Scaling Message Passing on Amazon Web Services with Elastic Fabric Adapter

Raghunath Raja Chandrasekar

© 2020, Amazon Web Services, Inc. or its Affiliates.

Agenda

- Overview of HPC on AWS
- Evolution of networking on AWS
- What is EFA?
- Scalable Reliable Datagrams (SRD) Protocol
- MPI implementations and support for EFA
- Application scaling case-studies

© 2020, Amazon Web Services, Inc. or its Affiliates.

HPC on AWS

© 2020, Amazon Web Services, Inc. or its Affiliates.

Time-to-results as the Metric for Success

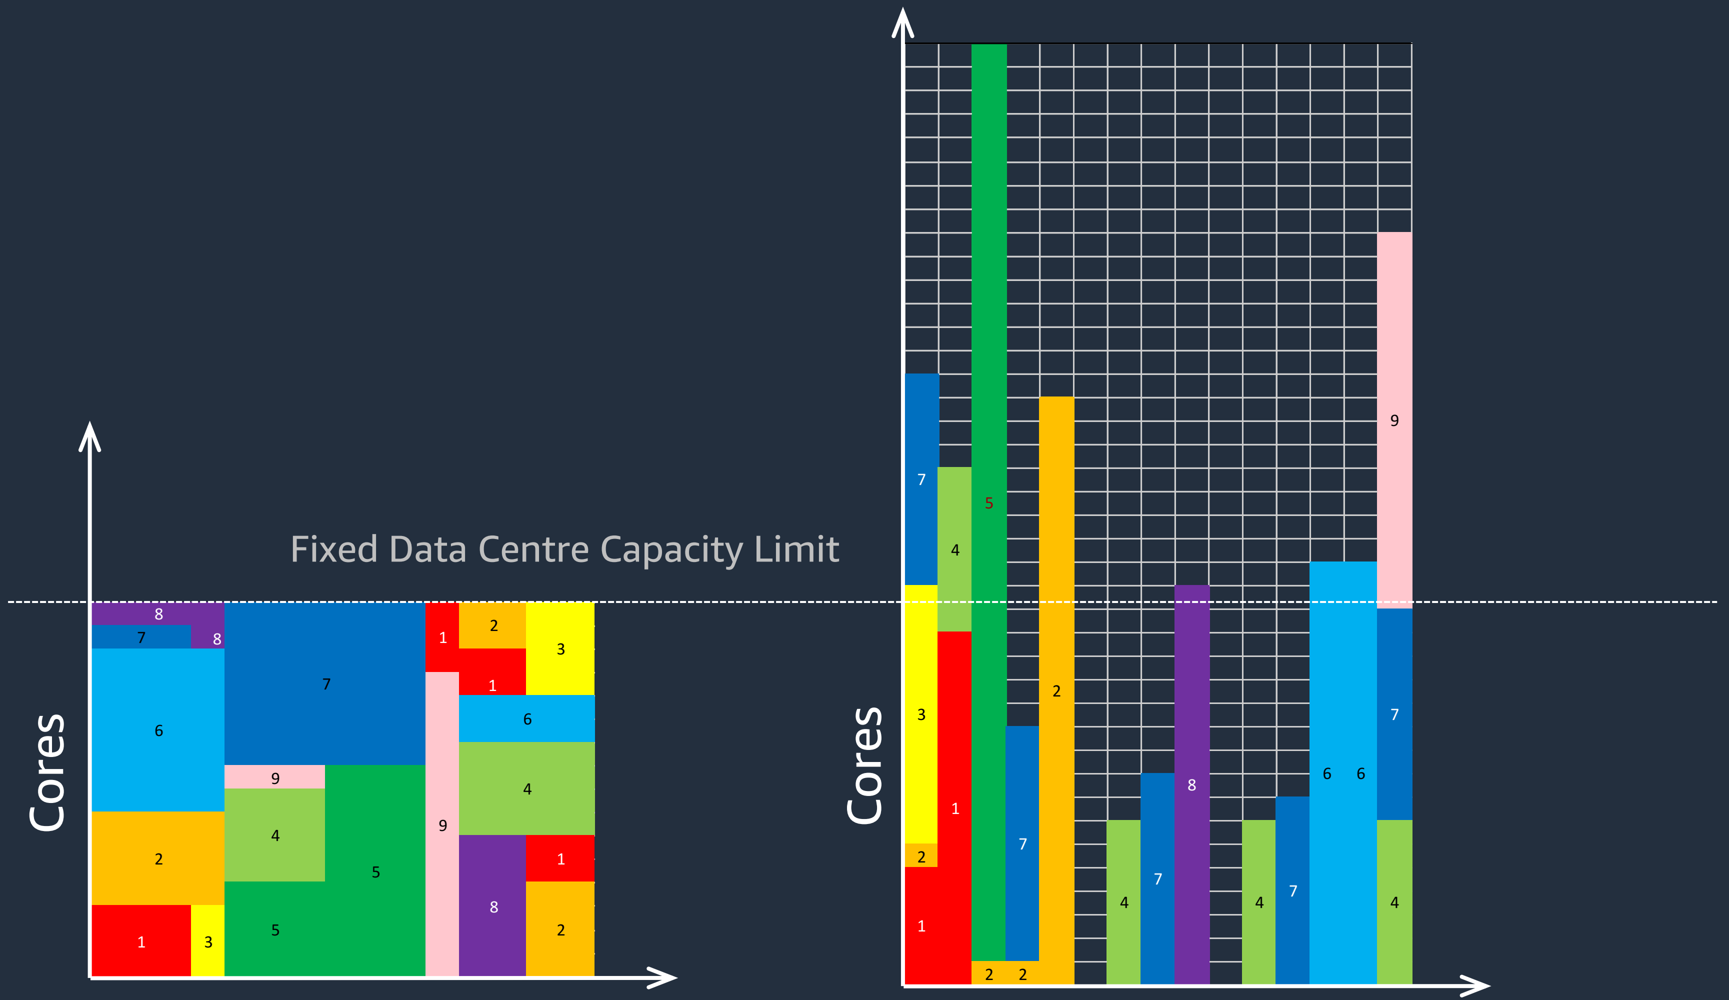

Finite capacity, usually with long queues to wait in.

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

© 2020, Amazon Web Services, Inc. or its Affiliates.

HPC on AWS: Solution Components

Automation and orchestration

- AWS Batch
- AWS ParallelCluster
- NICE EnginFrame

Storage

- Amazon EBS
- Amazon FSx for Lustre
- Amazon EFS
- Amazon S3

Compute

- Amazon EC2 instances (Compute and accelerated)
- Amazon EC2 Spot
- AWS Auto Scaling

Visualization

- NICE DCV
- Amazon AppStream 2.0

Networking

- Enhanced networking
- Placement groups
- Elastic Fabric Adapter

© 2020, Amazon Web Services, Inc. or its Affiliates.

Broad HPC Partner Community

Application partners

onscale

SIMSCALE

ANSYS

MAXELER Technologies
Maximum Performance Computing

CFD Direct

aneo
the other solution

Ridgeway Kite

rescale

DASSAULT SYSTEMES

RFD
Rock Flow Dynamics

TotalCAE

DNAnexus

SIEMENS
Ingenuity for life

illumina

ZENOTECH
SIMULATION UNLIMITED

Technology partners

intel

NVIDIA

AMD

arm

Consulting partners

CADFEM

DXC.technology

six nines

RONIN

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws

Global Infrastructure

We add the equivalent of an entire Fortune 500 company's compute capacity every day

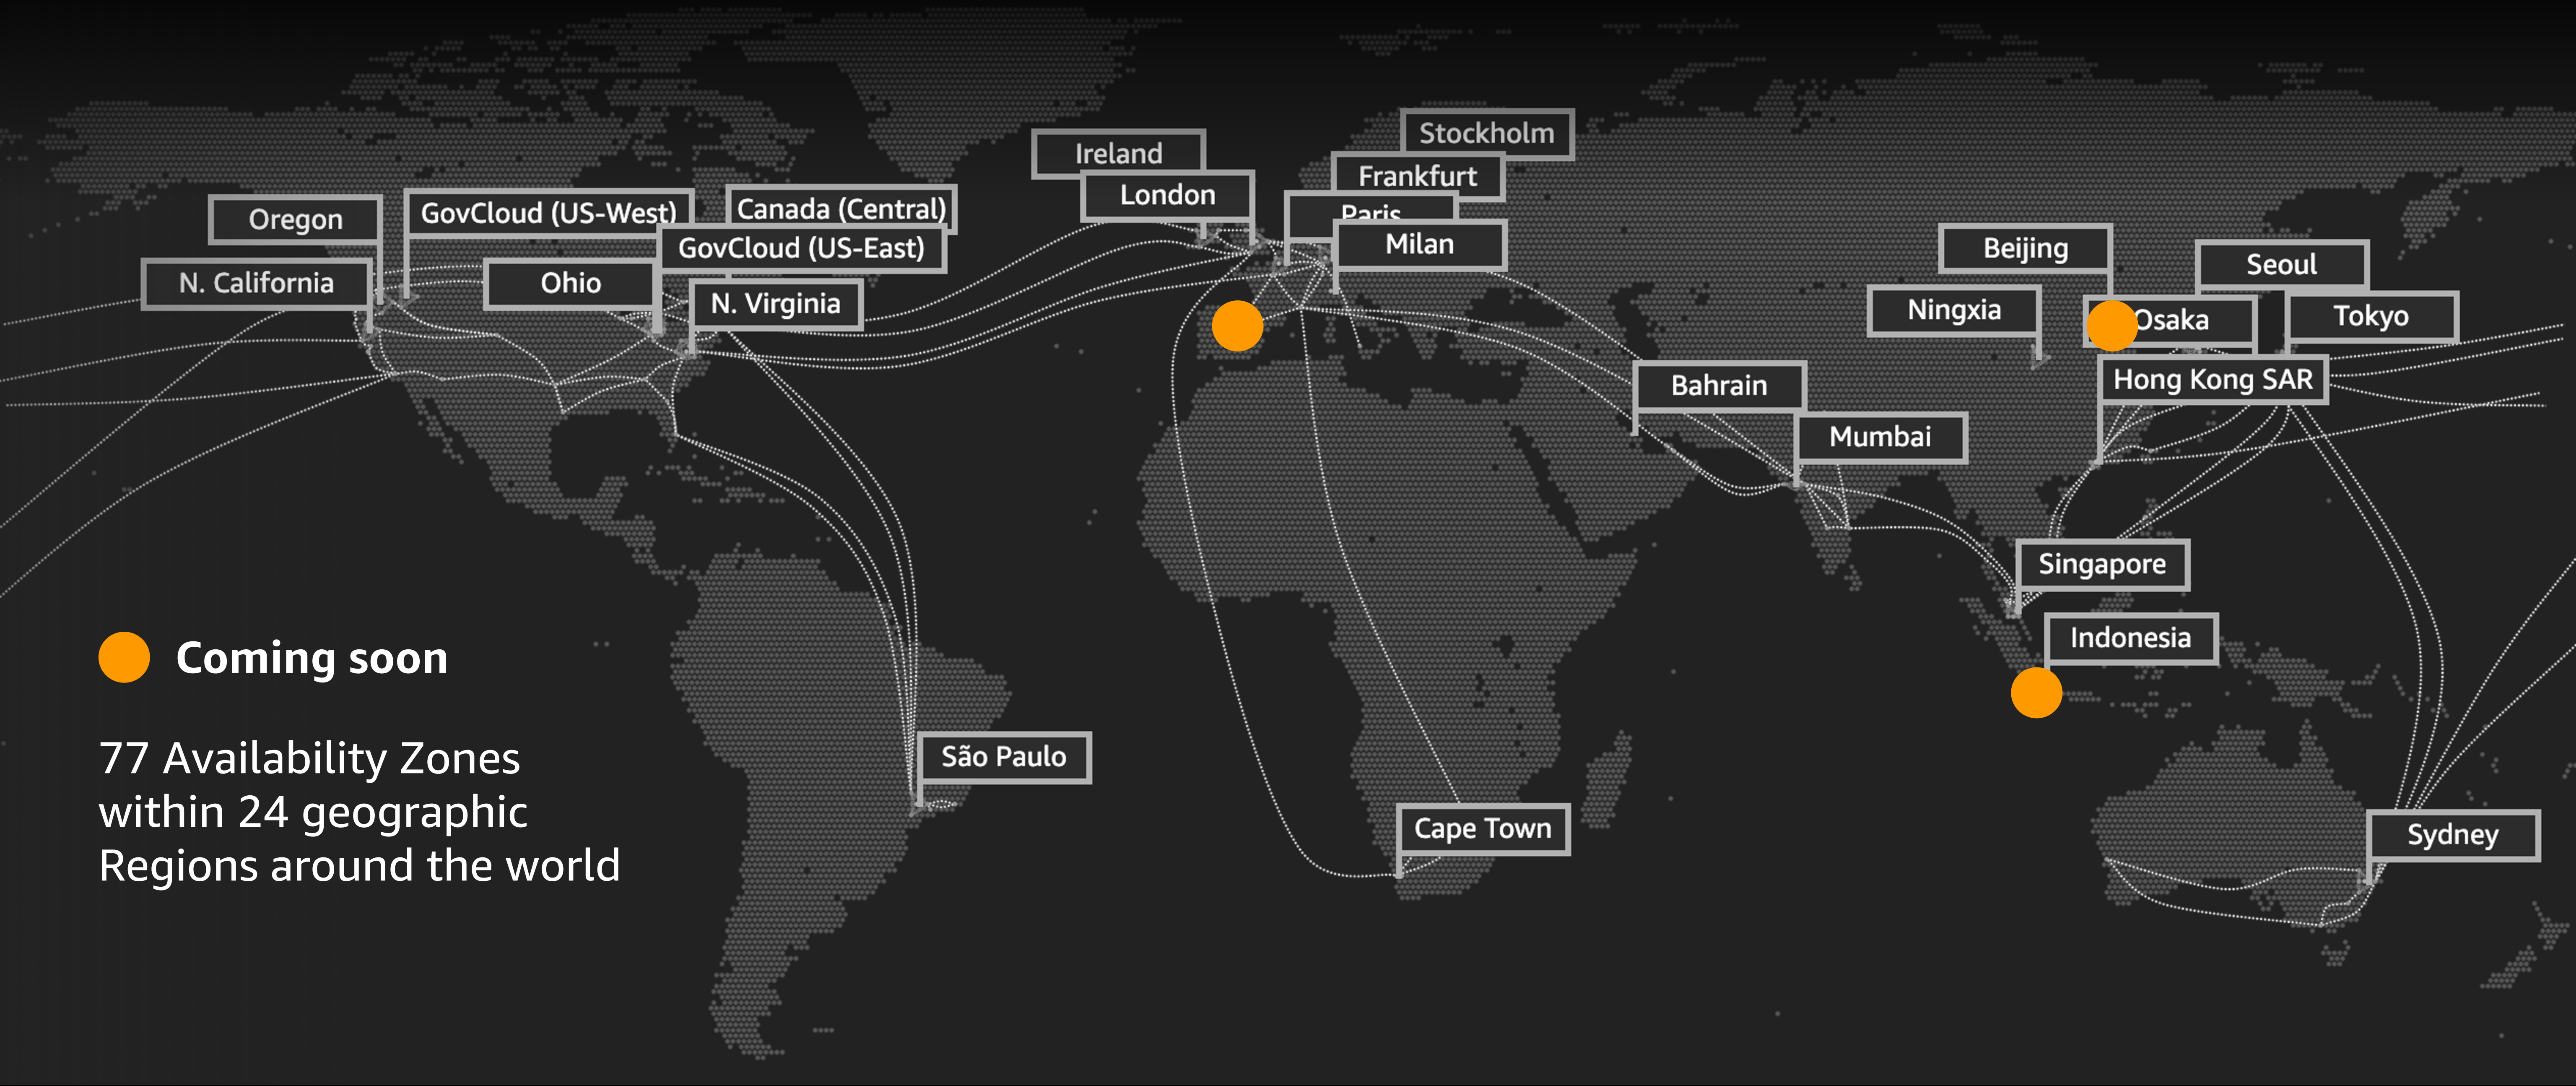

● Coming soon

77 Availability Zones within 24 geographic Regions around the world

Evolution of Networking on AWS

© 2020, Amazon Web Services, Inc. or its Affiliates.

