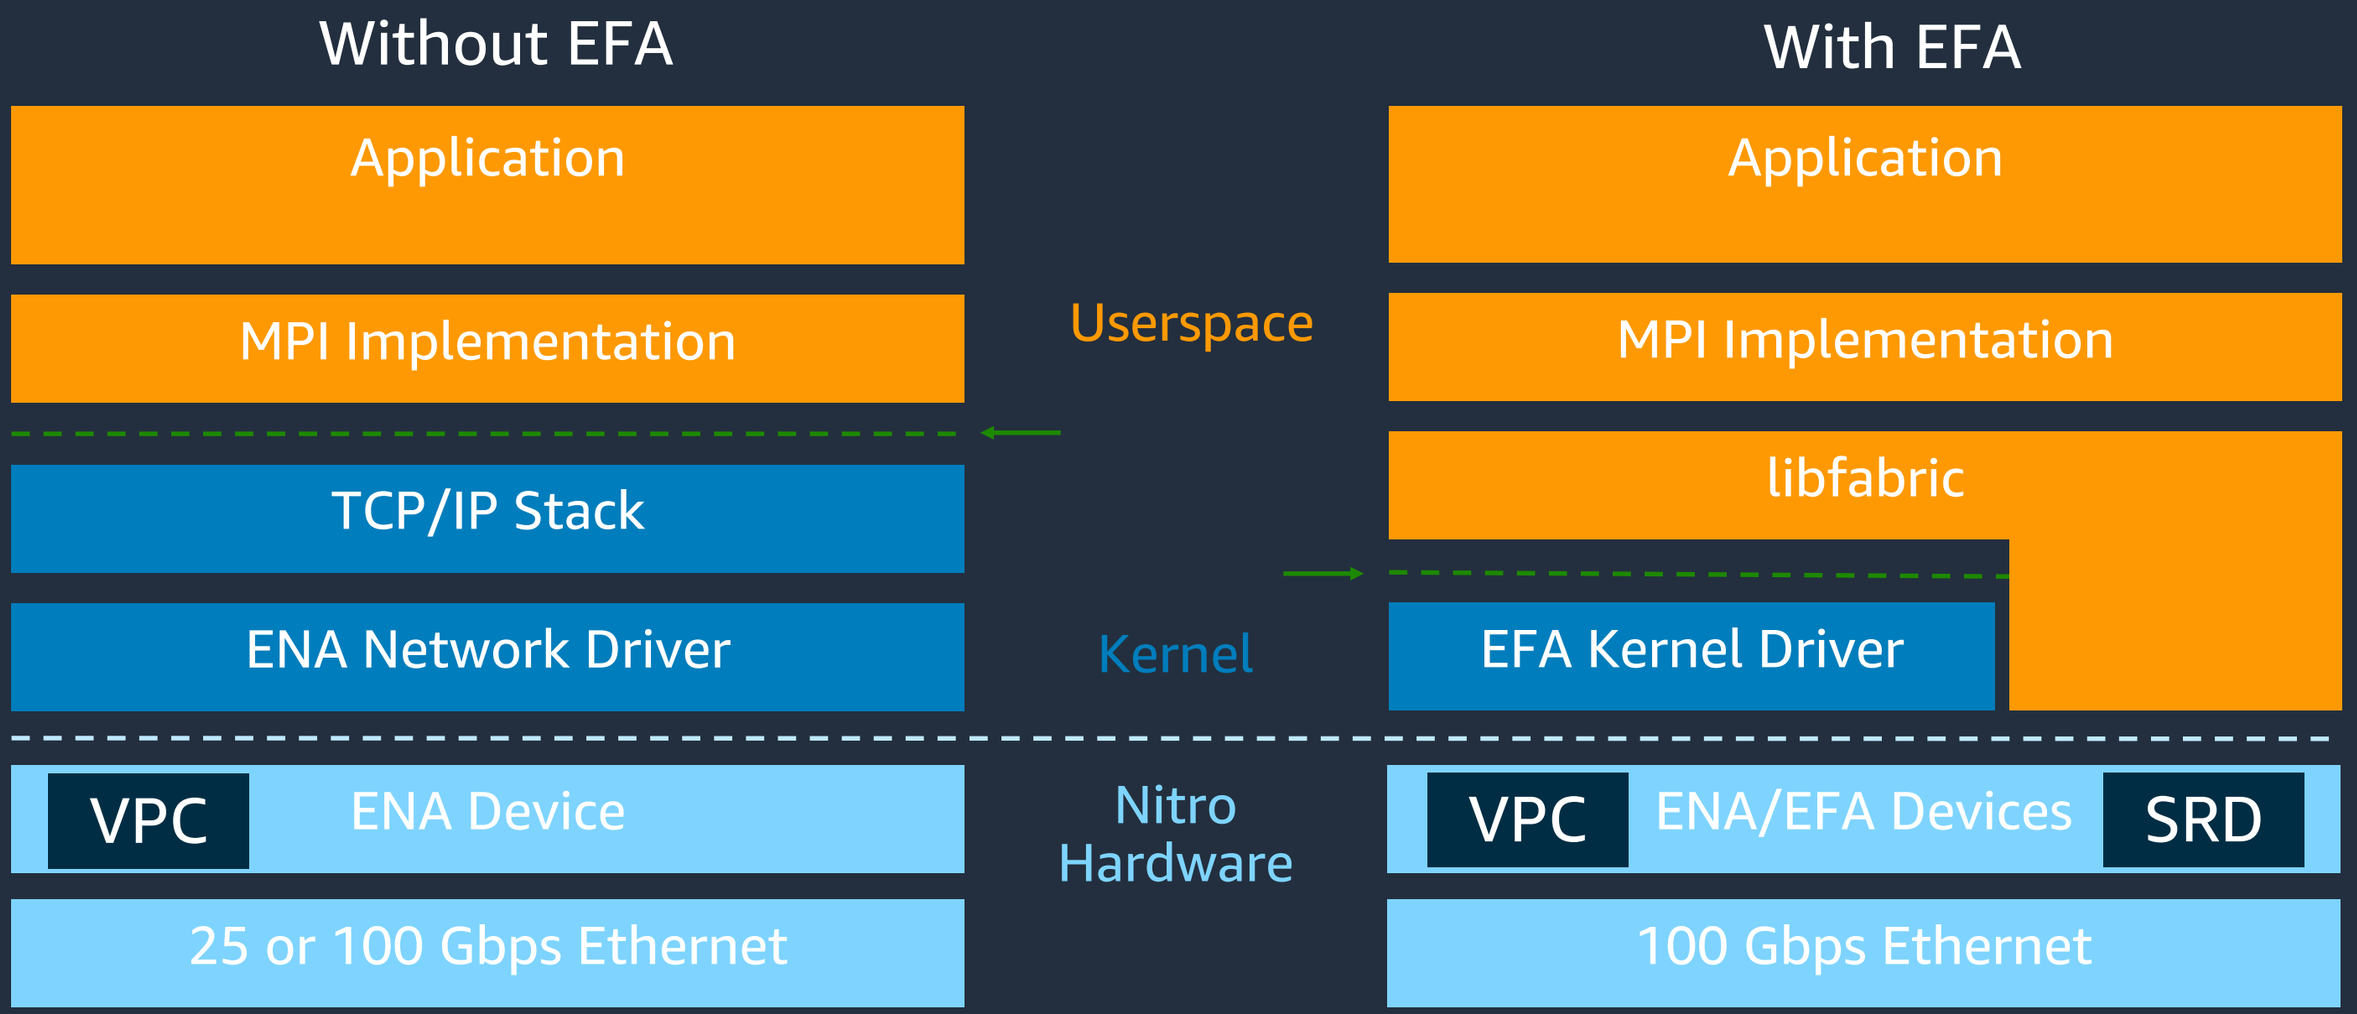

aws

Scalable Reliable Datagram (SRD)

A reliable high-performance lower-latency network transport

Inspired by Infiniband Reliable Datagram, without the drawbacks

- No limit on the number of outstanding messages per context

Out-of-order delivery – no head-of-line blocking

- Messages are independent: in many cases, application/middleware can restore ordering only if/when needed
- Same motivation as weak/relaxed memory ordering

Packet spraying over multiple ECMP paths

- Rapidly adapt to hot spots
- Fast and transparent recovery from network failures

Congestion control designed for large-scale cloud

- Maintains high throughput in the face of packet drops
- Minimize latency jitter

