

Getting data where it's needed

potential for CPU, GPU, accelerators, and smart networks

Sam Antao, Nick Malaya, Matthew McIntire, Chuck Gilbert

International Workshop on Smart Networks, Data Processing and
Infrastructure Units (DPU'23) – ISC'23
25/05/2023

AMD 
together we advance_

Agenda

-
1. Motivation and Data Processing Unit (DPU) definition
 2. Use cases
 3. DPU requirements

Motivation

- Network scale-out challenges – Frontier/LUMI
- Very tight integration/coupling
- Good for weak scaling, but worse for strong scaling
- Latency and bandwidth both major players
- Irregular network requirements
 - Inherent imbalance
 - Each rank with different problem sizes or workload
 - Workflows with different coupled apps



Data Processing Unit (DPU) definition

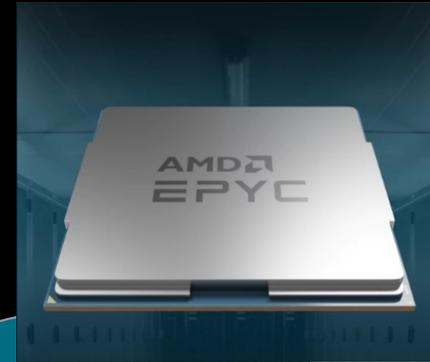
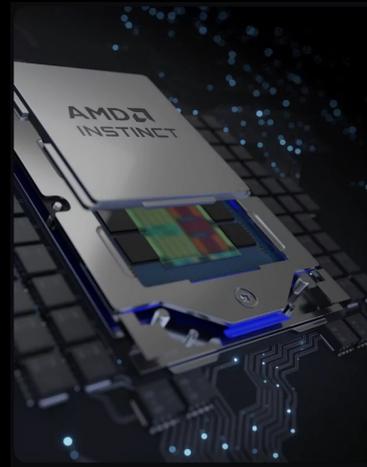
- DPU as a new type of accelerator or integration/coupling of different accelerators?
- Definition
 - Controlling data streaming and collective processing
 - Managing network traffic, protocols and storage
 - Abstraction of a DPU as a collection of devices
- Requirements clash and overlap: cloud vs. traditional HPC
 - Opportunities: what needs prioritizing for these different kinds of deployments
 - Cloud prioritizes data confidentiality and tenant isolation
 - HPC prioritizes low latencies
 - Both prioritize core count – more cores maps to more revenue (cloud) and performance (HPC)
 - Convergence of technologies: create technology that is transferable across HPC and cloud
 - Widespread modelling and simulation offerings
 - Deep learning frameworks
 - Variety of workloads in HPC which might be different from what is in the cloud
 - ML workloads showcasing different requirements than HPC apps
 - Requirement shift depending on the scale

Good use cases for DPUs

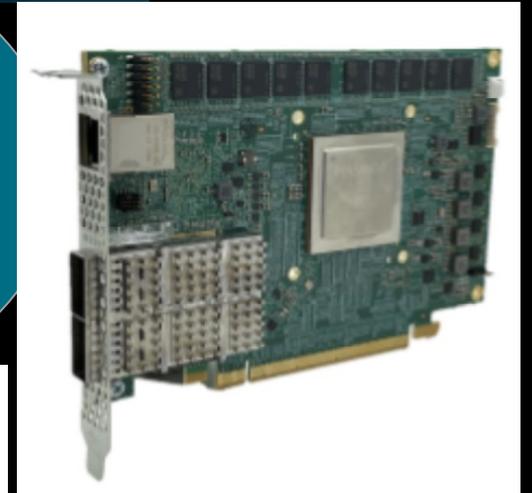
- Sparse Machine Learning communication challenges over the network
- Serialization/deserialization of data
- RDMA and shared memory abstractions for a collection of nodes
- Enabler for composability of resources
- High-performance distributed storage
- Offloading programming models (e.g., collectives or other distributed operations)
- Data-lake enabler
- Be smarter behind existing APIs
 - Evolve standards to accommodate user needs and HW limitations
 - Can we keep/simplify the same API and be smarter in their implementation?
 - Examples: MPI, RCCL/NCCL

Co-design across different classes of hardware

- Convergence of multiple flavors of technology
 - CPUs
 - GPUs
 - FPGAs
 - SW defined networks
 - Offloading model applied everywhere
 - From components to systems
 - From systems to a user abstraction
- Co-design
 - with system integrator
 - Within the SoC
 - Across a network
 - Toolchains – software libraries and tooling



Co-design



What's happening around the GPUs

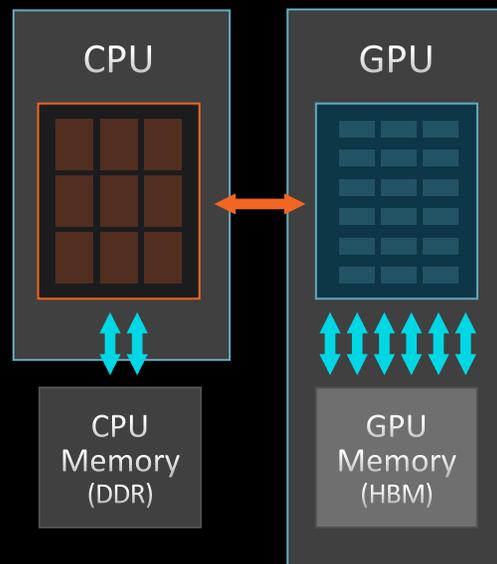
- CPU+GPU unification – MI300 APU
 - Keep data in the same place – move the compute
 - Abstraction lifts the need for user-defined copies
 - Can we expand the abstraction beyond the APU?

