




Arm HPC Solutions and Beyond

Brent Gorda, Senior HPC Executive, Arm


Arm HPC Platform Workshop @ PCCC | August 26th, 2021

Extreme Scale Compute System Design Objectives

CPU Per-Core
Performance



Performance
Efficiency



Memory Bandwidth
& IO Flexibility



Technical
Sovereignty

Arm Is the Driving Force Behind Innovation for HPC and AI



NVIDIA GRACE



SIPEARL
The Silicon Pearl

SIPEARL RHEA



Ministry of Electronics &
Information Technology,
Government of India

MEITY SOC

ETRI

ETRI K-AB21

FUJITSU

FUJITSU A64FX



ARM NEOVERSE V-SERIES CORES




SPECIALIZED PROCESSORS



ARM SCALABLE VECTOR EXTENSION (SVE)



HIGH BANDWIDTH MEMORY




HPC

arm NEOVERSE

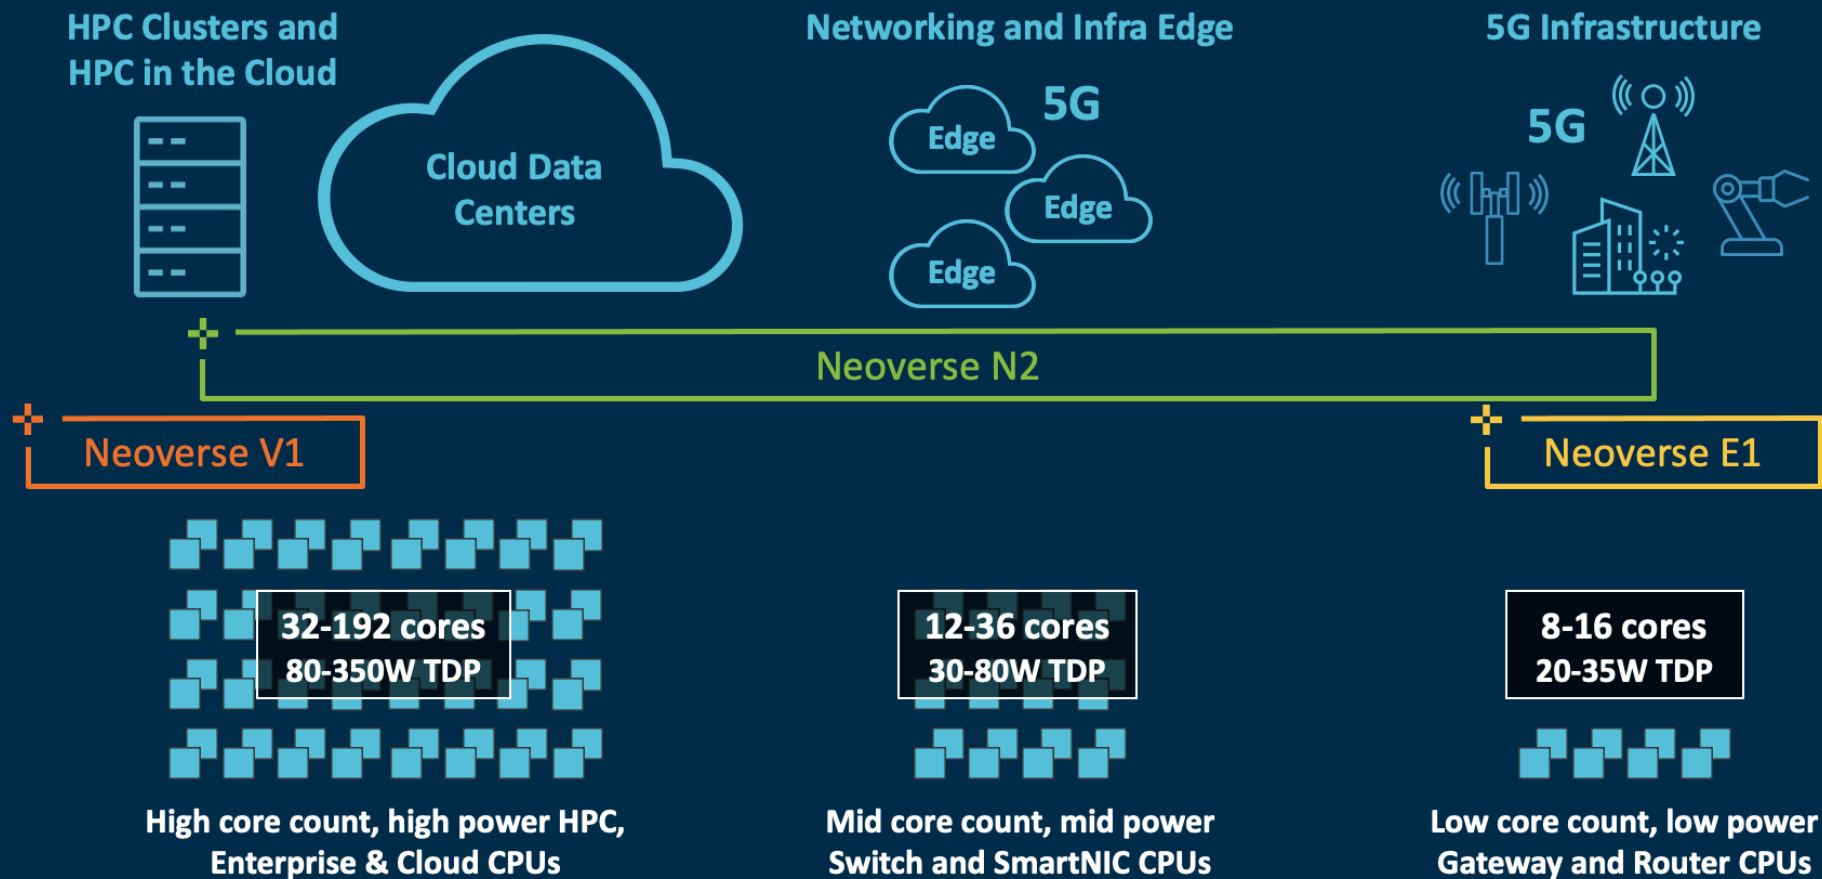
arm CORTEX

New infrastructure for support of 1T Arm devices

... will be designed from the edge in



Today's Server Silicon Design with Arm Neoverse



Cloud Native Software Development on Arm



100s of OSS Projects
Native Build



Performance
Scalable Neoverse CPU Cores




100+ ISVs
Commercial Support



Efficiency
Best \$/throughput in the industry



100K+
Docker Hub Images



Platform Diversity
Widest possible choice of Edge Platforms



1M+
CI/CD Build Minutes
per Month




Optimization
Purpose built use cases

AWS Graviton2 is Quickly Expanding its EC2 Footprint

Steady growth and regional expansion since mid-2020

2020 AWS Instance Type Additions



49% of AWS EC2 instance additions in 2020 are based on AWS Graviton2



NETFLIX



nielsen

flickr

hotelbeds

redbox.

