Arm in HPC

Jonathan Beard
Staff Research Engineer

arm NEOVERSE

Arm HPC Ecosystem

Arm IP

- Neoverse IP roadmap
- SVE Scalable Vector ISA Extension
- Arm v8.x ISA

With a growing software ecosystem

Operating Systems	   
Container & Virtualization	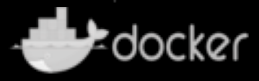  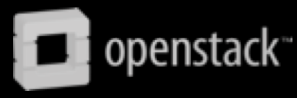 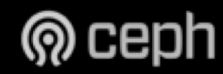 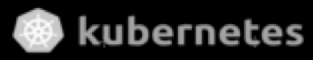  
Language & Libraries	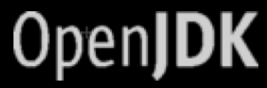  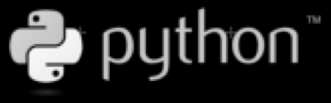 
Dev Tools & Resources	   
Open Source Projects Server & networking	  

Arm HPC Ecosystem

Arm IP

- Neoverse IP roadmap
- SVE Scalable Vector ISA Extension
- Arm v8.x ISA

Si Partners

FUJITSU
XILINX
HUAWEI
AMPERE
HISILICON
amazon

Integrators

Hewlett Packard Enterprise
Atos
CRAY

Deployments

Sandia National Laboratories
University of BRISTOL
cea
epcc
Stony Brook University
RIKEN

arm NEOVERSE

Recent Announcements

First Public Disclosure of Isaml
Performance Results

February 5, 2018 by Rich Brueckner Leave a Co

FUJITSU
shaping tomorrow with you

Supercomputer "Fugaku"
Formerly known as Post-K

ARM IS THE NNSA'S NEW SECRET WEAPON

November 7, 2018 Nicole Hemsoth

HPC Wire
Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

Home
Technologies

Cray ARM-based 'Ookami' to Serve as Testbed for Computational Studies at Stony Brook
August 16, 2019

STONY BROOK, N.Y., August 16, 2019 – A $5 million grant from the National Science Foundation (NSF) to the Institute of Advanced Computational Science (IACS) at Stony Brook University will enable researchers nationwide to test future supercomputing technologies and advance computational and data-driven research on the world's most pressing challenges.

CRAY

GW4, the Met Office, and Cray F
Supercomputer in Europe

Email

November 11, 2018 13:00 ET | Source: Cray Inc; The G

Data Centre ▸ HPC

HPE donates 3 mini-supercomputers to UK universities boning up on Arm

Muscling in Arm supers message on road to exascale

By Chris Mellor 16 Apr 2018 at 15:35 12 SHARE

Atos

Home / Newsroom / CEA acquires BullSequana supercomputer from Atos equipped with Marvell ThunderX2 Arm based processors

CEA acquires BullSequana supercomputer from Atos equipped with Marvell ThunderX2 Arm-based processors

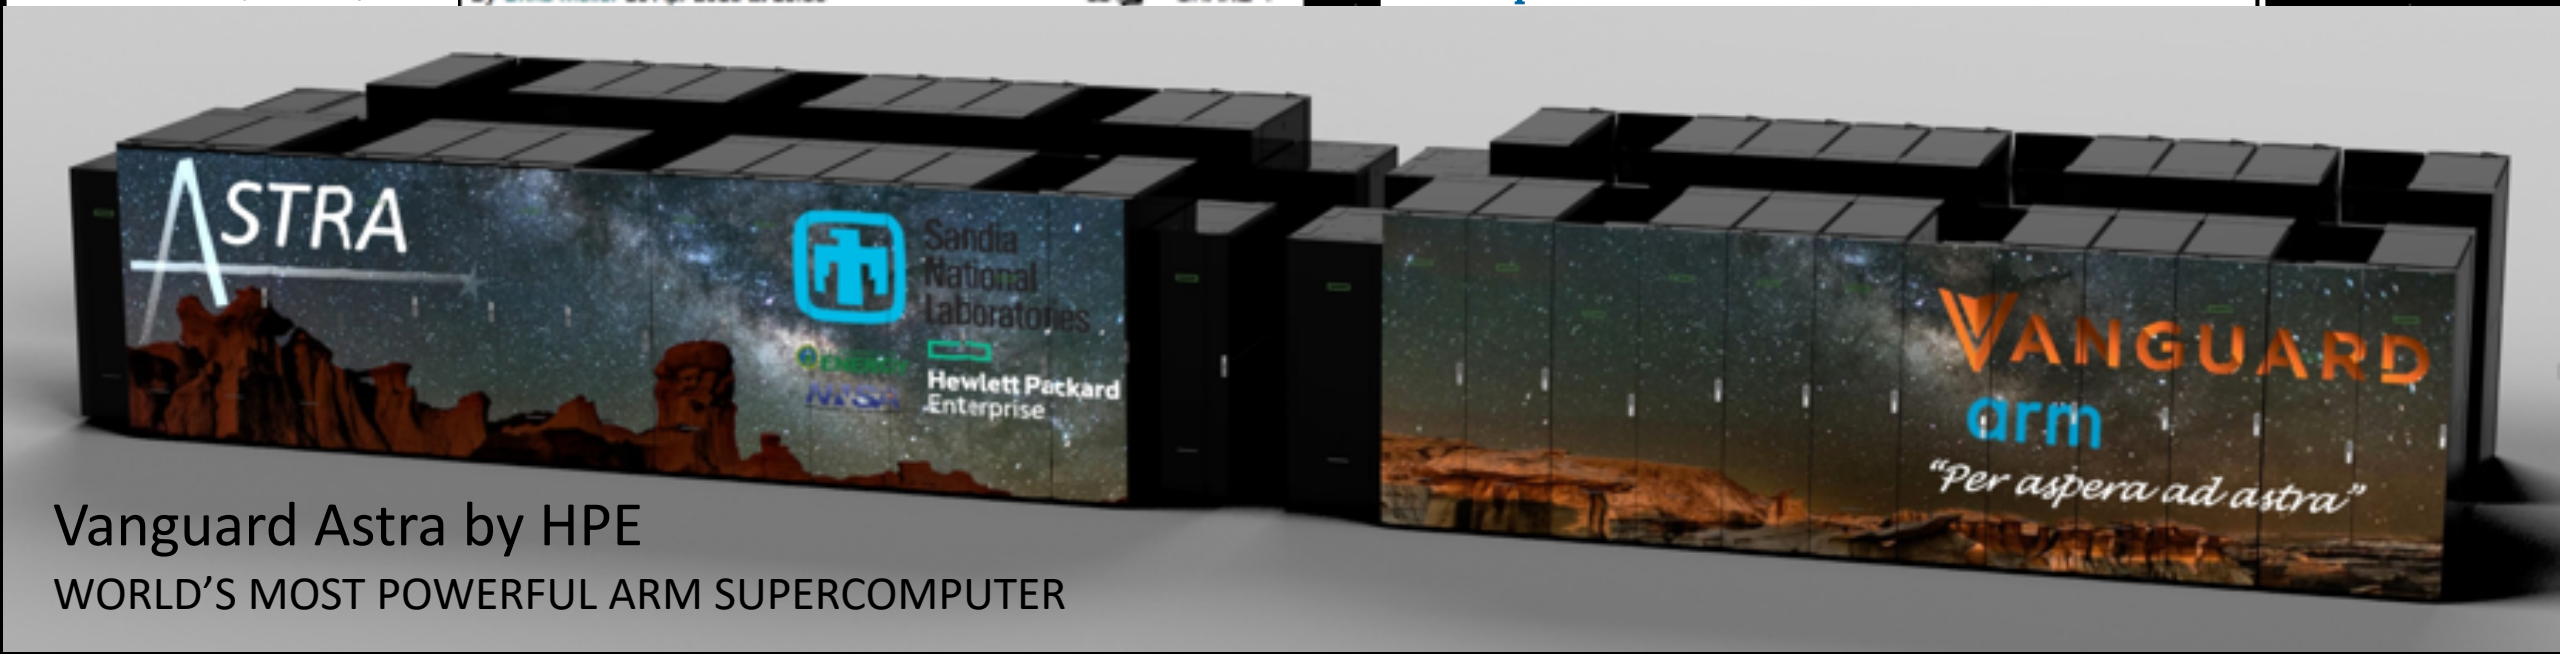

Vanguard Astra by HPE

WORLD'S MOST POWERFUL ARM SUPERCOMPUTER

Arm Neoverse IP

Cloud server
Control plane
Max efficiency

Cosmos Platform

Cortex-A72, A75
CMN-600
GIC, MMU

DDR4, 16nm

Ares Platform

Ares CPU
CMN + CCIX
GIC, MMU

7nm

Zeus Platform

Zeus CPU
CMN,
GIC, MMU

DDR5, 7nm+

Poseidon Platform

Poseidon CPU
CMN,
GIC, MMU

5nm

Maximum efficiency
Data plane

Cortex-A55

Helios Platform

Silicon in market

Today

2020

2021

arm NEOVERSE

Edge to data center

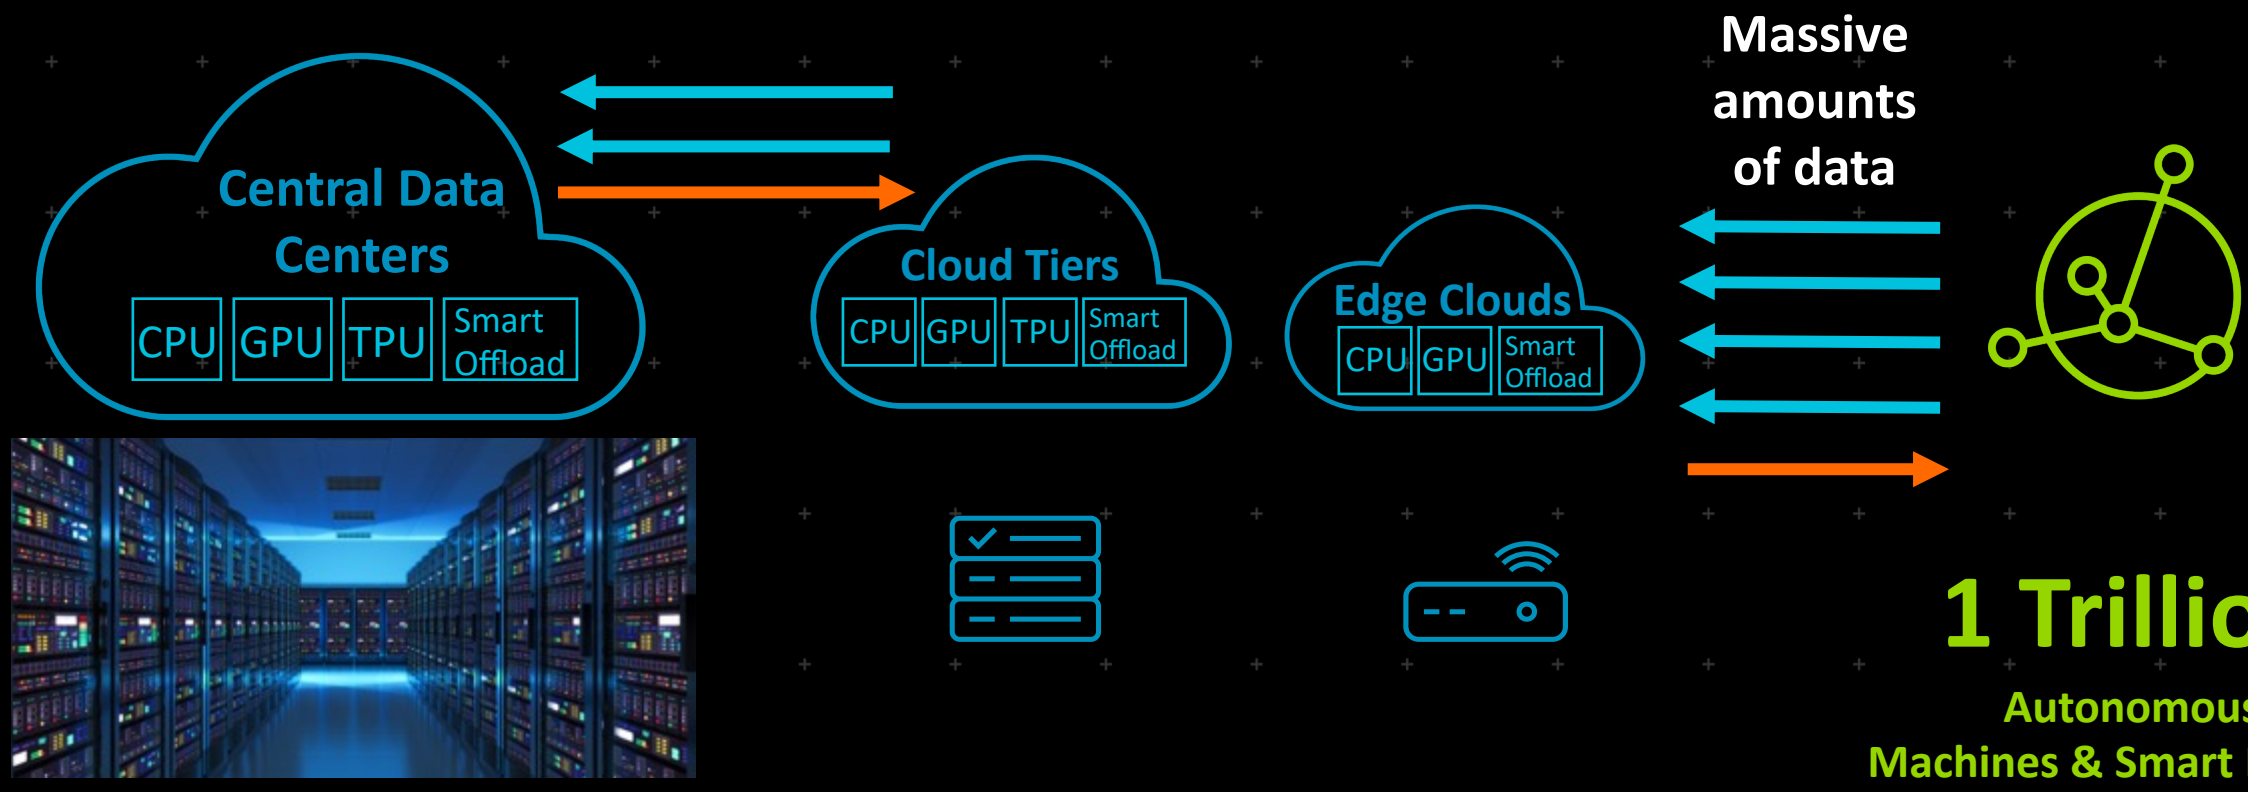

Infrastructure Will Be Designed from the Edge In

arm NEOVERSE

Potential long-term research directions

(The more fun stuff, for me at least)