New network transport protocol

Designed for AWS's unique datacenter network

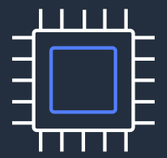

Part of AWS's 3rd generation Nitro Chip

Exposed as reliable datagram interface in EFA

https://ieeexplore.ieee.org/document/9167399

aws

EFA improvement – US Navy NAVGEM model

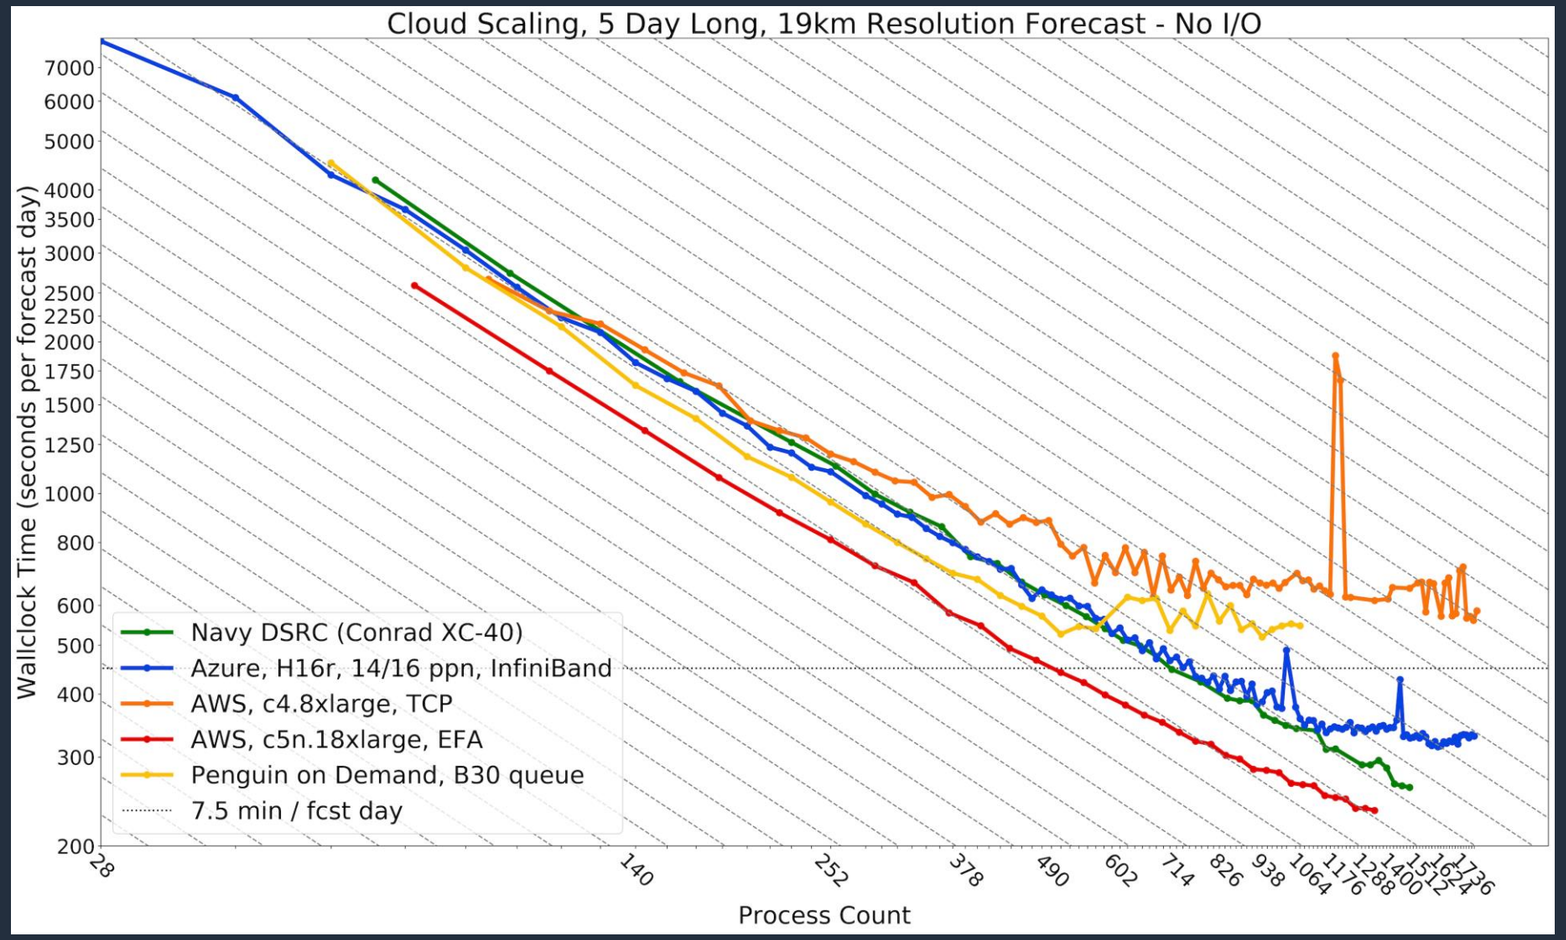

Performance improvement for C5n with EFA:

- 139% faster than previous generation AWS instance (C4)
- 6% faster than Navy DSRC (Conrad XC-40)

And lower cost per run (on-demand rates):

