Porting and Benchmarking GRAF on AWS

Timothy Brown

Principal Solutions Architect, AWS

Global High-Resolution Atmospheric Forecasting System

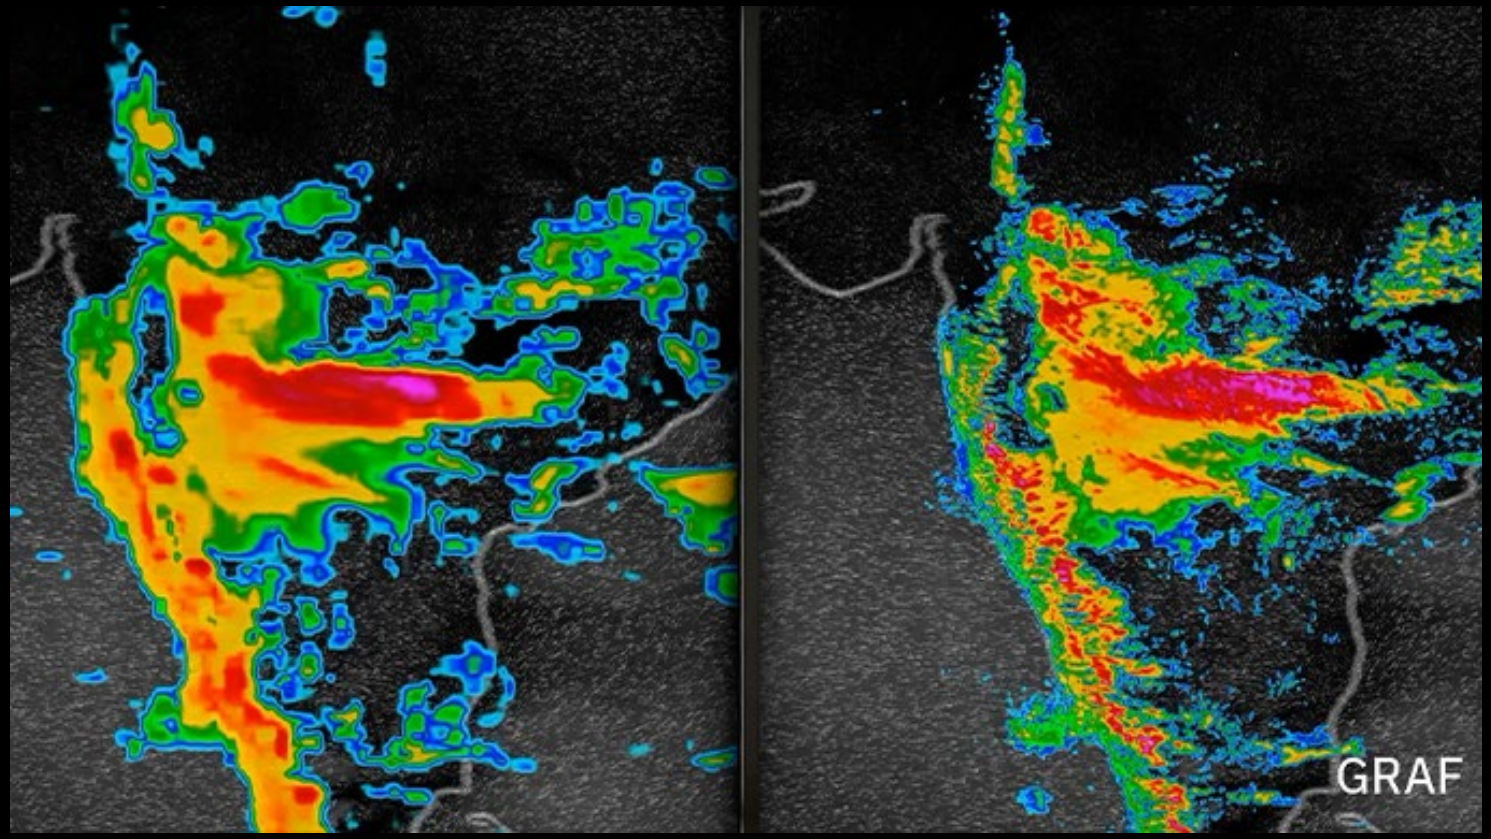

The Weather Company

- Hourly forecasts
- 5 min output for US and EU
- 3.5 km over US and EU
- GPU optimized
- Modelling system

Global High-Resolution Atmospheric Forecasting System

- MPAS Atmosphere
  - Version 7
  - OpenACC
  - RRTMG scheme is on CPUs

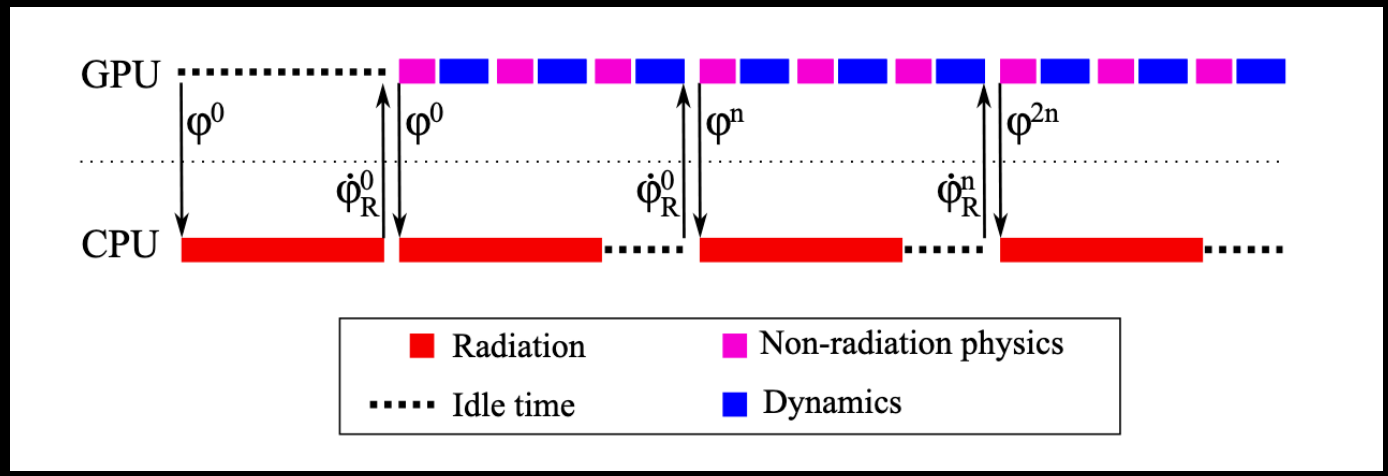

https://www.mmm.ucar.edu/models/mpas

Reference AWS HPC Stack

Remote Visualization	Amazon DCV /Research and Engineering Studio
Cluster Management / Job Scheduling	AWS Parallel Computing Service / AWS Batch
CPU Compute + GPU Acceleration	EC2: Hpc7g, Hpc7a, Hpc6a, C8gn, C6id, P6, P5, P4
High Performance Parallel File System	Amazon FSxfor Lustre
High BW + Low Latency Networking	EFA (SRD Transport Protocol)
High Throughput Block Storage	Amazon EBS
Scalable & Durable Object Storage	Amazon S3