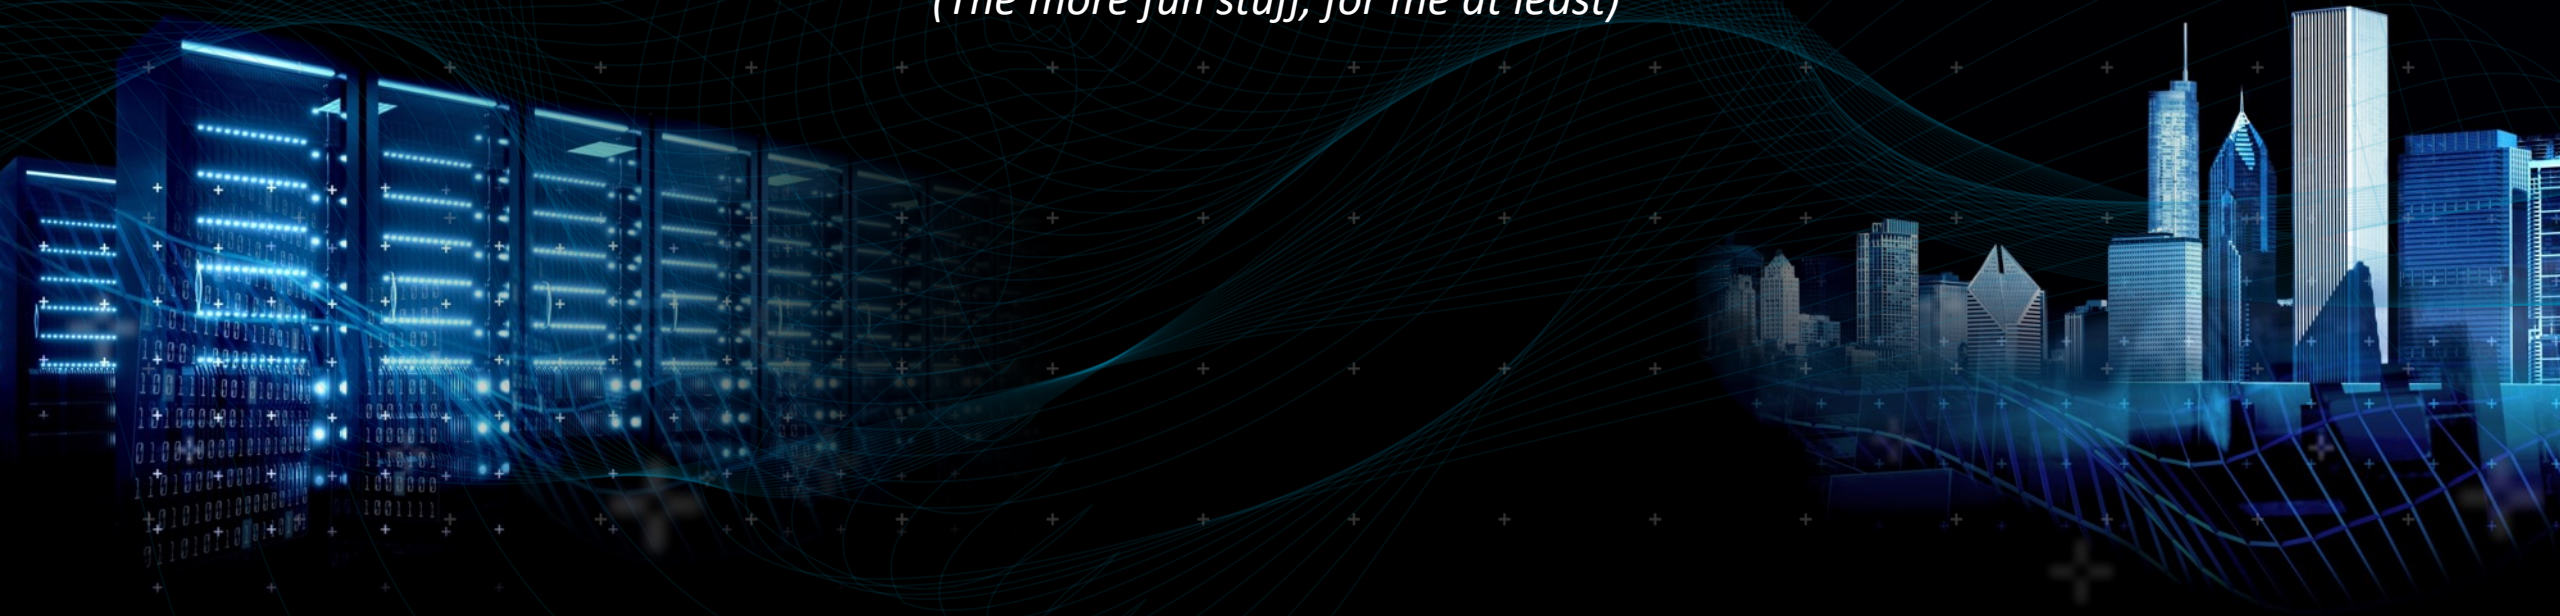

Where is the edge drawn?

CPU

Where is the edge drawn?

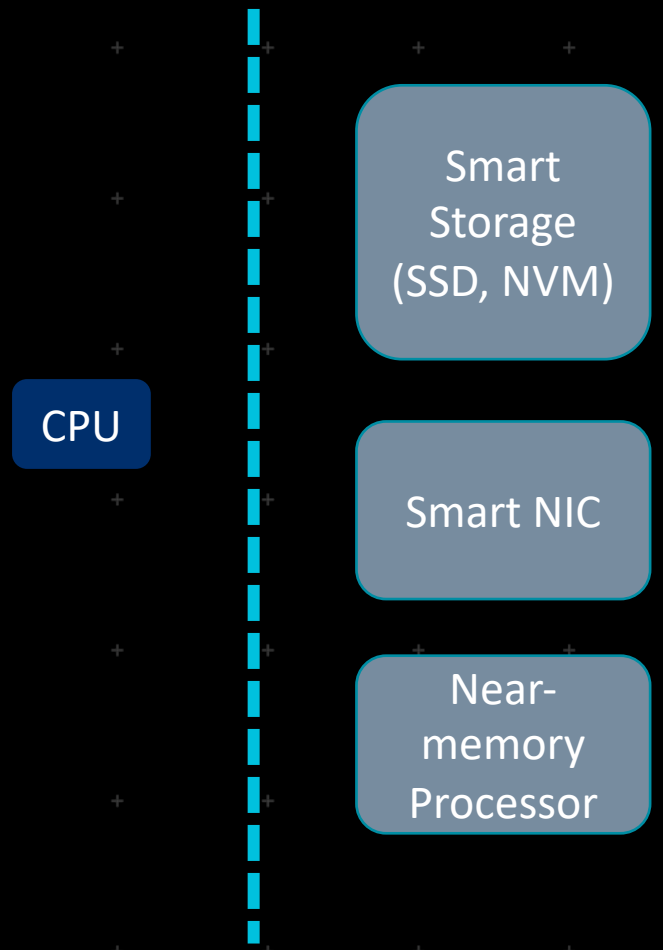

Where is the edge drawn?

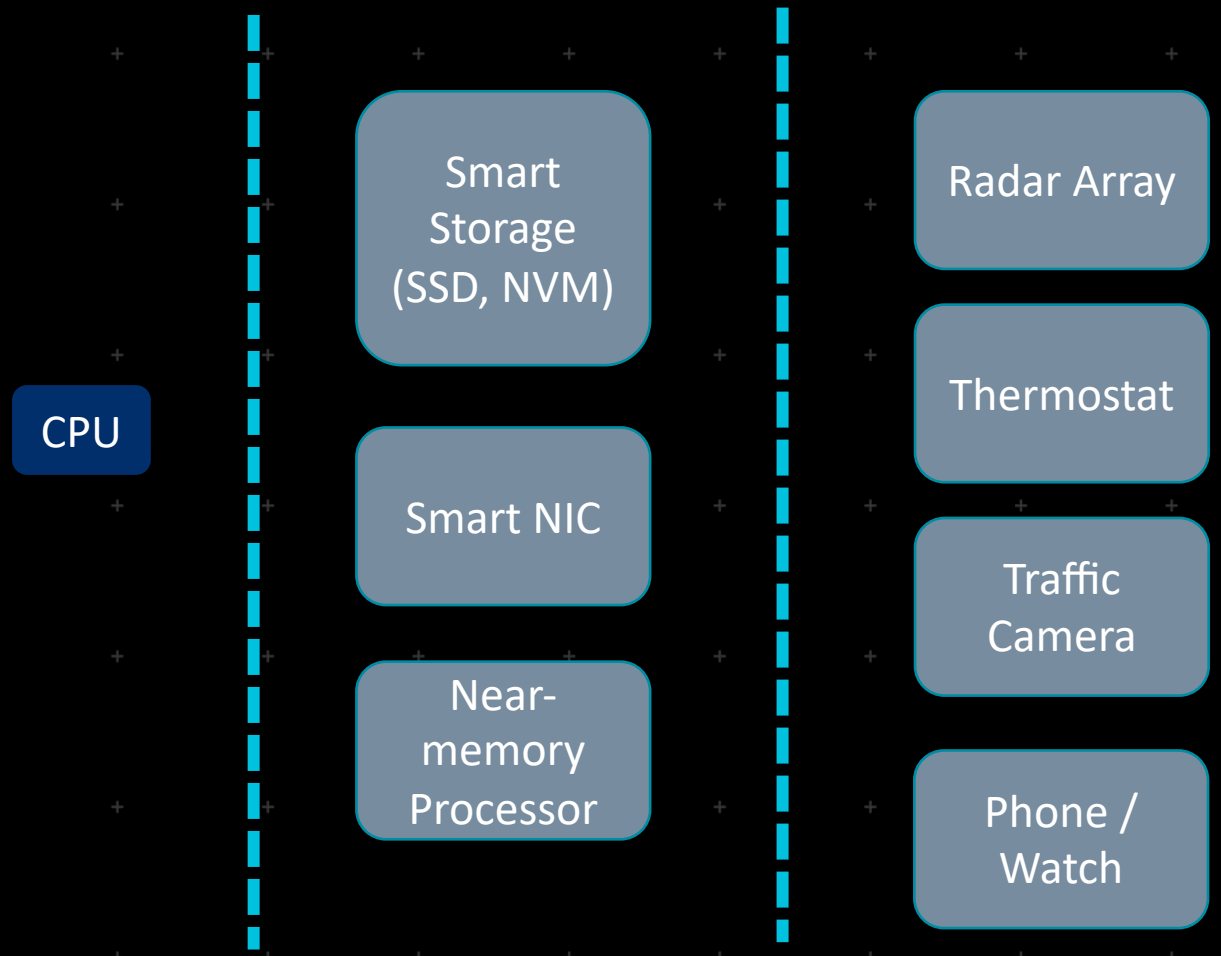

Where is the edge drawn?

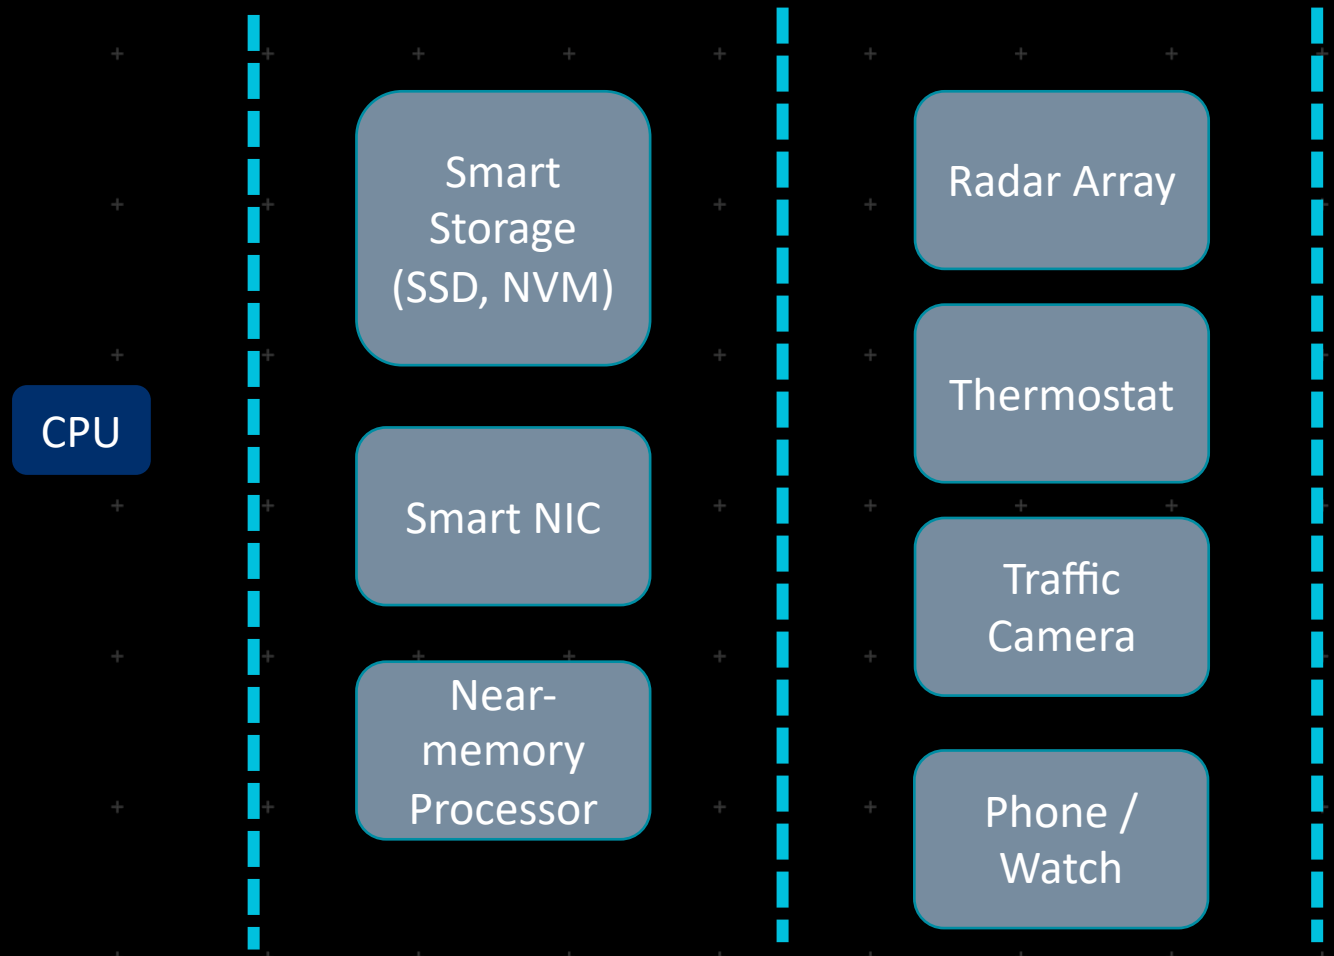

- Each of these likely has a core capable of running a container
- Sometimes similar security concerns (can we abstract this?)
- Can we make processing look the same in each of these devices? Or at least as easy to use?
- Goal is same in most cases, reduce data movement…

The magic 8-ball says

Adapted/modified from original figure courtesy of Dilip Vasudevan (LBL)