- C5n with EFA: USD $13.76
- Previous generation: USD $18.65

https://youtu.be/GTHWf0OVGrw

EC2 UltraClusters with P4 instances

400Gbps, NVIDIA GPUDirect, FSx Lustre storage

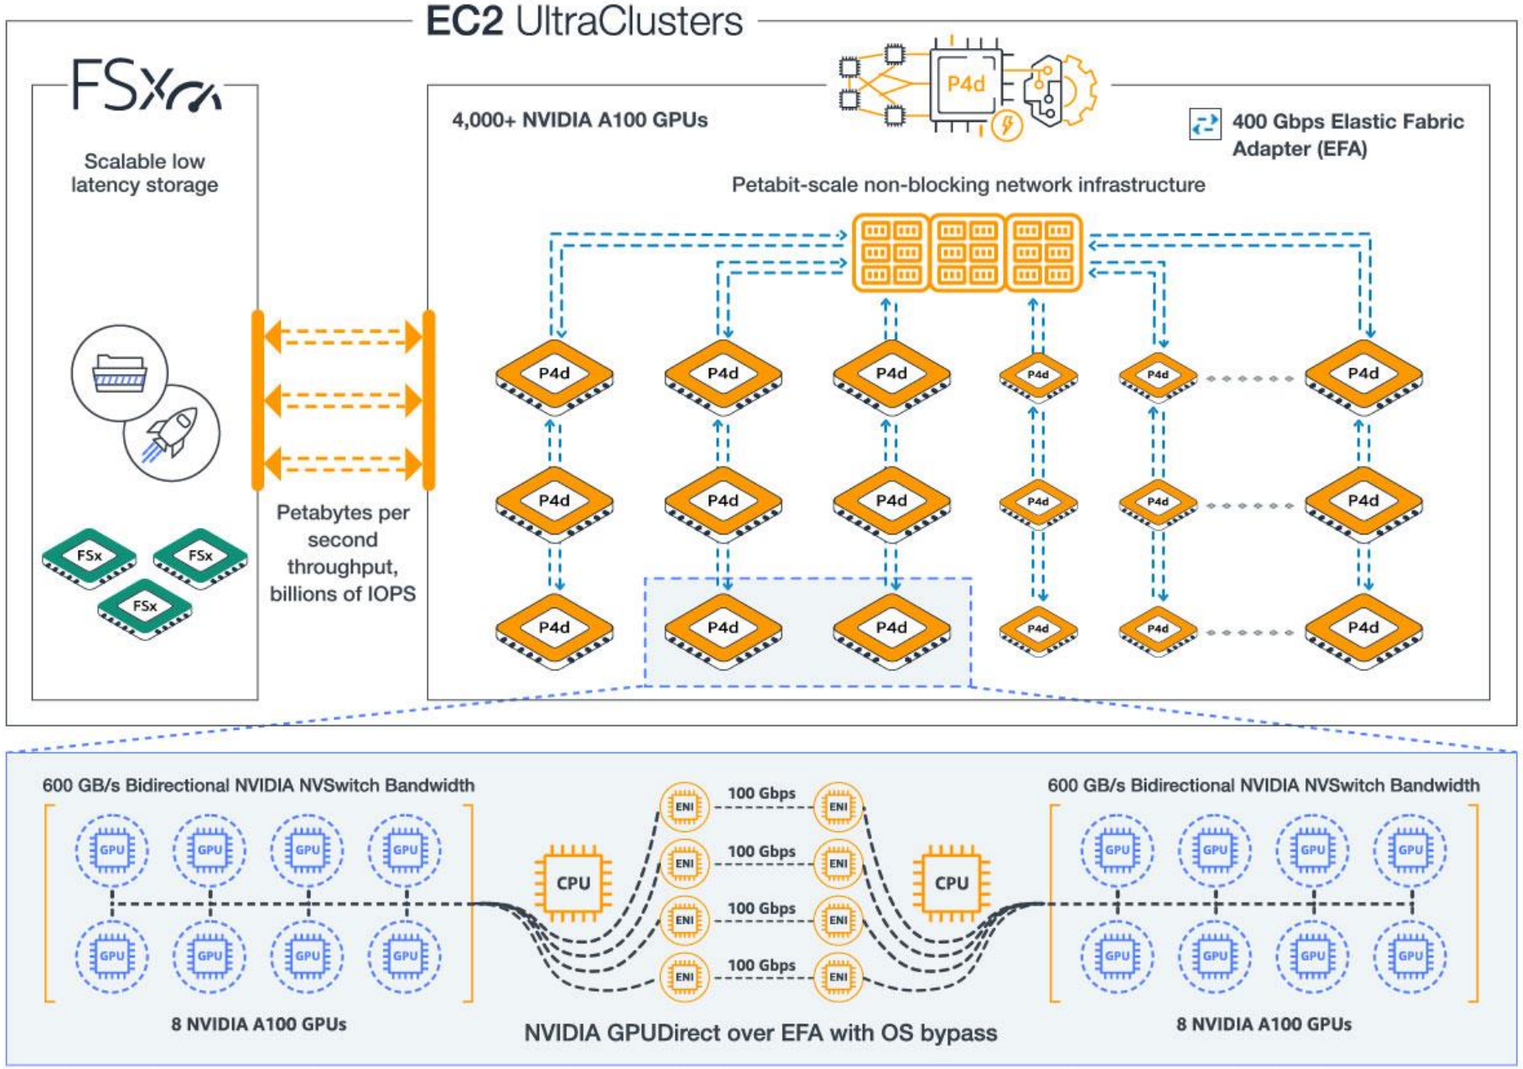

Each P4 instance has:

- 96 vCPUs (48 cores)
- 1152 GB memory
- 8 x 1TB NVMe SSD
- 8 x NVIDIA A100 GPUs
- 400Gbps EFA network

Compared to previous P3 generation (2019):