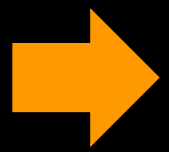

Reference Architecture – AWS Parallel Computing Service

Configuration required resources

- AMI
- GPU instance
- Storage
- Software Stack

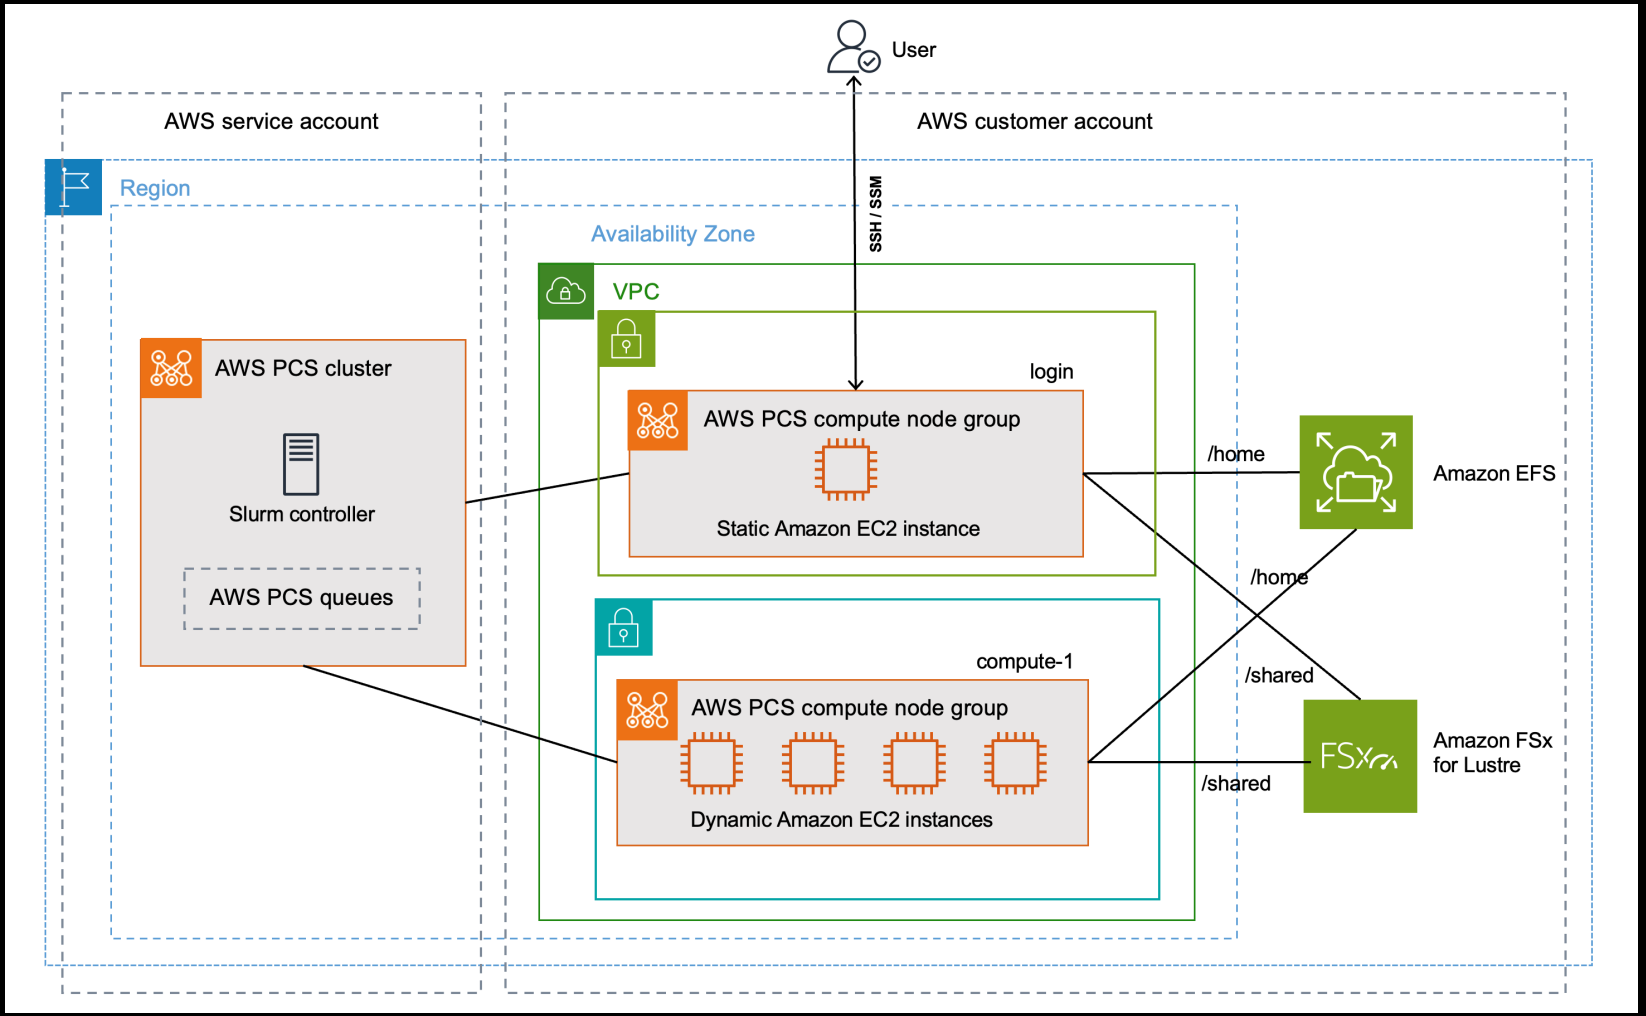

Building a network building block

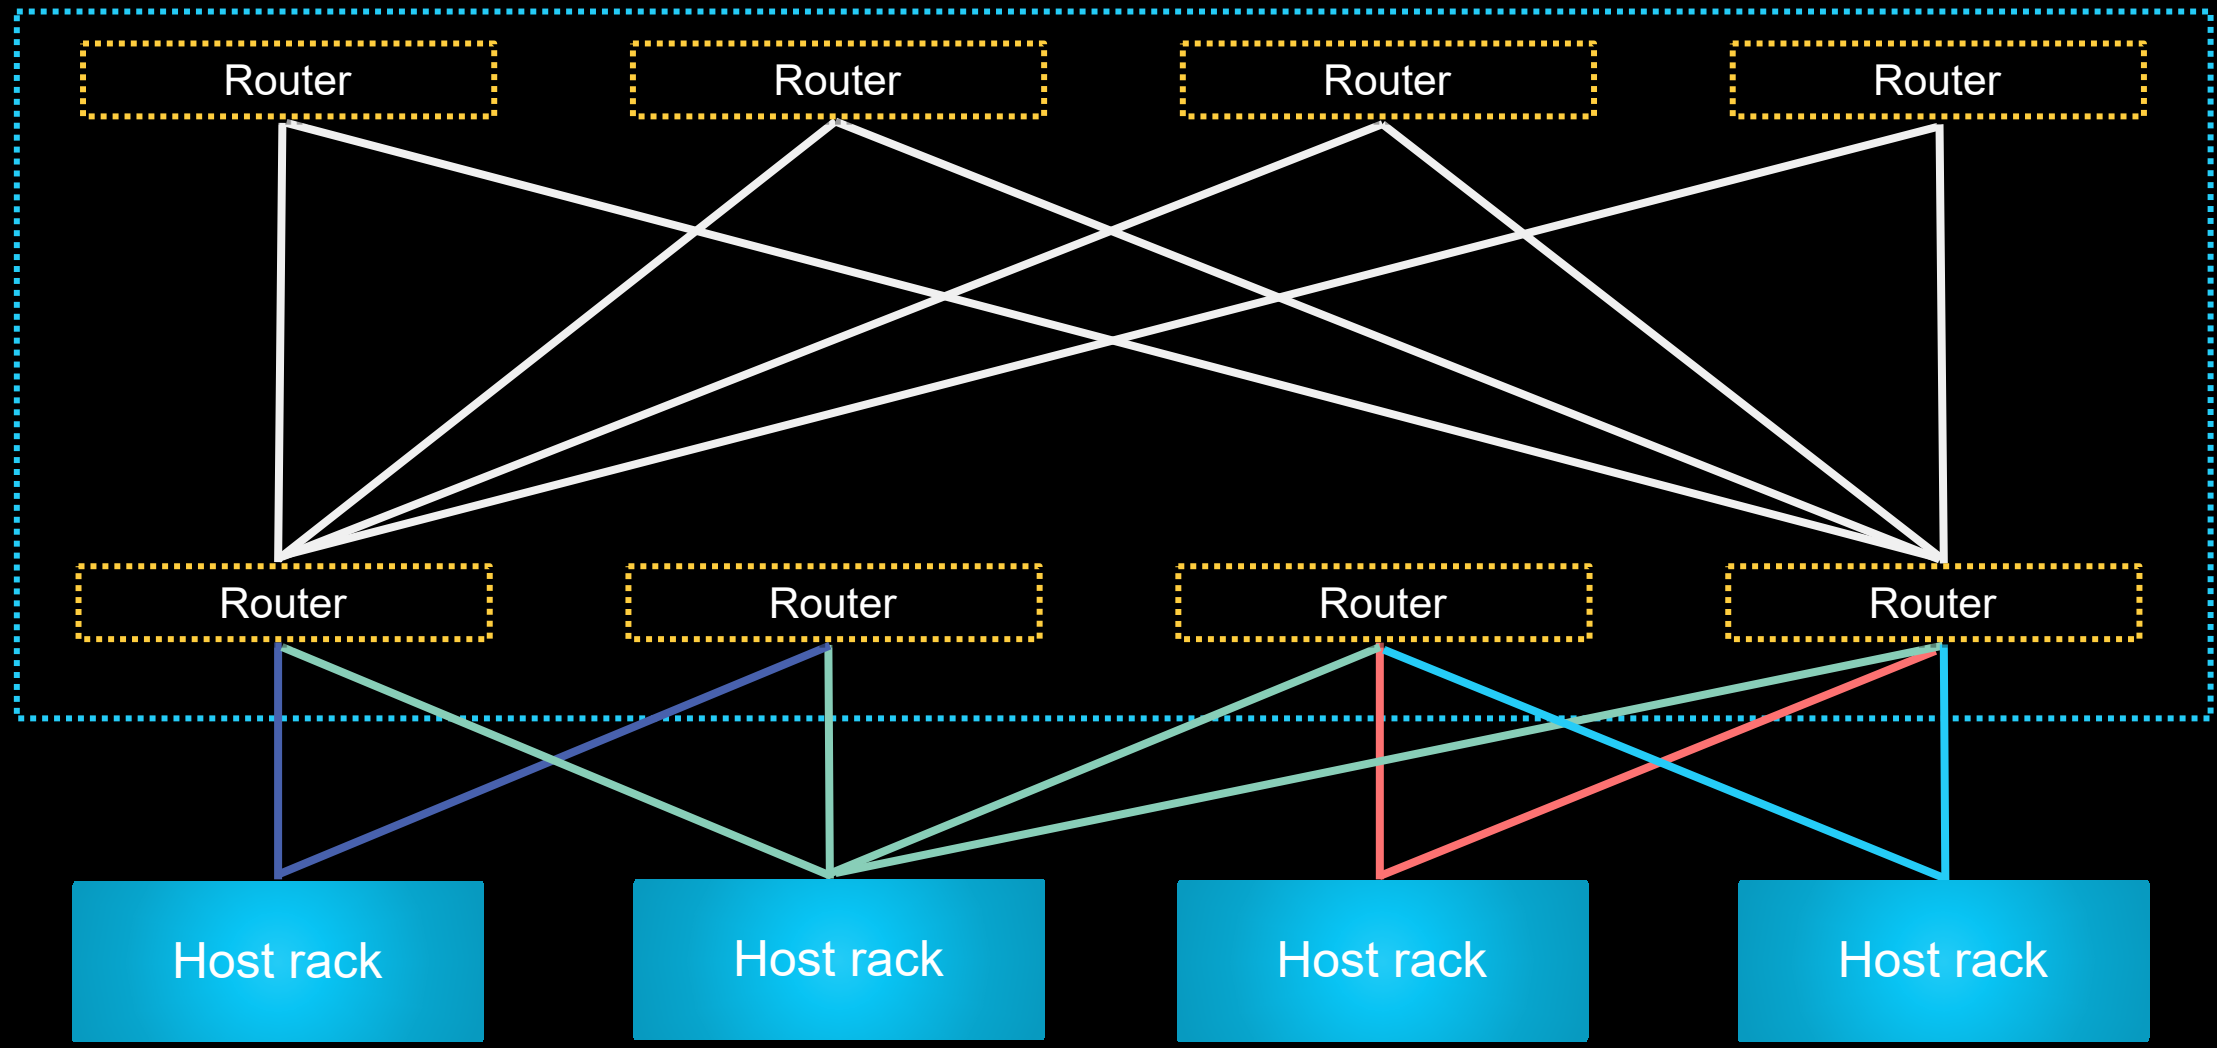

Source: AWS re:Invent 2019: Innovation and operation of the AWS global network infrastructure (NET339)