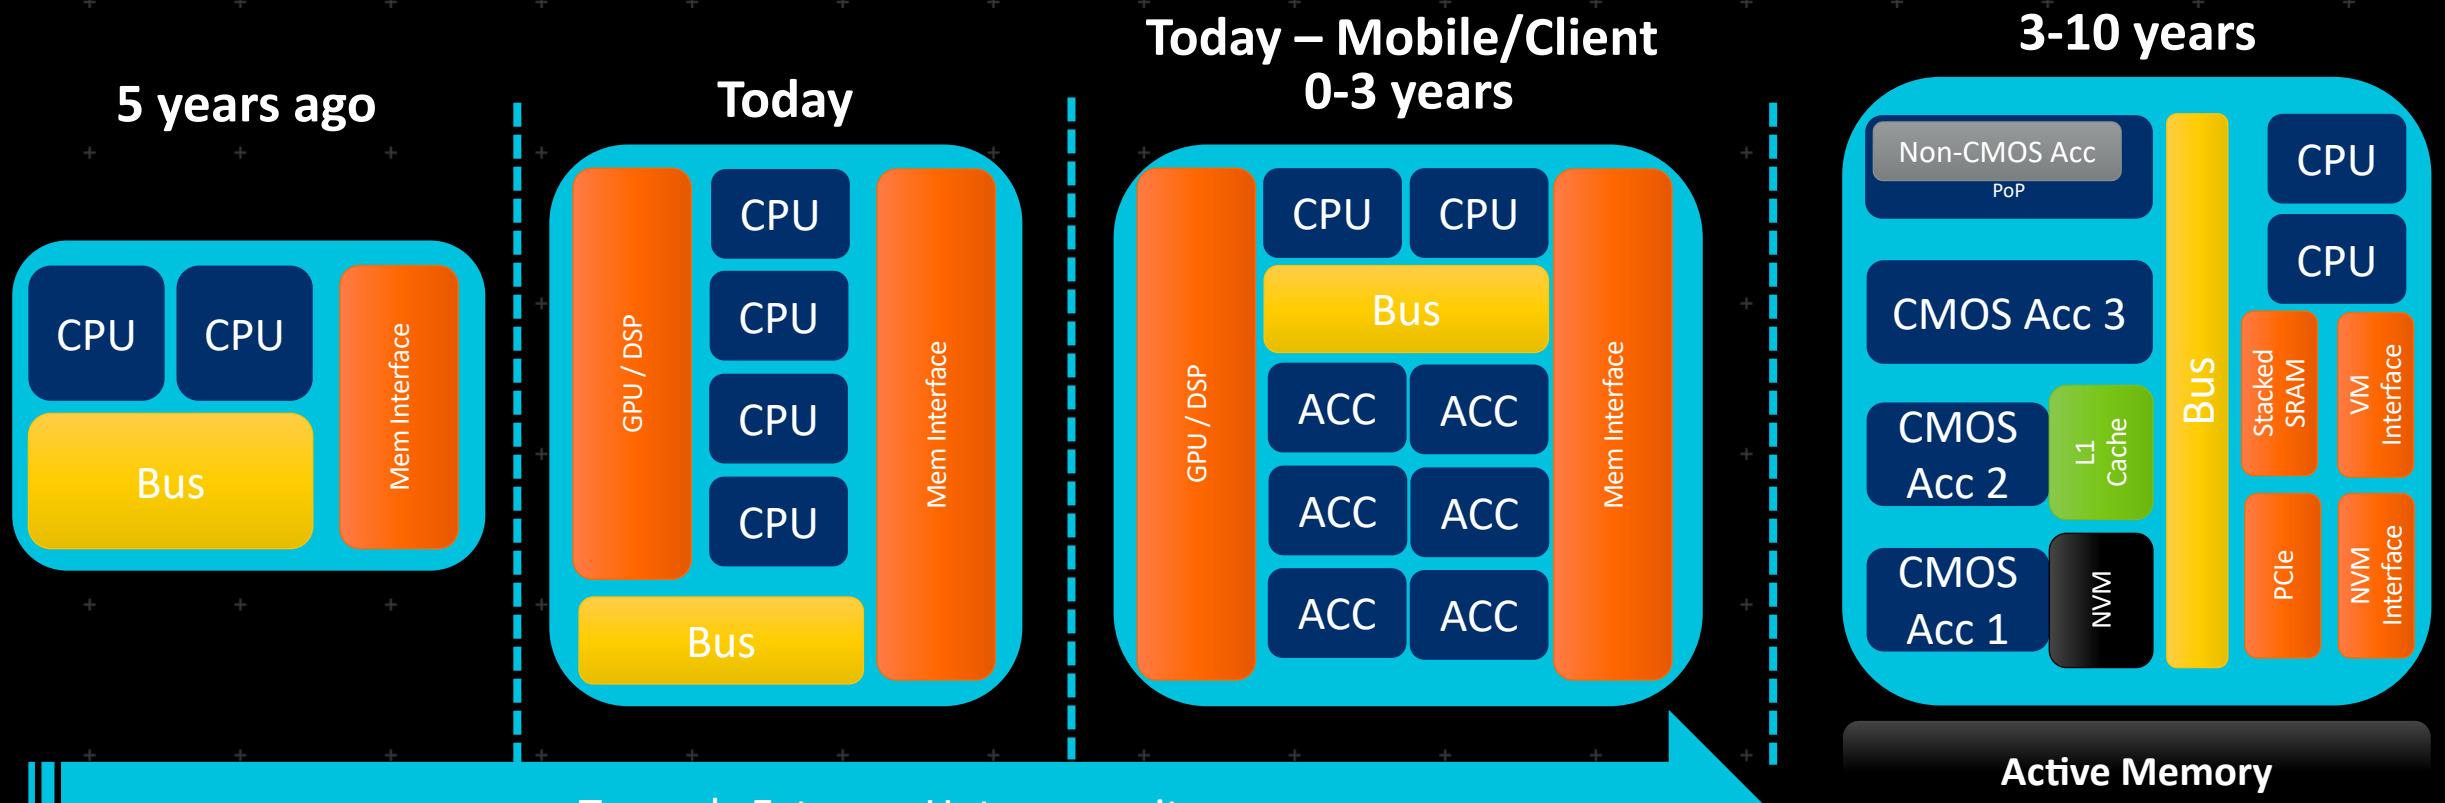

Towards Extreme Heterogeneity

Low — **Cost to build/adopt/run** — **Extreme**

Low — **Opportunity for Disruption** — **High**

arm NEOVERSE

Making extreme heterogeneity a reality

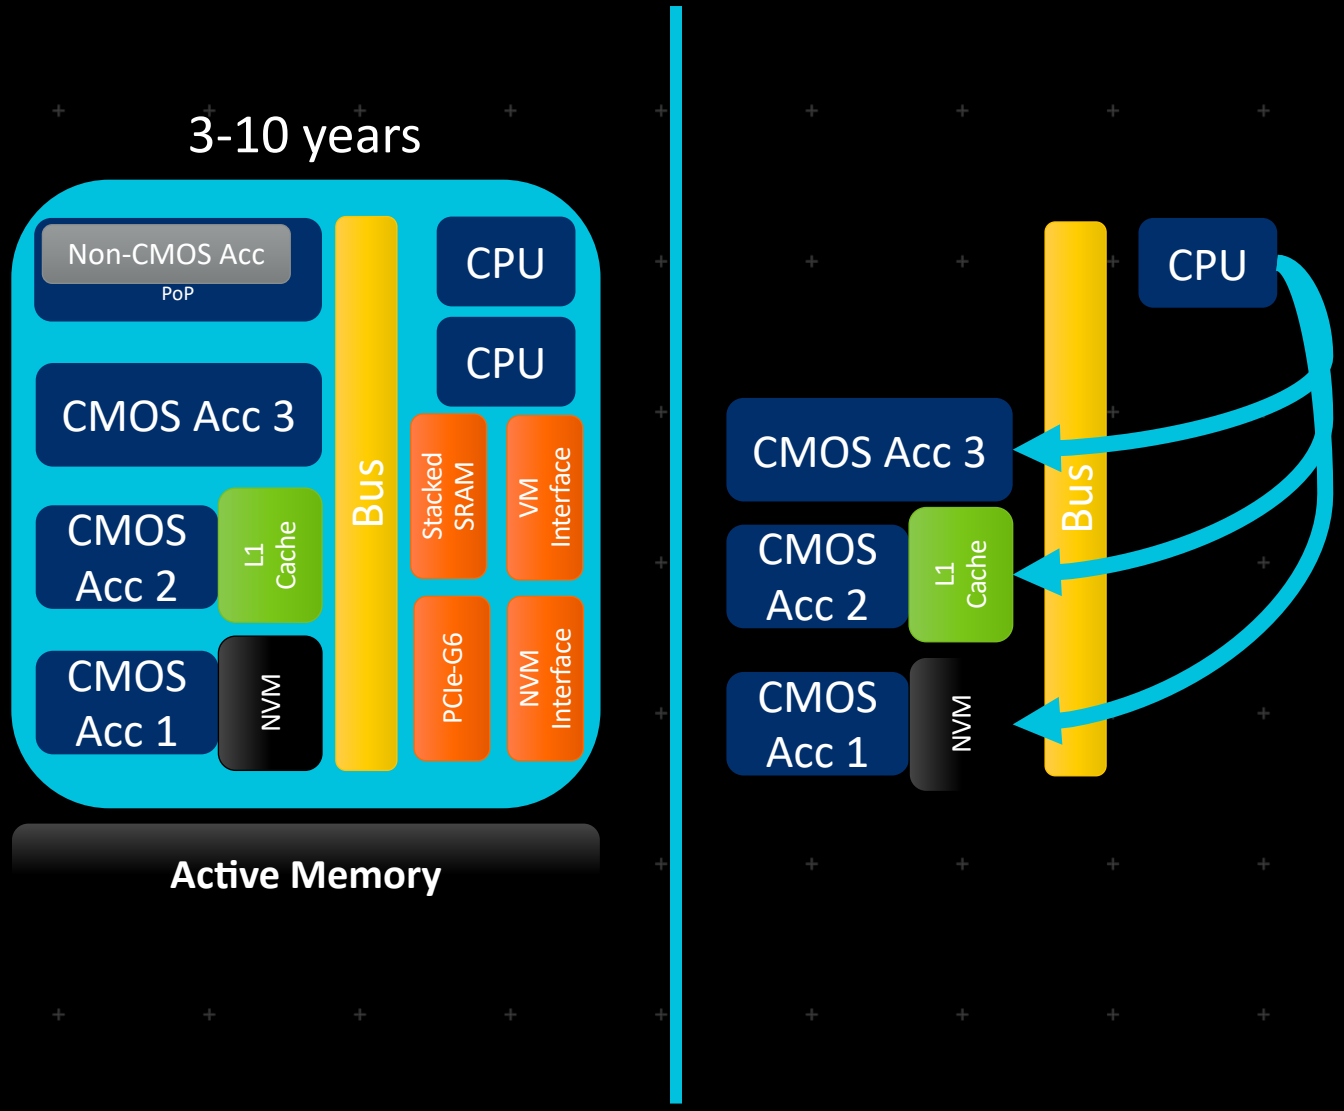

1: Software adoption (or lack of) kills most novel accelerators

- **Standardizing on the interface layer would be useful, can we do this?**

2: Efficiency of data movement

- **logic is cheap(-ish), data movement is**

3: Coherence scales only so far. Dataflow graph execution works well, can we virtualize it to be transport medium agnostic?

4: Virtualization and translation for accelerators is an afterthought at the moment, can we do more?

Enabling an ecosystem

Unlock innovation on both sides of interface! – Minimize software disruption, maximize innovation pace

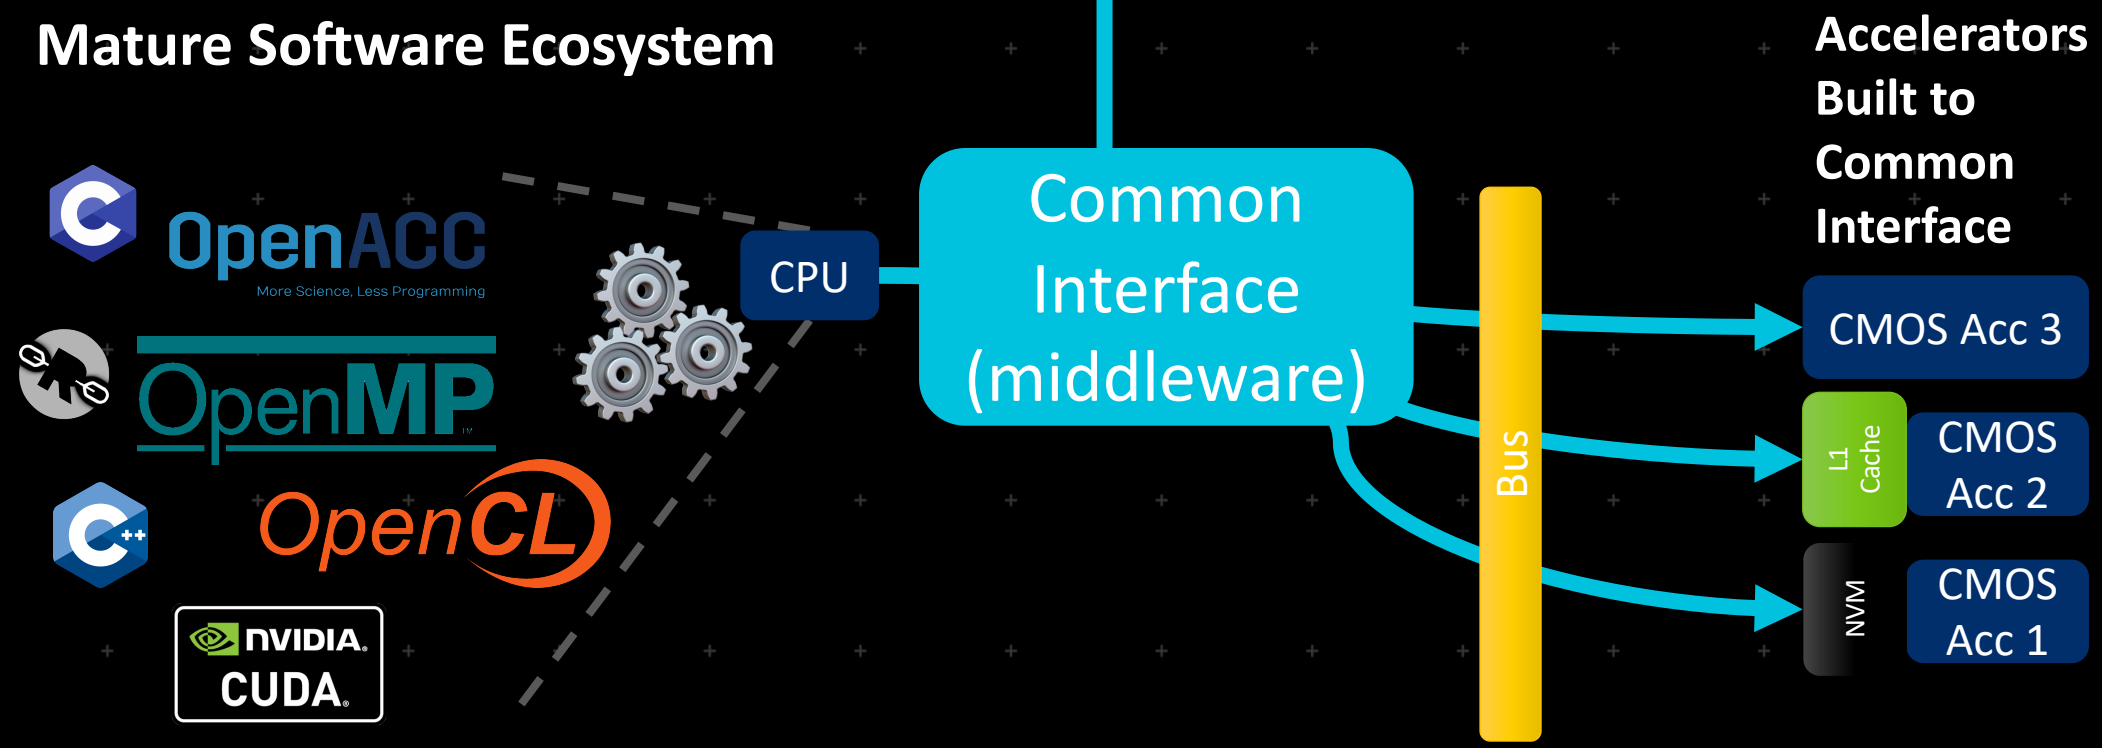

arm NEOVERSE

Compute efficiency post-Moore

Data movement dominates, parallelism is critical

- Reduce latency to initiate a heterogeneous-parallel task
- Decrease communications overhead (reduce state transference)
- Increase data locality
- Reduce programmer effort for heterogeneous systems