- 4x network bandwidth
- 2.5x GPU performance

https://aws.amazon.com/ec2/instance-types/p4/

Getting started with AWS ParallelCluster

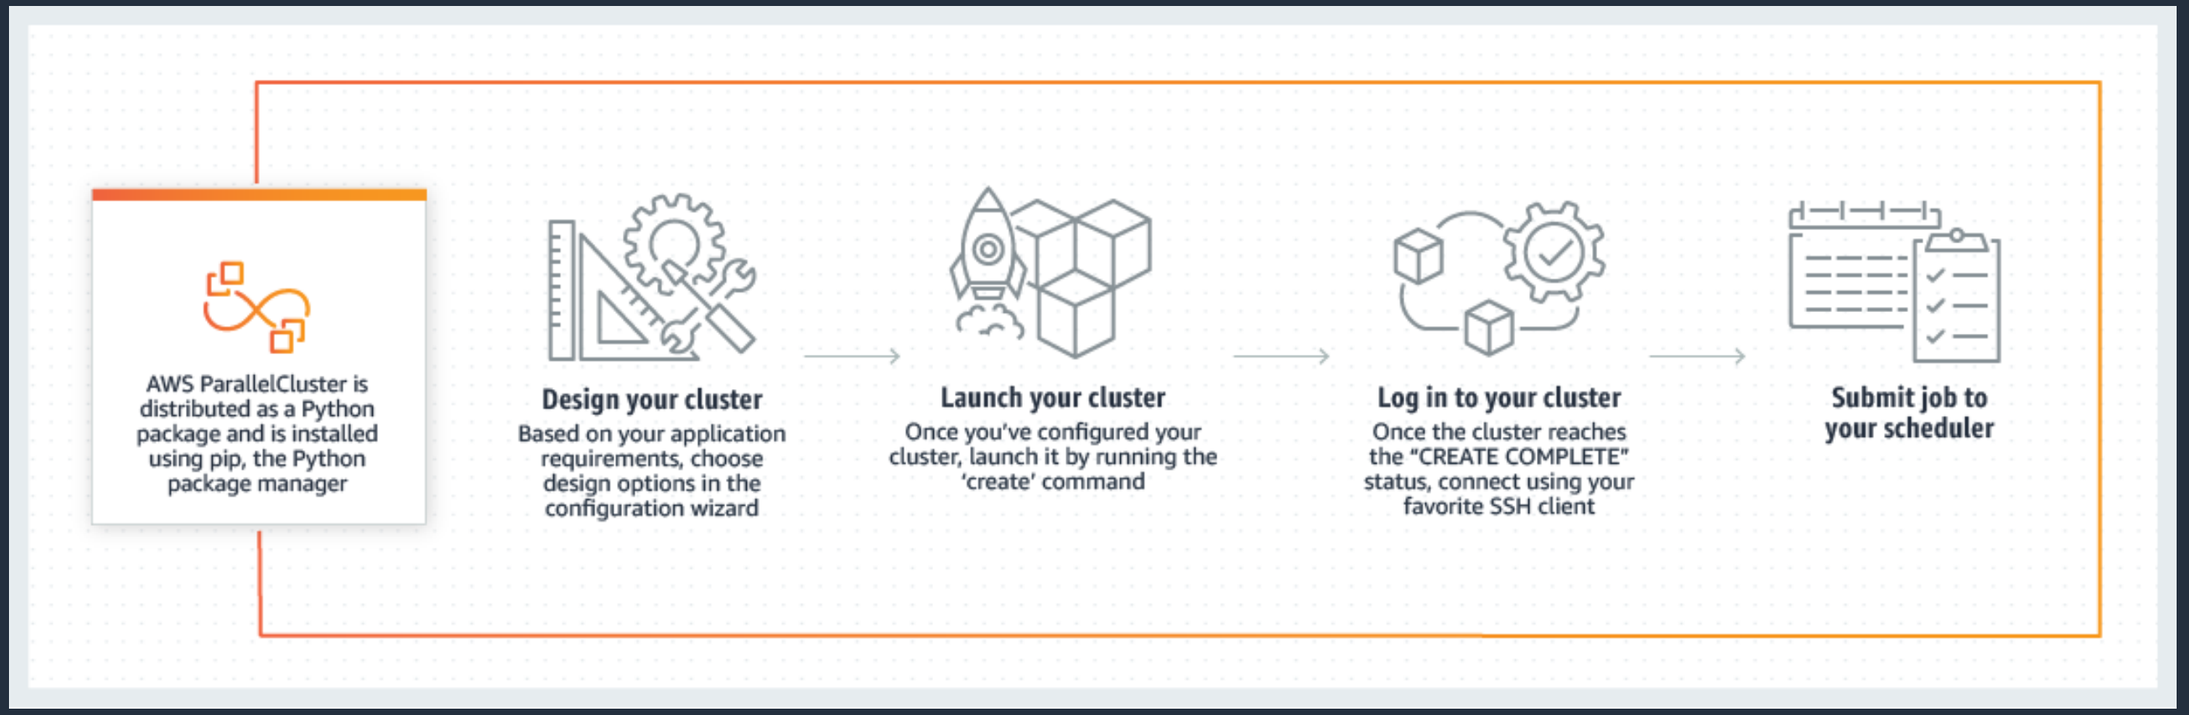

AWS ParallelCluster automates the deployment of your HPC cluster, and manages the auto-scaling of compute nodes which will run only when there are submitted jobs.

https://aws.amazon.com/hpc/parallelcluster/

aws

Collaborating with shared data in the cloud

Bring your algorithms to the data

Traditional approach:

Move data to computing resources

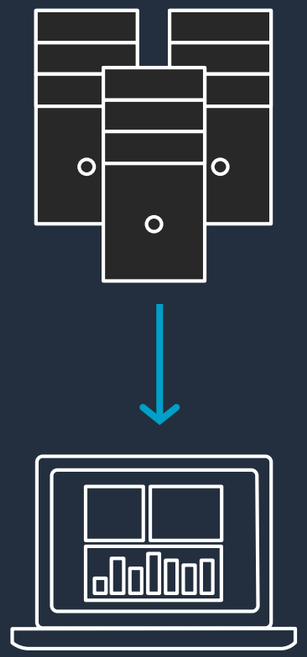

Cloud approach:

Move computing resources to data

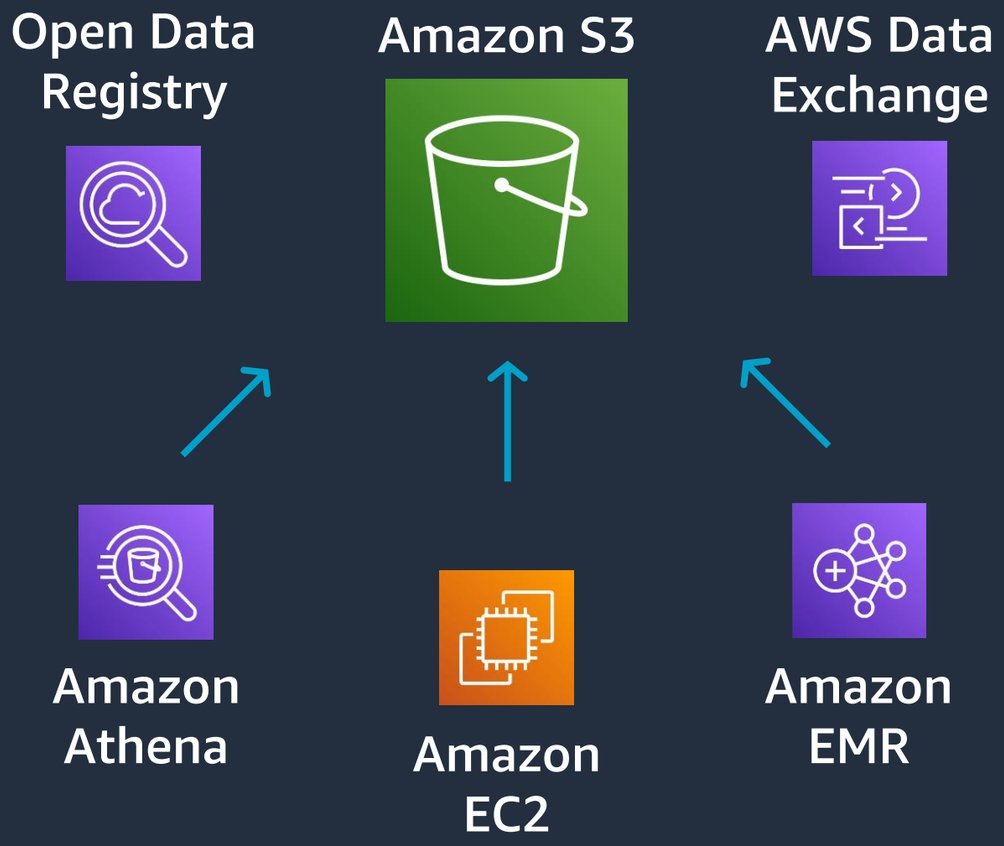

aws

AWS Open Data Registry

https://registry.opendata.aws/

NASA Space Act Agreement

https://registry.opendata.aws/collab/nasa/

"***Multi-Scale Ultra High Resolution Sea Surface Temperature***"

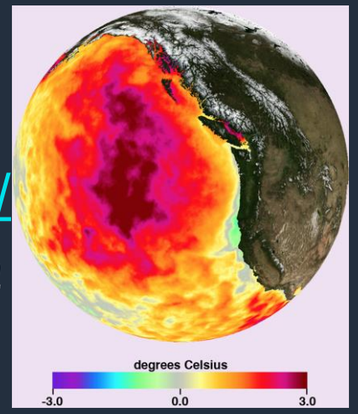

Dr. Chelle Gentemann, sr. scientist at Farallon Institute: "Something that took me three months and 3,000 lines of code now takes me 10 minutes with 20 lines of code. Now you don't have to have a big supercomputer and a system administrator to figure out how to download, store, and access the data. This is a transformative technology that's paving the way for the democratization of science."