Cellular network pattern

Scale out by adding cells

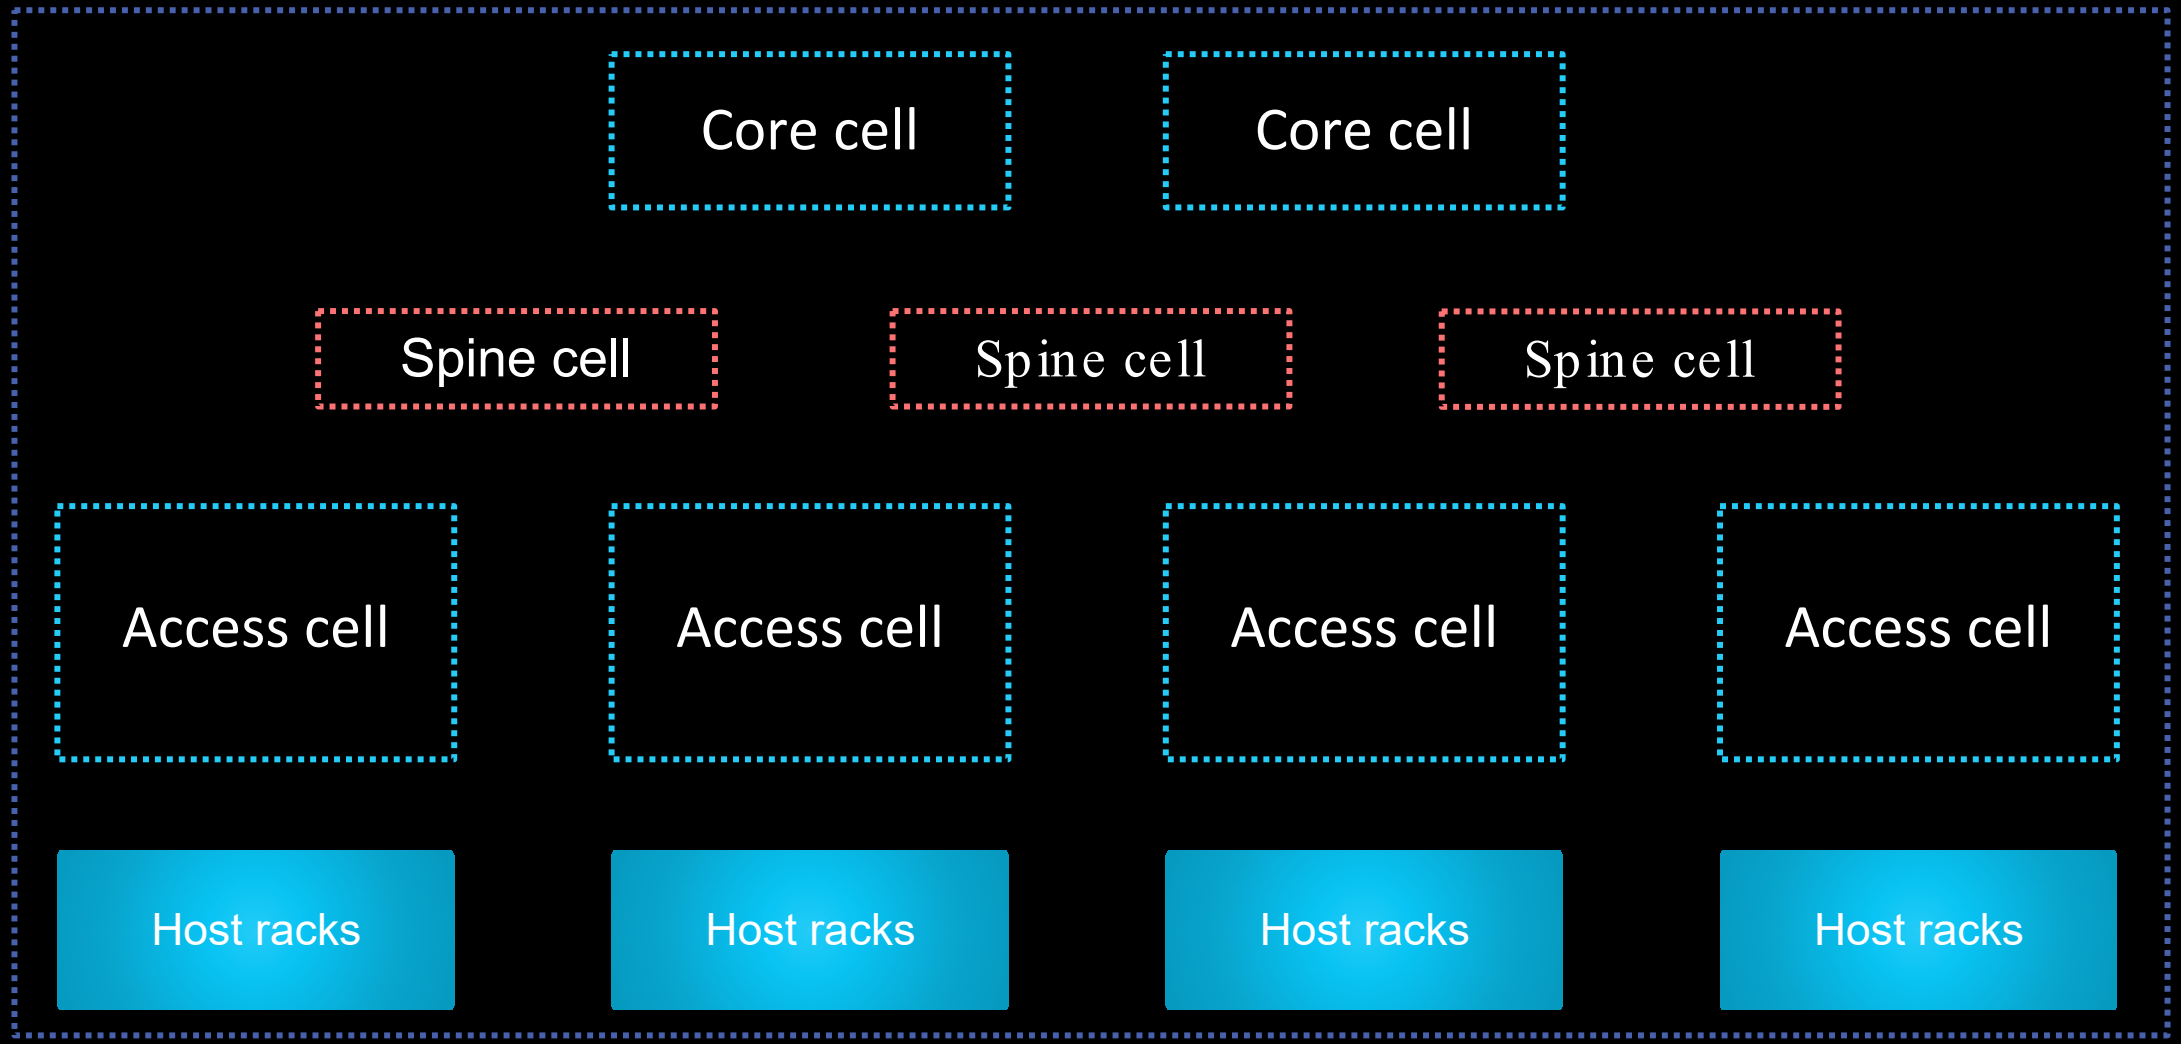

Source: AWS re:Invent 2019: Innovation and operation of the AWS global network infrastructure (NET339)

Amazon EC2 Capacity Blocks

- Reserve a start time up to 8 weeks in advance
- Set a reservation duration of one to 14 days or a multiple of 7 days, up to 182 days
- Configure up to 64 instances per Capacity Block
- Configure up to 256 instances across multiple Capacity Blocks
- Capacity Blocks are in a single Placement Group