AMD CDNA™ 2 Coherent Memory Architecture

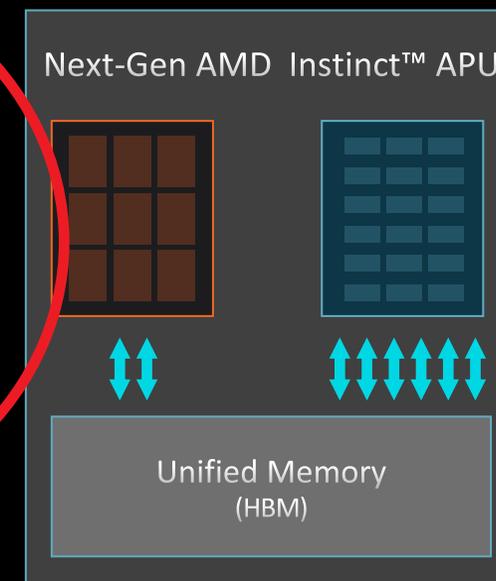


AMD CDNA™ 3 Unified Memory APU Architecture

- Simplifies Programming
- Low Overhead 3rd Gen Infinity Interconnect
- Modular System Design



- Eliminates Redundant Memory Copies
- High-Efficiency 4th Gen AMD Infinity Architecture
- Low TCO with Unified Memory APU Package



