Maximizing heterogeneous system performance with ARM interconnect and CCIX

ARM

Neil Parris, Director of product marketing
Systems and software group, ARM

Teratec June 2017

Intelligent flexible cloud to enable new use cases

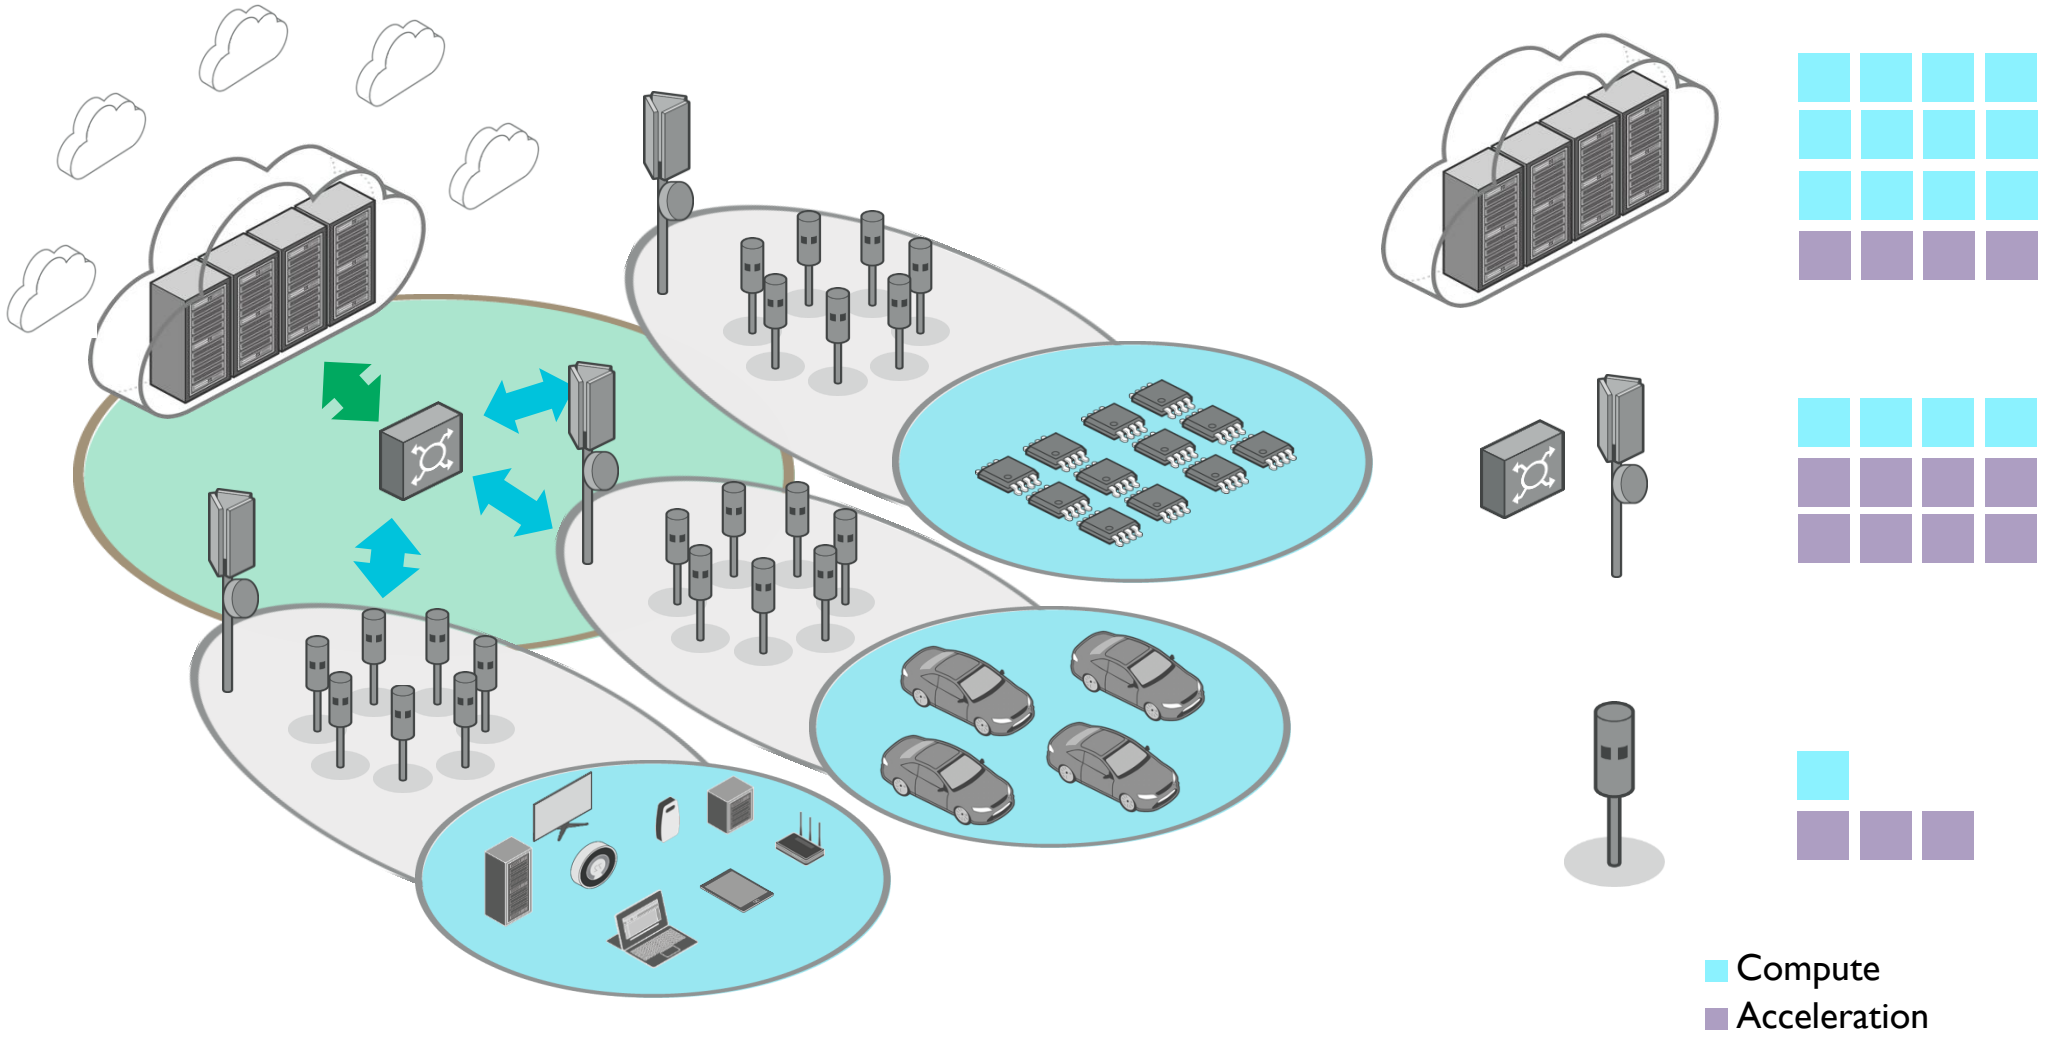

Distributing intelligence efficiently at scale

Cortex-A75: Ground-breaking performance

Cortex-A55: High-efficiency redefined

Unveiling DynamIQ – New single cluster design

- New application processors
  - Cortex-A75 and Cortex-A55 CPUs
  - Built for DynamIQ technology
- Private L2 and shared L3 caches