Anatomy of an EC2 Instance – Circa 2011

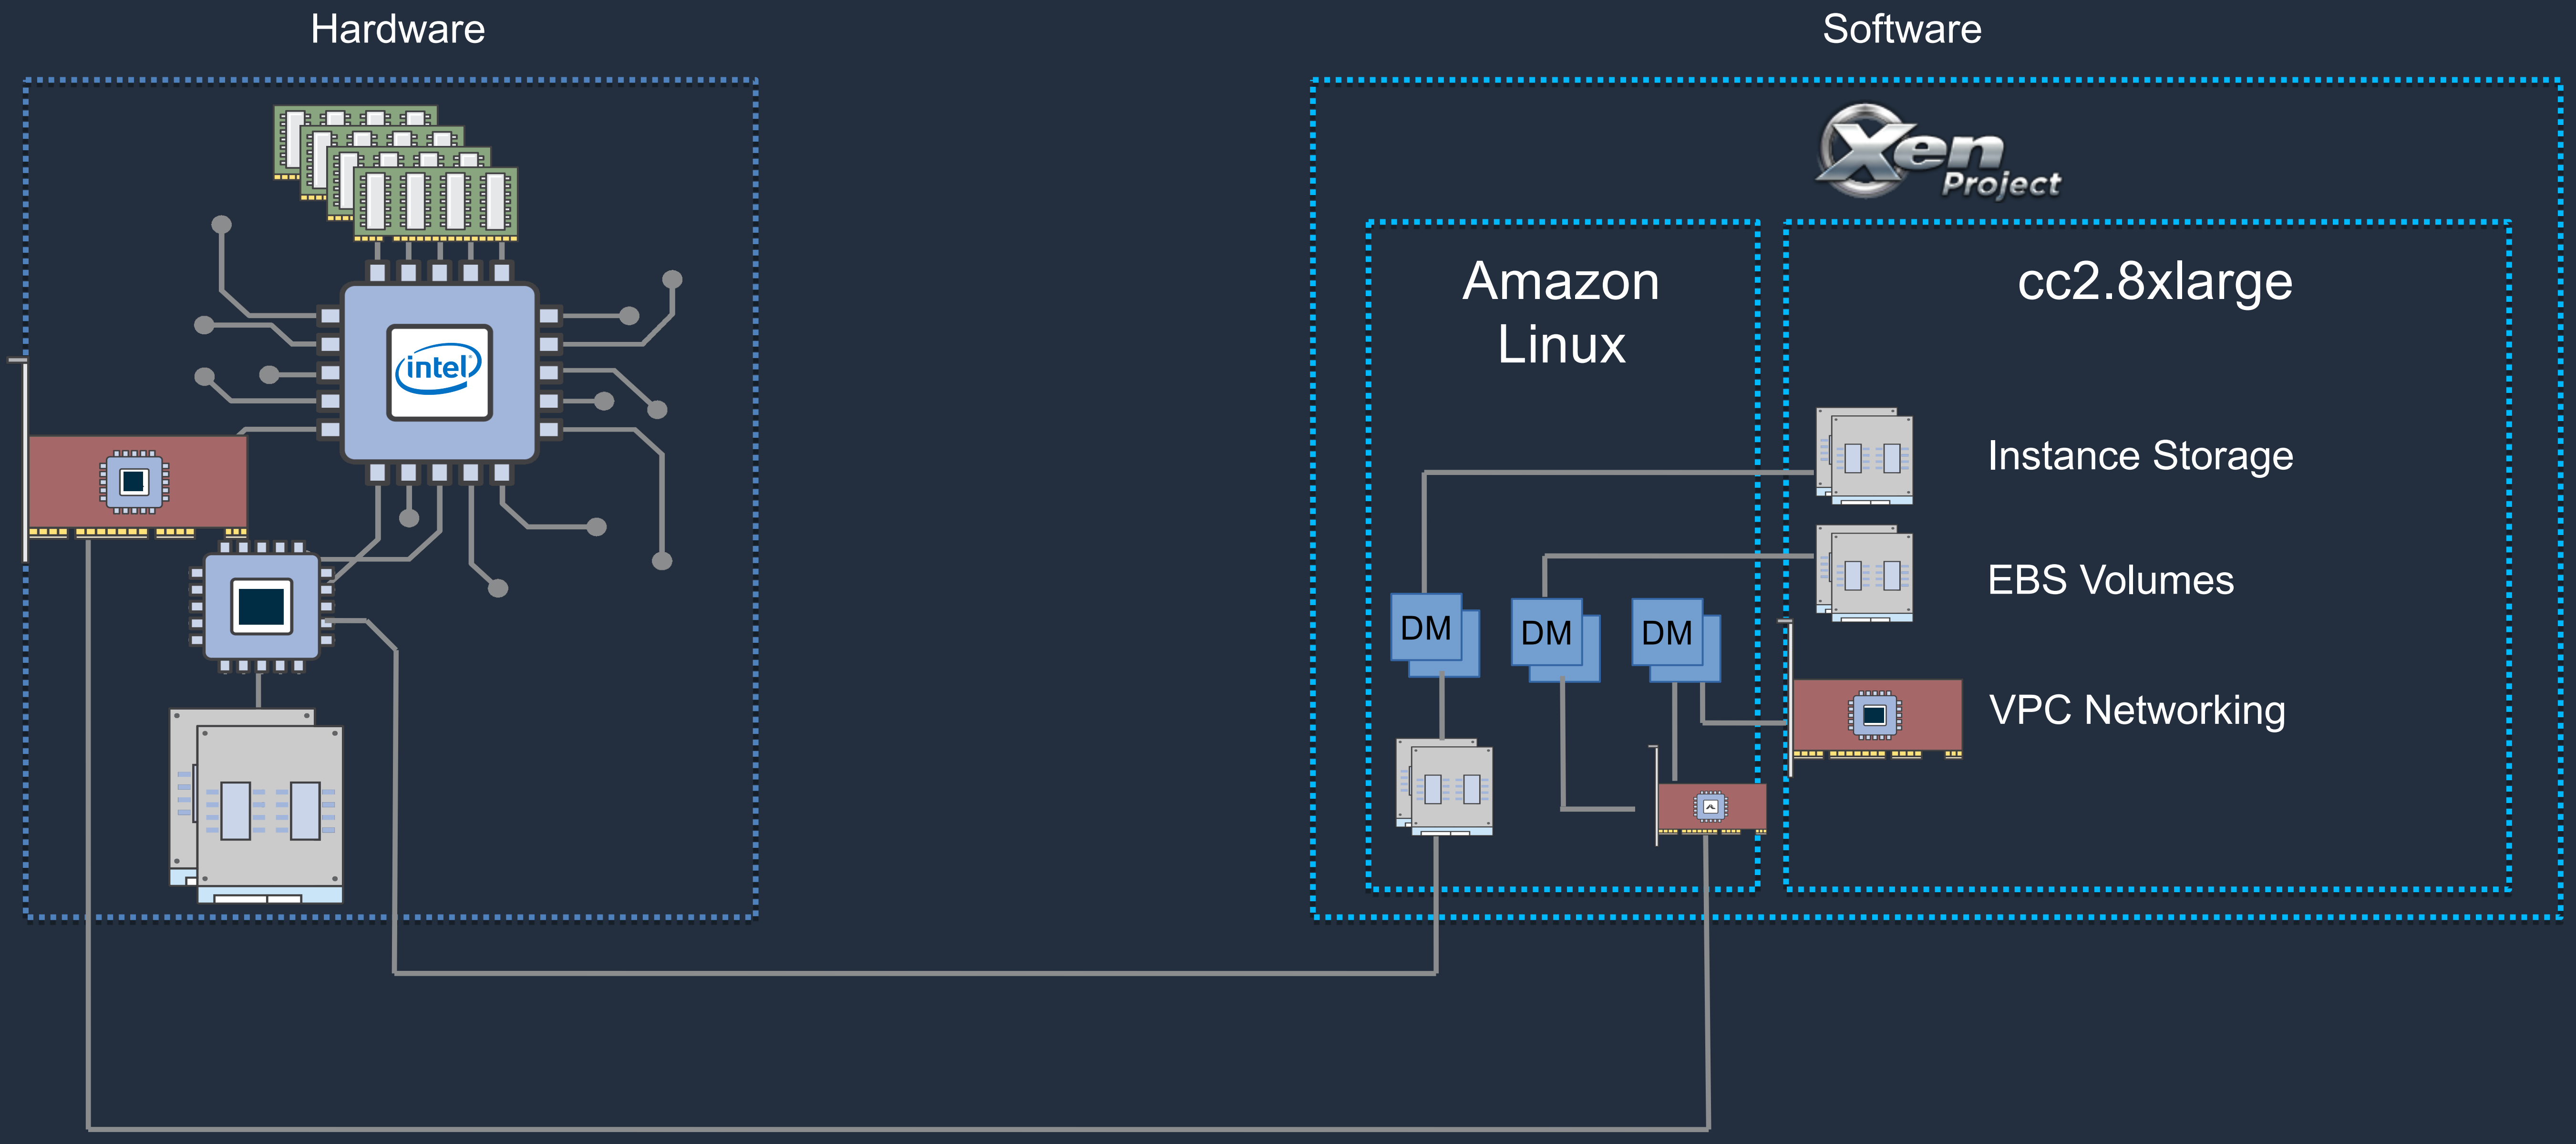

© 2020, Amazon Web Services, Inc. or its Affiliates.

What EC2 Instances Look Like Today

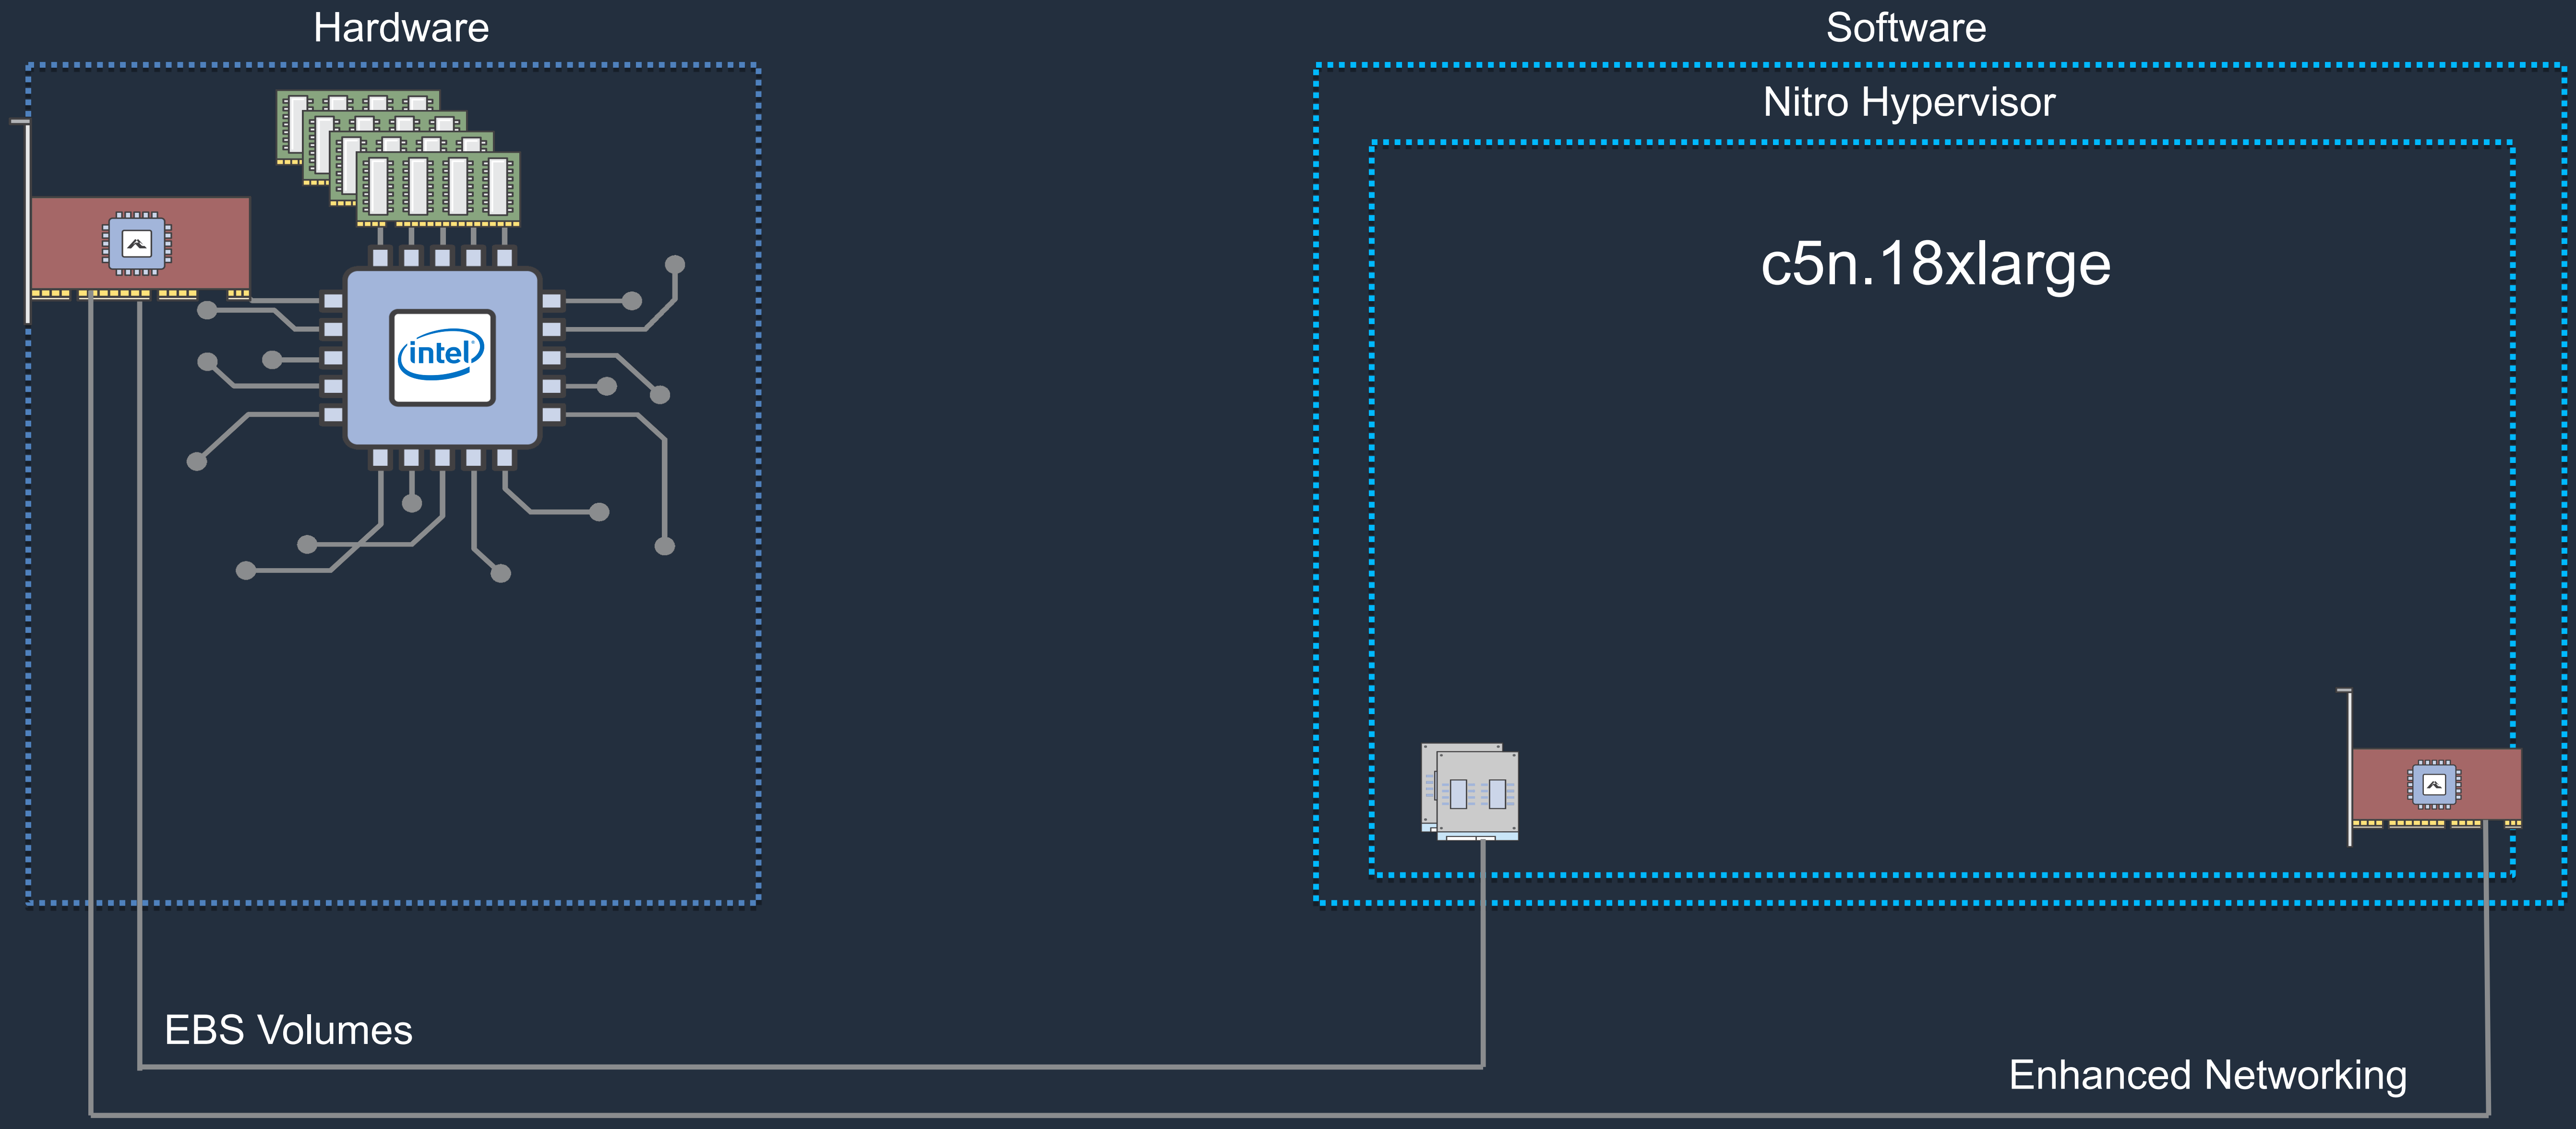

© 2020, Amazon Web Services, Inc. or its Affiliates.