LexisNexis®
RISK SOLUTIONS




SmugMug


honeycomb.io

RAYGUN


AWS Graviton2 Customer Adoption





NEW



+ R6gd



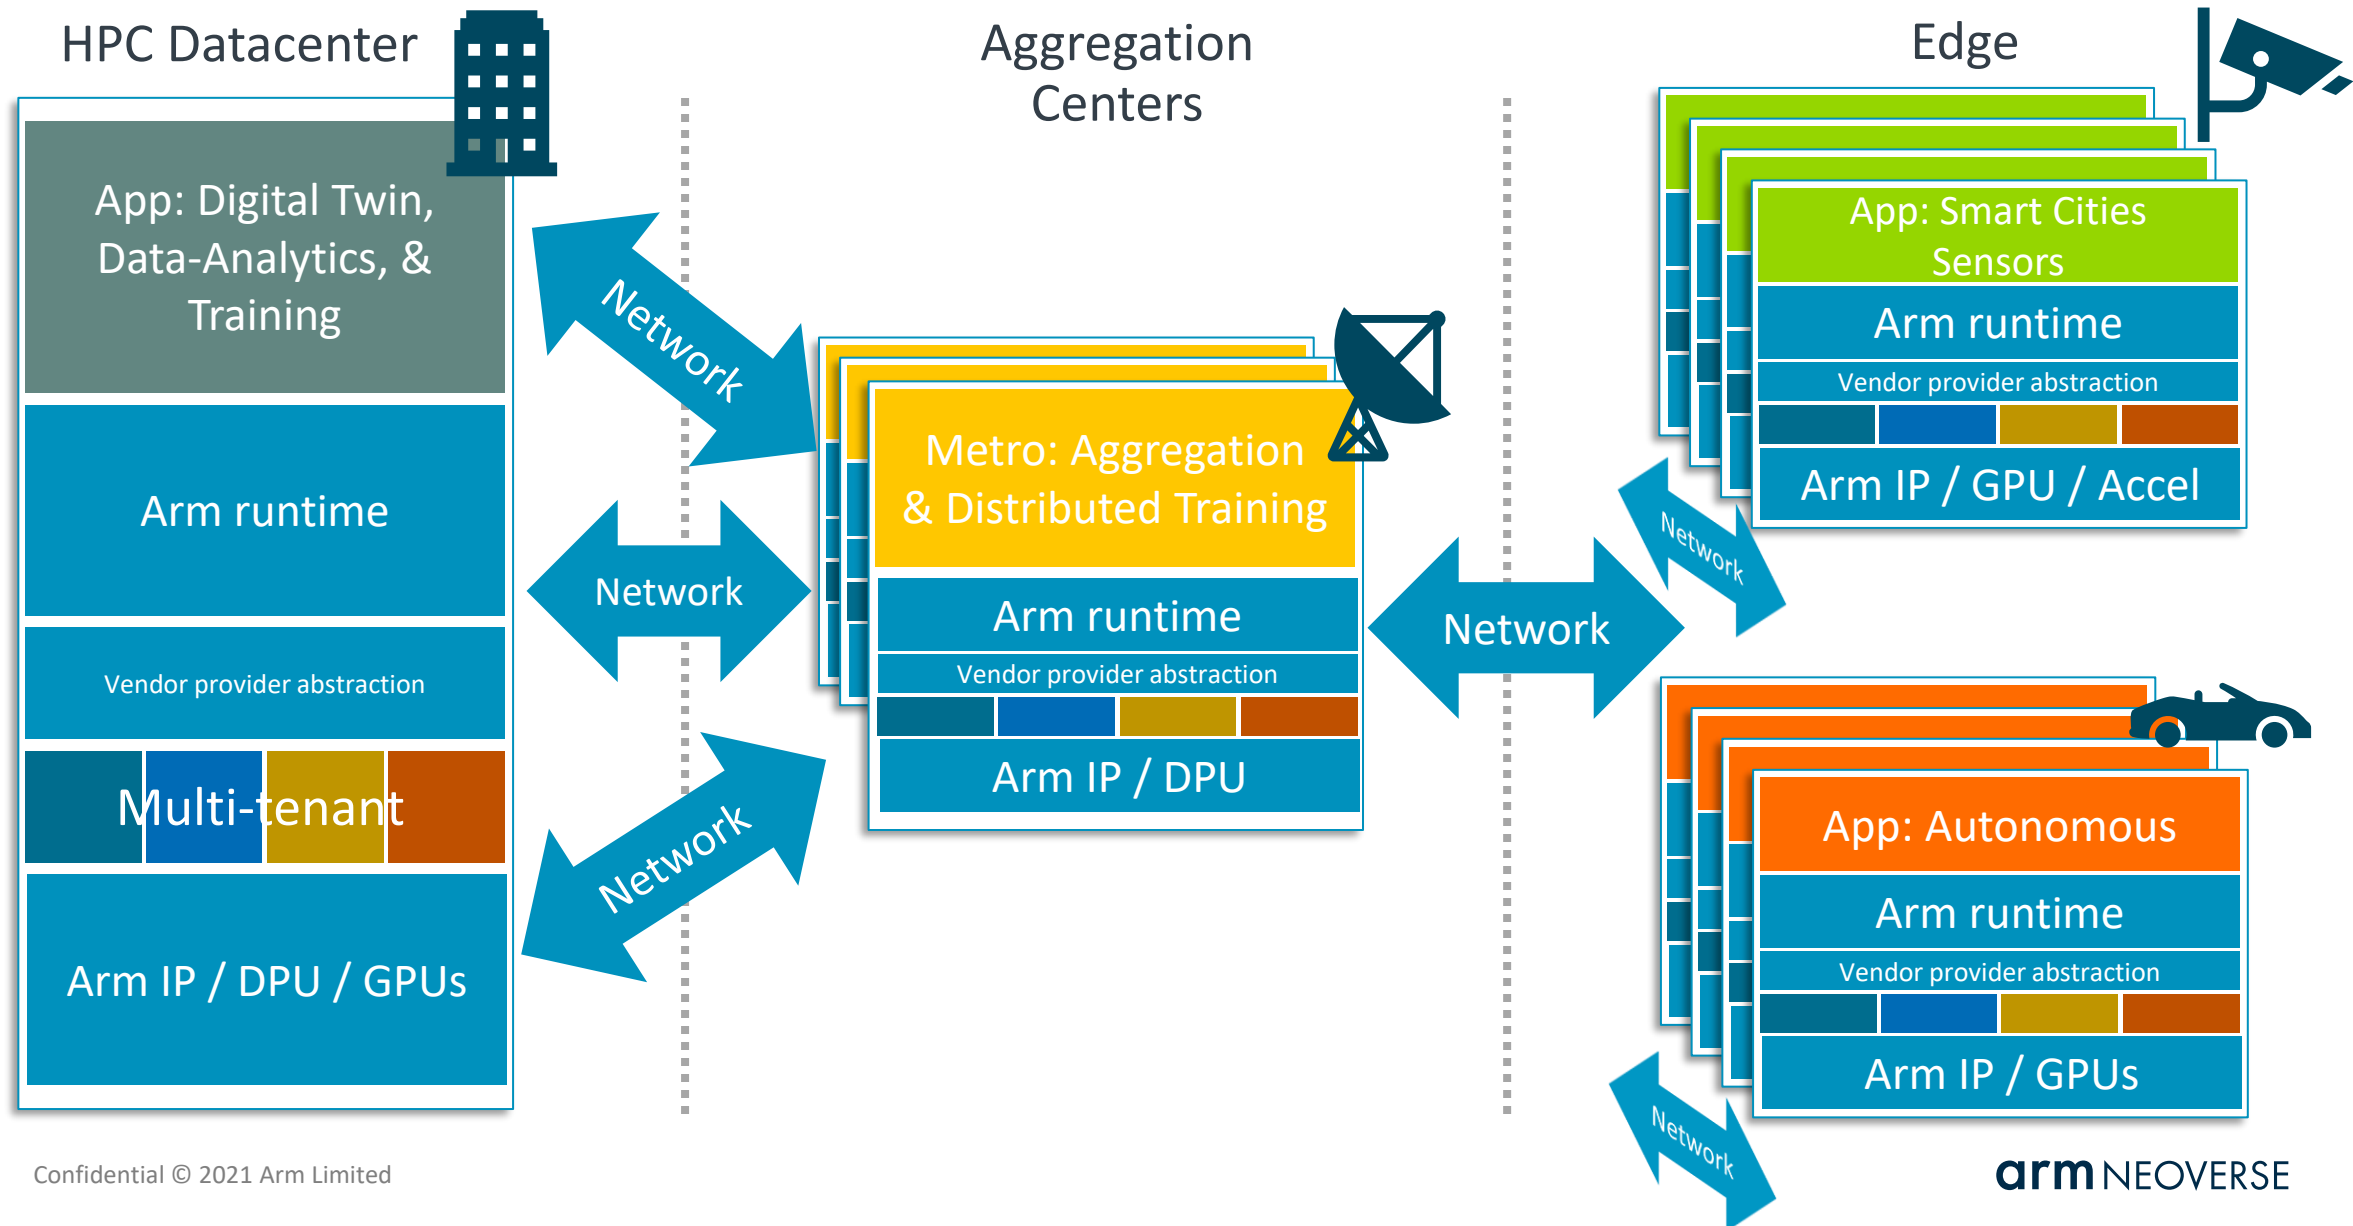
+ C6gd, C6gn




+ M6gd

Source: Litr Insights

Securely scaling ML application across heterogenous HW




Sensors



LIDAR

Software
Defined
Radios




Hyperspectral
Imaging

Facilities



Actuators




Dynamic
adaptation

Servos



AI @ Edge

Powerful
Parallel Edge
Computing



Artificial Intelligence
Deep Learning Inference
Lightweight Training

Edge computing and deep learning
with feedback for continuous
improvement



Reduced, Compressed data

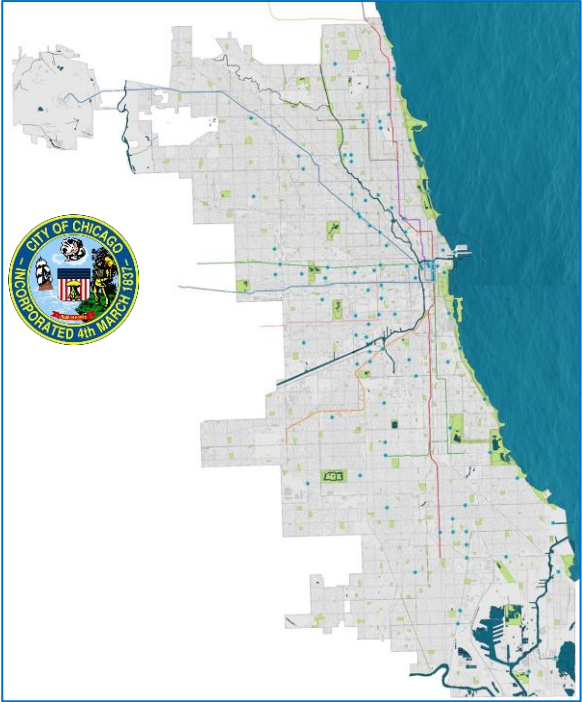
New inference (model)
Adaptive steering

HPC




Deep Learning Training
Simulation / Forecast