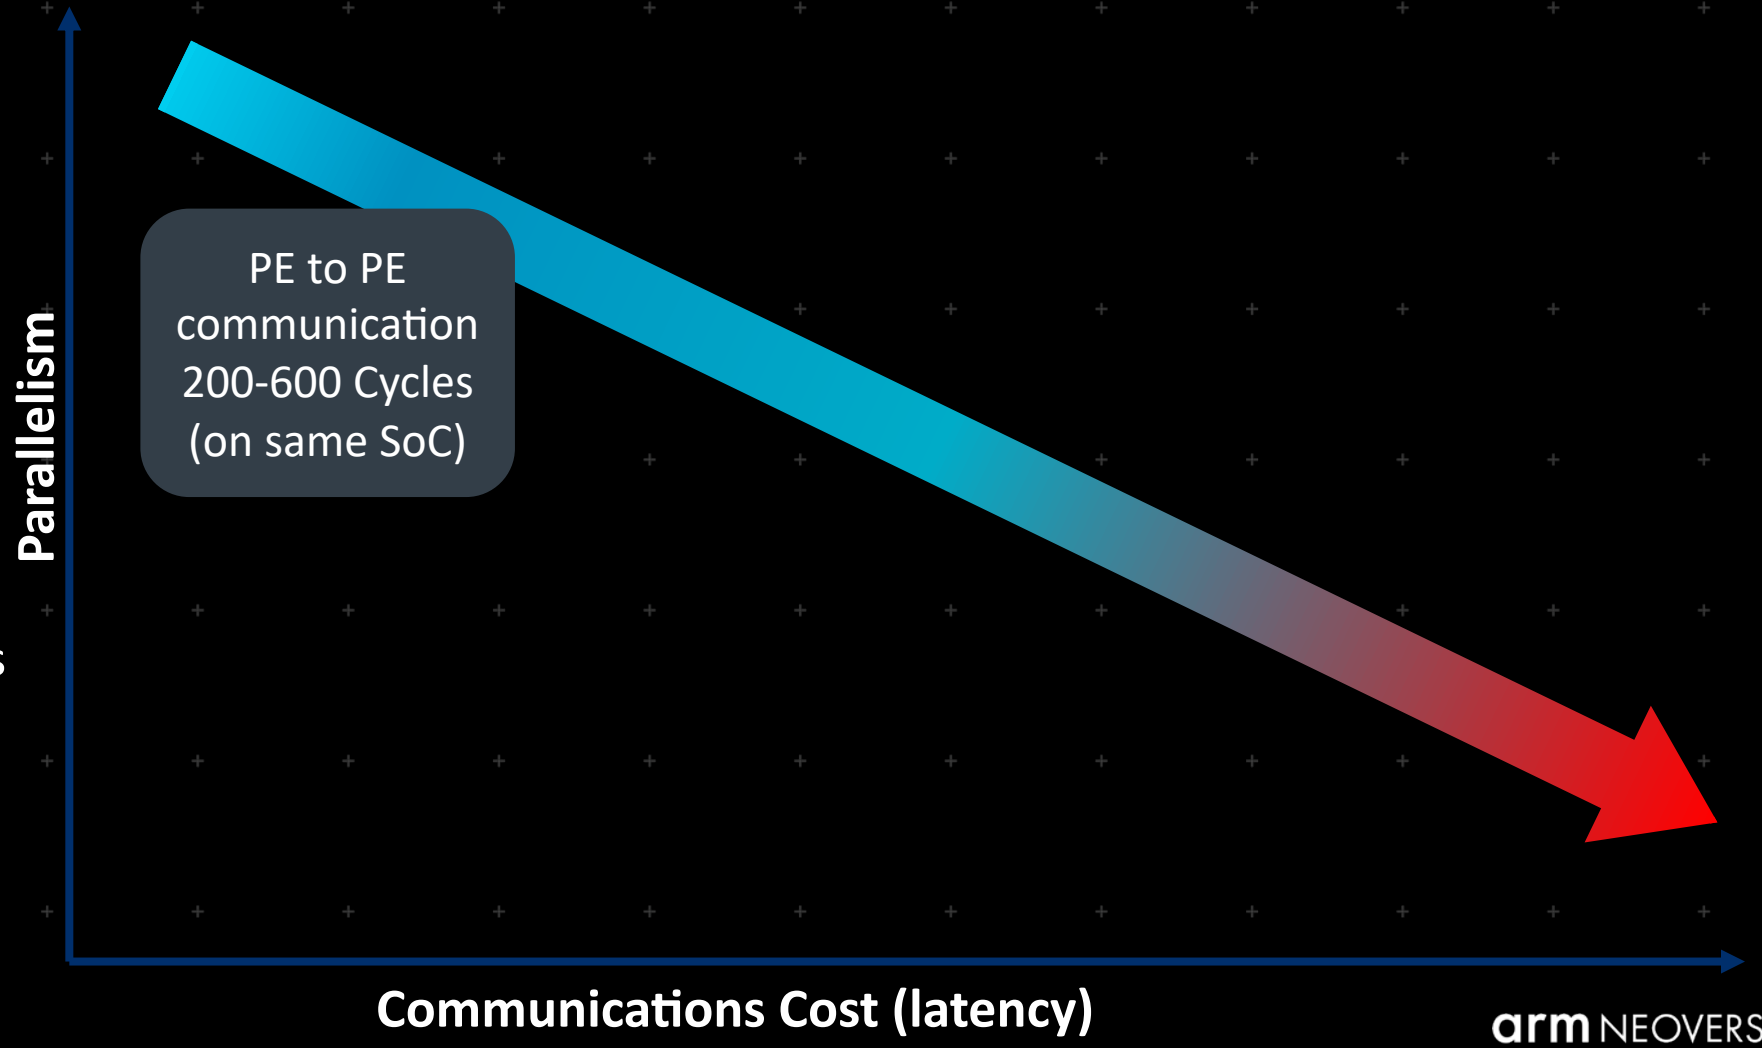

arm NEOVERSE

Compute efficiency post-Moore

Data movement dominates, parallelism is critical

- Reduce latency to initiate a heterogeneous-parallel task
- Decrease communications overhead (reduce state transference)
- Increase data locality
- Reduce programmer effort for heterogeneous systems

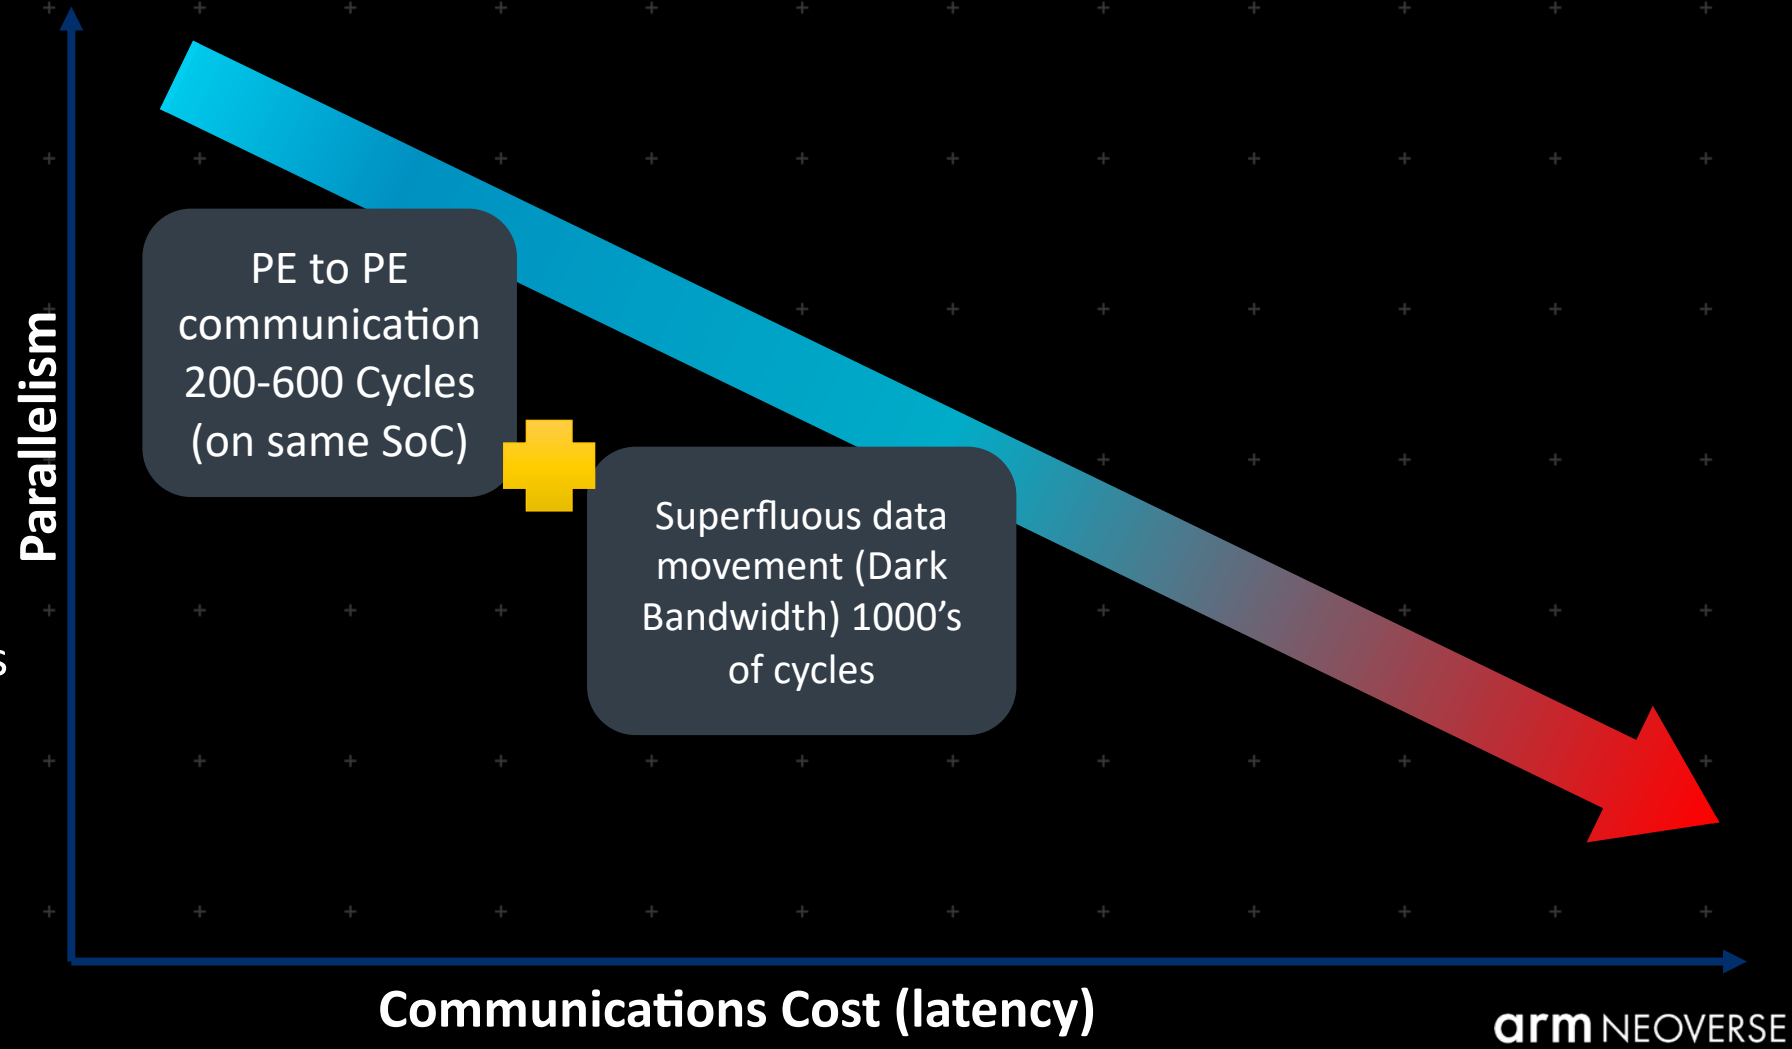

arm NEOVERSE

Compute efficiency post-Moore

Data movement dominates, parallelism is critical

- Reduce latency to initiate a heterogeneous-parallel task
- Decrease communications overhead (reduce state transference)
- Increase data locality
- Reduce programmer effort for heterogeneous systems

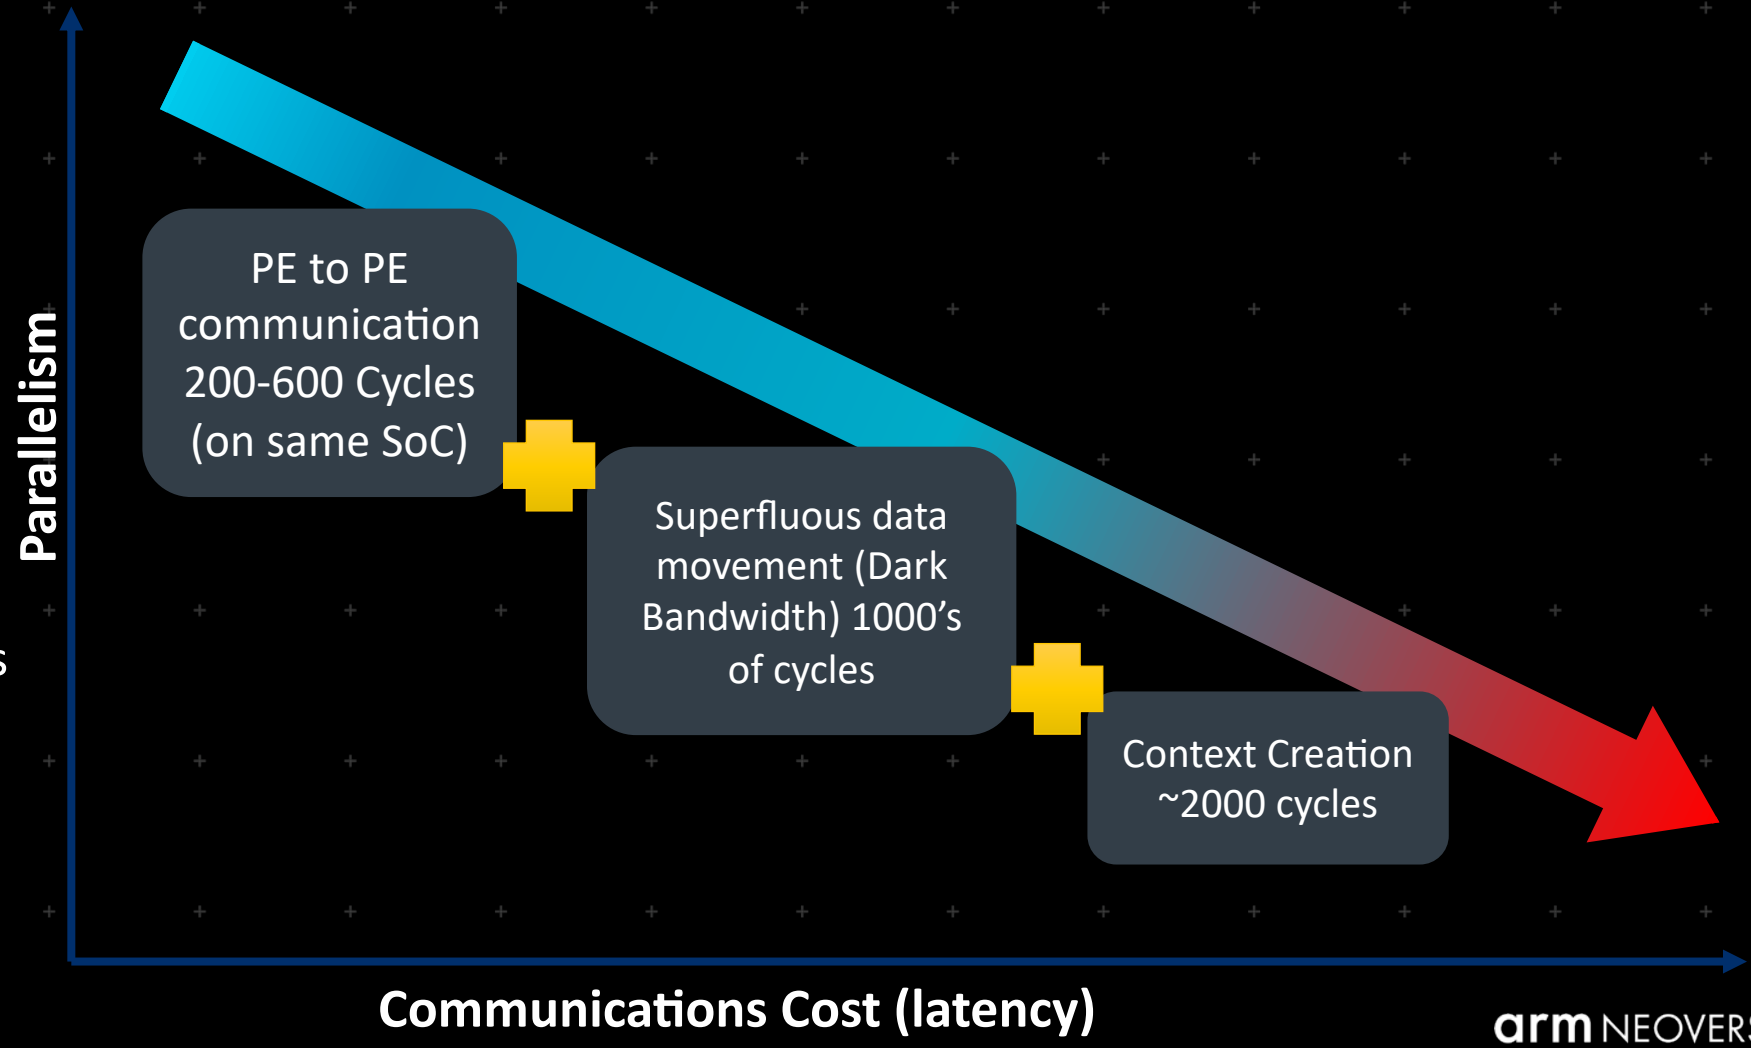

arm NEOVERSE

Compute efficiency post-Moore

Data movement dominates, parallelism is critical

- Reduce latency to initiate a heterogeneous-parallel task
- Decrease communications overhead (reduce state transference)
- Increase data locality
- Reduce programmer effort for heterogeneous systems