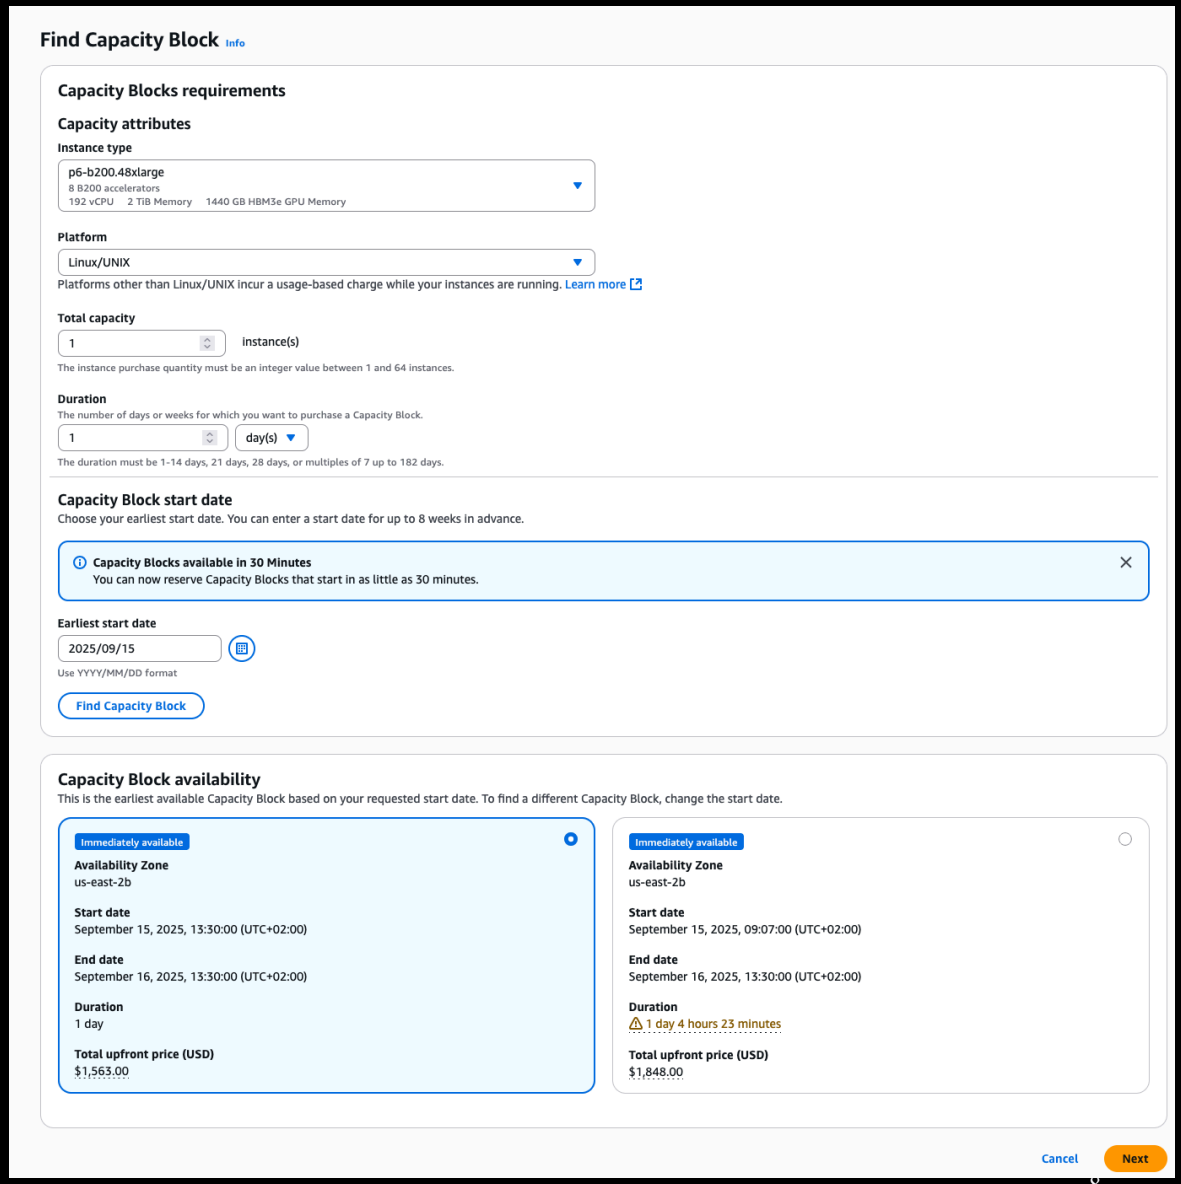

EC2 Topology

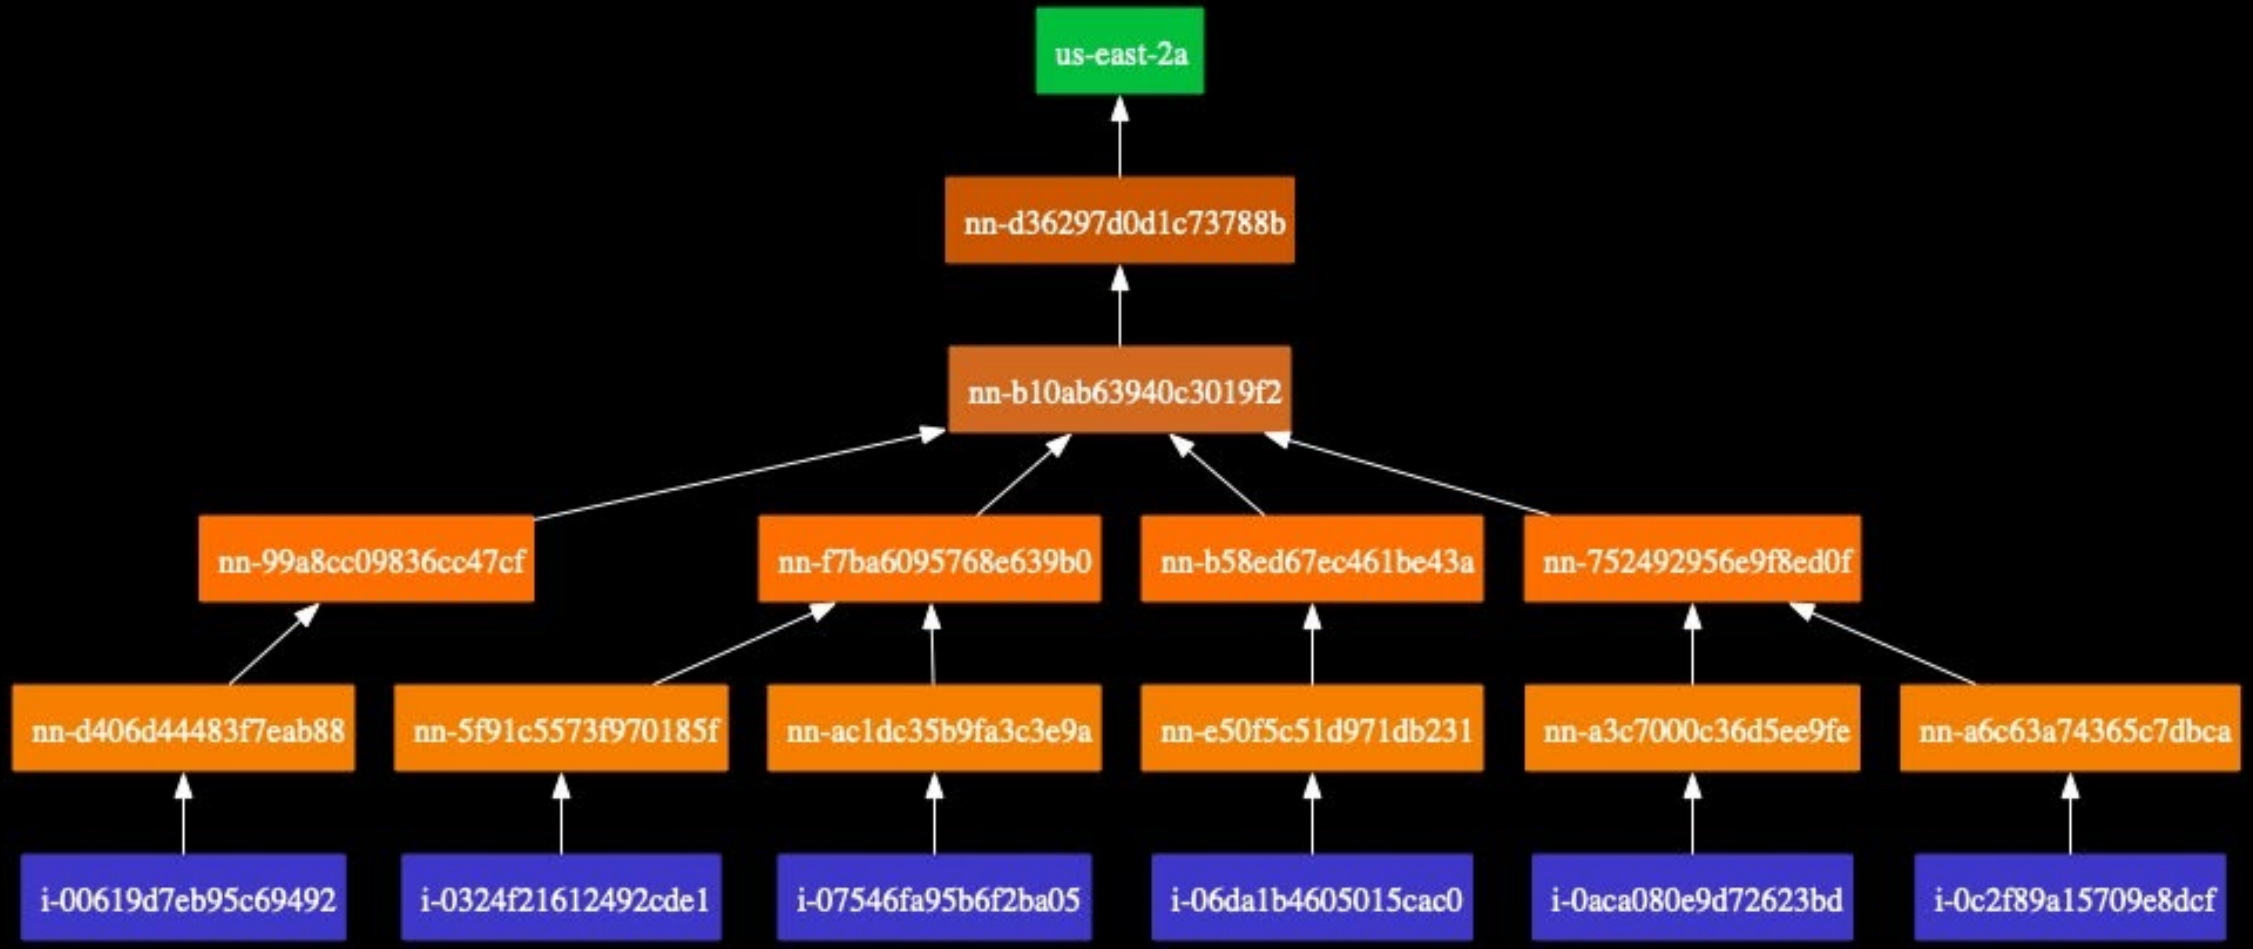

https://github.com/aws-samples/ec2-topology-aware-for-slurm

AWS P5 Instance

P4.48xlarge
- Intel
- 8 NVIDIA A100
- 400 Gbps EFA
- 8 TB NVMe

P5.48xlarge
- Intel
- 8 NVIDIA H100
- 3,200 Gbps EFA
- 30 TB NVMe