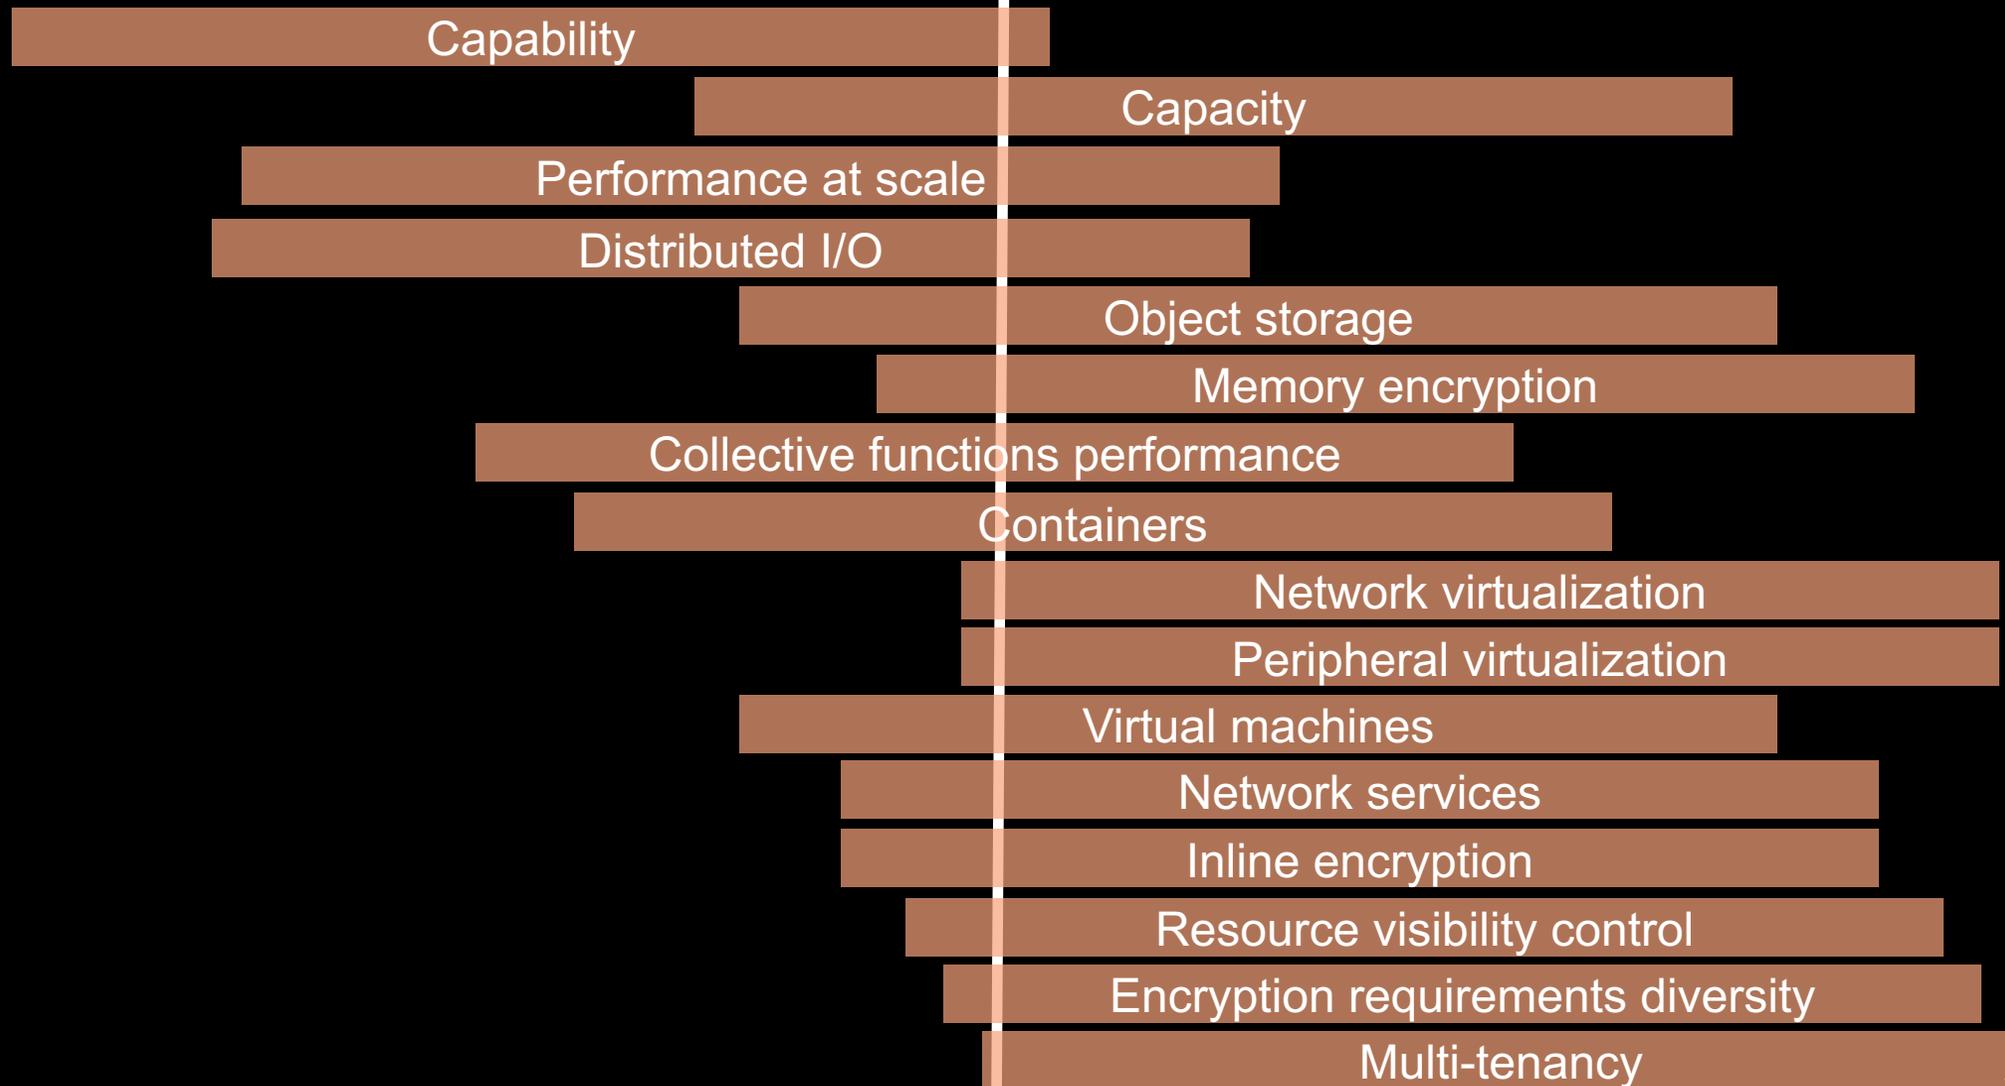
A global shared memory abstraction

- Reflective memory
- Stop thinking about nodes but a collection of resources
- Composability of resources
- The enablers
 - RDMA around compute units
 - Flexible/adaptable (SW defined) network
 - Smart NICs – configurable data processors