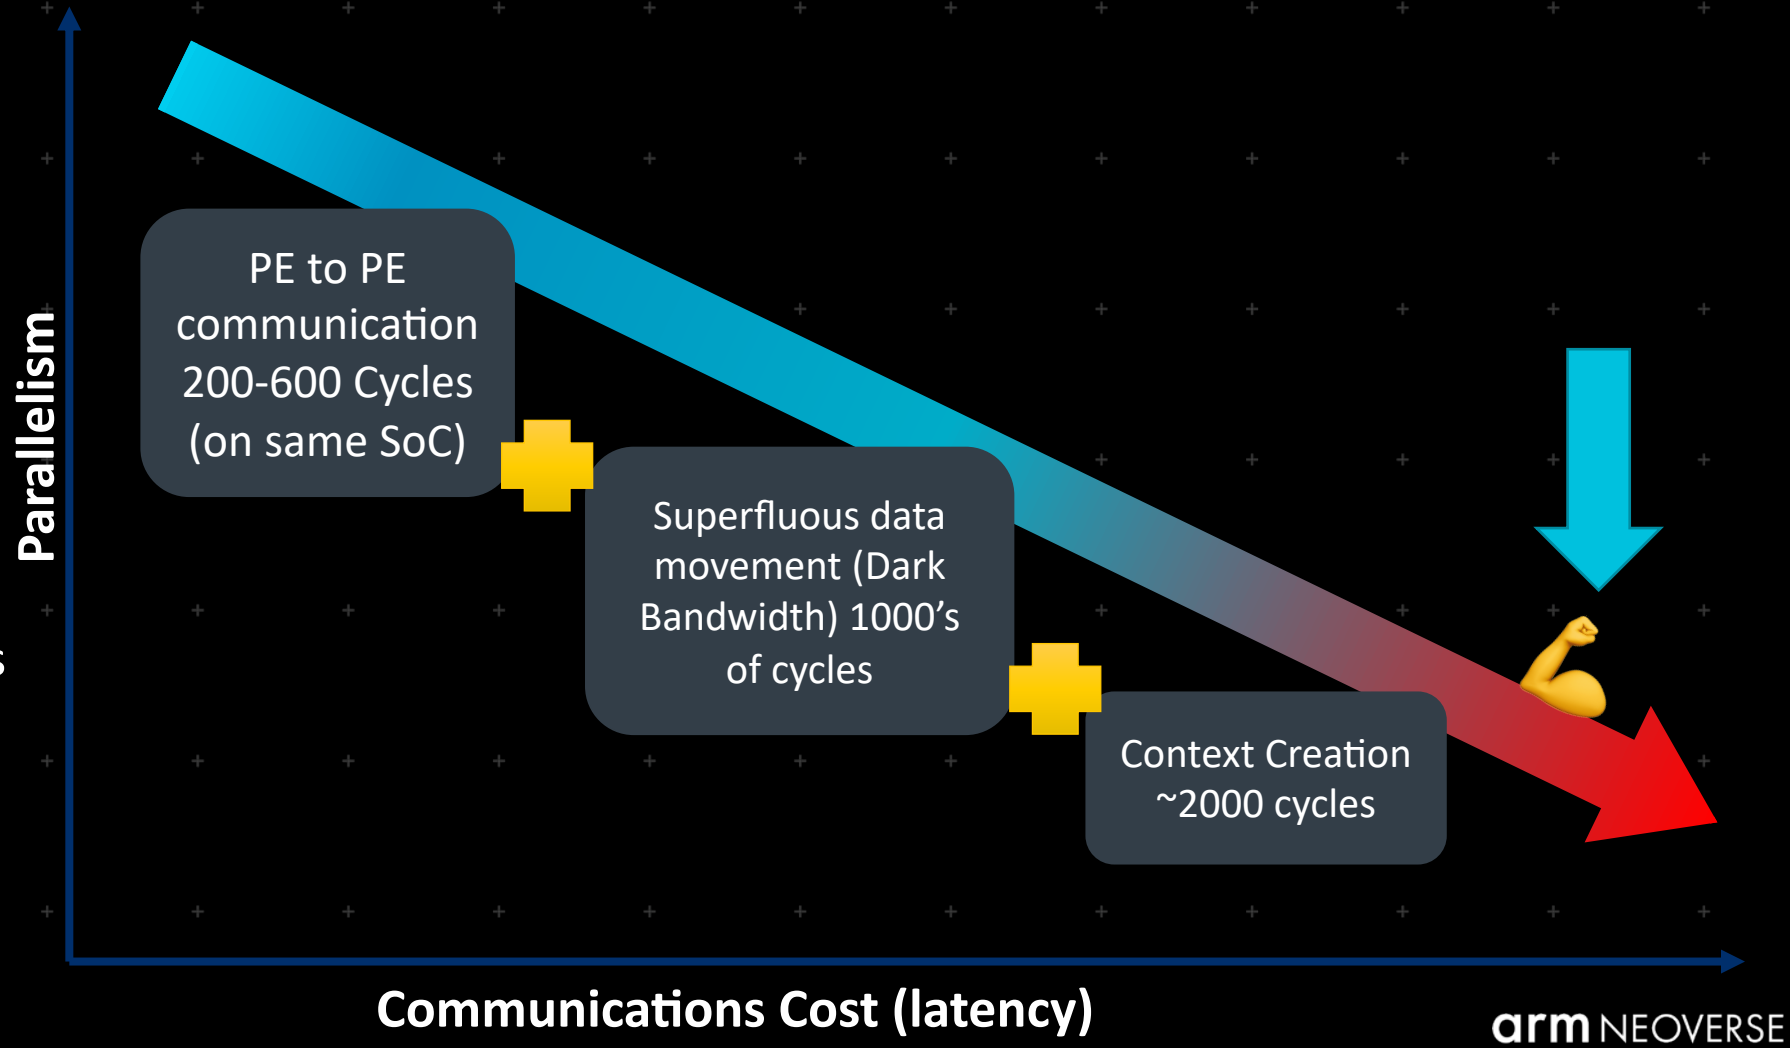

arm NEOVERSE

Enabling an ecosystem

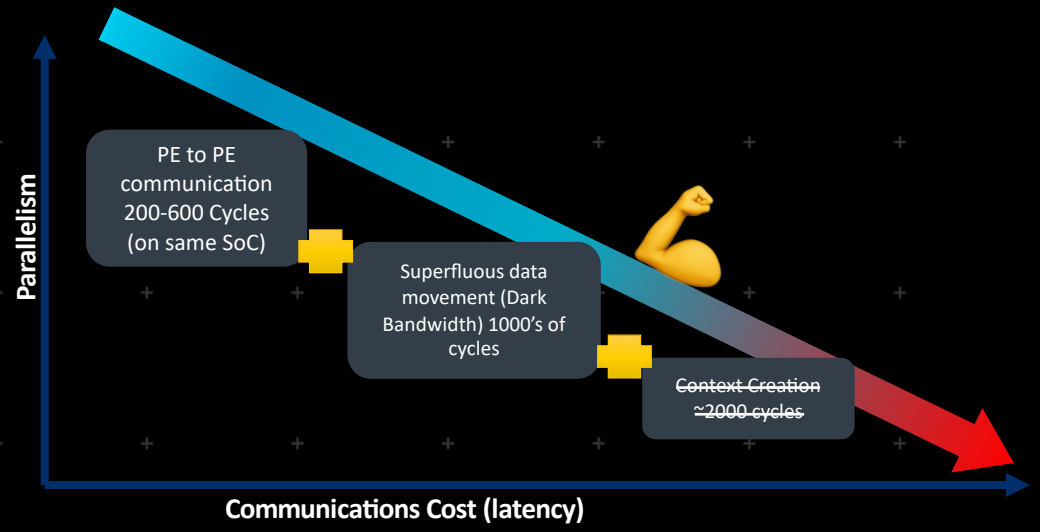

Optimize for latency and speed

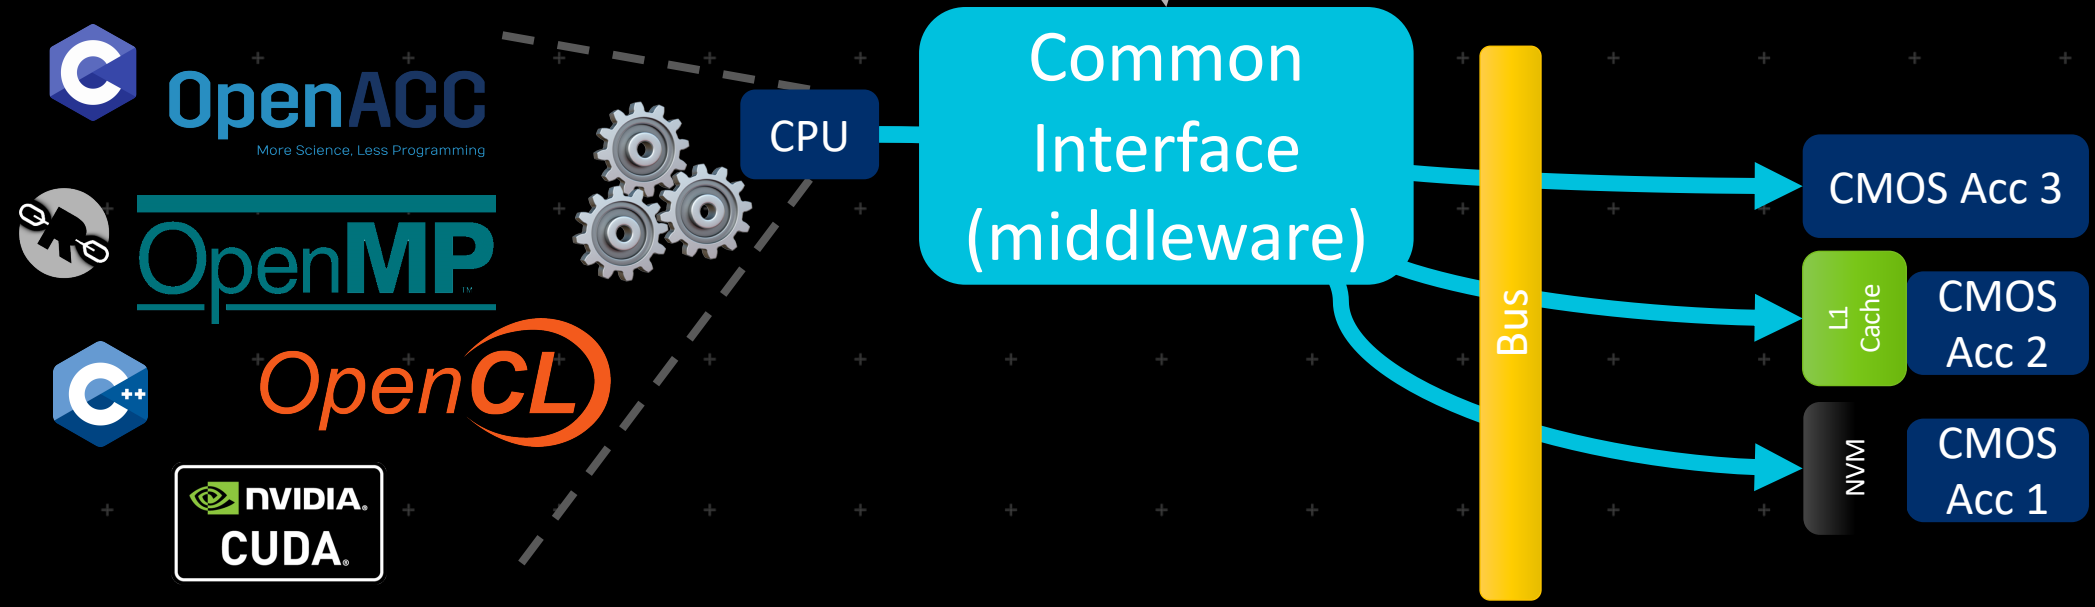

arm NEOVERSE

Enabling scalability and heterogeneity

- *Data in cache lines are often underutilized* (0-80%) before eviction, median ~42%.
- Synchronization takes extra cycles (~30ns between cores)
- Presented differently based on where you are: OoO Core, GPGPU, FPGA, etc.

Enabling scalability and heterogeneity

- *Data in cache lines are often underutilized* (0-80%) before eviction, median ~42%.
- Synchronization takes extra cycles (~30ns between cores)
- Presented differently based on where you are: OoO Core, GPGPU, FPGA, etc.

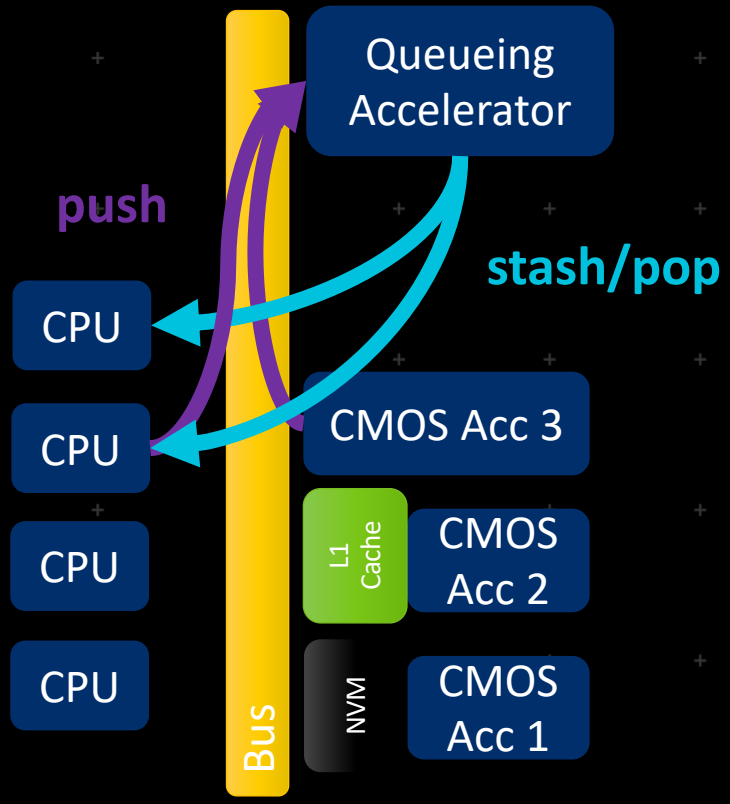

- *Dense transfers of packed cache lines*
- *Keep transfers inside coherence bus, instead of memory write-back*
- Can we virtualize underlying mechanism?

Enabling scalability and heterogeneity

- *Data in cache lines are often underutilized* (0-80%) before eviction, median ~42%.
- Synchronization takes extra cycles (~30ns between cores)
- Presented differently based on where you are: OoO Core, GPGPU, FPGA, etc.

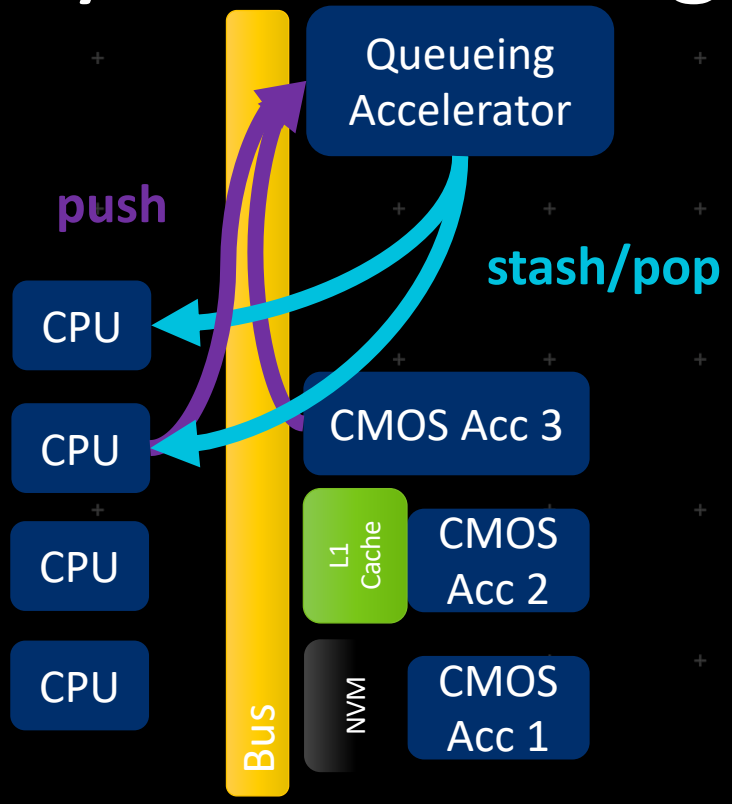

- *Dense transfers of packed cache lines*
- *Keep transfers inside coherence bus, instead of memory write-back*
- Can we virtualize underlying mechanism?