What EC2 Looks Like Today… when you don't want hypervisors

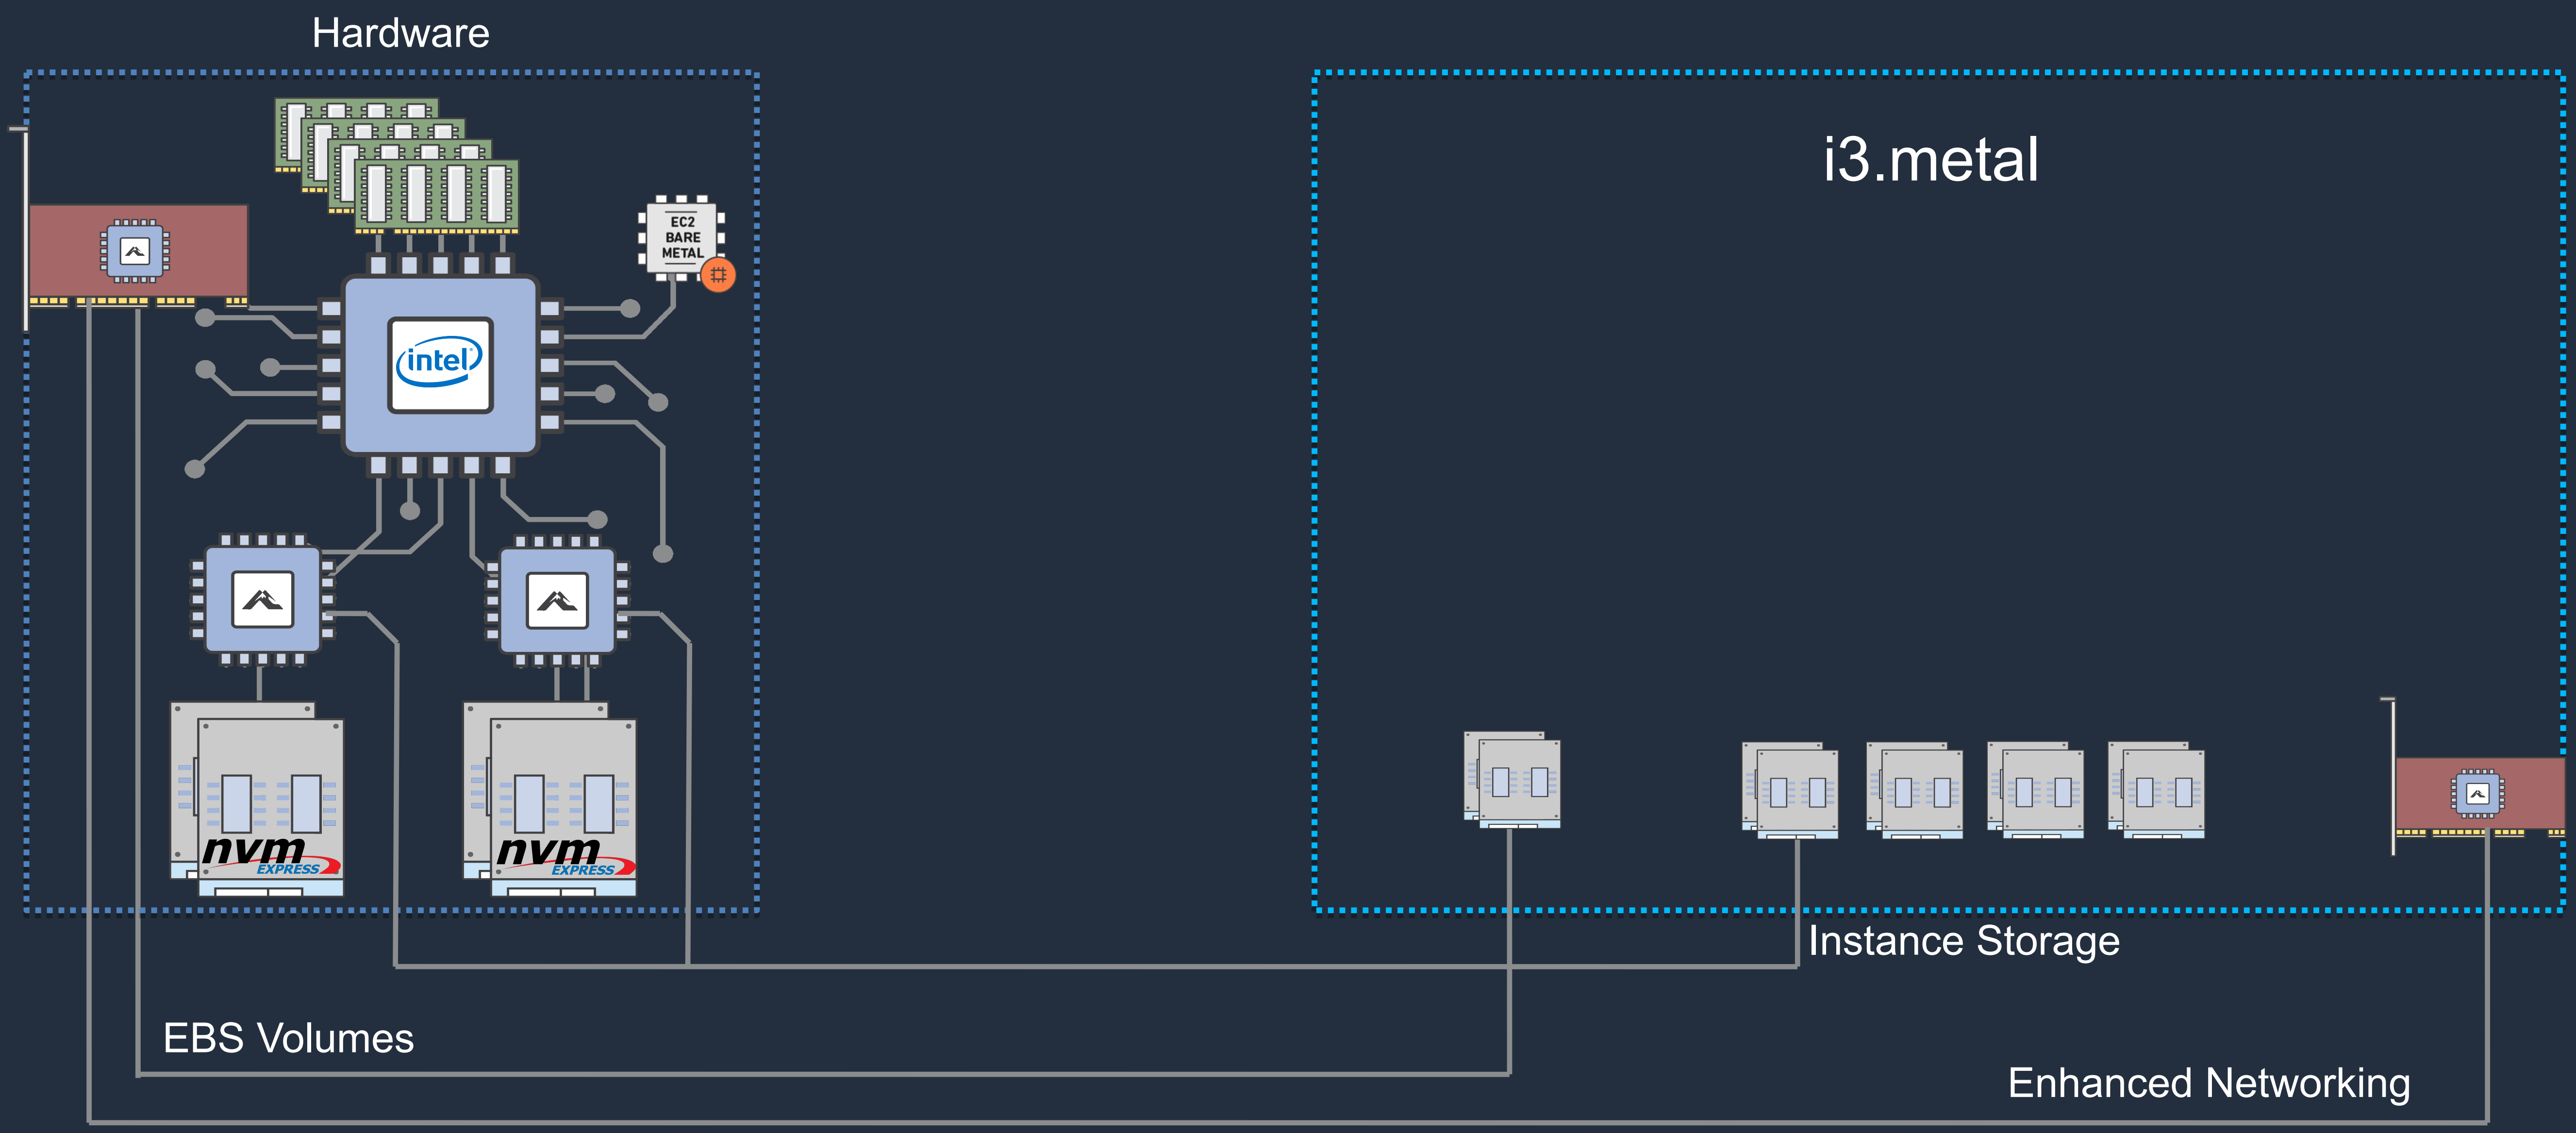

© 2020, Amazon Web Services, Inc. or its Affiliates.

What AWS Looks Like Today… with EFA

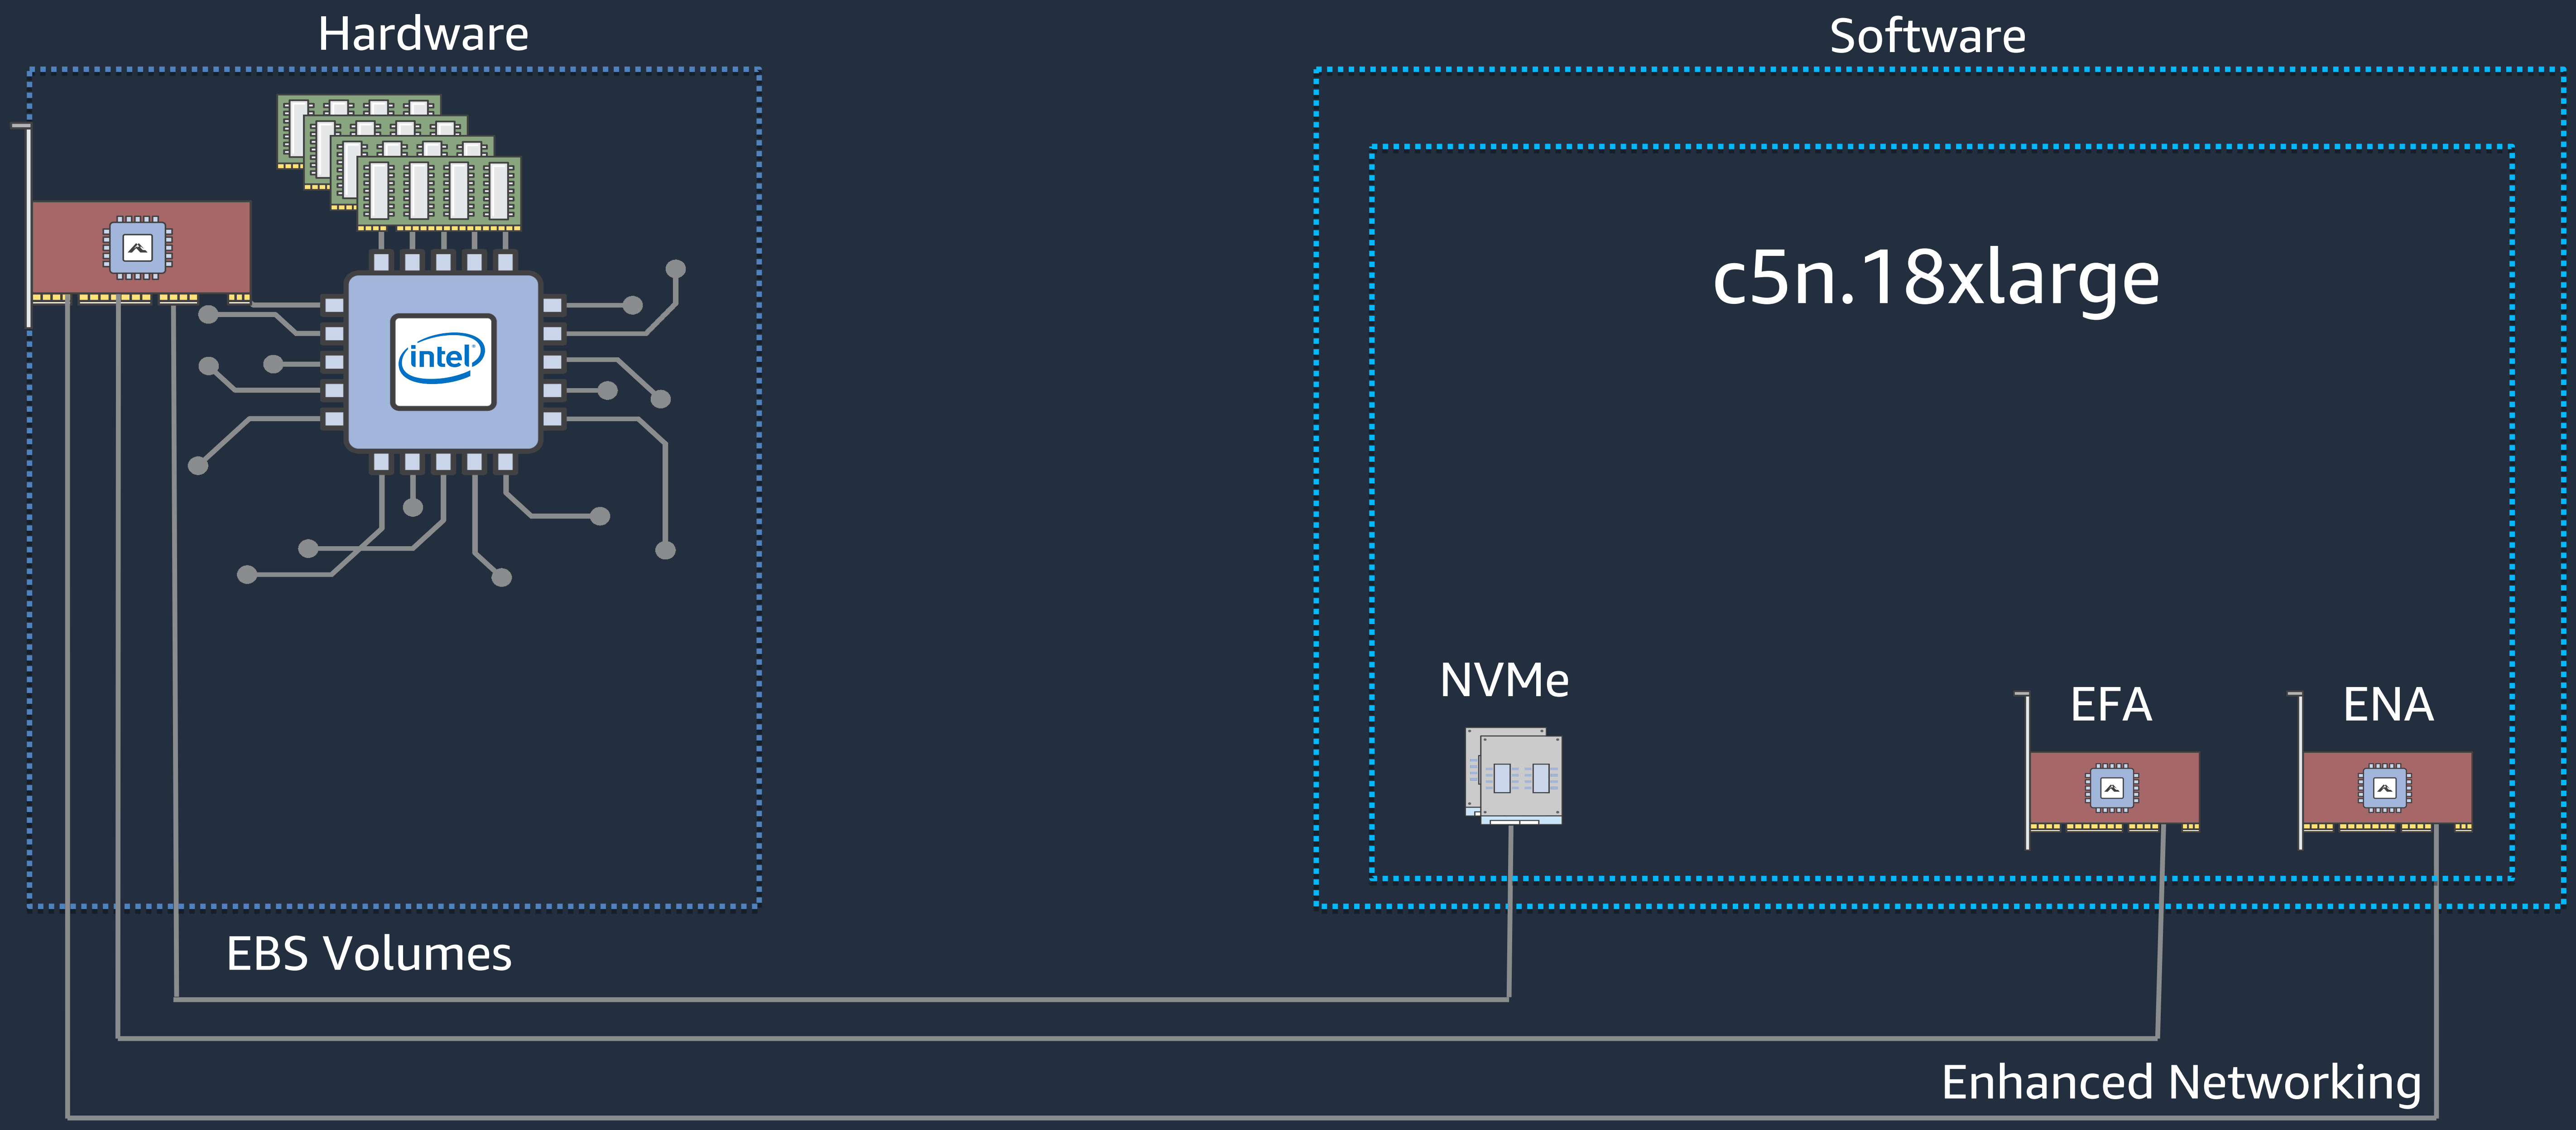

© 2020, Amazon Web Services, Inc. or its Affiliates.

EFA – For HPC and Machine Learning

Scale **tightly-coupled**
HPC applications on AWS

C5n | P3dn | I3en | M5(d)n | R5(d)n | G4dn Metal* | C5a(d)n Metal*

intel | NVIDIA | AMD

* **Coming soon**

EFA

Elastic Fabric Adapter, best for large HPC workloads

High data throughput

100 Gbps network bandwidth

Congestion control for cloud scale and rapid packet loss recovery.

Lower latency for message passing and more effective application-layer comms.

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws

Scalable Reliable Datagrams (SRD)

© 2020, Amazon Web Services, Inc. or its Affiliates.