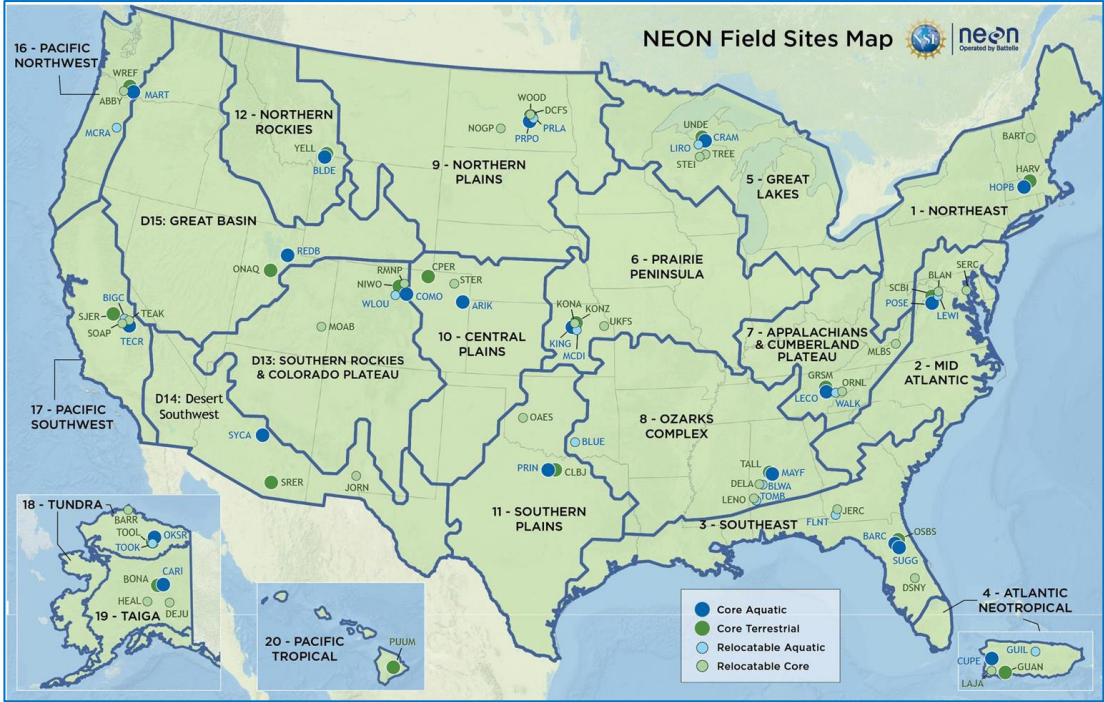
SAGE Partner Instruments




AoT: *Neighborhood* scale urban environment and activity.



HPWREN/WIFIRE: *Regional* Environmental Conditions and Events.



NEON: *Continental* scale ecology and environment.



Closing Remarks

- Arm brings mature SW stack (community driven) and competitive HW (licensees)
- Arm's business model empowers customers to design solutions for their markets
 - The A64FX is designed by Fujitsu to focus on RIKEN applications
 - Takes advantage of the software ecosystem (redhat, gnu, etc) from the community
 - Graviton2 & Ampere Altra leverage Arm IP: quick-to-market solutions that compete (+ use software)
- Arm has a very healthy roadmap and will continue to create ISA and IP solutions for the data center. Indeed growth in the hyperscalar market is a significant part of our future
- Arm is growing into cloud & hpc, from the device where we enjoy a healthy business
- The need of hyperscalers to innovate provides ever-more powerful Arm options for HPC
- We continue to invest in software in collaboration with the community
 - A-HUG.org is the HPC community

The background is a dark blue, monochromatic graphic. It features a perspective view of rows of server racks, with each rack displaying a grid of lights or indicators. Overlaid on this are diagonal bands of binary code (0s and 1s) that create a sense of depth and digital connectivity.

arm NEOVERSE

The Cloud to Edge Infrastructure Foundation
for a World of 1T Intelligent Devices

Thank You!