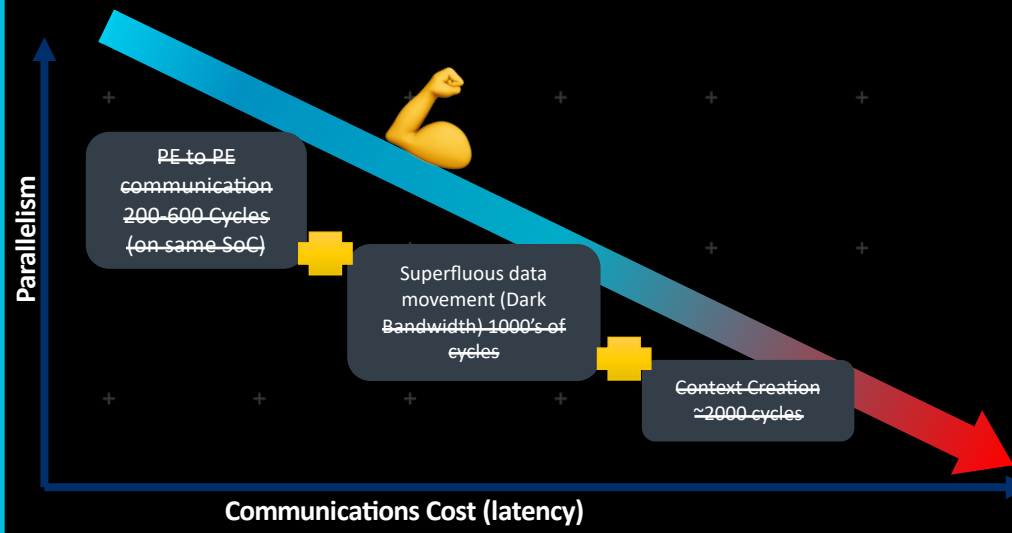

arm NEOVERSE

Accelerator Rich Systems

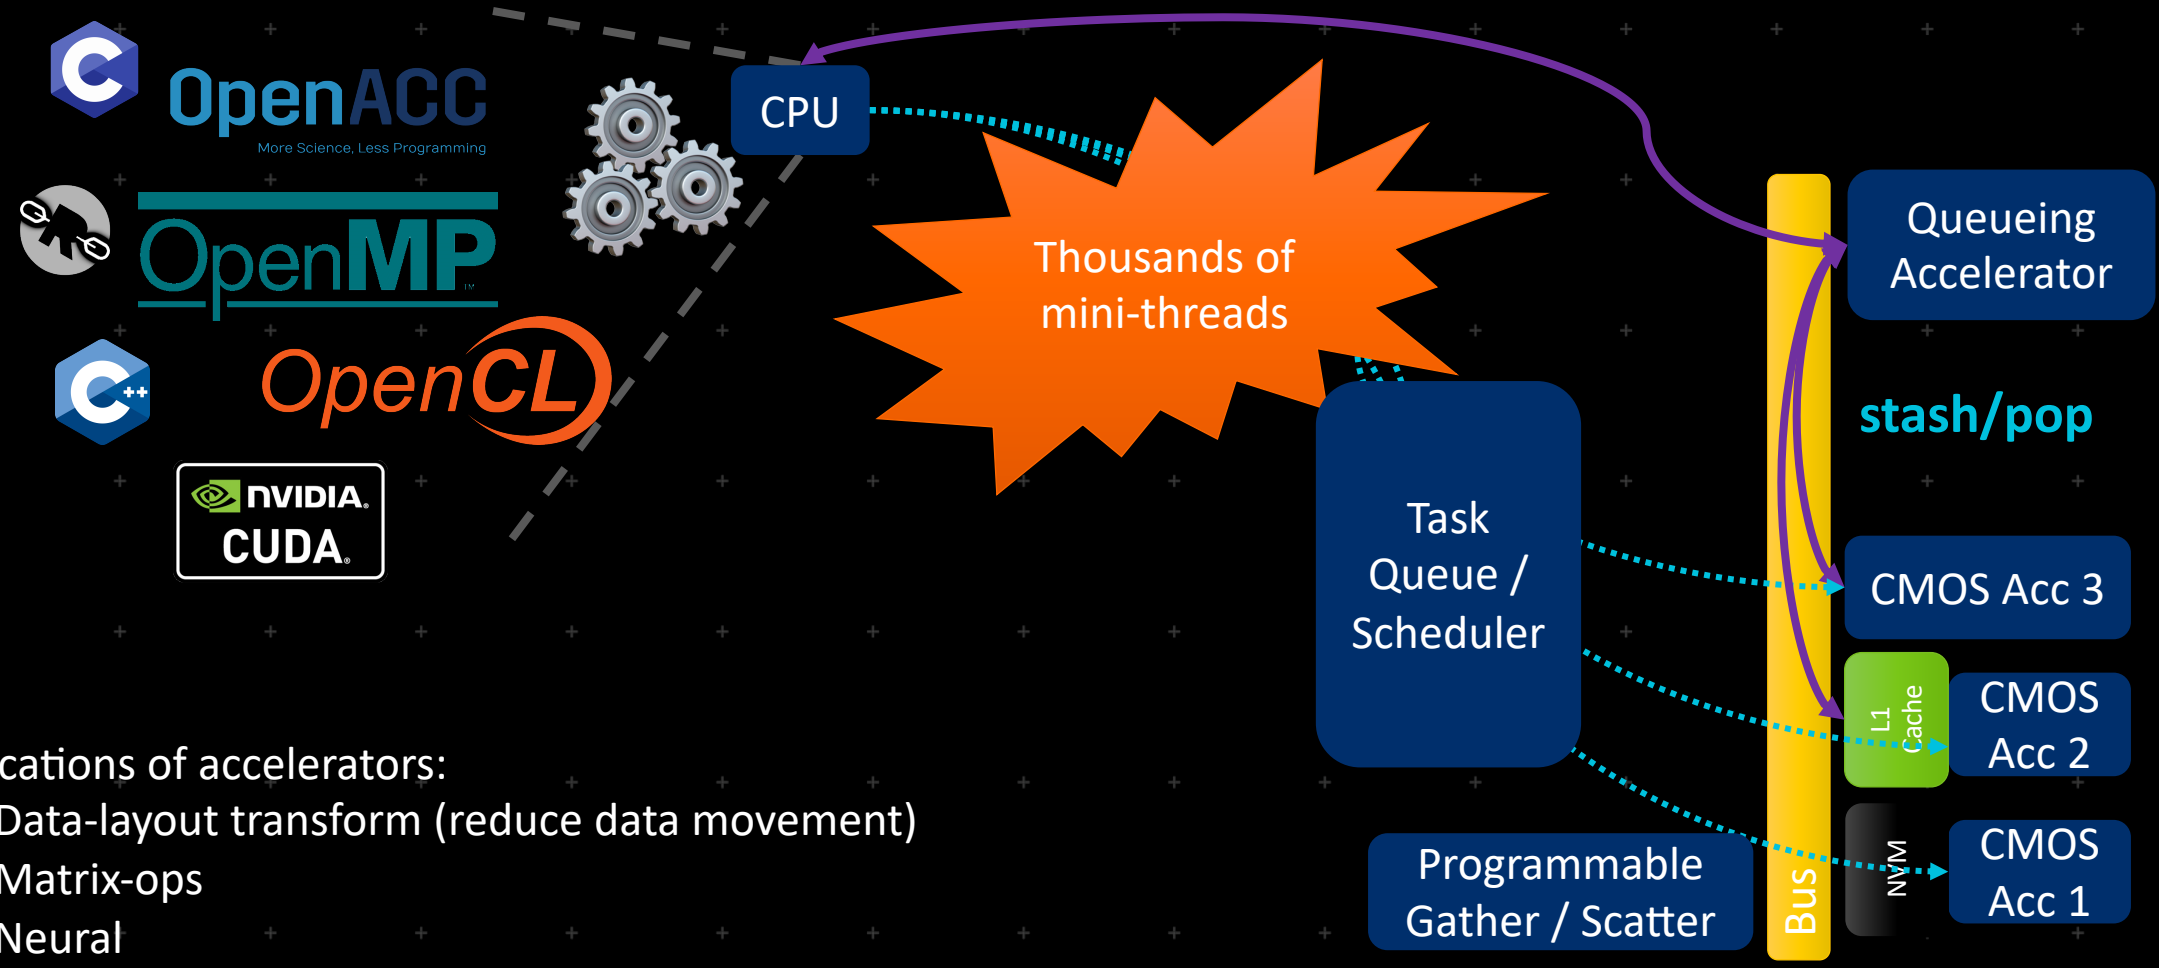

- Applications of accelerators:
  - Data-layout transform (reduce data movement)
  - Matrix-ops
  - Neural
  - Near-data
  - *Edge*
  - Etc.