Scalable Reliable Datagram Protocol

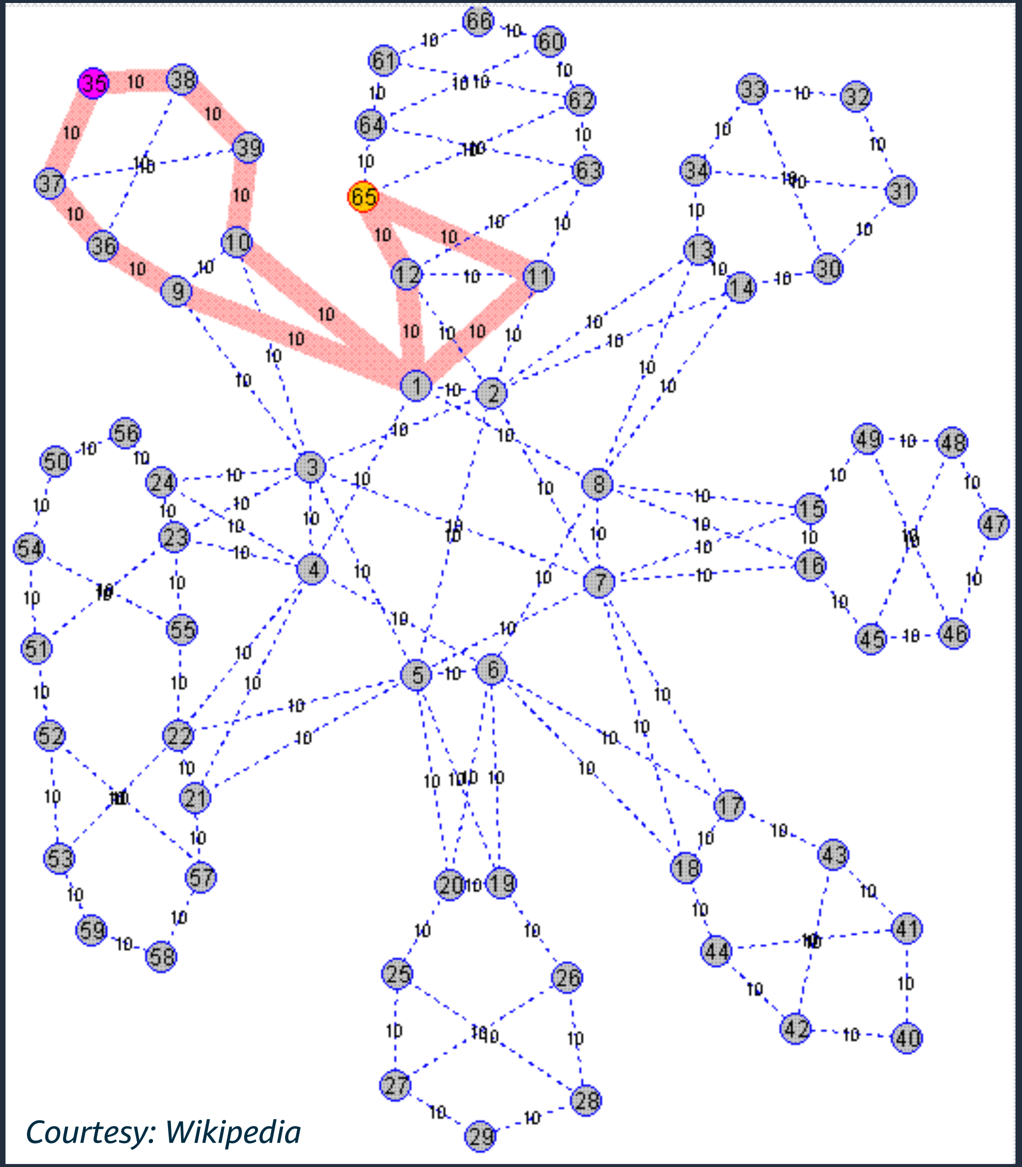

Courtesy: Wikipedia

- New protocol designed for AWS's unique datacenter network
- Implemented as part of our 3^{rd} generation Nitro chip
- EFA exposes SRD as a reliable datagram interface
- Inspired by IB Reliable Datagram, without the drawbacks
- Packet spraying over multiple ECMP paths
- Out-of-order delivery – no head-of-line blocking
- Congestion and flow control designed for large-scale cloud

L. Shalev, H. Ayoub, N. Bshara and E. Sabbag, "Supercomputing on Nitro in AWS Cloud," in IEEE Micro 2020

© 2020, Amazon Web Services, Inc. or its Affiliates.

SRD Link Failure Handling - Throughput

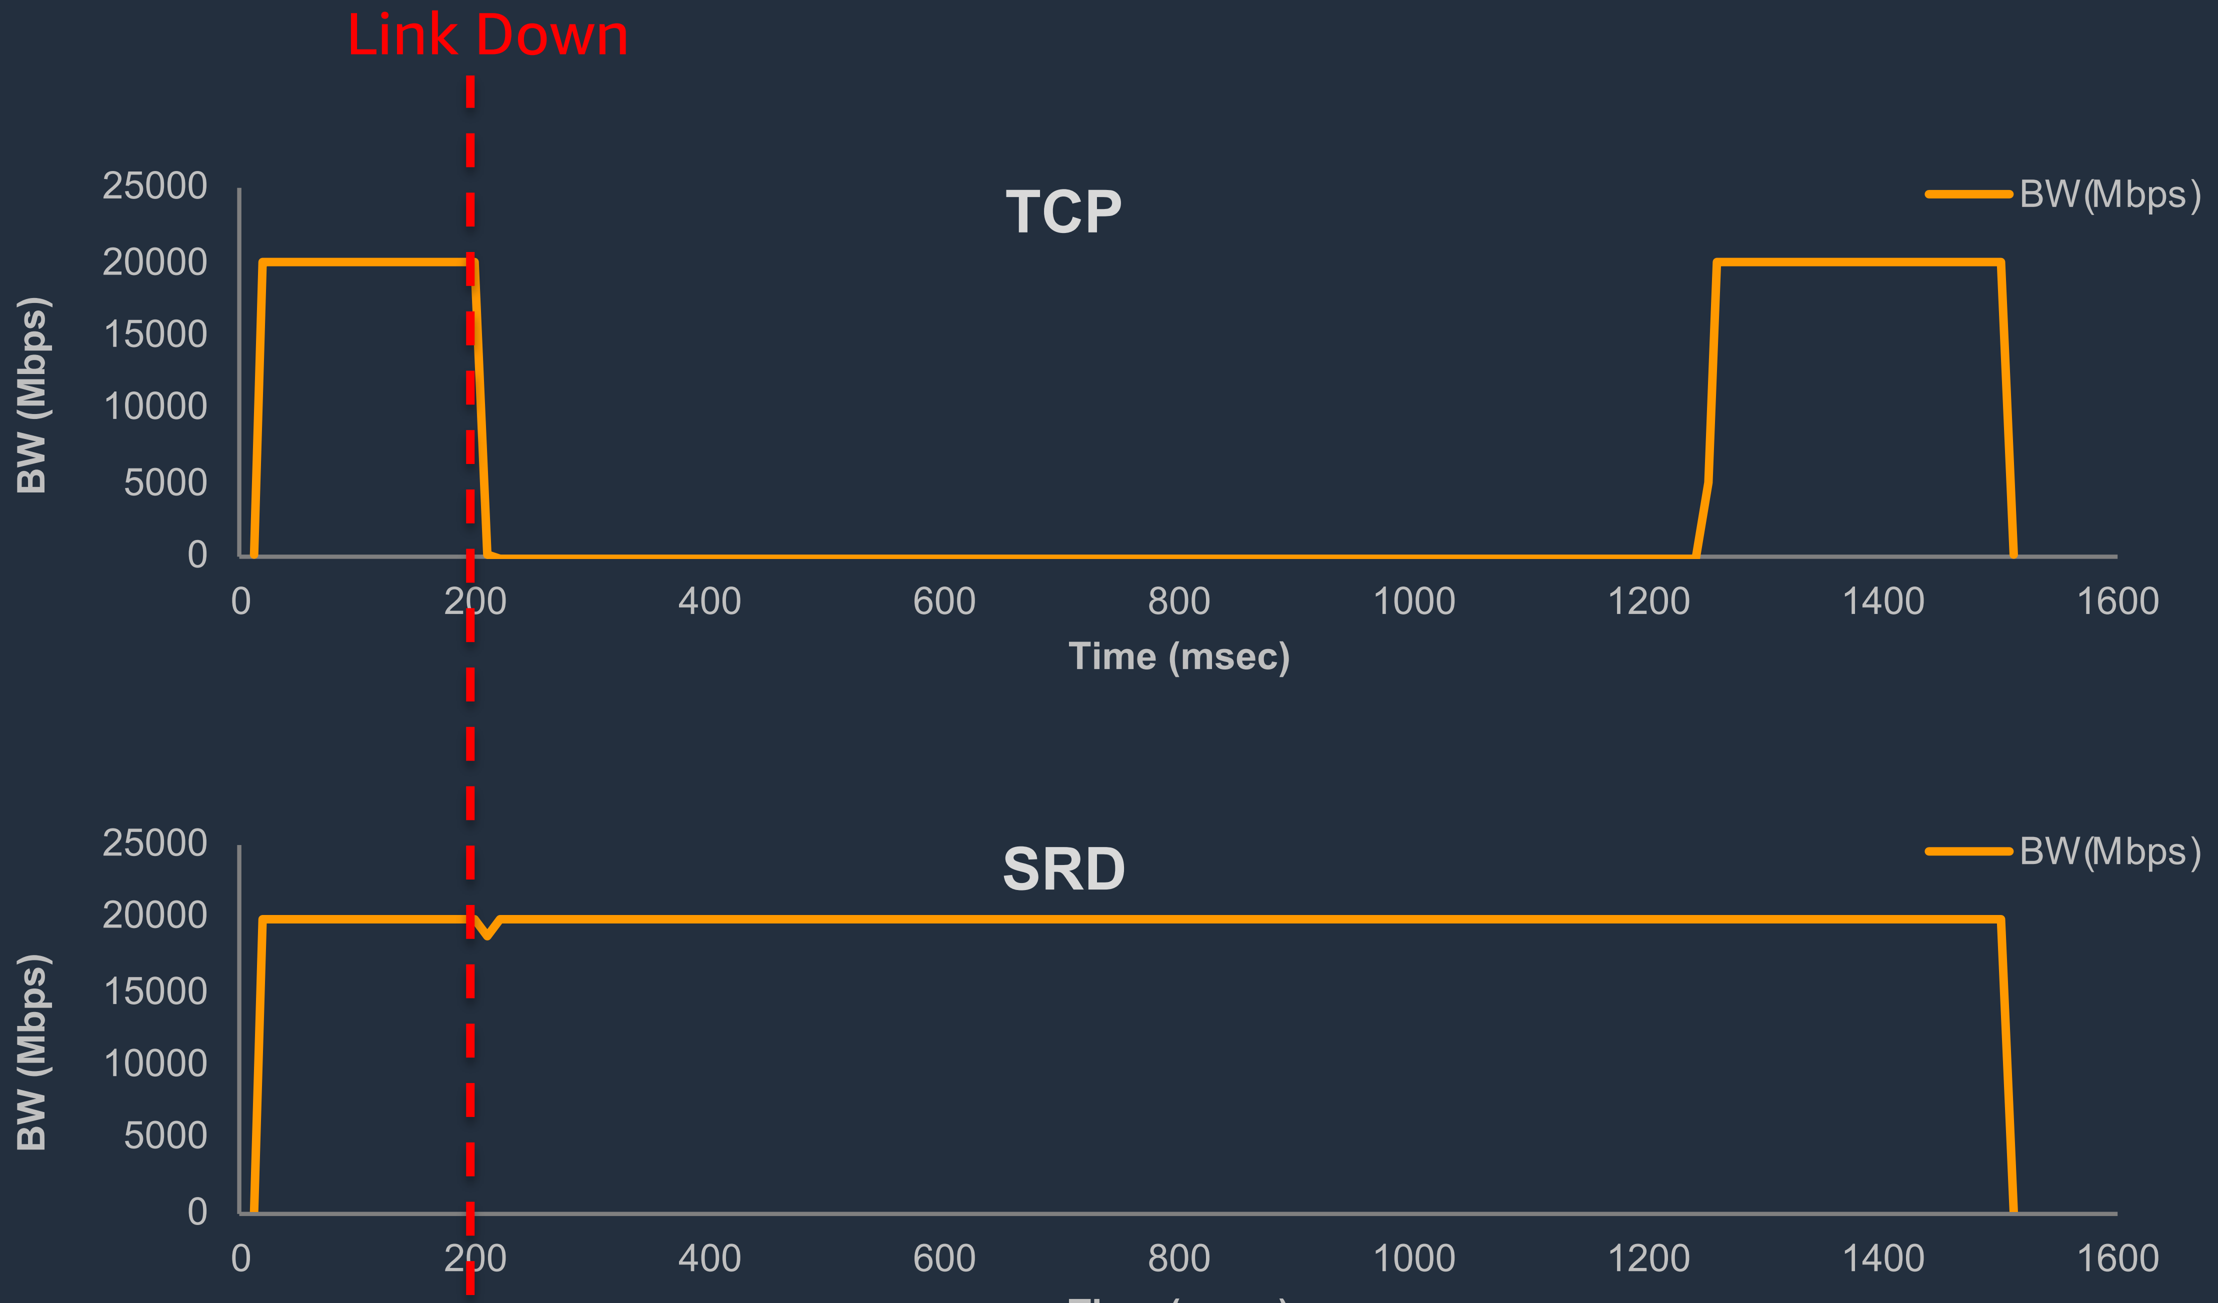

© 2020, Amazon Web Services, Inc. or its Affiliates.