HPC/Cloud requirements/priorities

On-premise HPC

Cloud

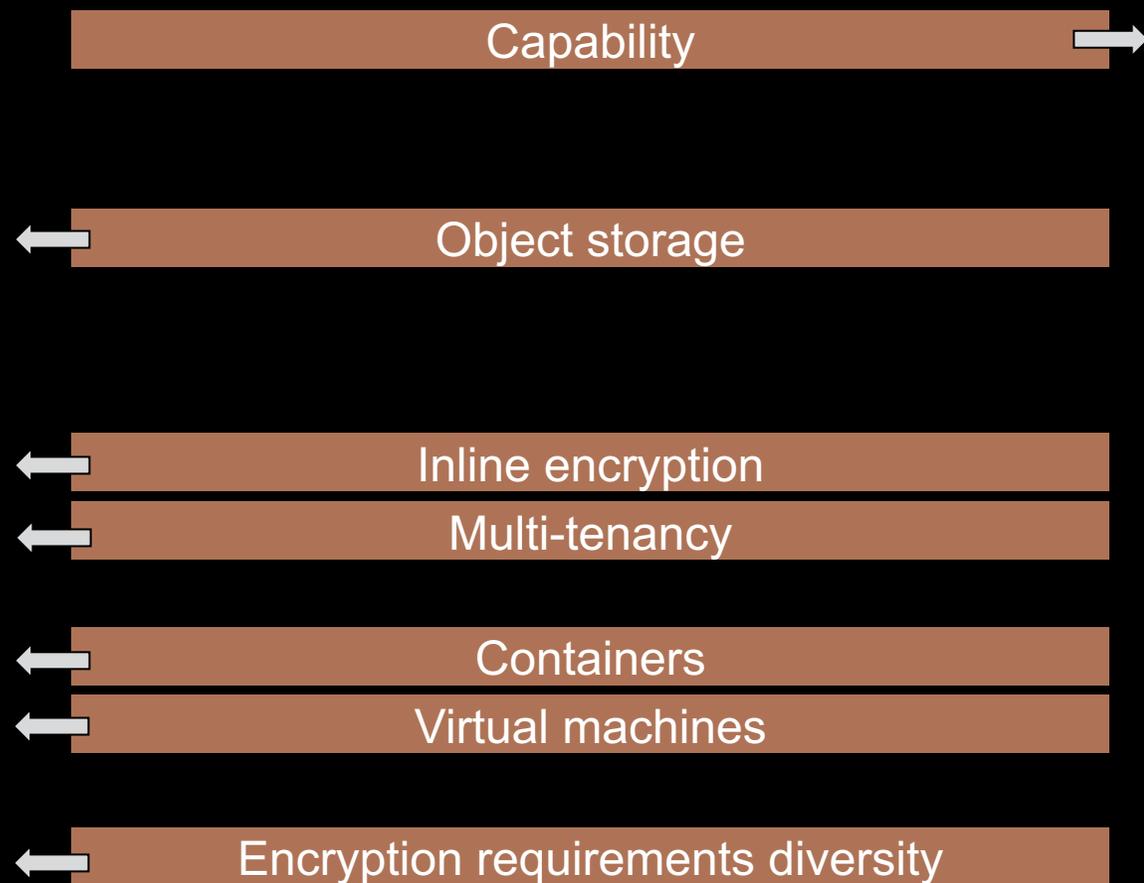


How are these bars moving?

- More capability runs attempted on the cloud
 - Distributed deep learning
 - Modeling and simulation offerings
- Object storage implementations in-house
 - Data organization beyond the file abstraction
 - Extra metadata annotations
 - Data streaming from object store services
- Sensitive data used everywhere
 - General Data Protection Regulation – GDPR
 - Data quality as competitive advantage
 - Encryption desired in many workloads
- Containers and other sort visualization first class citizens
 - Most HPC sites support container engines
 - Dodge native filesystem – in-memory overlay file system
- Flexible deployments
 - No static/immutable functions
 - Accommodate new standards and workloads

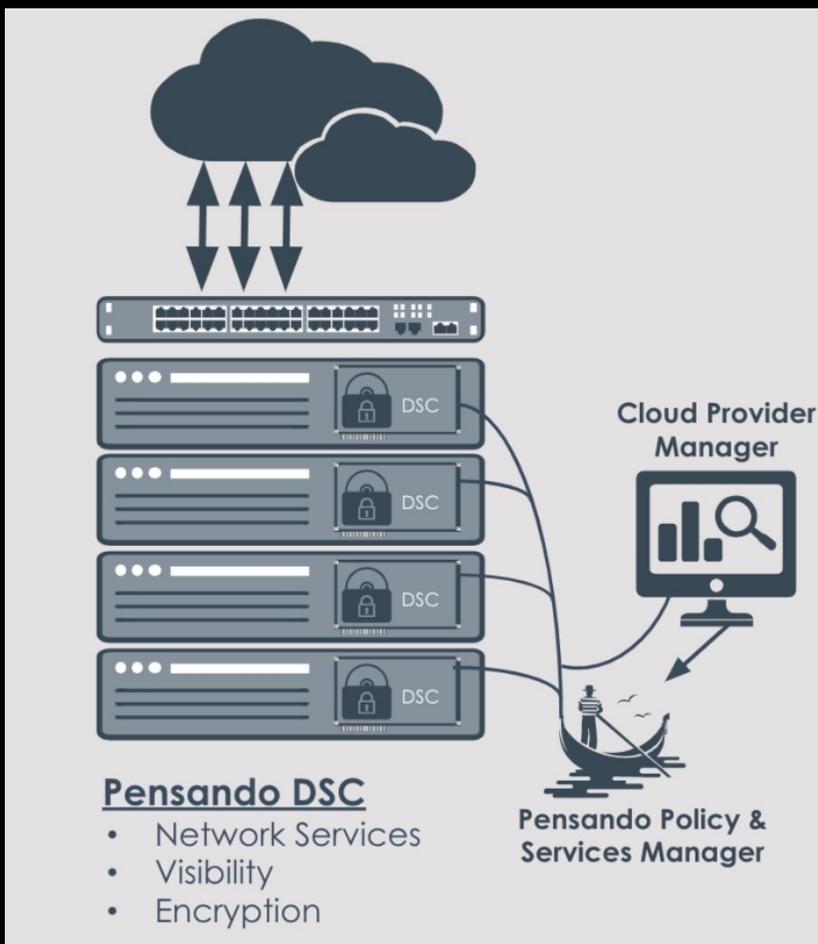
On-premise HPC

Cloud



DPU and Smart NICs – complementing capabilities

Distributed Services Card (DSC)