DAG Execution

Programmer / Runtime / Compiler Provided Kernels

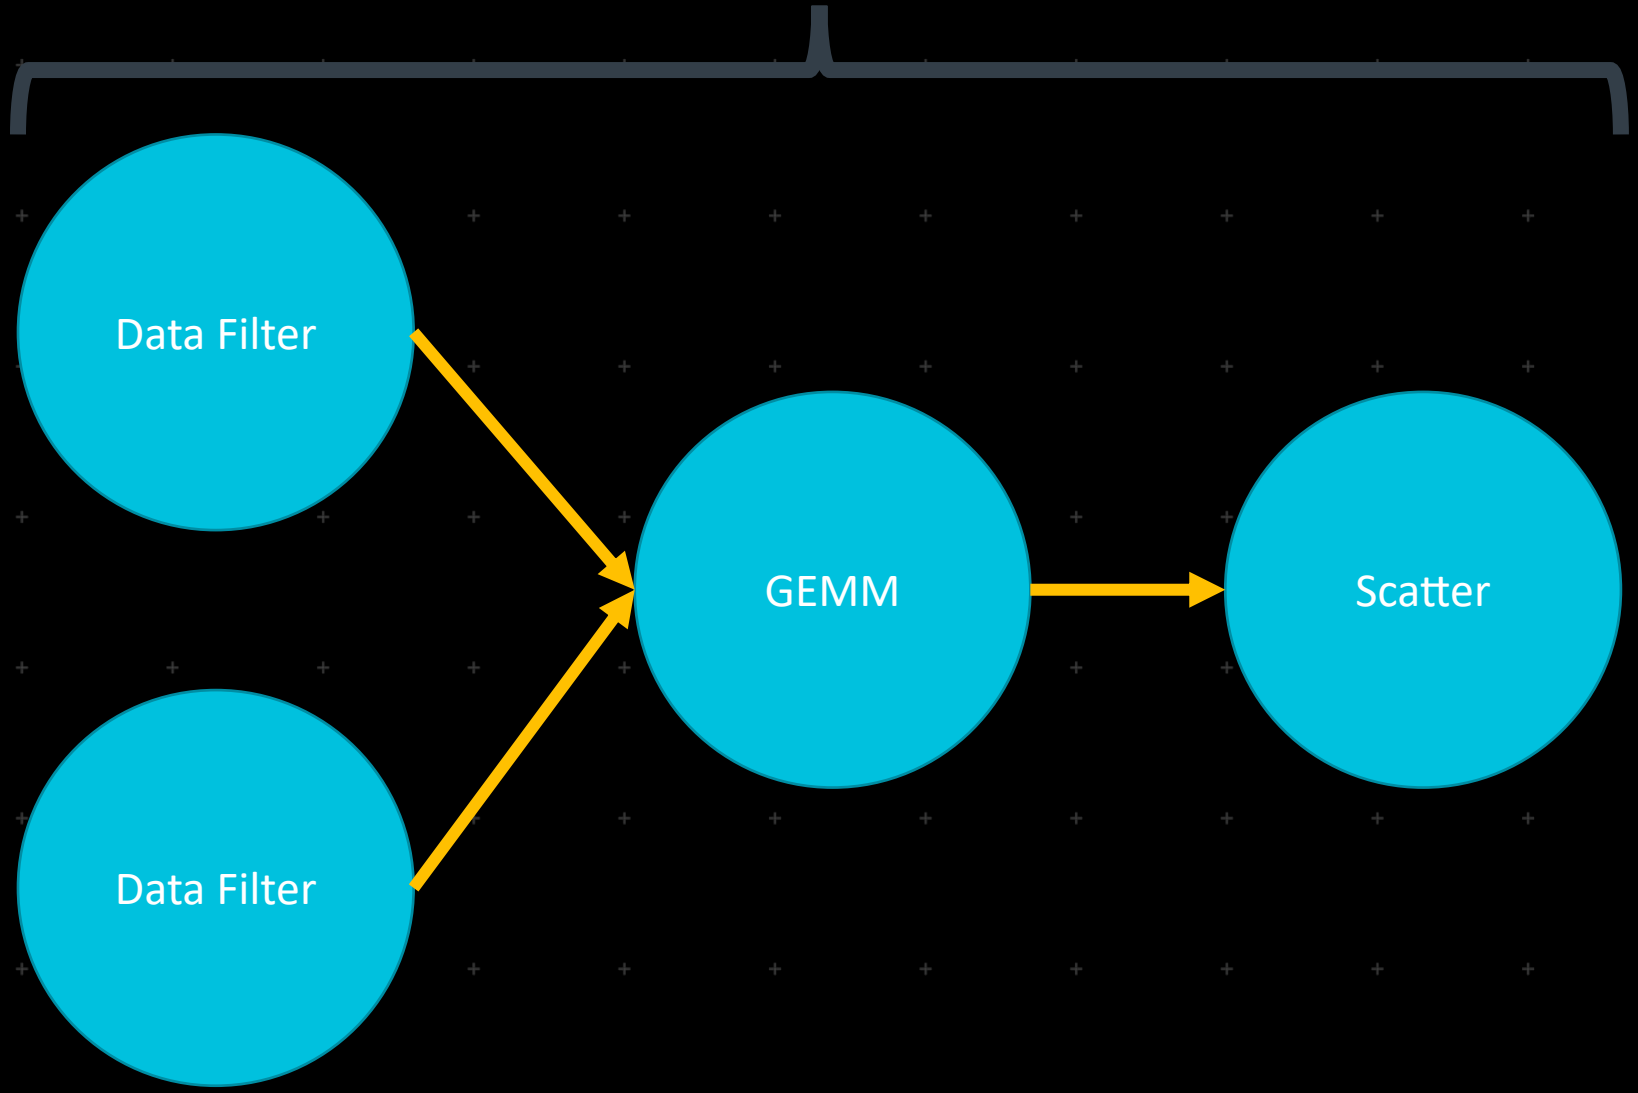

arm NEOVERSE

DAG Execution

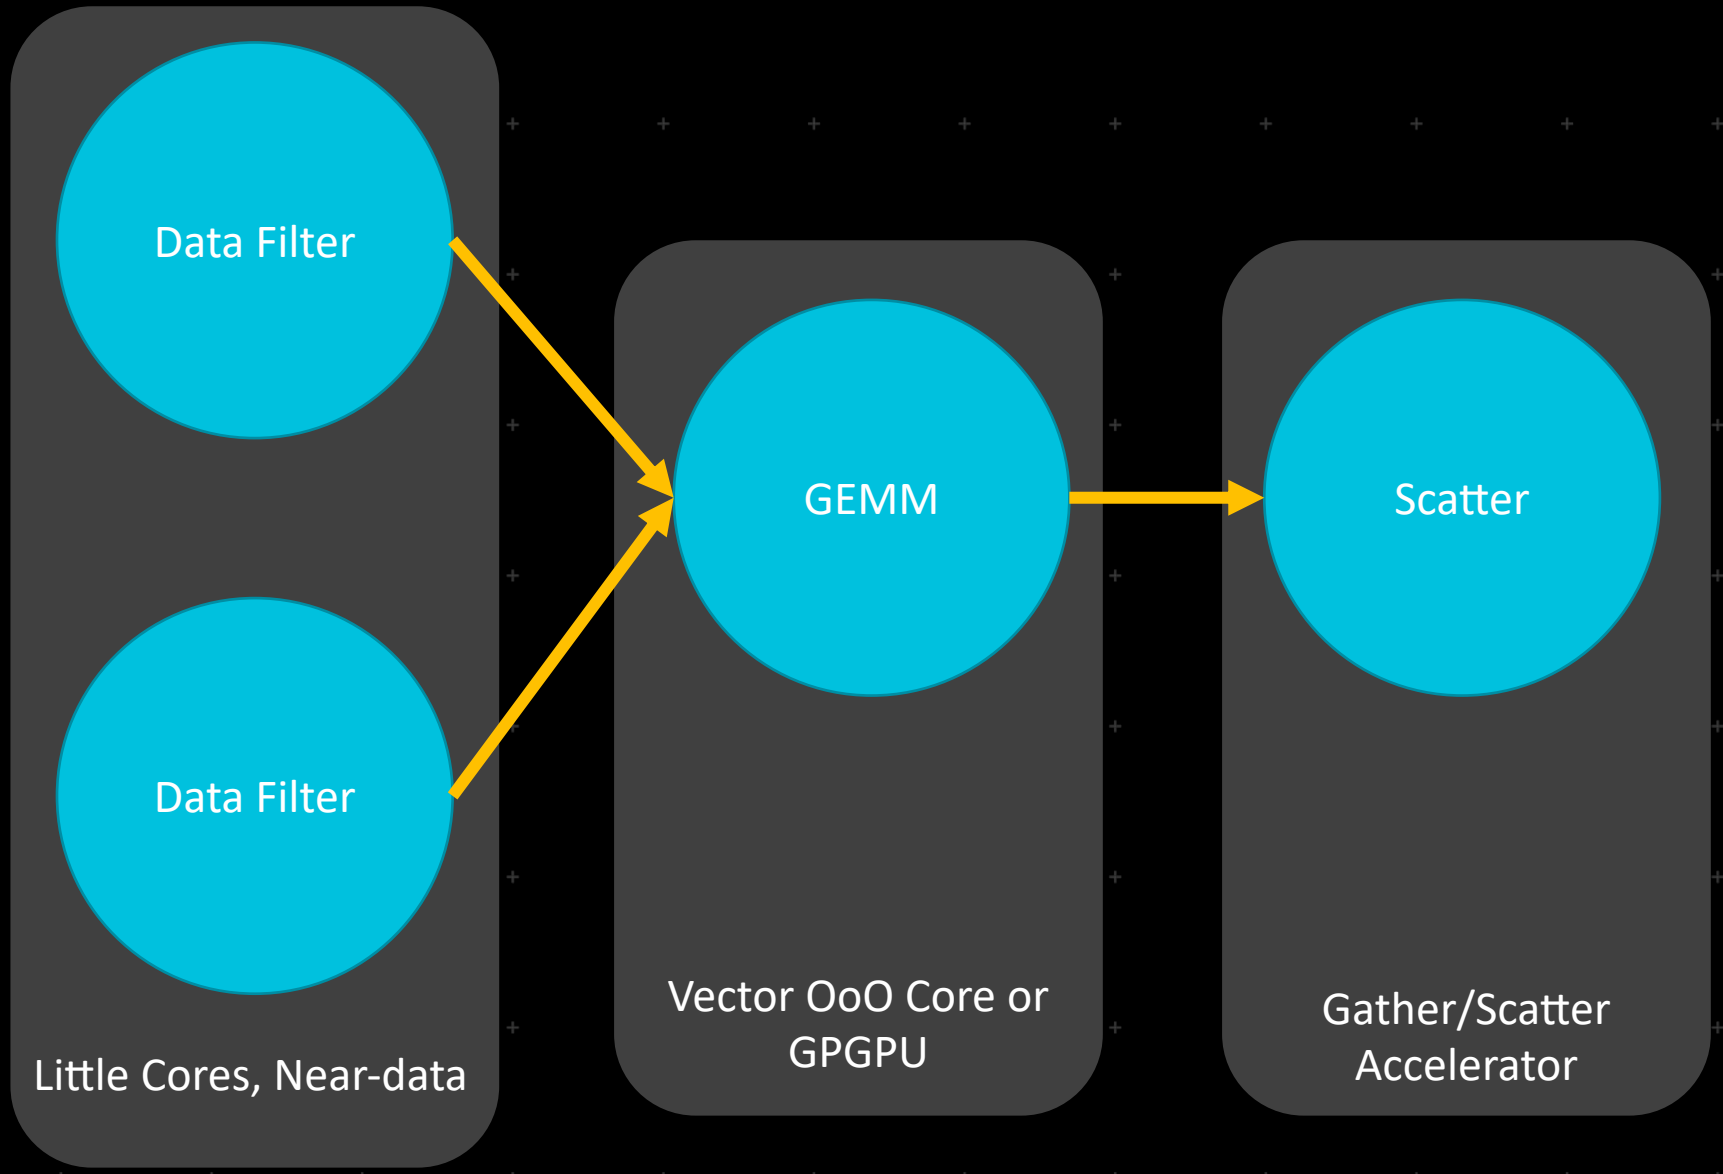

arm NEOVERSE

DAG Execution

- Programmer shouldn't have to care what medium these links are (e.g., MPI, simple FIFO, RDMA, etc.)
- Filter data from edge to data center (maybe just *center*??)