SRD Link Failure Handling - Latency

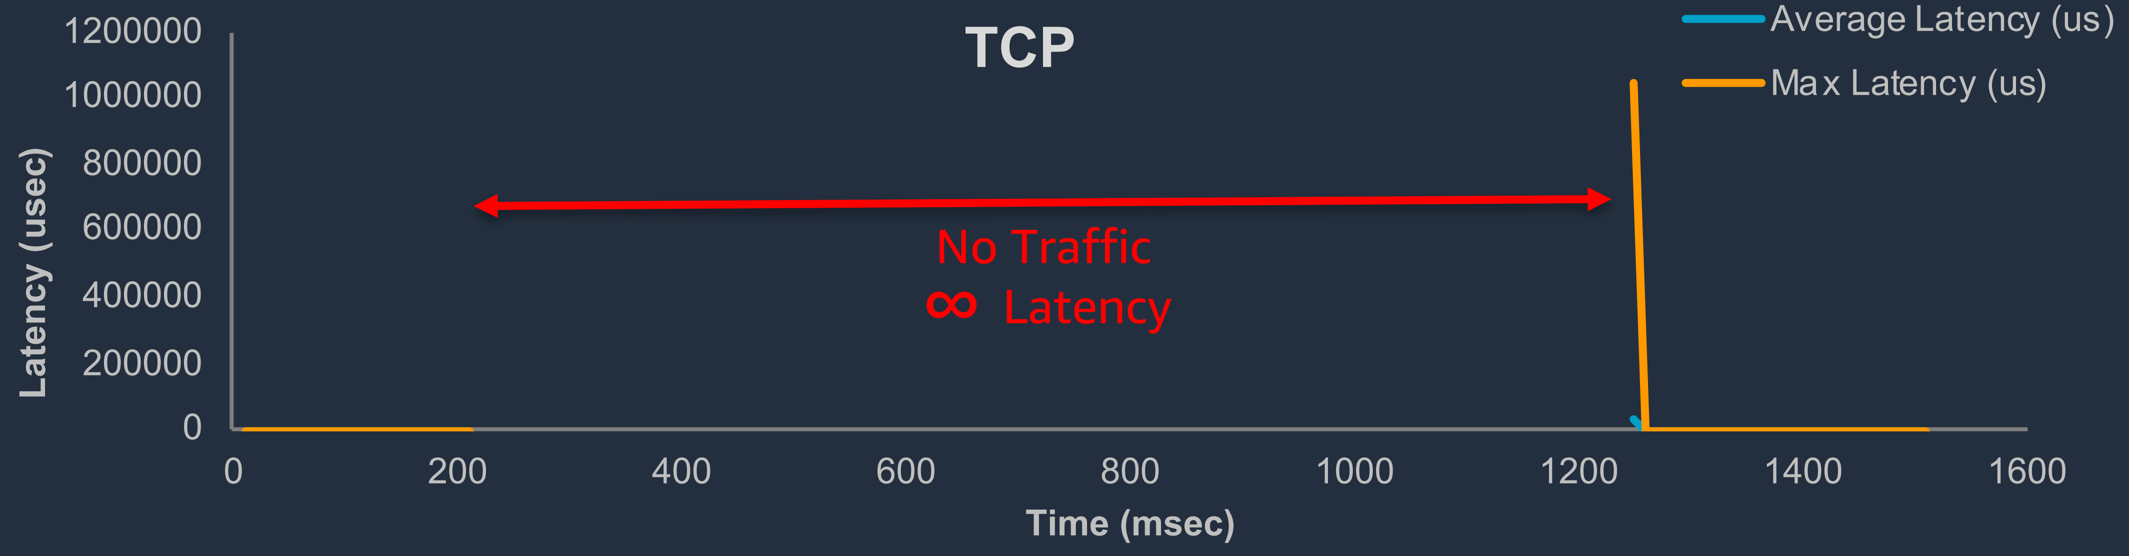

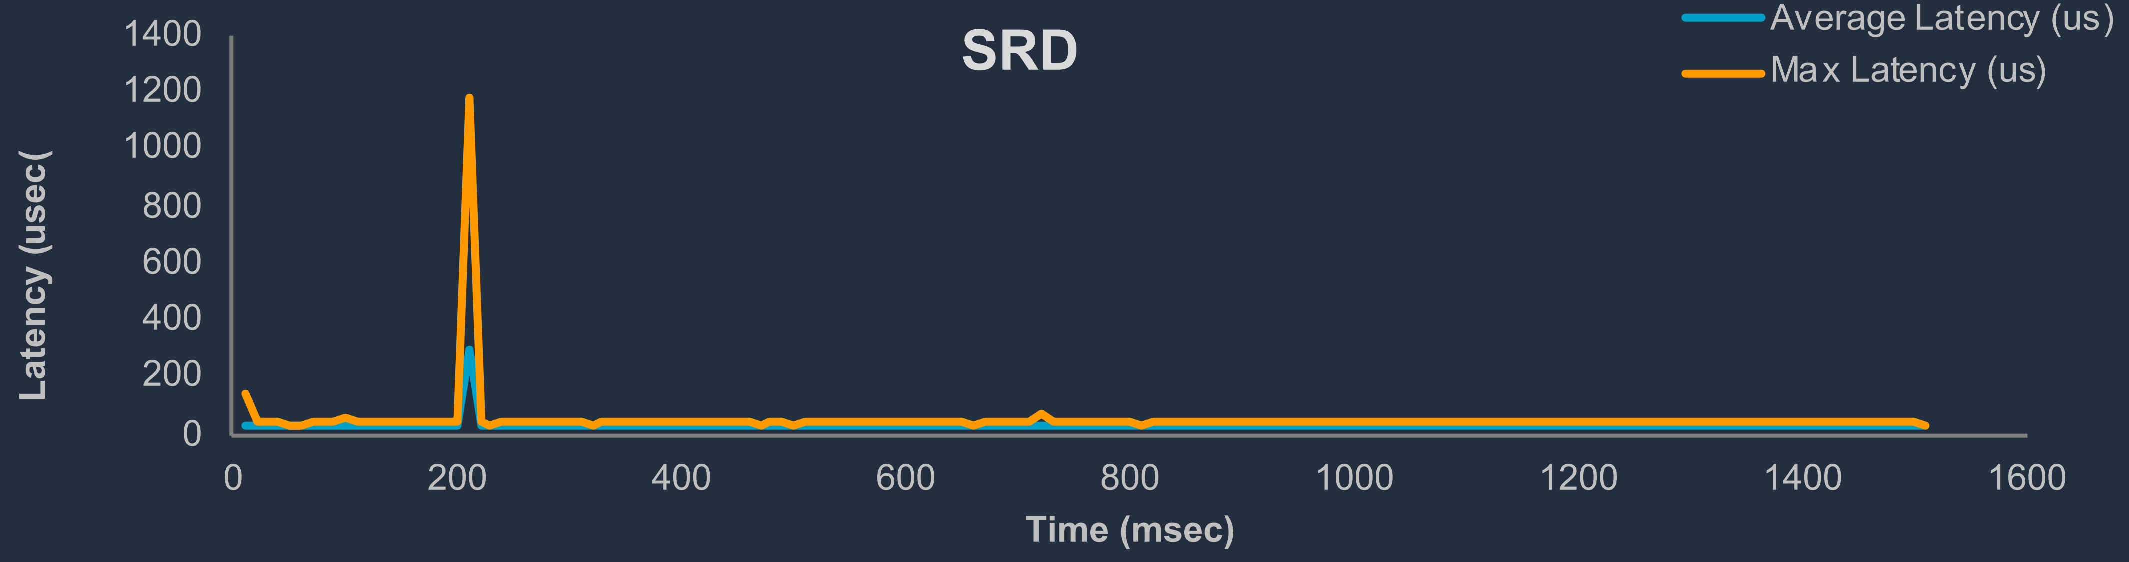

© 2020, Amazon Web Services, Inc. or its Affiliates.

SRD Congestion Control on Incast

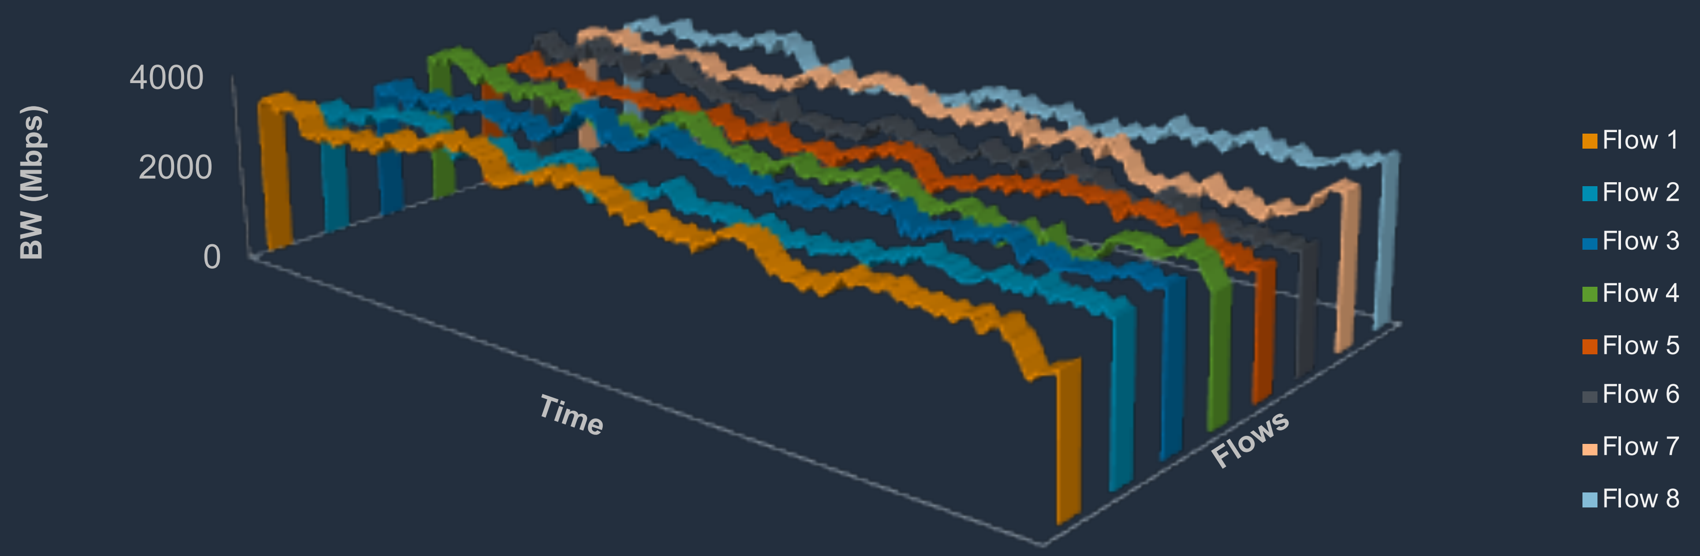
SRD

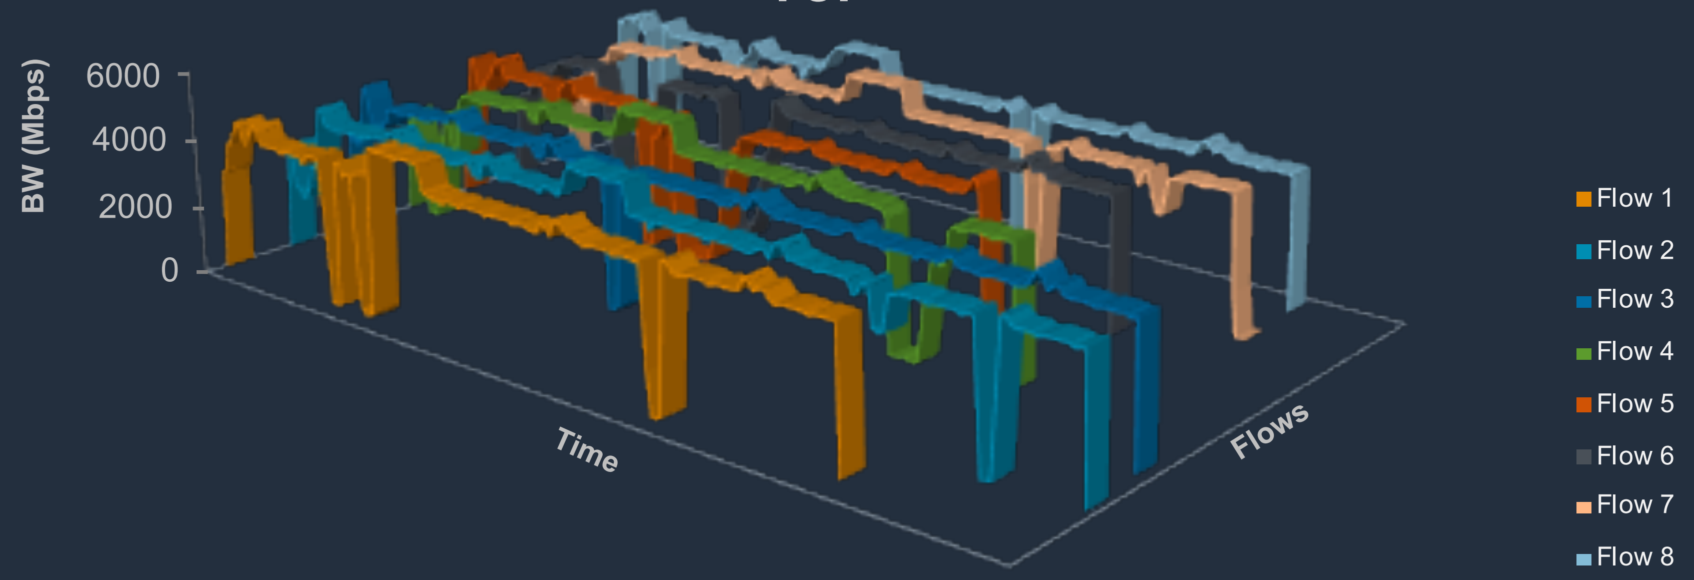
TCP

© 2020, Amazon Web Services, Inc. or its Affiliates.

EFA Software Ecosystem

© 2020, Amazon Web Services, Inc. or its Affiliates.

EFA Software Ecosystem

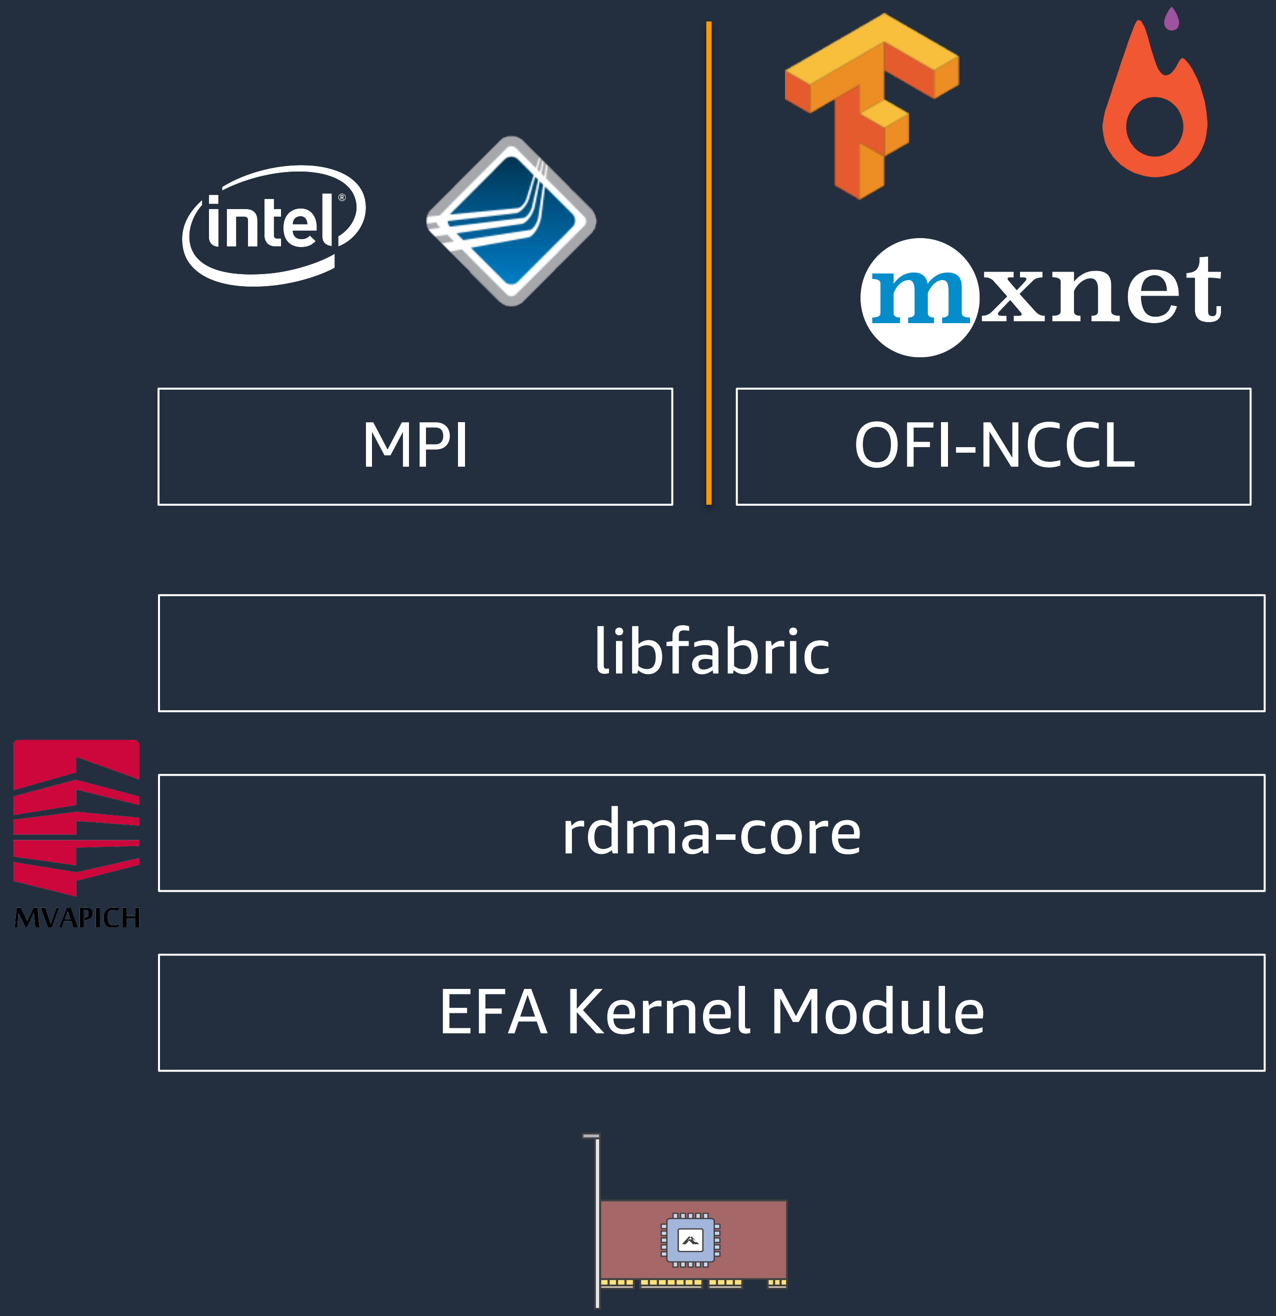

MPI

OFI-NCCL

libfabric

rdma-core

EFA Kernel Module

© 2020, Amazon Web Services, Inc. or its Affiliates.

MVAPICH2-X-AWS

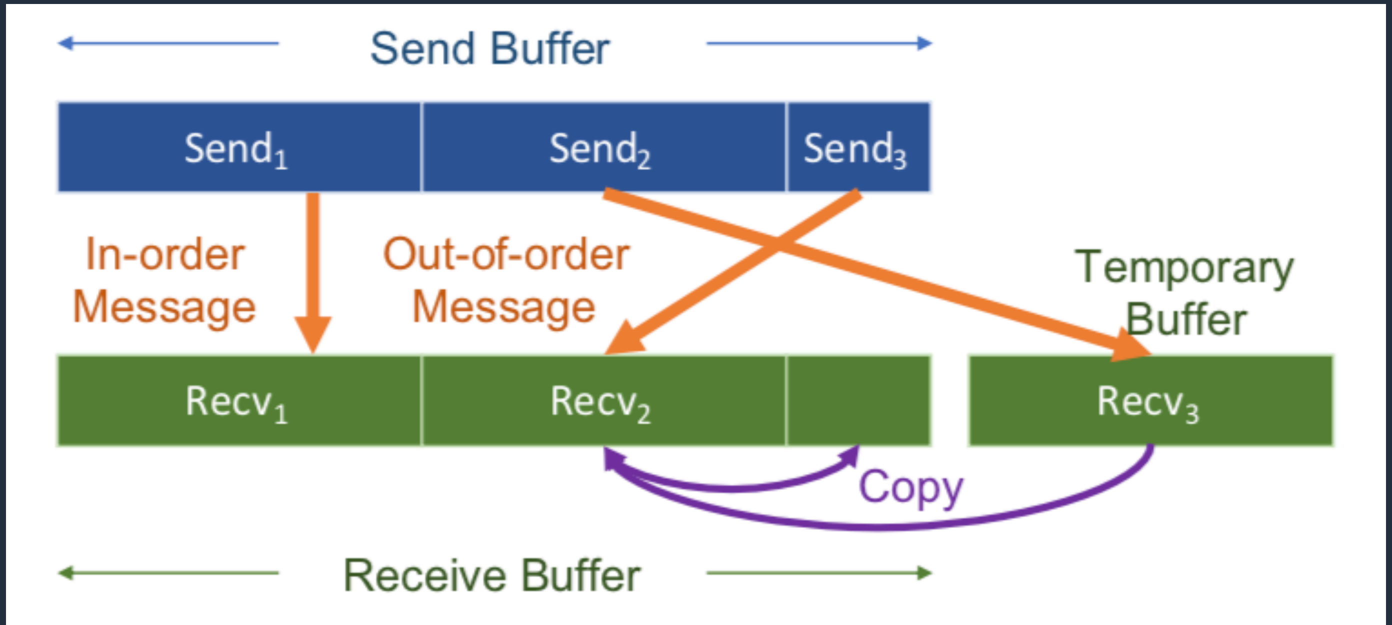

- Directly programs to rdma-core
- Reordering with copy-out
- Use of immediate data for seq ID
- Long message packetization
- Tag matching
- Intra-node path with XPMEM

S. Chakraborty, S. Xu, H. Subramoni and D. K. Panda, Designing Scalable and High-Performance MPI Libraries on Amazon Elastic Adapter, Hot Interconnect, 2019

© 2020, Amazon Web Services, Inc. or its Affiliates.