Highly manageable and fully programmable

AMD SmartNIC Accelerator

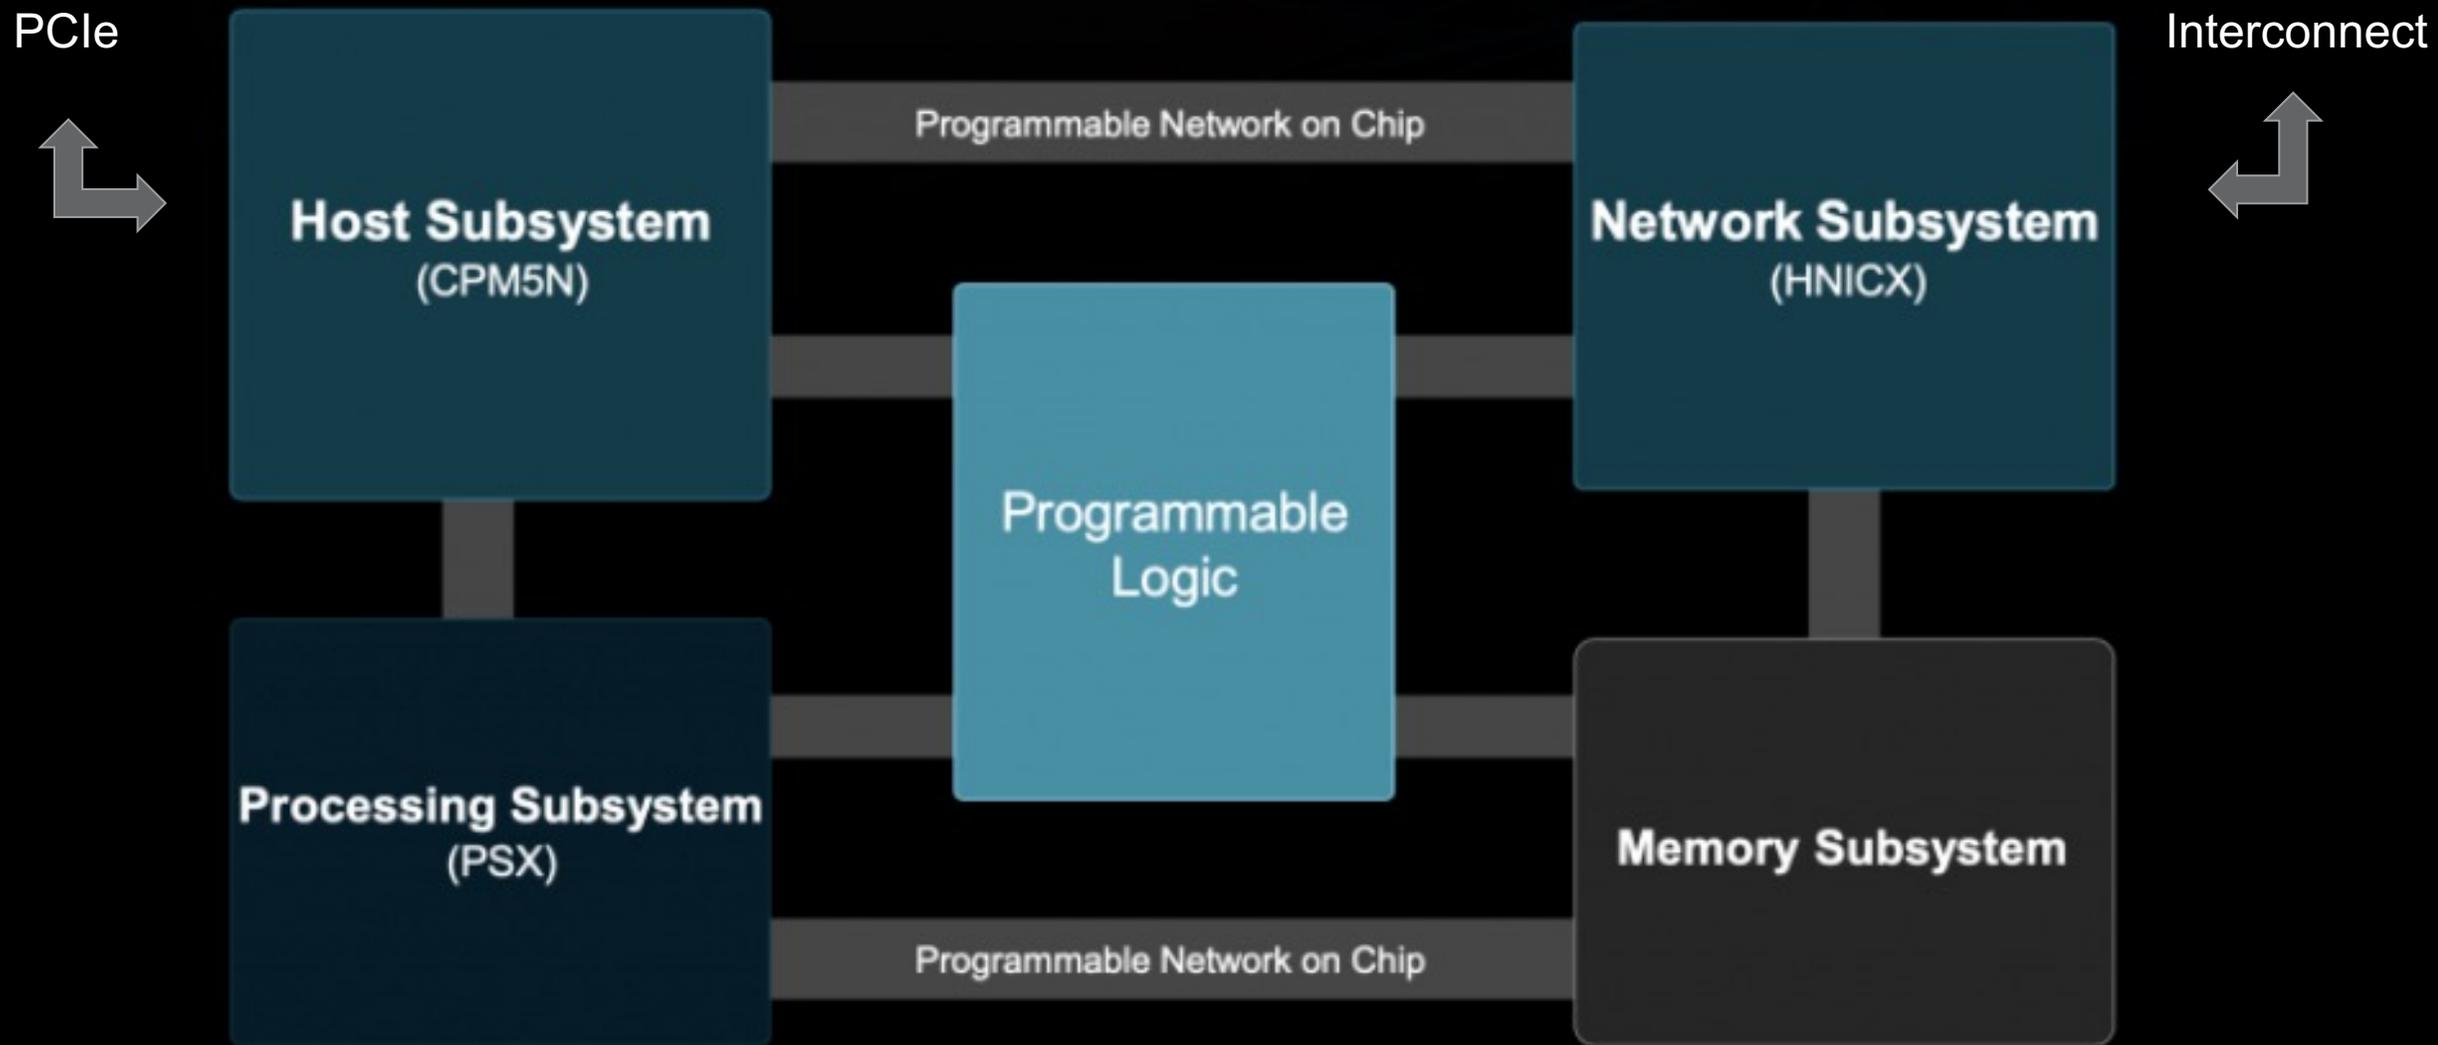


Complete function set tightly coupled with programmable logic (up to 400G)