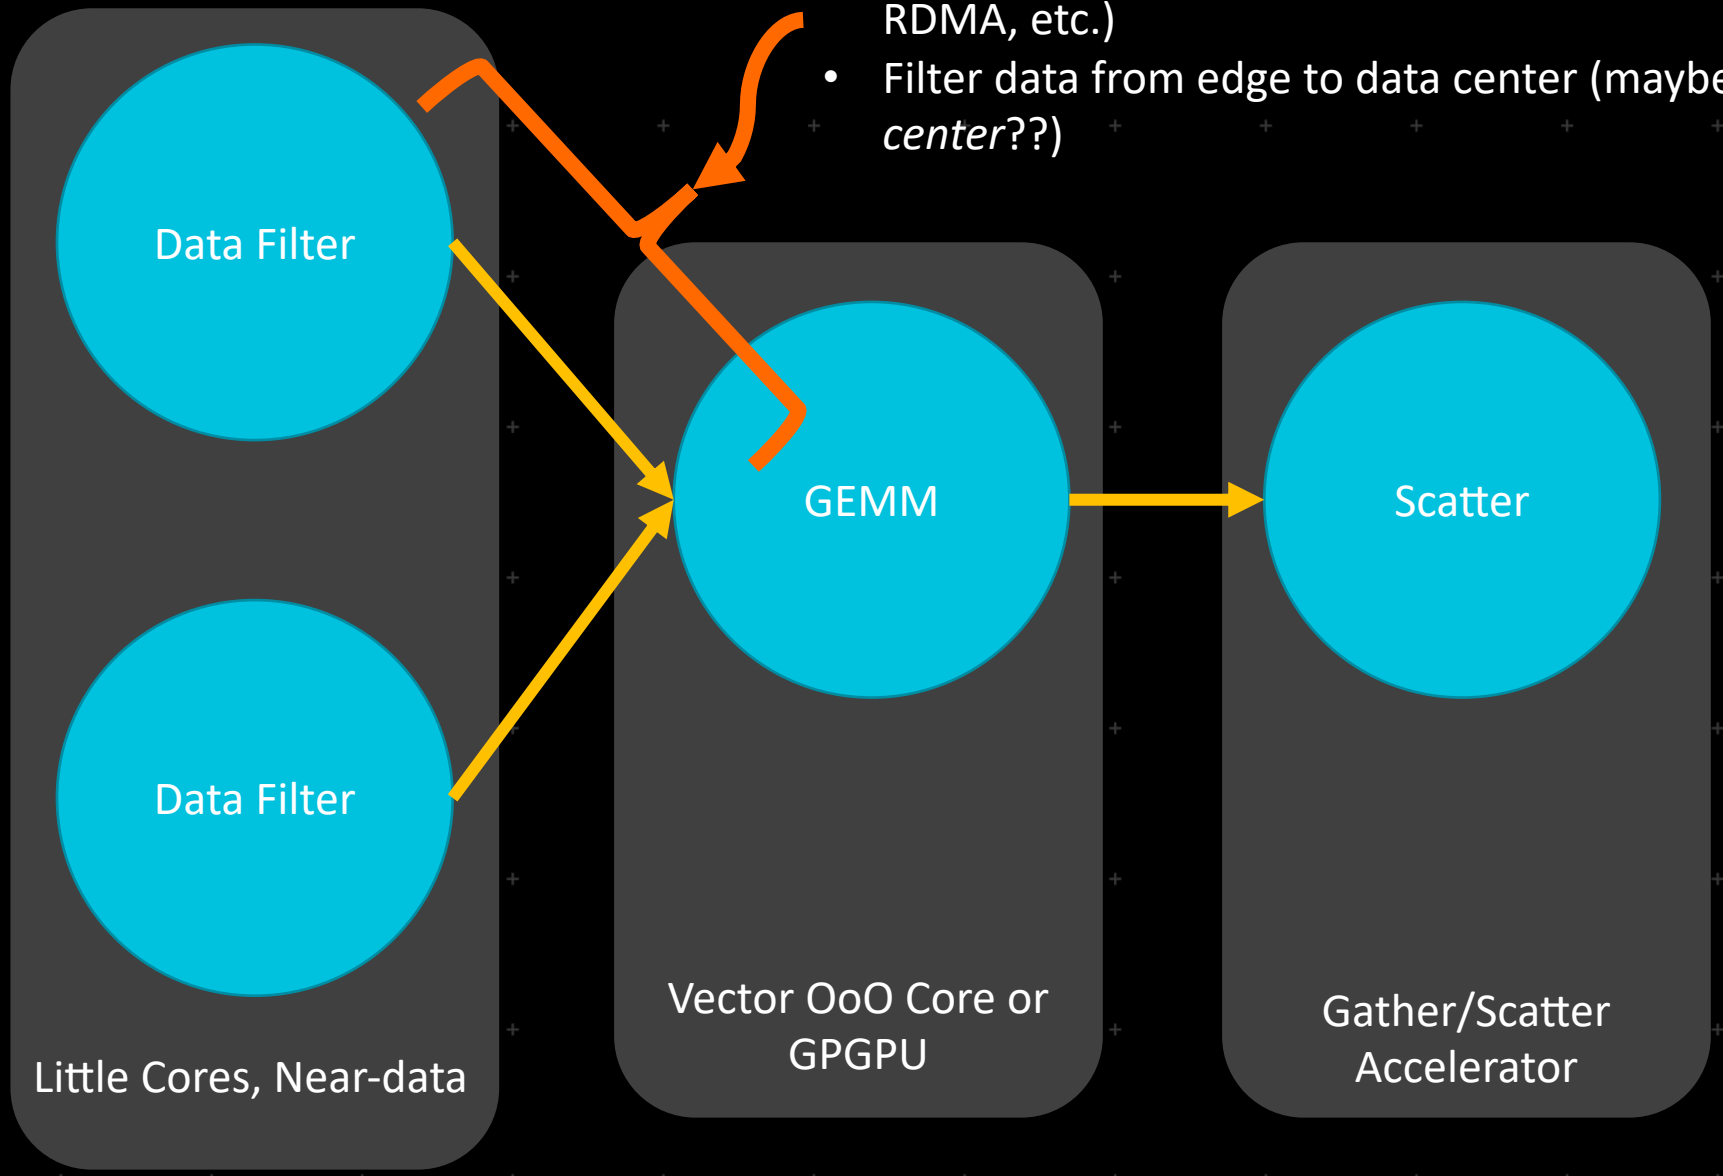

arm NEOVERSE

Productive Accelerator Rich Systems

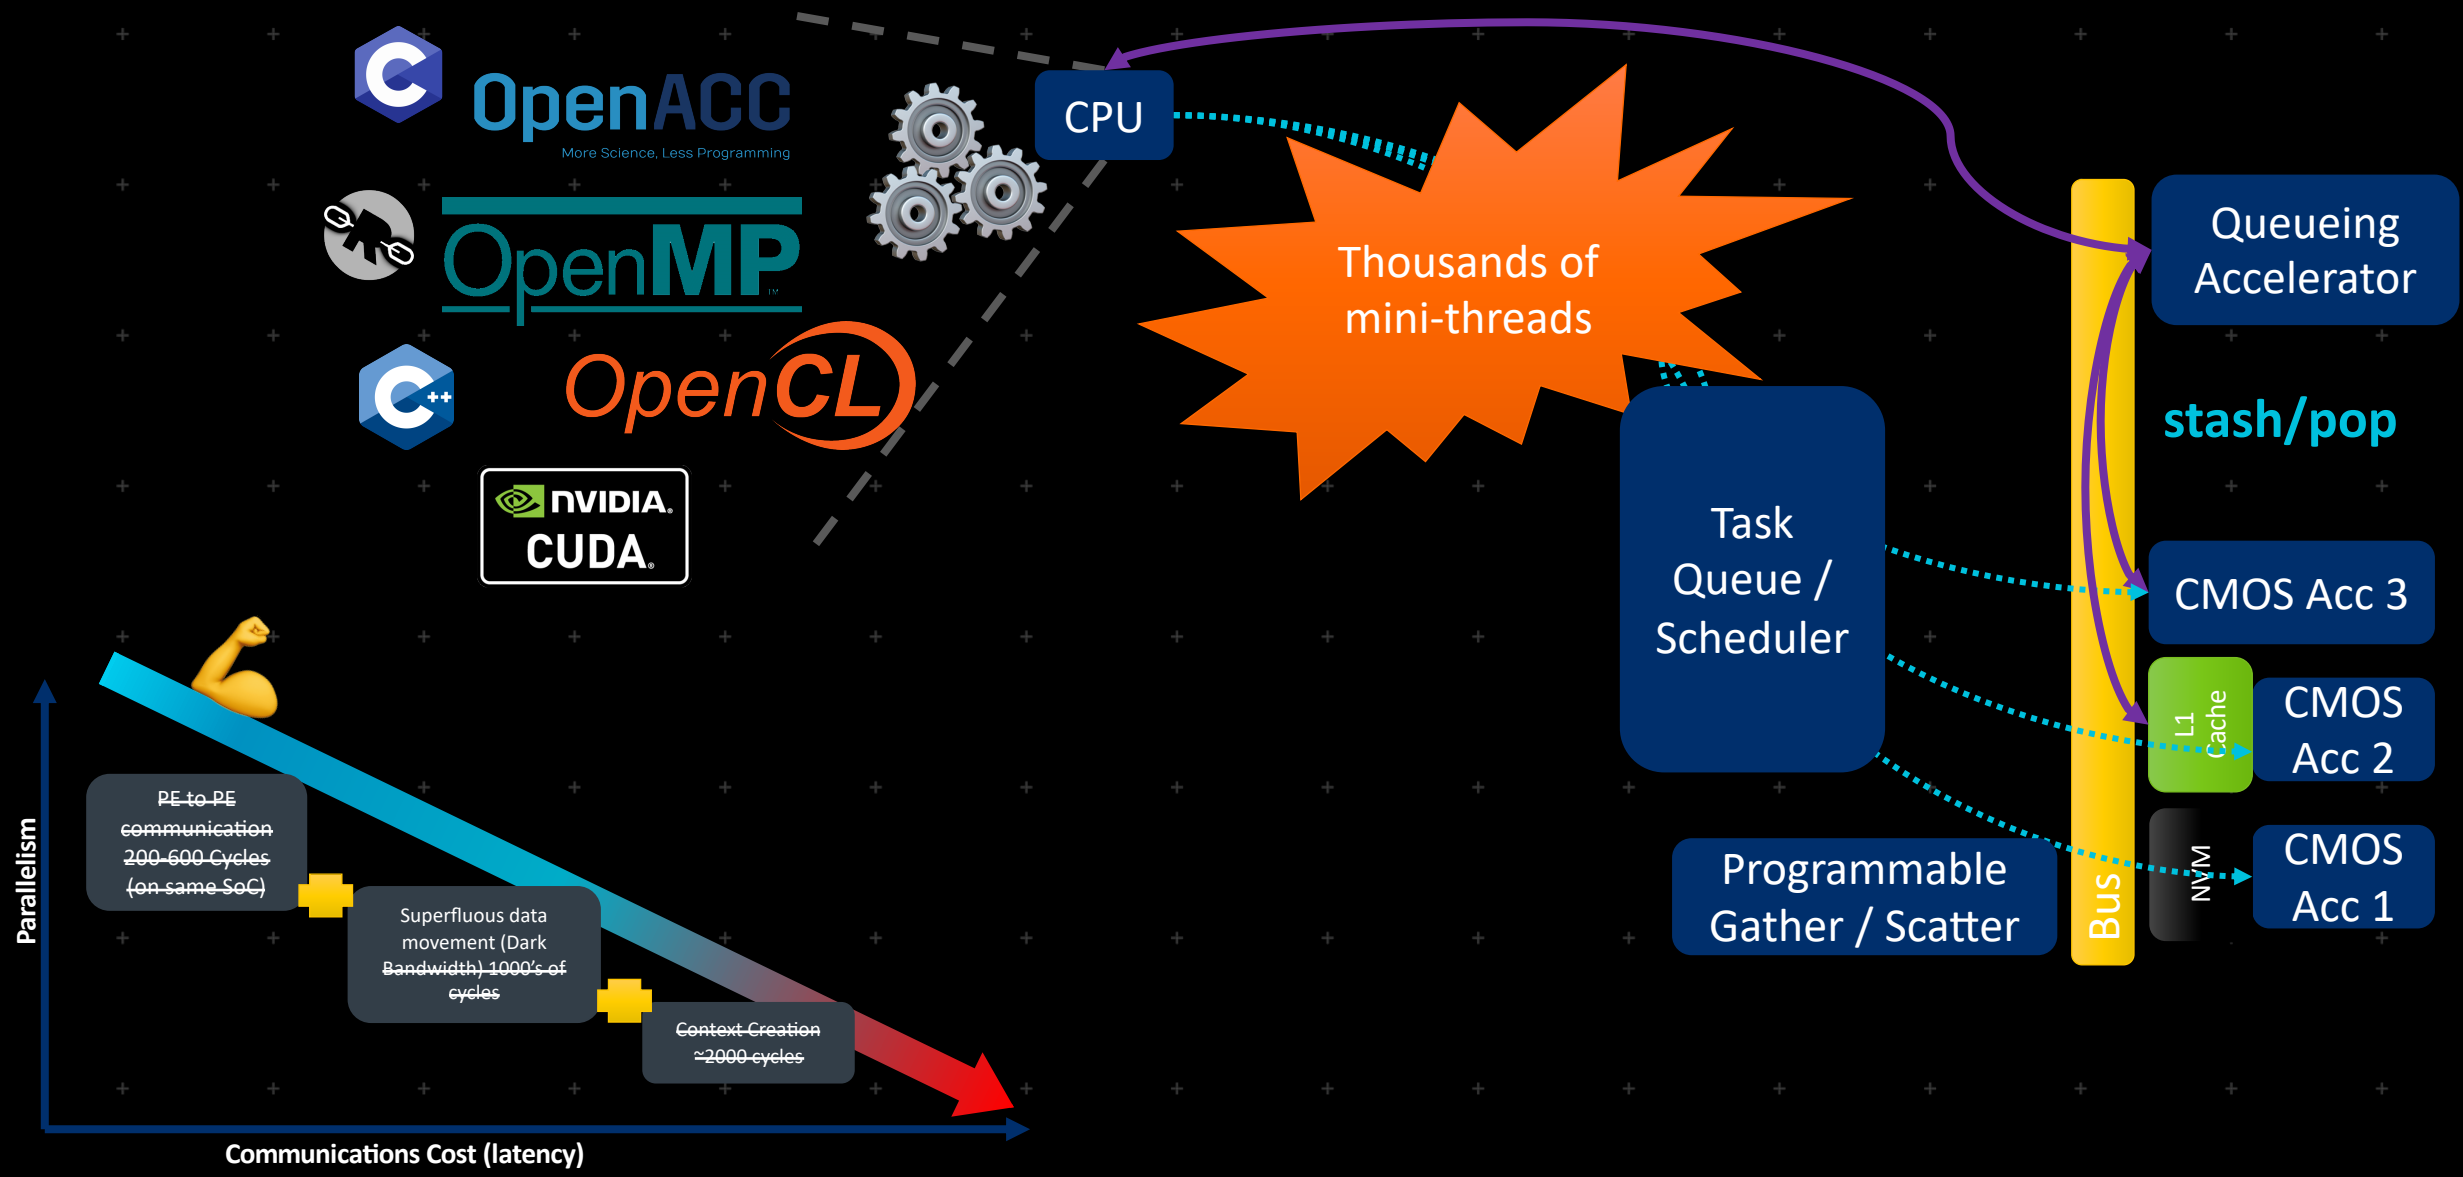

arm NEOVERSE

Productivity from Edge to HPC

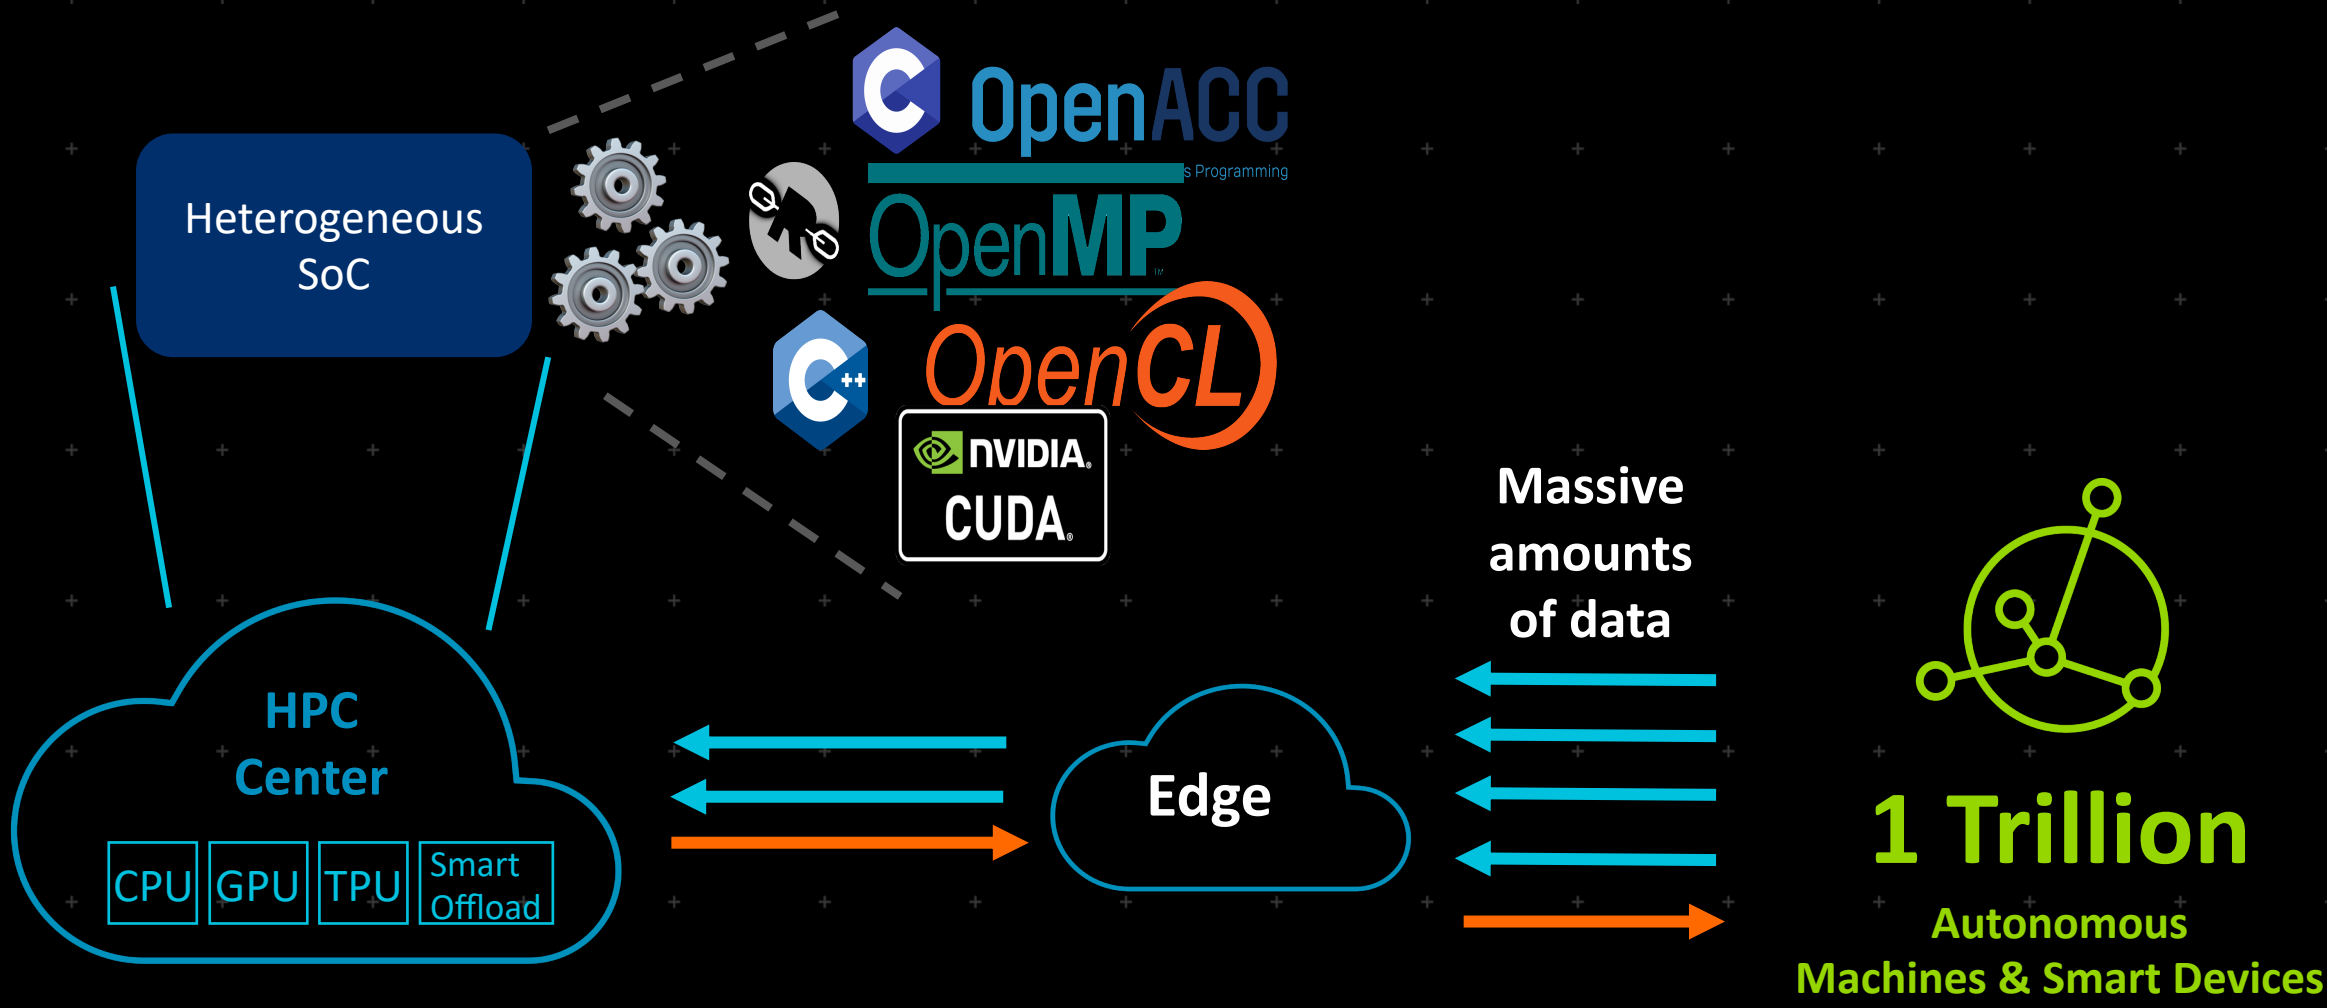

arm NEOVERSE

Research summary / parting thoughts

- Scientists should be able to focus on the science, programmers on the algorithm…often juxtaposed to reducing time to solution.
- Is *edge* really just an extended form of near-data processing?
- The *edge* complicates an already complicated world for application programmers.
- Same old problems, still no “sticky” solutions…yet.

Closing Comments

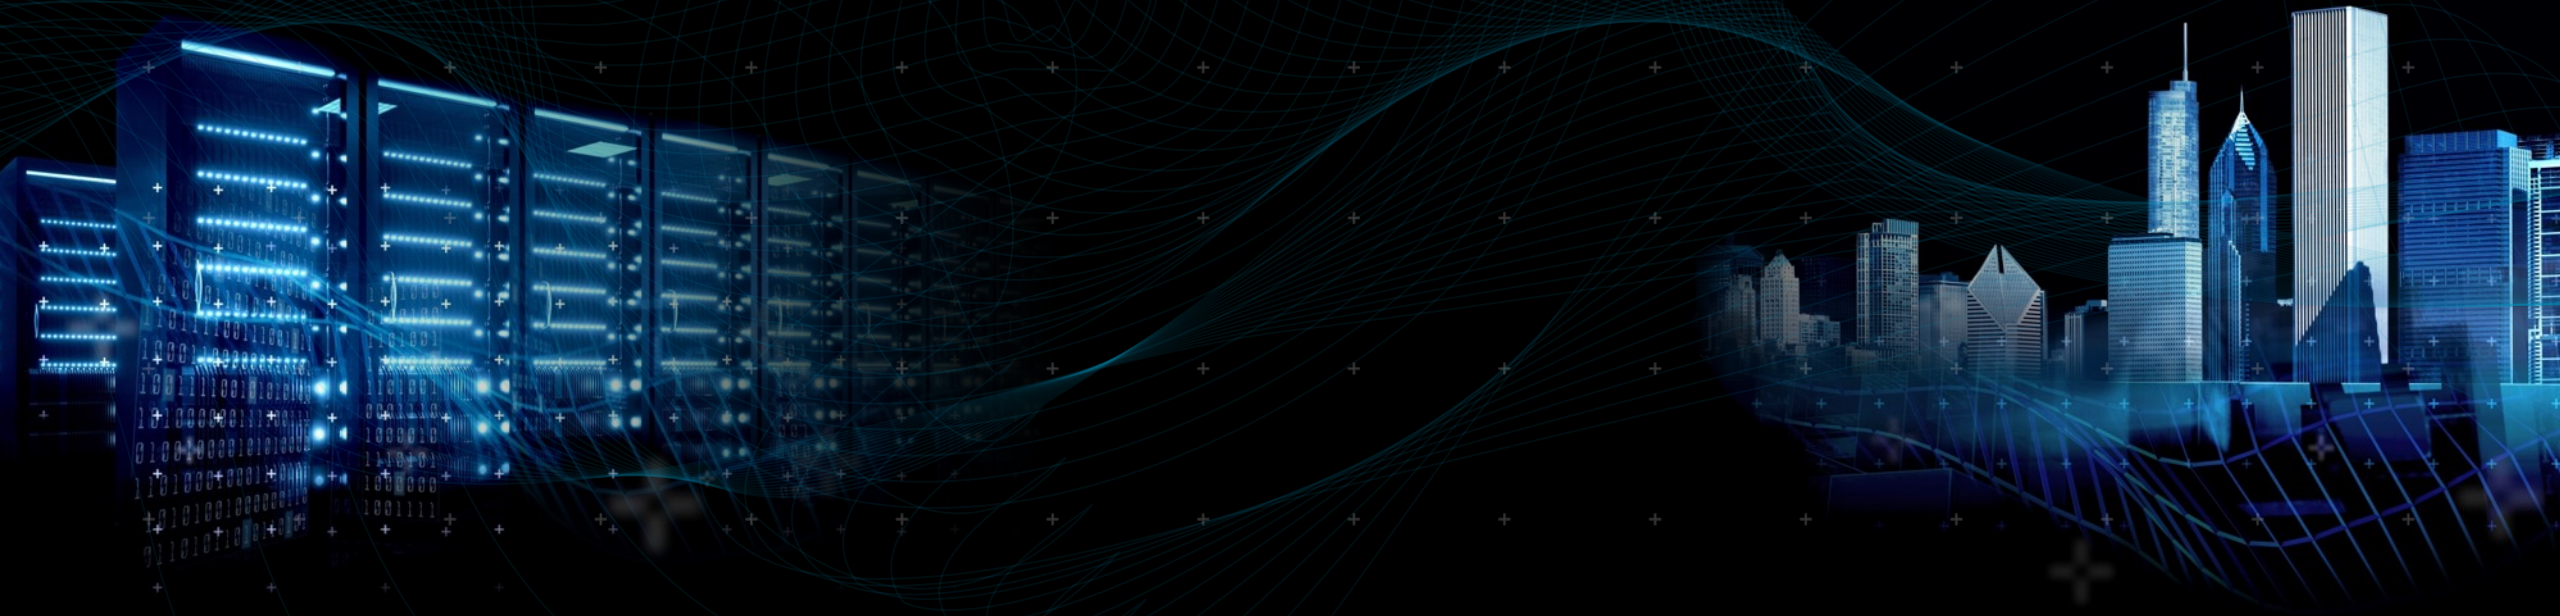

arm NEOVERSE

Arm HPC Community – Arm.com/hpc

Communication Portals

- Arm.com HPC resources
- developer.arm.com/HPC (HPC Ecosystem Landing page)
- community.arm.com/tools/HPC (HPC Blogs, Forum)

Arm HPC User Group Community

- Gitlab HPC Packages Wiki (software ecosystem)
- Arm-HPC @ Groups.IO (<=NEW)

Supporting Arm HPC Community end-users and developers.

Who you gonna call?

Arm Professional Services!

- Application performance engineering!
  - Compiler not vectorizing?
  - Performance not what it should be?
- New system tuning!
  - Do you have the right SMT mode on your ThunderX2?
  - Does your InfiniBand need tweaking?
- Hackathons and tutorials!
  - Does a team need a mentor at your hackathon?
  - Looking for a jumpstart with Arm HPC?

arm Research Summit

Austin, Texas, September 15-18, 2019

Last Call for Registration:

https://www.arm.com/company/events/research-summit

*Discount room block closes 30 August 2019


arm NEOVERSE
The Cloud to Edge Infrastructure Foundation
for a World of 1T Intelligent Devices
Thank You!