Open MPI and EFA

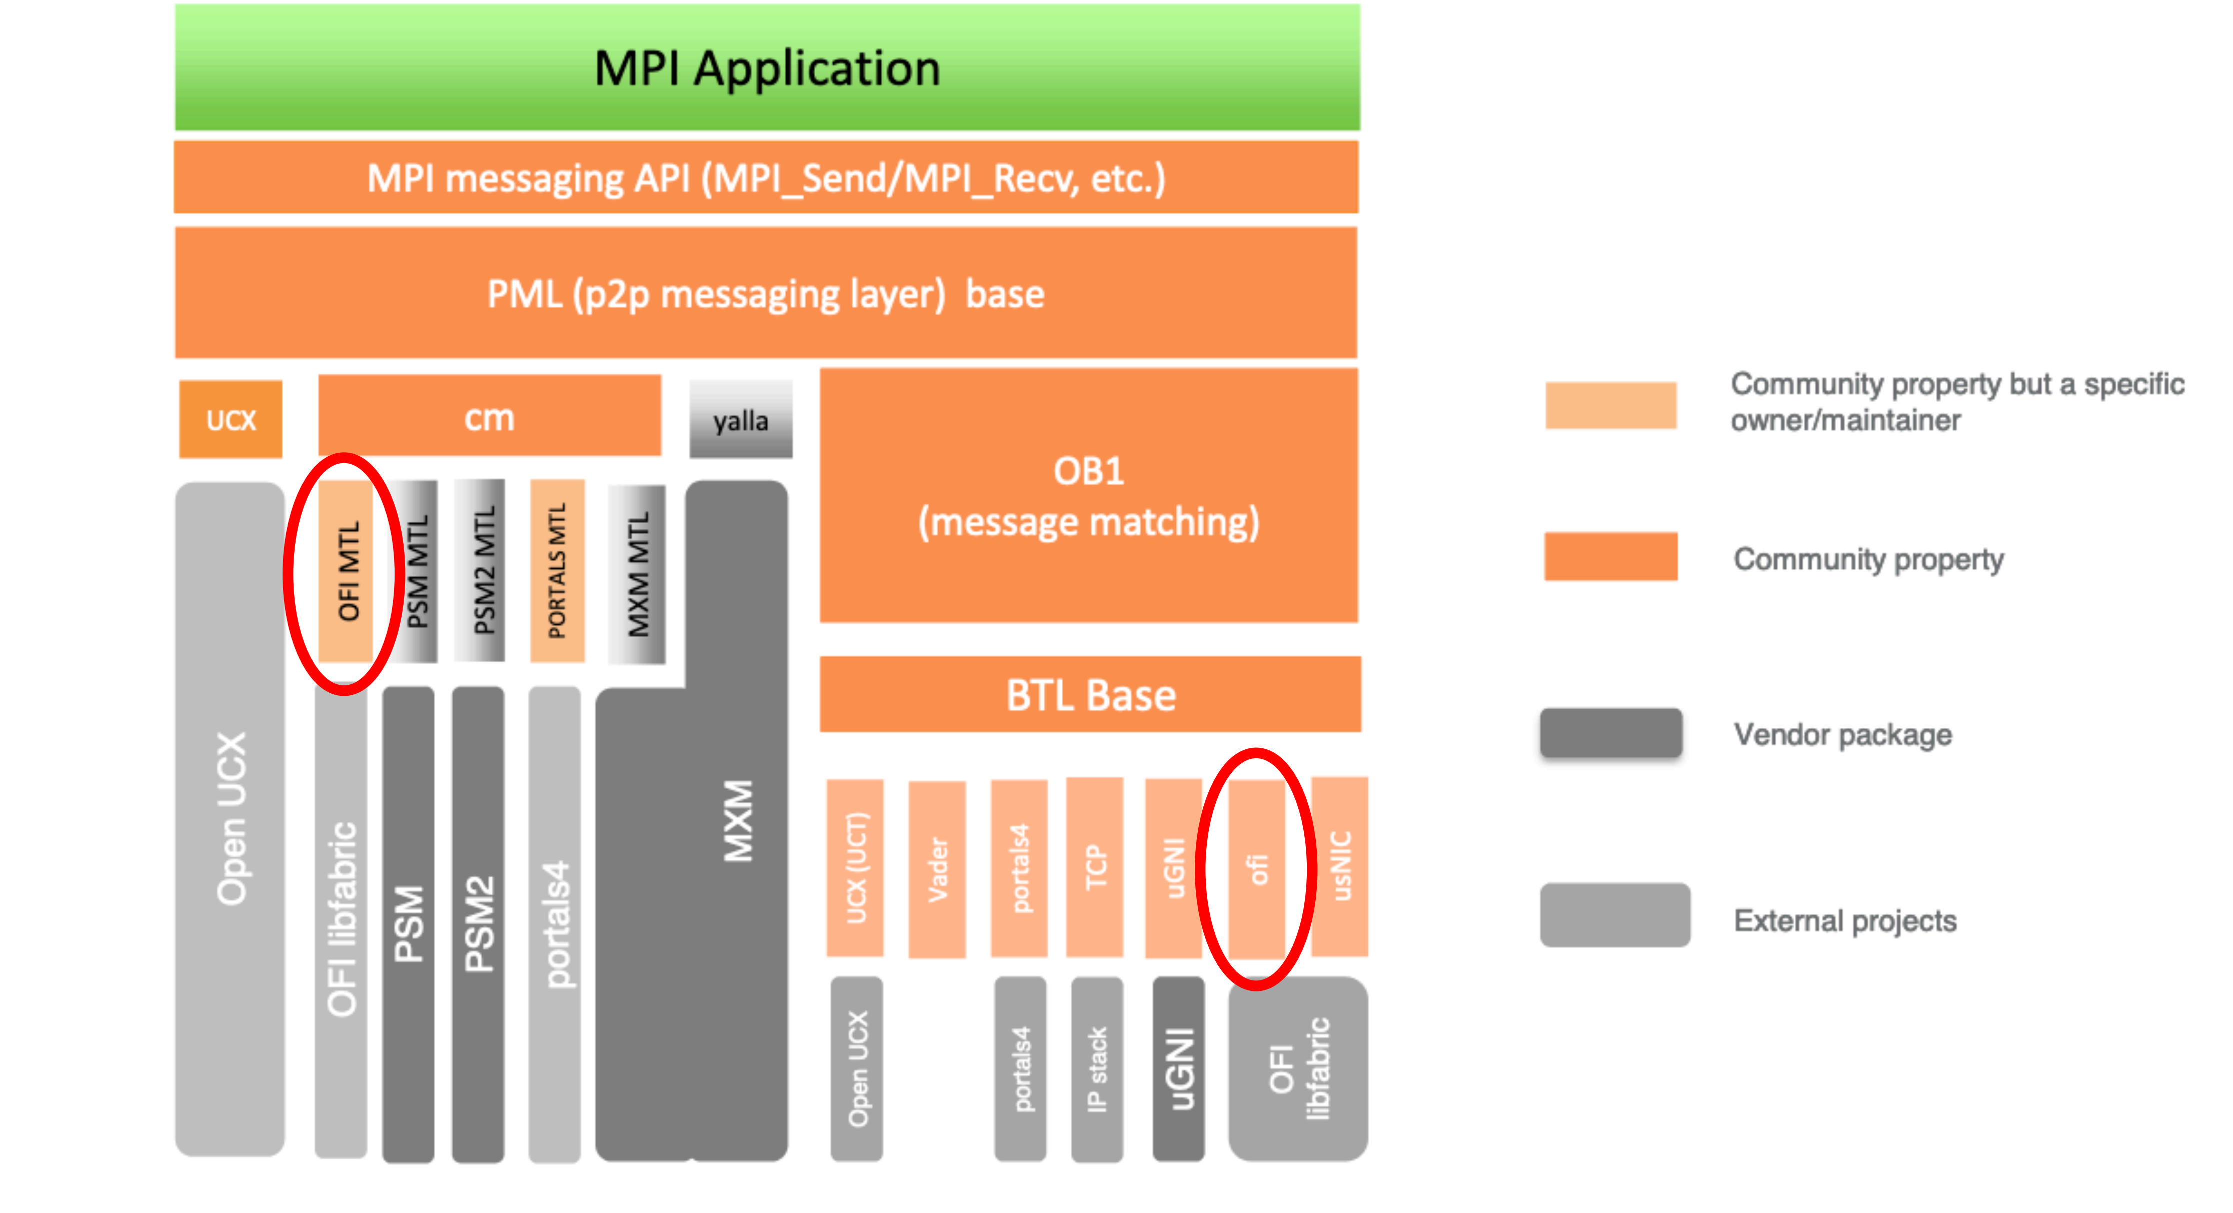

© 2020, Amazon Web Services, Inc. or its Affiliates.

Intel MPI and EFA

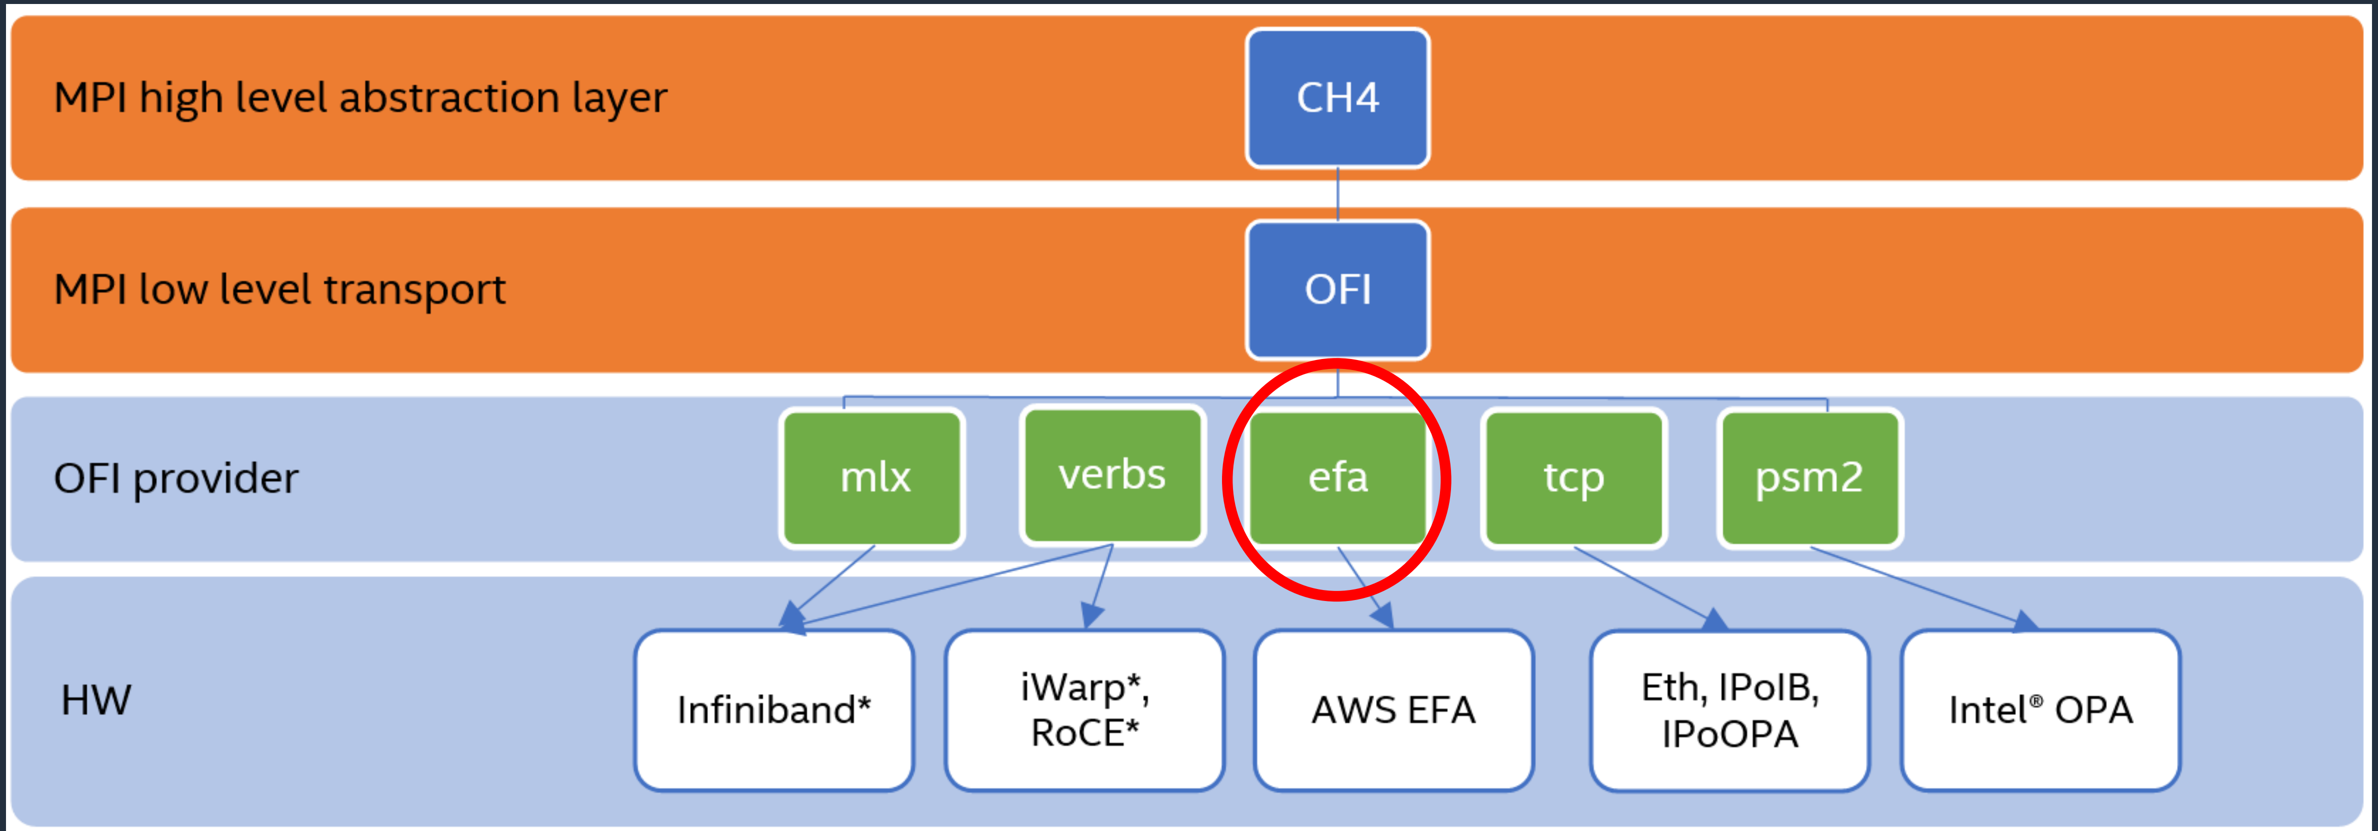

MPI high level abstraction layer

CH4

MPI low level transport

OFI

OFI provider

mlx

verbs

efa

tcp

psm2

HW

Infiniband*

iWarp*,
RoCE*

AWS EFA

Eth, IPoIB,
IPoOPA

Intel® OPA

https://software.intel.com/content/www/us/en/develop/articles/intel-mpi-library-2019-over-libfabric.html20

© 2020, Amazon Web Services, Inc. or its Affiliates.

EFA Installer

```
$ curl -O https://s3-us-west-2.amazonaws.com/aws-efa-installer/aws-efa-installer-latest.tar.gz
$ tar -xf aws-efa-installer-1.4.1.tar.gz
$ cd aws-efa-installer
$ sudo ./efa_installer.sh -y
= Starting Amazon Elastic Fabric Adapter Installation Script =
= EFA Installer Version: 1.4.1 =

== Installing EFA dependencies ==
<snip>
== Writing EFA profile.d configuration ==
== Configuring system limits for EFA ==
Limits for Elastic Fabric Adapter configured.
== Testing EFA device ==
Starting server...
Starting client...
<snip>
===================================================
EFA installation complete.
- Please logout/login to complete the installation.
- Libfabric was installed in /opt/amazon/efa
- Open MPI was installed in /opt/amazon/openmpi
===================================================
```

© 2020, Amazon Web Services, Inc. or its Affiliates.

Performance Results

© 2020, Amazon Web Services, Inc. or its Affiliates.