Pensando DSC

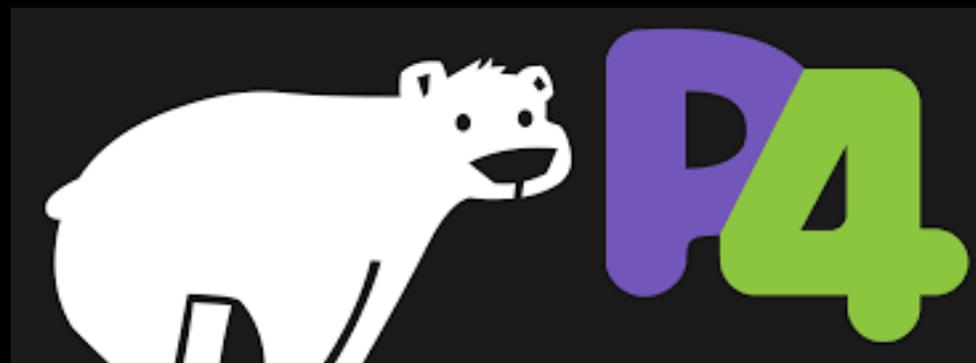


FPGA Smart NICs structure



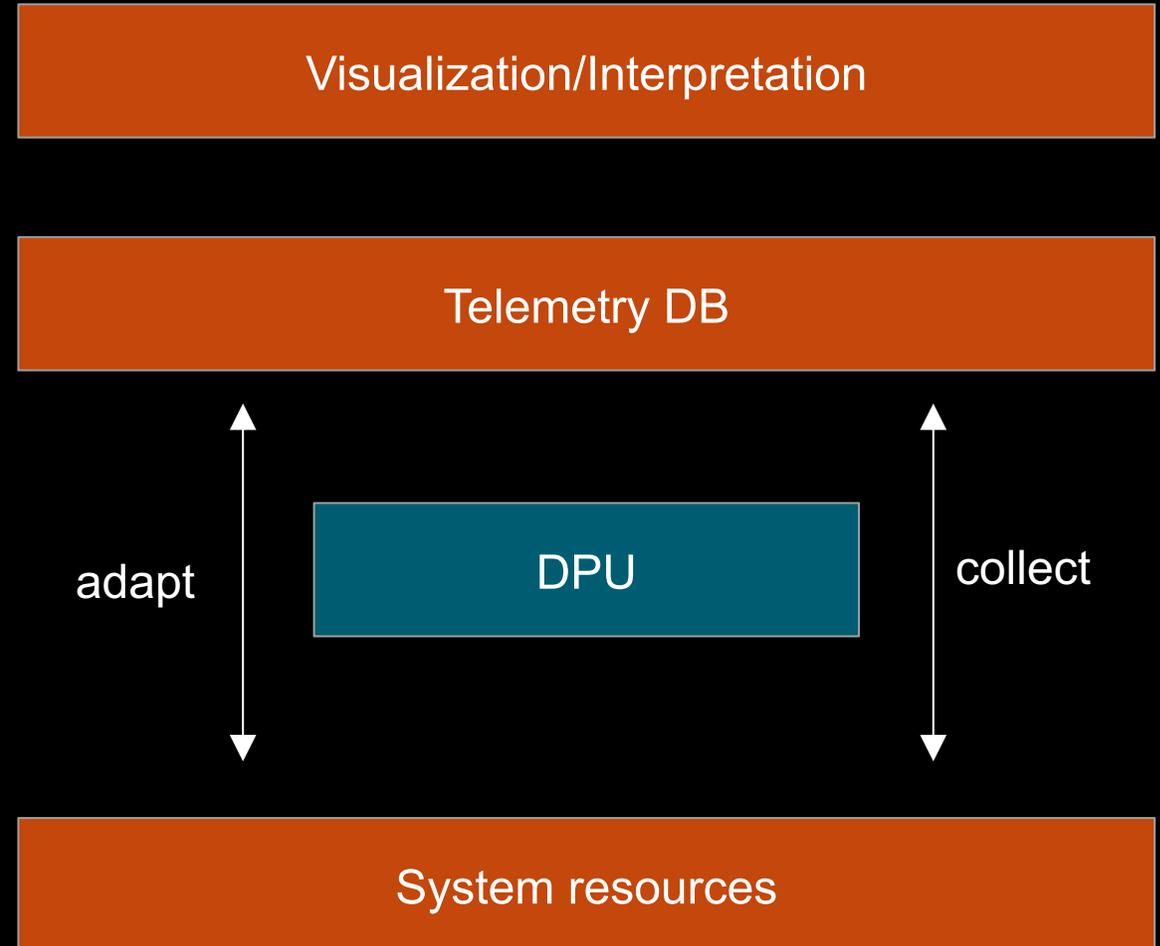
Programmability across all devices

- Bespoke solutions
 - Accommodate diversity of use cases
- Ease of use
 - User-friendly abstractions
 - Enable developers and end users
 - Facilitate training and knowledge transfer
- Evolve solutions over time
 - New security standards
 - New communications patterns
 - Address lessons learned
- AMD DPUs support
 - P4 packet domain specific language
 - C programming languages
- FPGA
 - Feature rich high-level-synthesis tooling
 - Many ready-to-deploy configurable IP instances