https://earthdata.nasa.gov/learn/articles/tools-and-technology-articles/mur-sst-in-the-cloud

Estimating hurricane wind speeds with ML

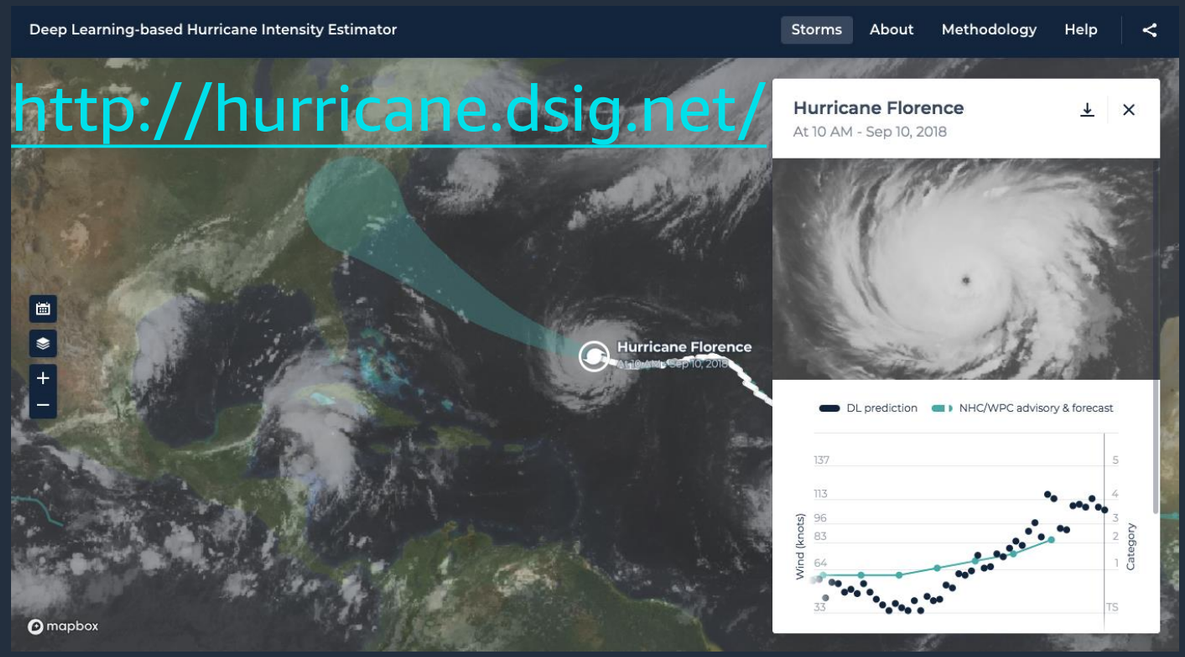

Uses the GOES16 satellite dataset on AWS, updated every 15 minutes. More accurate and much faster. (6h → 15 min)

https://aws.amazon.com/blogs/publicsector/estimating-hurricane-wind-speeds-with-machine-learning/

aws

Research Flywheel

Teach using data sets and similar provisioning

Educate

Better links between research and teaching

Teach from the same resources

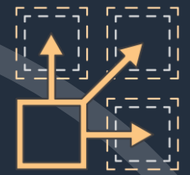

Metadata

Key metadata programmatically generated and passed between all stages

Gain better insights

Meta analysis, trends of success

Central Catalogue

Data sets are tagged, discoverable and sharable

Control through visibility

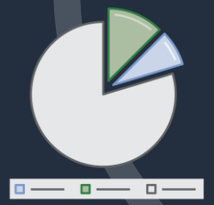

Unified Reporting & Analysis

Guardrails

Microservices approach to provide guard rails rather than rigid controls

Cloud delivery

Cloud delivery of resources for easy access drives more efficient use

Speed to Research

Balance control of spend with ability to innovate

aws

Thank You!

HPC on AWS: https://aws.amazon.com/hpc

Open data on AWS: https://opendata.aws/

Steve Gillard
Senior Solutions Architect – Government & Research
Amazon Web Services

gillards@amazon.com

aws