MPI Benchmarks

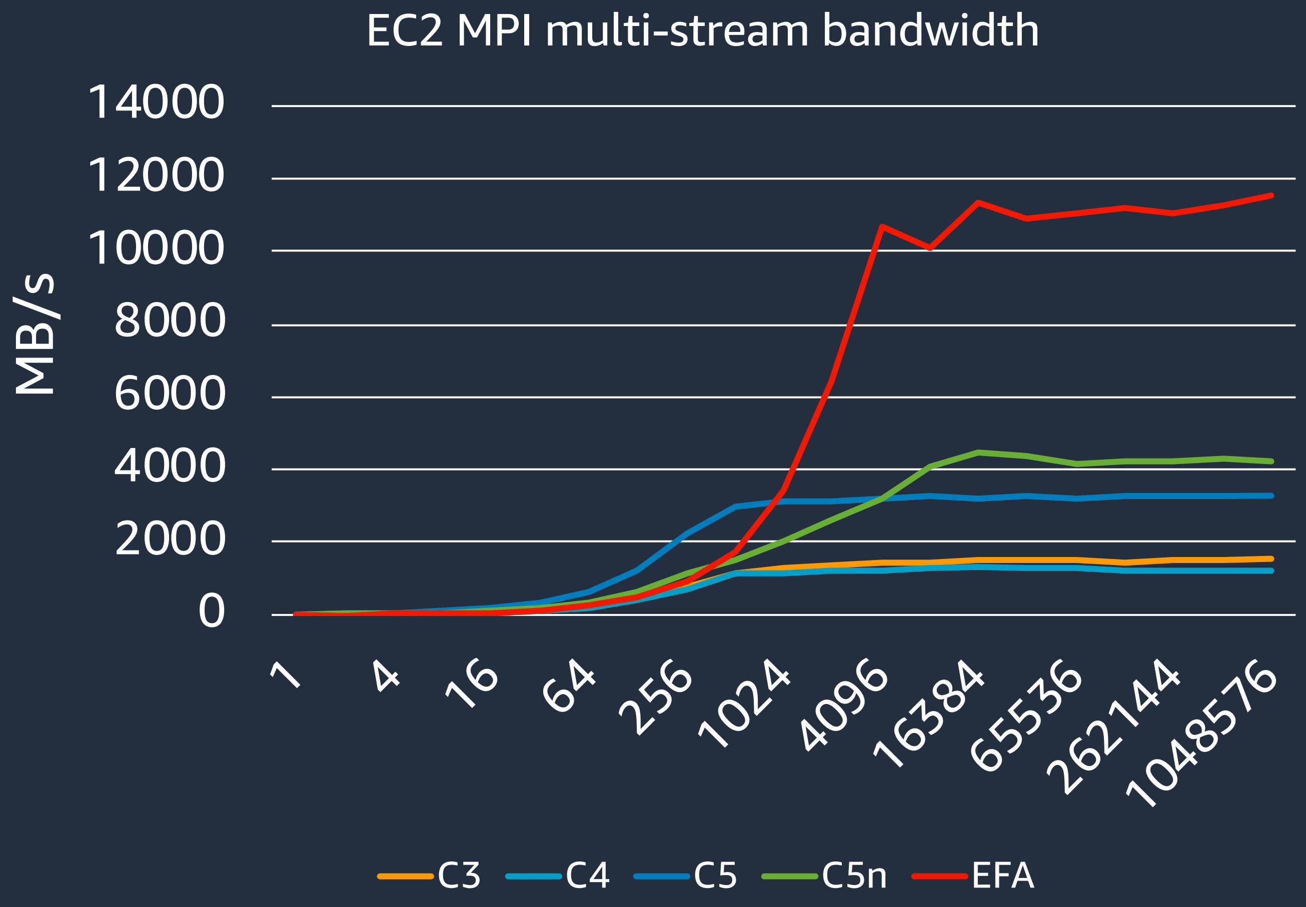

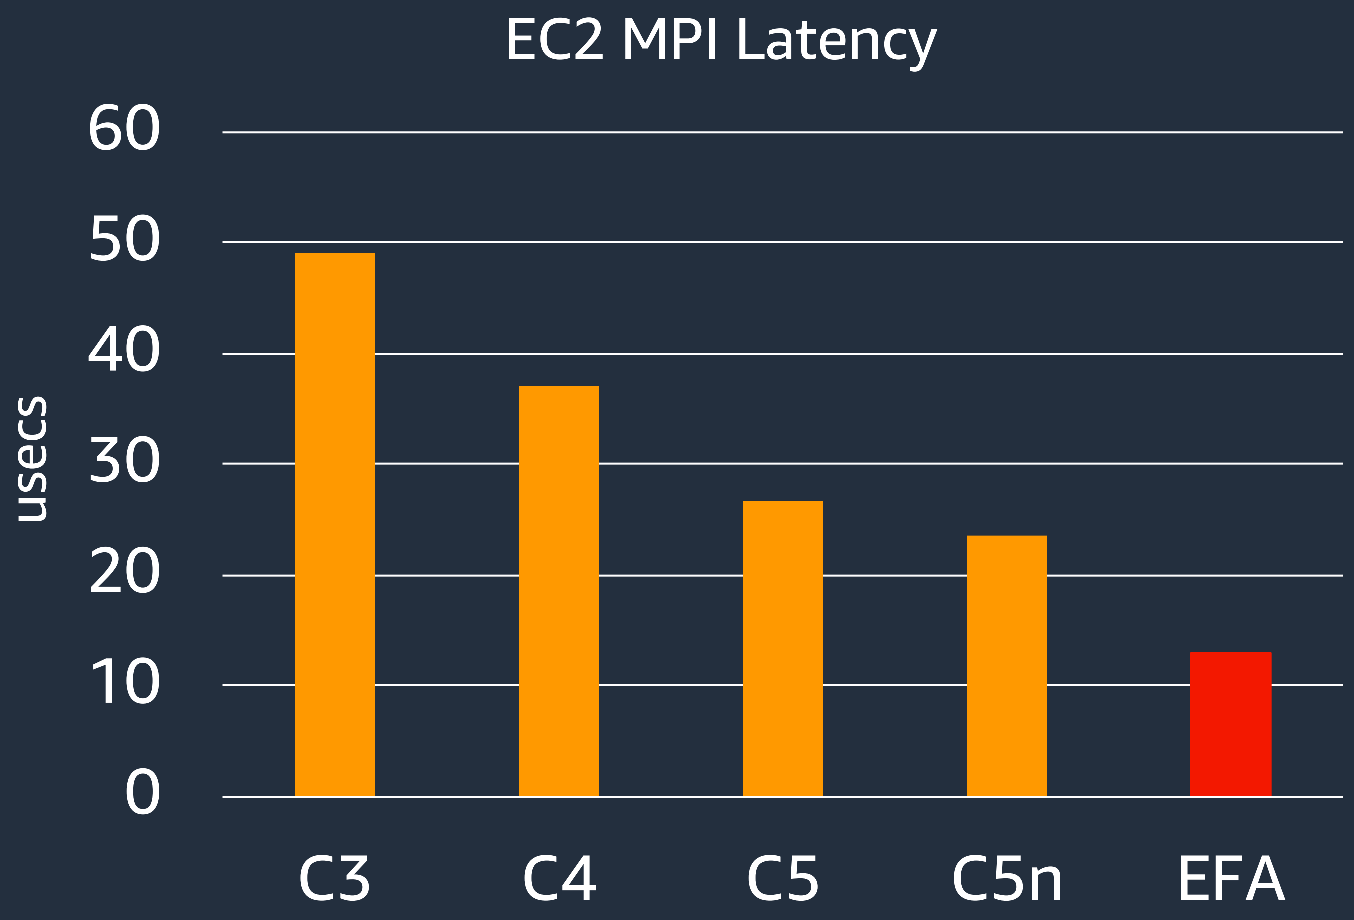

© 2020, Amazon Web Services, Inc. or its Affiliates.

Scaling with Applications: Fluent

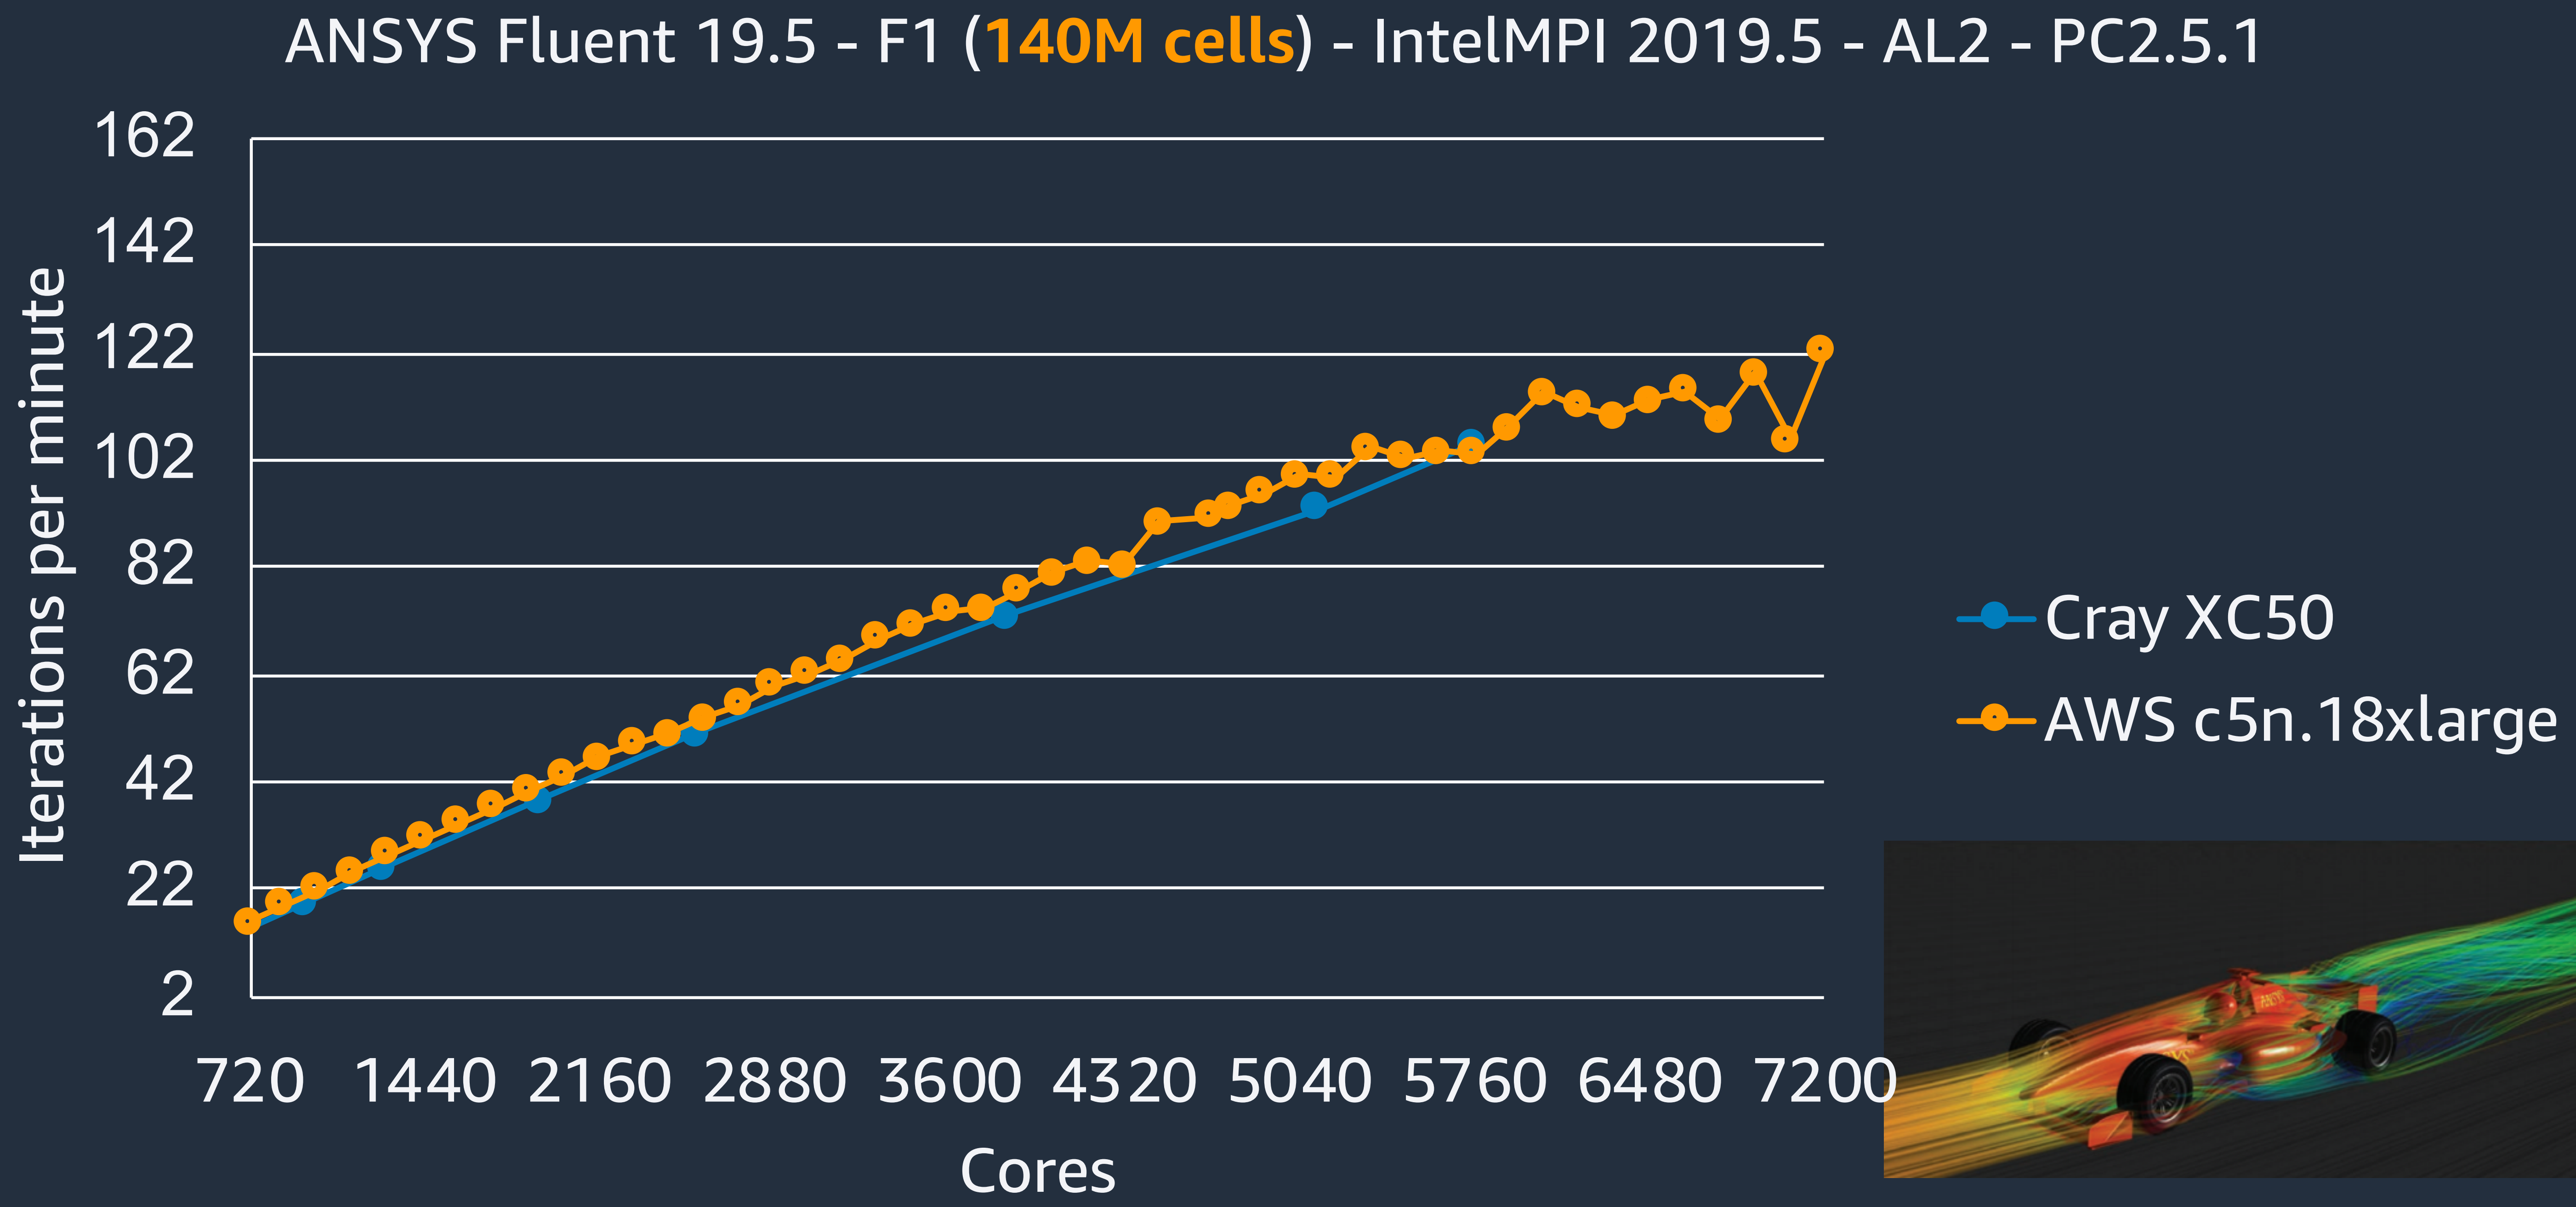

© 2020, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Scaling with Applications: zCFD