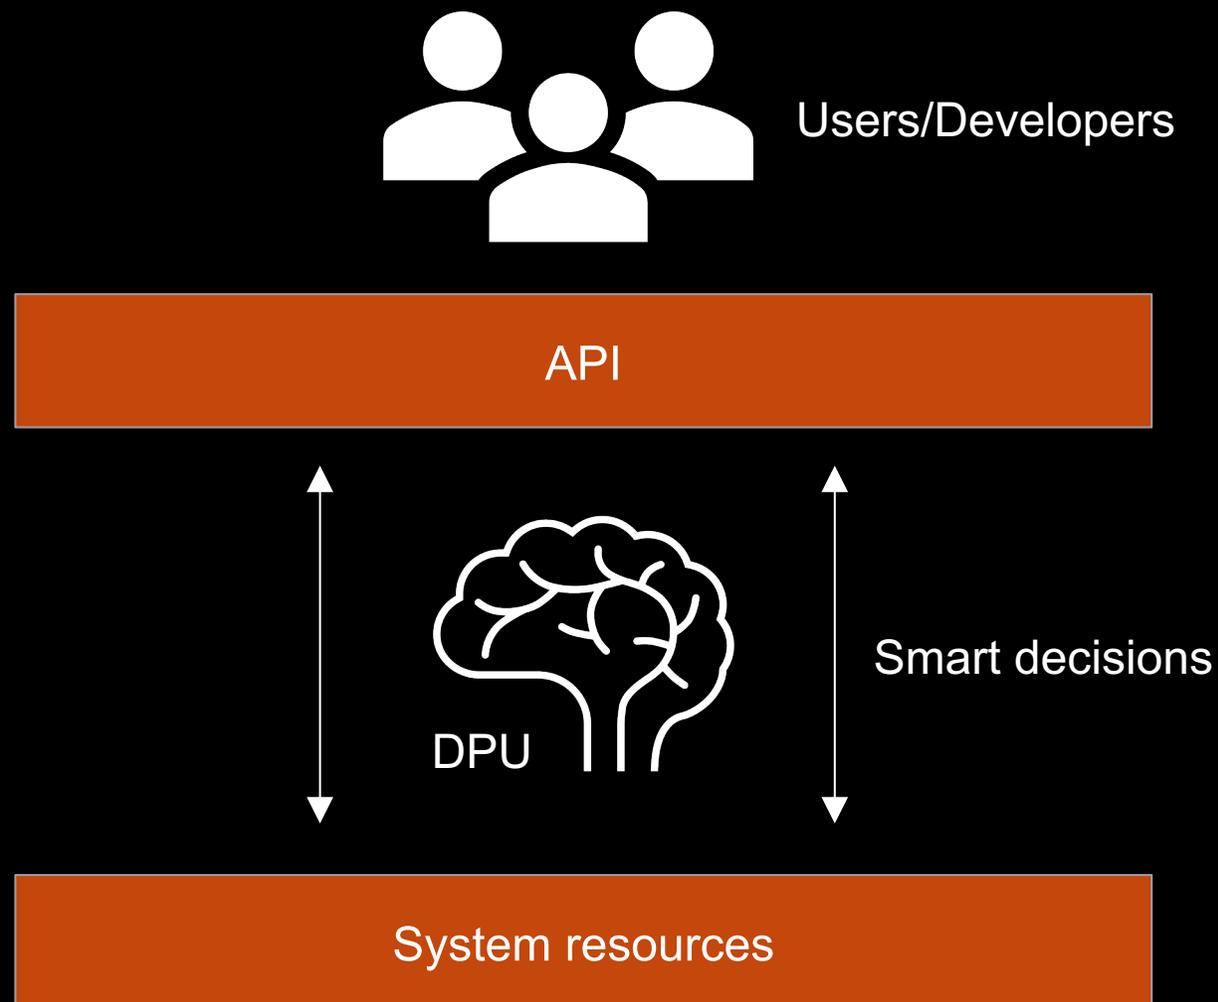
Tooling and telemetry

- Performance is hard to assess
 - Collect data seamlessly
 - No workload disruption
 - Minimum user intervention
 - Low overhead performance sampling
- Integrate with popular profiling tools
- Visualization of traces
 - Realtime view
- Find what is actionable
 - How far off from peak empirical performance
 - Self calibration
- Translate findings into the network configuration
- Telemetry for different workloads
 - Learn what is latency bound
 - Learn what is bandwidth bound
 - Learn when to save power
 - Link cooling down



Making smarter APIs

- Users (would) love
 - Stick with legacy APIs
 - Have simple and high-level application-level API
 - Ignore hardware intricacies
 - Simple abstractions
 - Focus on their science
- API developers (would) love
 - Digest all possible application-level information
 - Offload parameterization to the user
 - Have the user pick sensible defaults
- Can we meet in between?
- DPUs to make API implementation smarter
 - Collect telemetry
 - Realize sub-optimal decisions
 - Tune on-the-fly
 - Maintain data-base of good practices
 - Identify application signatures and map it to better defaults



Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Third-party content is licensed to you directly by the third party that owns the content and is not licensed to you by AMD. ALL LINKED THIRD-PARTY CONTENT IS PROVIDED "AS IS" WITHOUT A WARRANTY OF ANY KIND. USE OF SUCH THIRD-PARTY CONTENT IS DONE AT YOUR SOLE DISCRETION AND UNDER NO CIRCUMSTANCES WILL AMD BE LIABLE TO YOU FOR ANY THIRD-PARTY CONTENT. YOU ASSUME ALL RISK AND ARE SOLELY RESPONSIBLE FOR ANY DAMAGES THAT MAY ARISE FROM YOUR USE OF THIRD-PARTY CONTENT.

© 2023 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ROCm, Radeon, Radeon Instinct, DSC, Alveo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.

The P4 name and the P4 logo are registered trademarks of the Open Networking Foundation.

Thank you!

AMD 