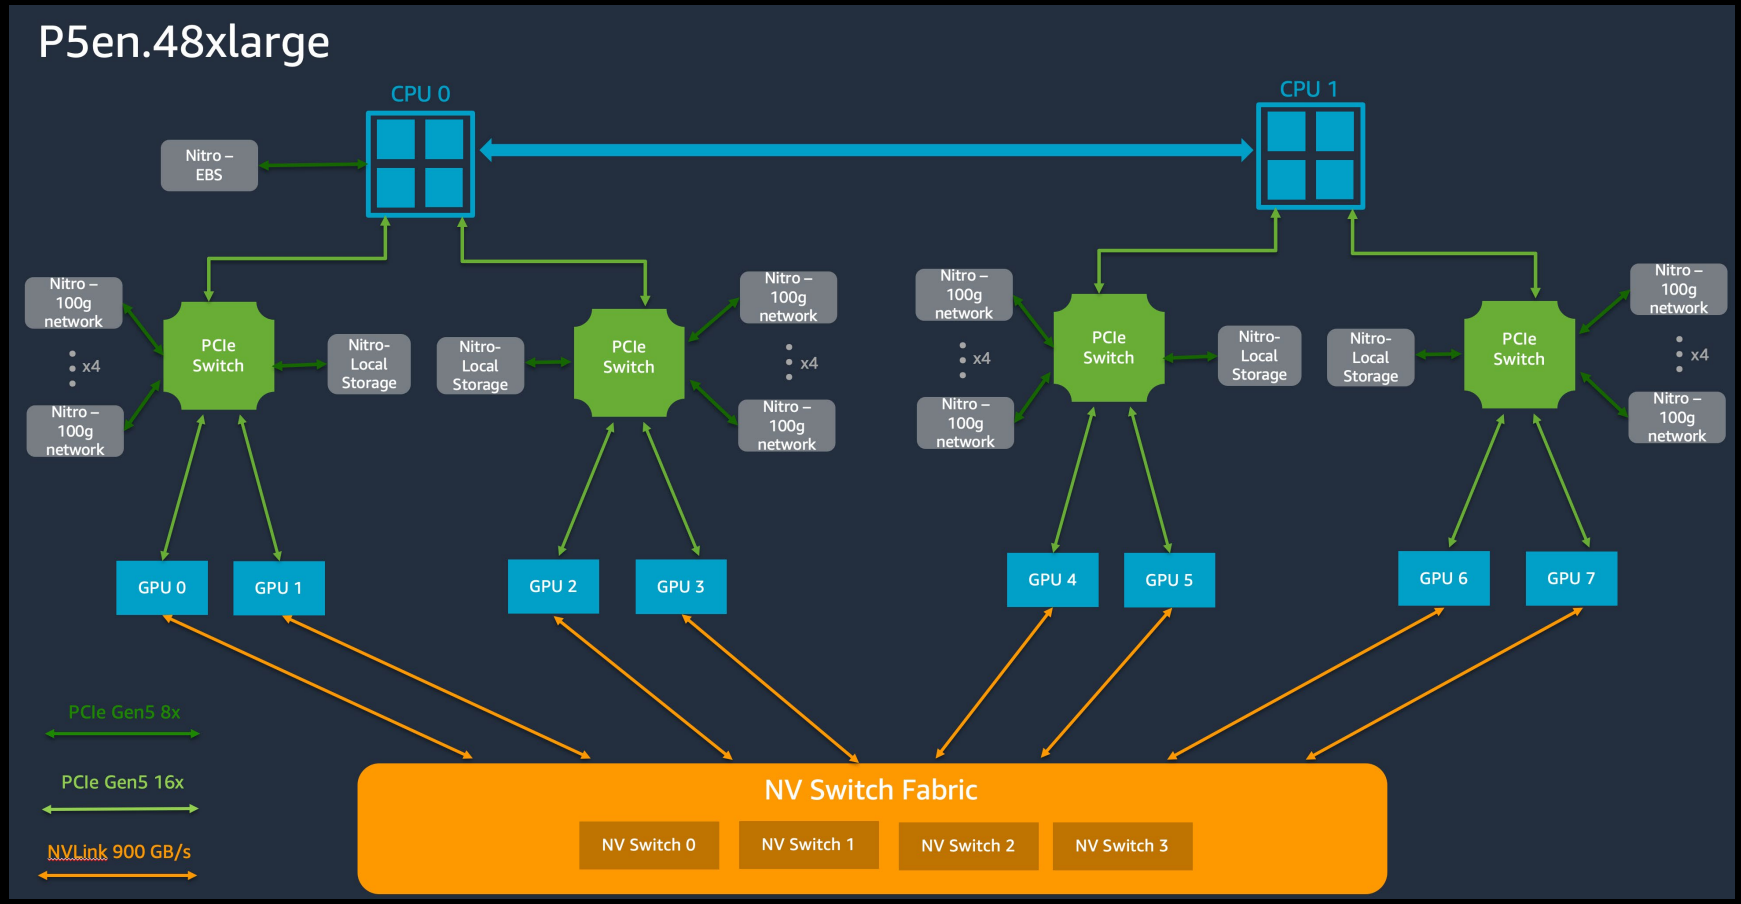

AWS P5 Instance

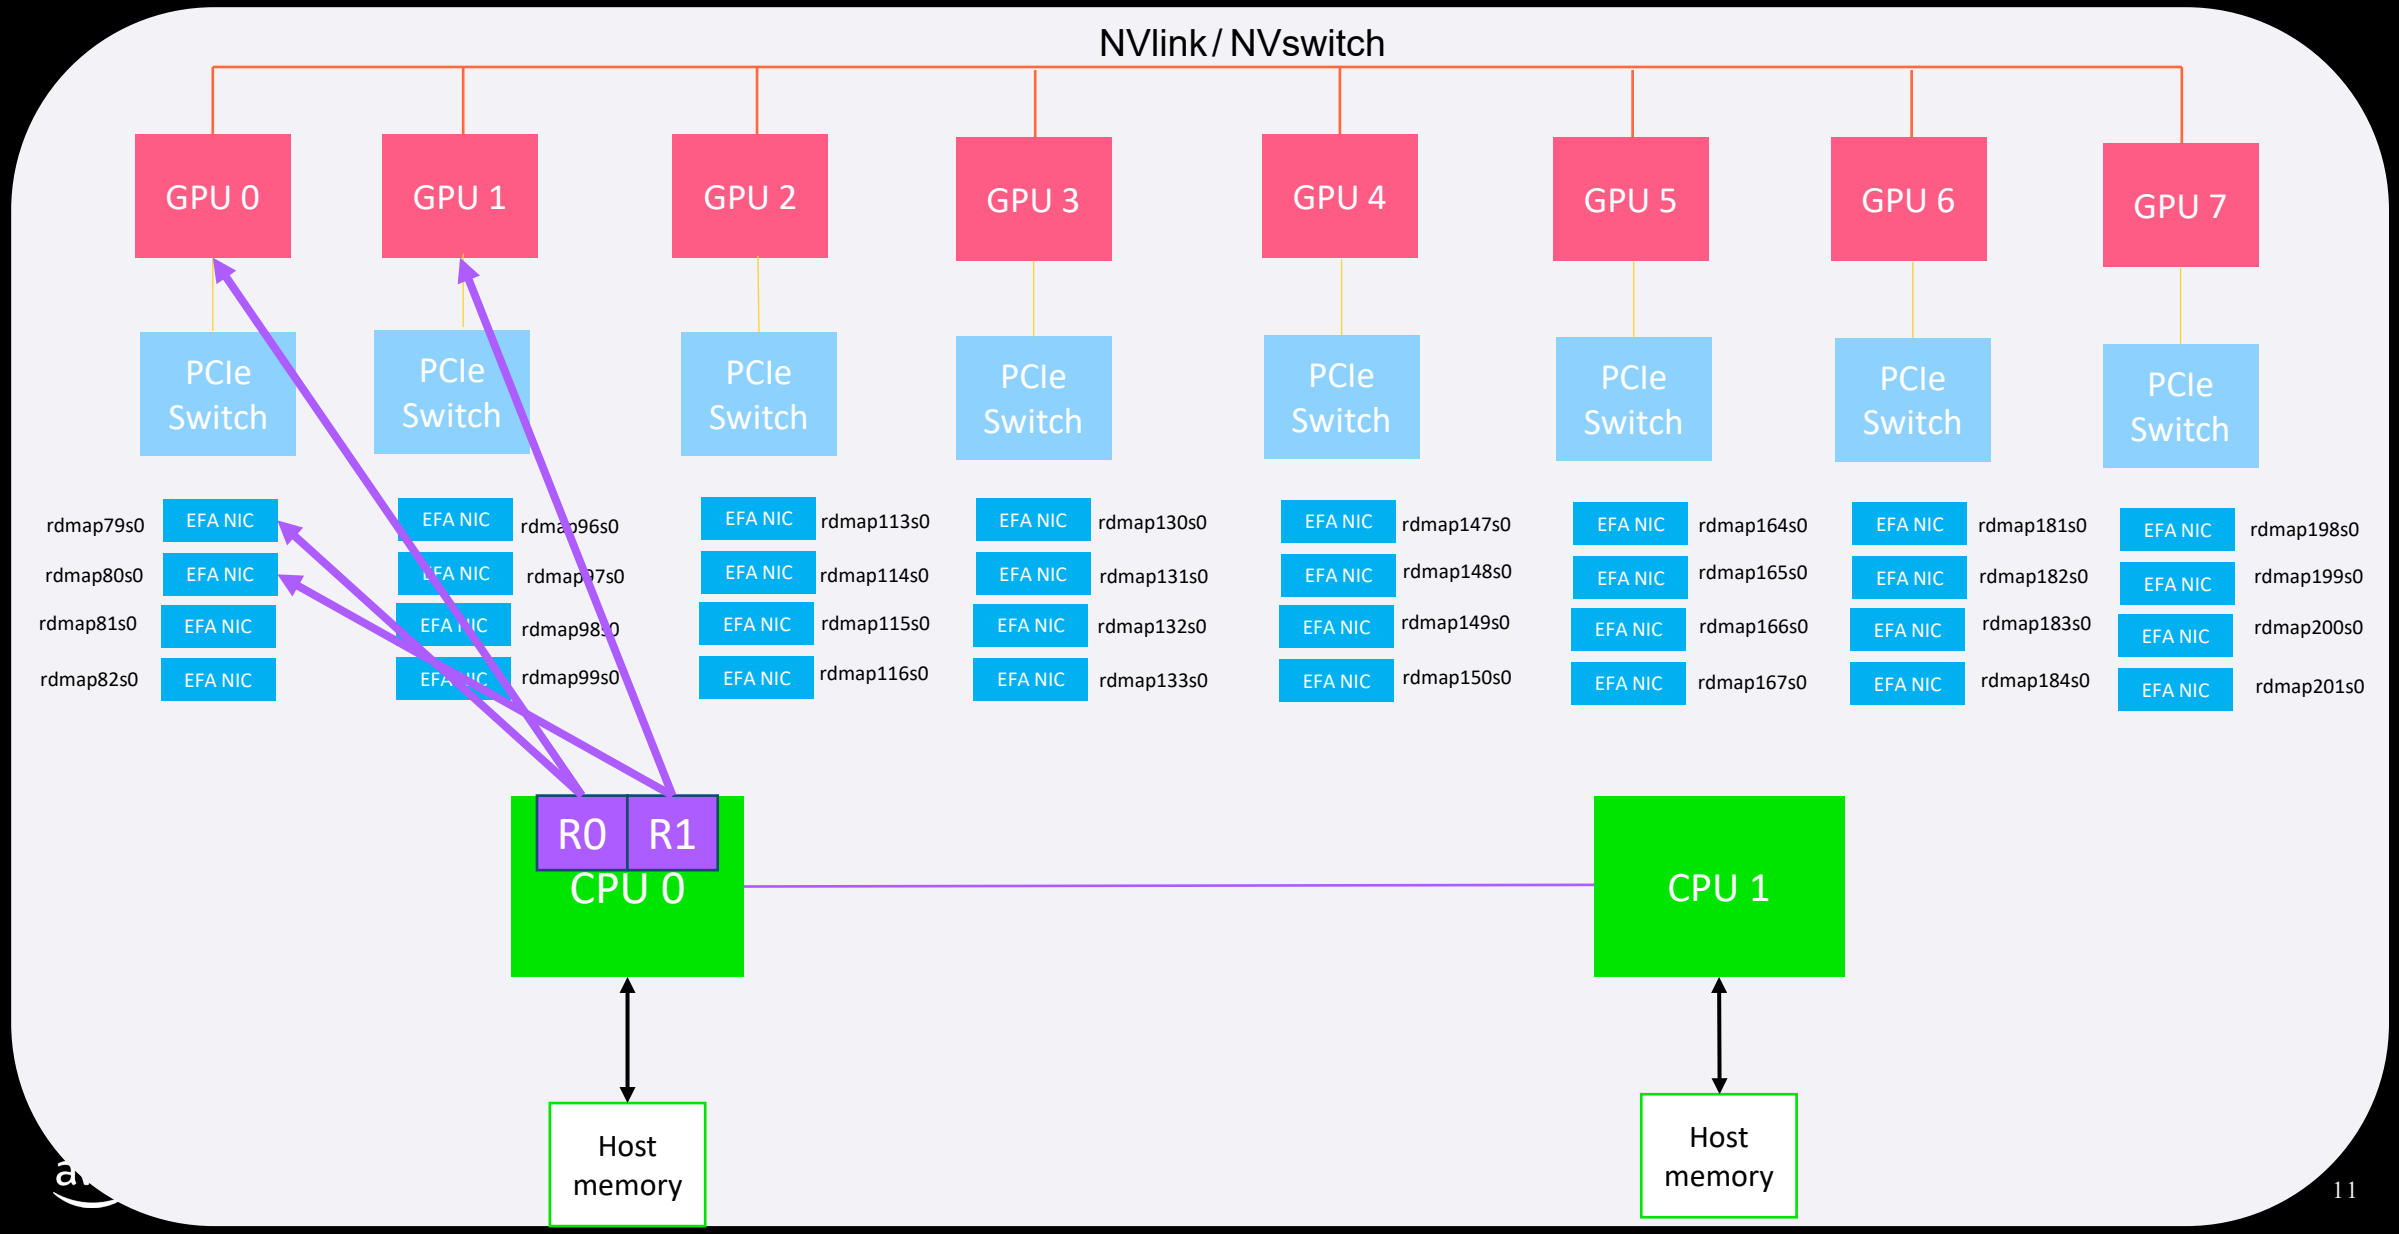

Flamegrph

- With bad NIC bindings
- Exorbitant time is spent in libfabric

NEW!!

12/9/25 Extended EFA metrics in /sys

https://aws.amazon.com/about-aws/whats-new/2025/09/aws-efa-metrics-improved-observability-networking/