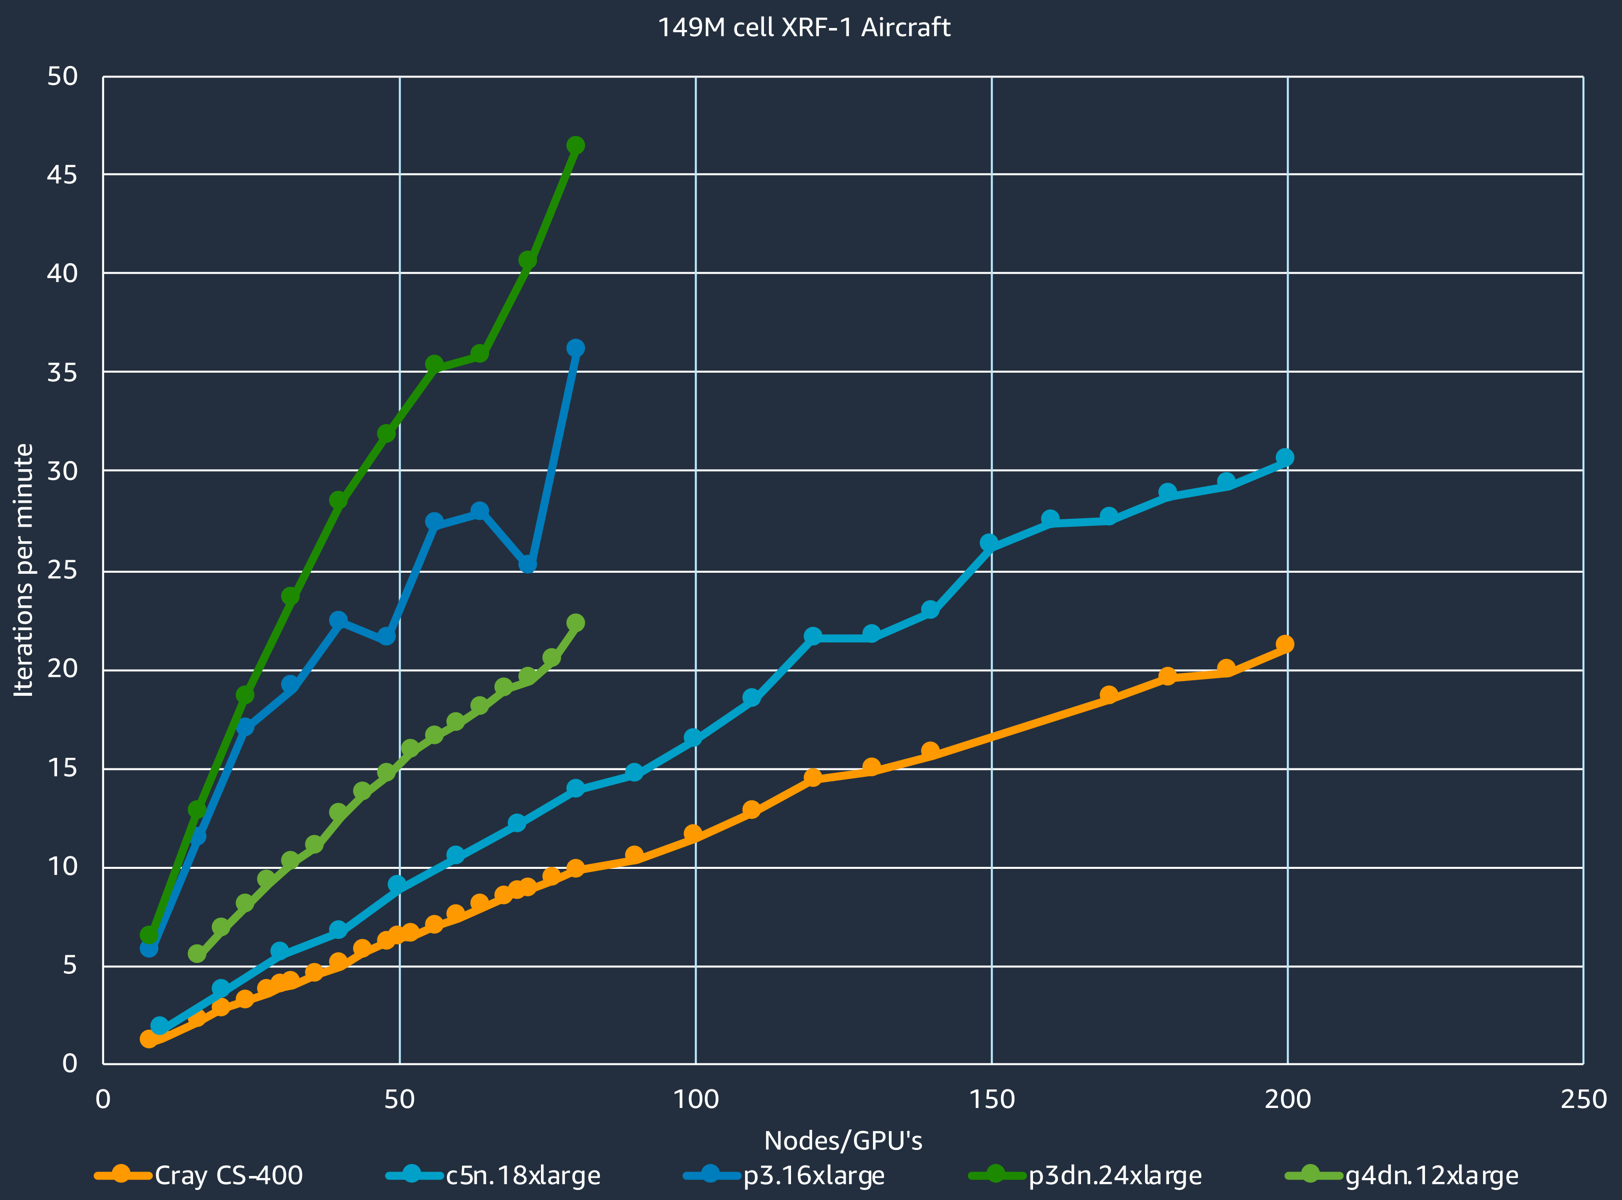

- Recent work with Zenotech (zCFD)
- GPUs (p3.24xlarge and g4dn.16xlarge Amazon EC2 instances i.e Nvidia v100 and T4's) deliver results faster for a lower cost

ZENOTECH SIMULATION UNLIMITED

© 2020, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Customer case study: Formula 1 on AWS

INTERNAL ONLY

- No access to wind-tunnels for 12-24 months, only CFD
- **60hrs** to **10hrs** thanks to AWS.
- 192 cores to 1440 cores (EC2 c5n.18xlarge+ AWS ParallelCluster)
- On-demand + now smaller spot jobs (exploring other instance types)

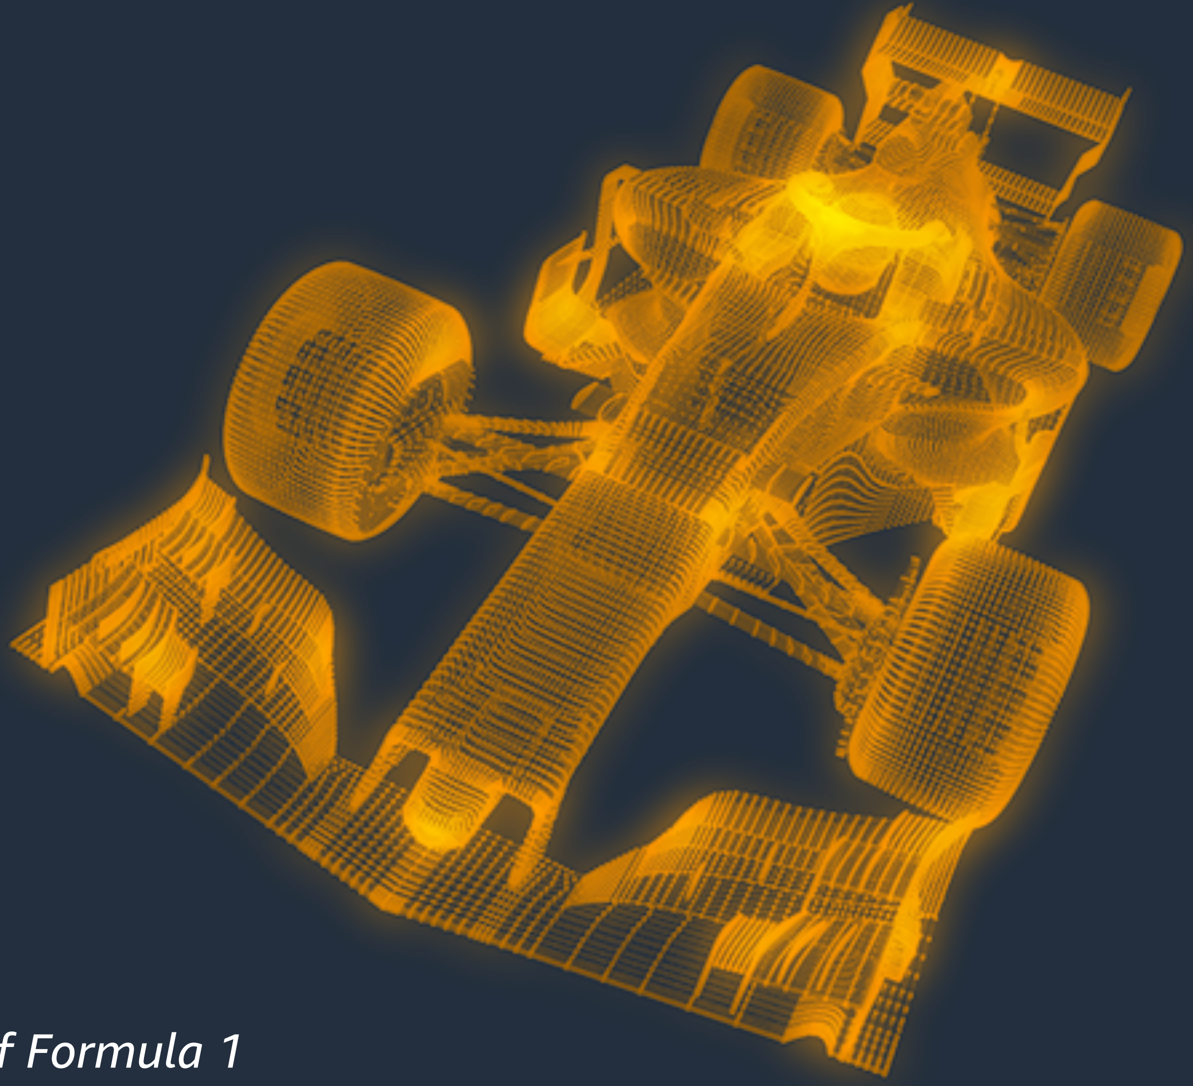

"This project with AWS was one of the most revolutionary in the history of Formula 1 aerodynamics," said Pat Symonds, Chief Technical Officer of Formula 1

© 2020, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Scaling with Applications: NAVGEM

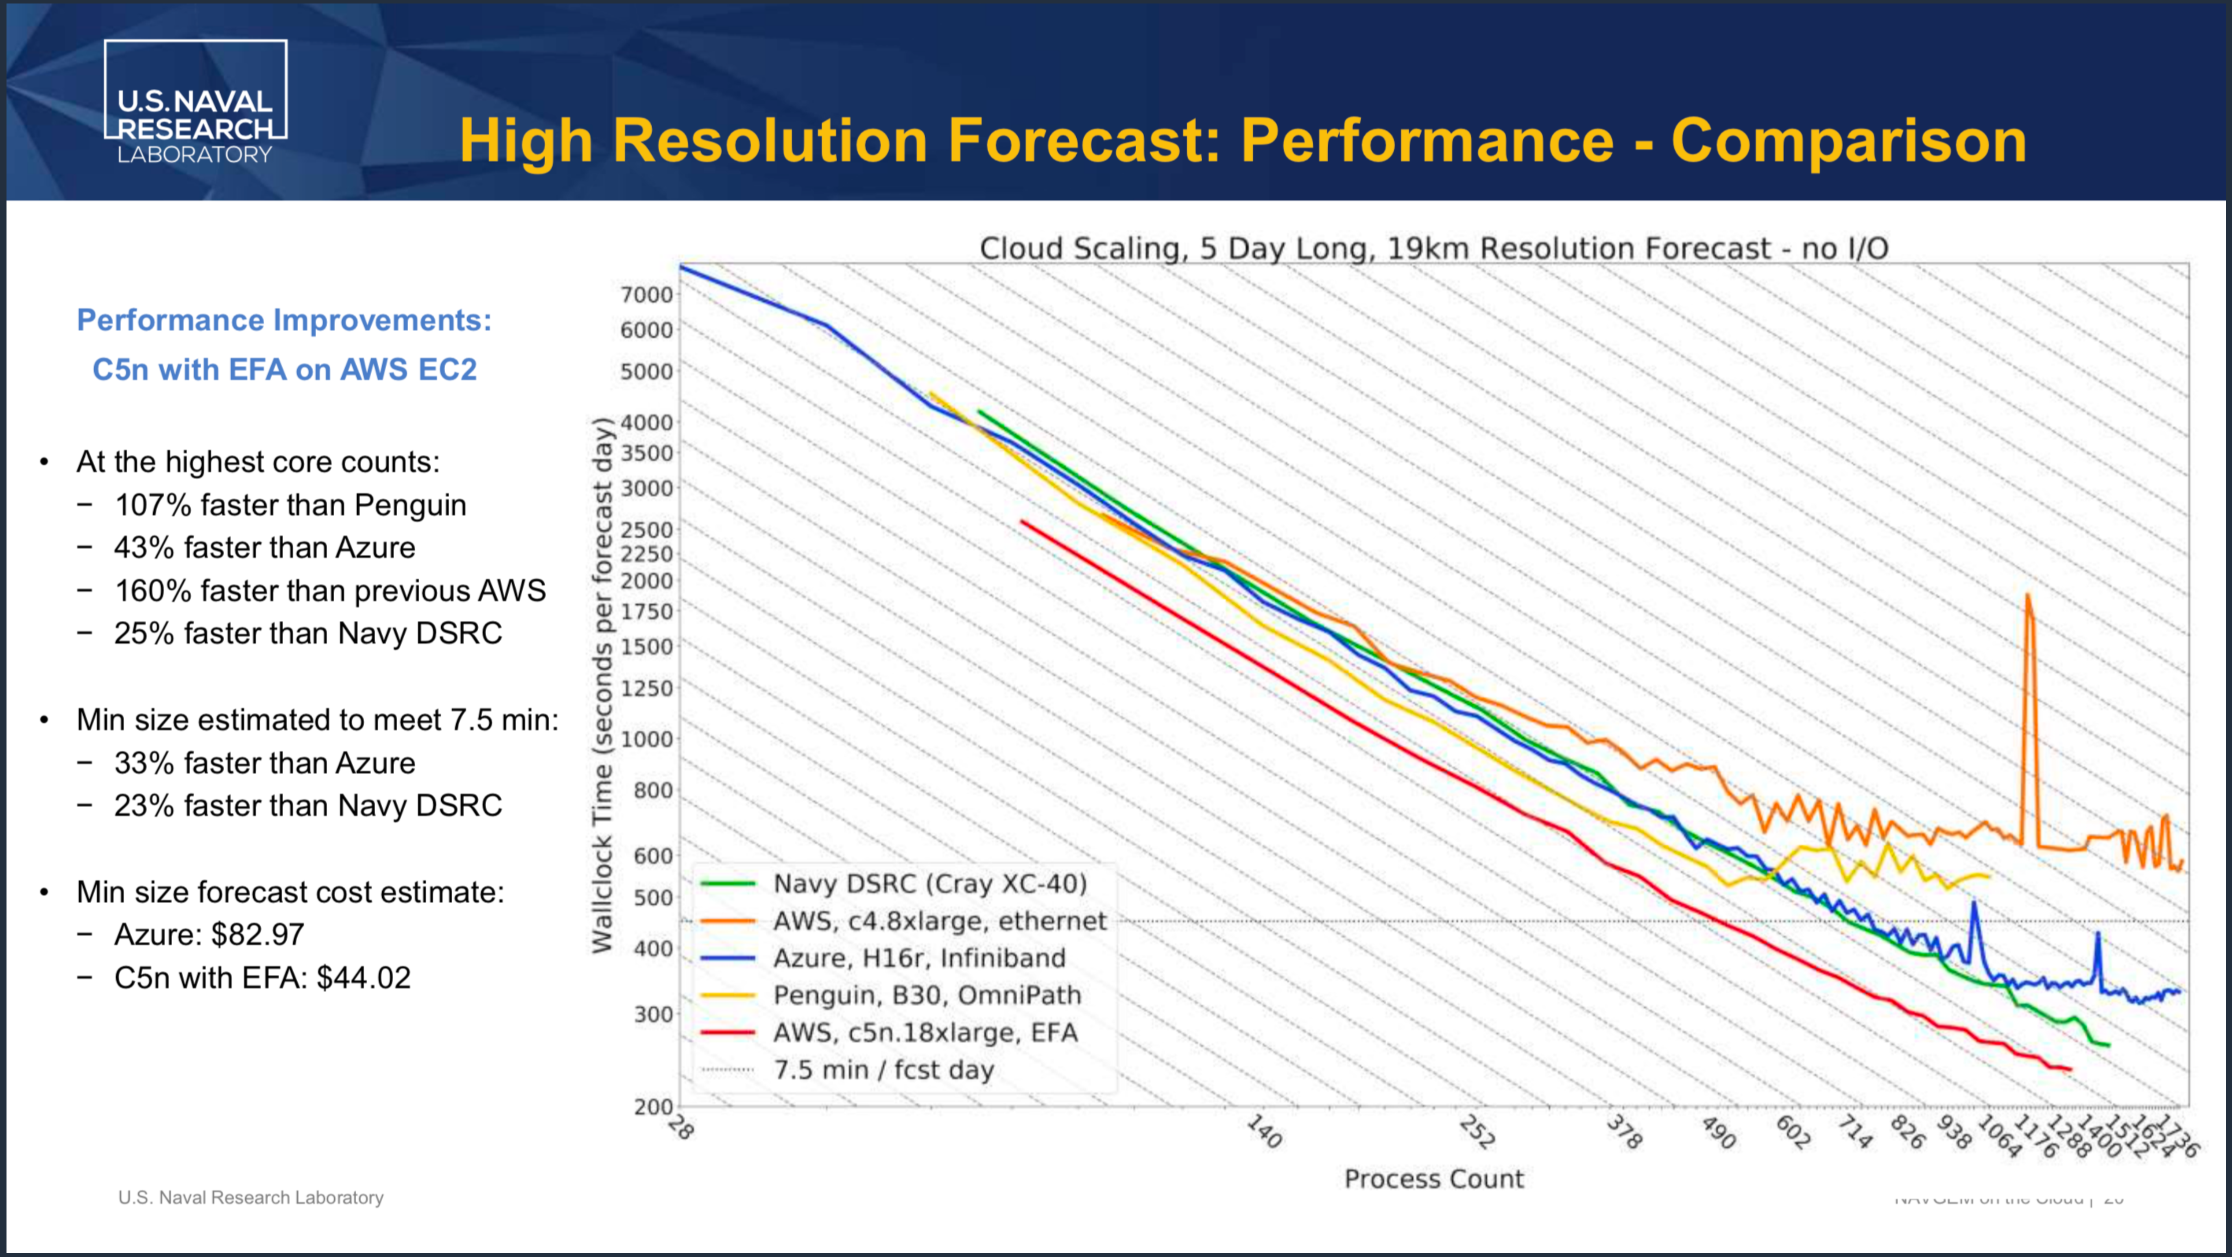

U.S. NAVAL RESEARCH LABORATORY

High Resolution Forecast: Performance - Comparison

Performance Improvements:
C5n with EFA on AWS EC2

- At the highest core counts:
  - 107% faster than Penguin
  - 43% faster than Azure
  - 160% faster than previous AWS
  - 25% faster than Navy DSRC
- Min size estimated to meet 7.5 min:
  - 33% faster than Azure
  - 23% faster than Navy DSRC
- Min size forecast cost estimate:
  - Azure: $82.97
  - C5n with EFA: $44.02

U.S. Naval Research Laboratory

Courtesy: U.S. Naval Research Laboratory https://www.youtube.com/watch?v=GTHWf0OVGrw&t=1177s

© 2020, Amazon Web Services, Inc. or its Affiliates.

Thank you!

Raghu Raja (craghun@amazon.com)

@rrcsraghu

© 2020, Amazon Web Services, Inc. or its Affiliates.