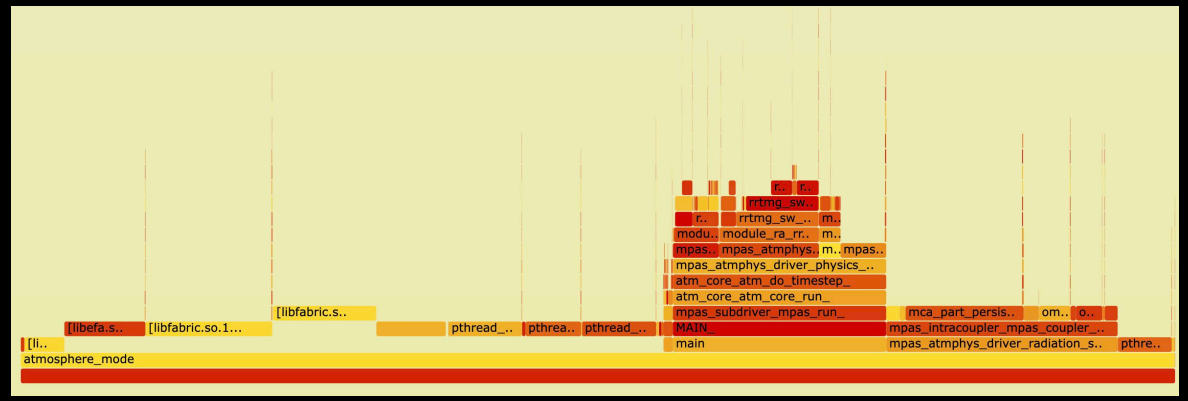

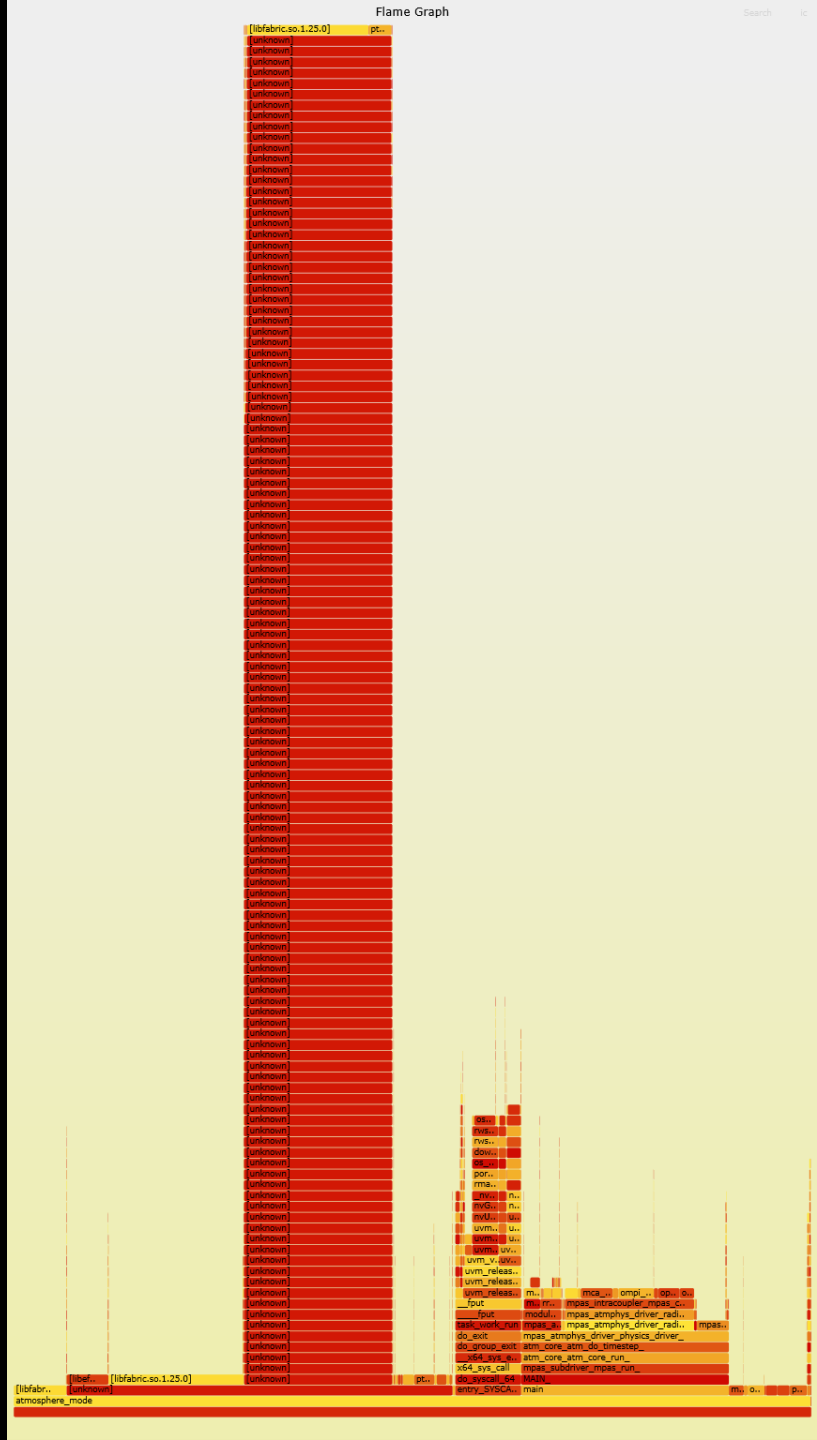

Performance Requirements & Results

- Hourly global forecast
- MPAS-A maximum runtime 51 minutes
- Pre and post can have overlapping running time – to an extent

Conclusion & Future Work

- Able to meet production run -time requirements of 51 minutes
- With 14 p5 instances, 112 H100. Vs 840 V100.

- Benchmark on p5en instances.
- Investigate optimizations for p6 -b200 instances (Blackwell GPUs)
- Investigate IO performance and the possibility of async IO

Thank you

tpbrown@amazon.com