  - Local cache close to processors
  - L3 cache shared between all cores
- DynamIQ shared unit
  - Contains L3, Snoop Control Unit (SCU) and all cluster interfaces
  - New AMBA 5 CHI bus interface

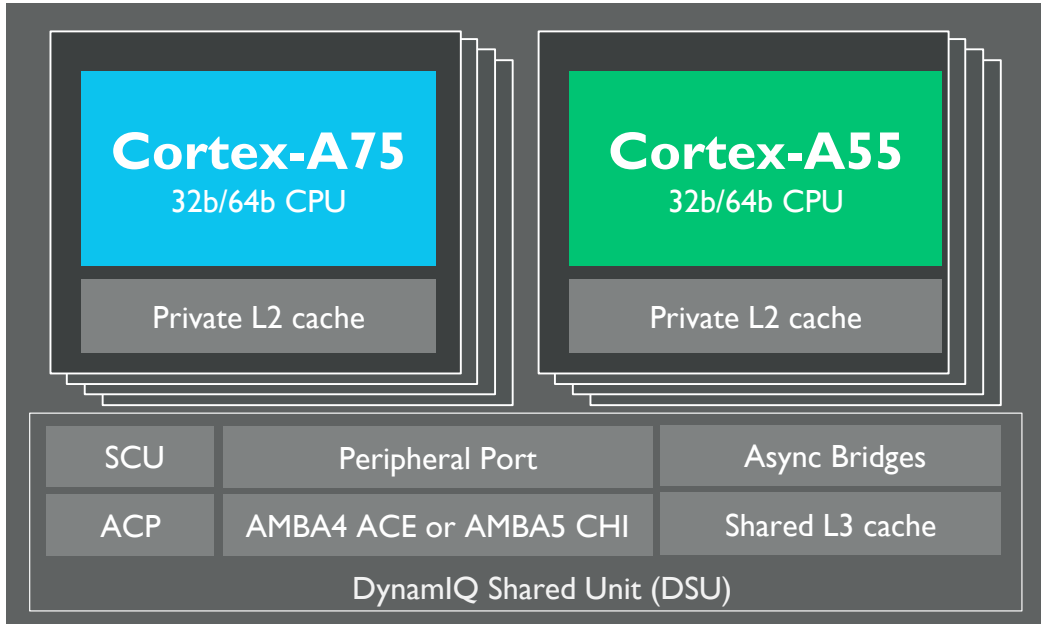

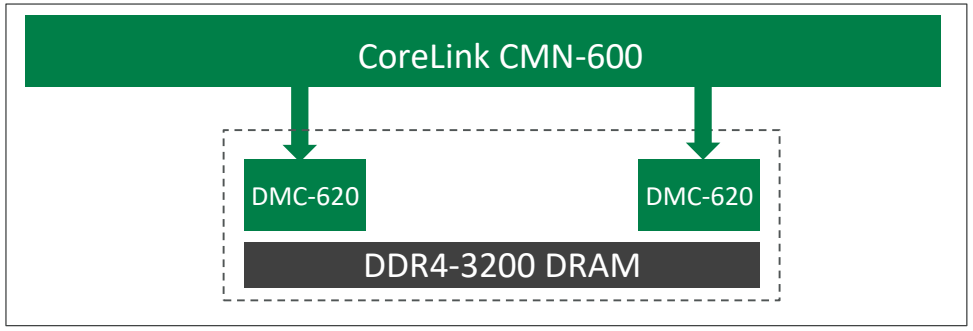

Build higher performance systems

- Boost performance
  - Up to 7.5x more compute
- Innovations to increase throughput
  - Up to 3x more packets per second
- Tailor solutions from edge to cloud
  - Scale from 1 to 256 CPUs on a single chip
- Multichip connectivity supporting CCIX
  - Scale to 1000+ CPUs

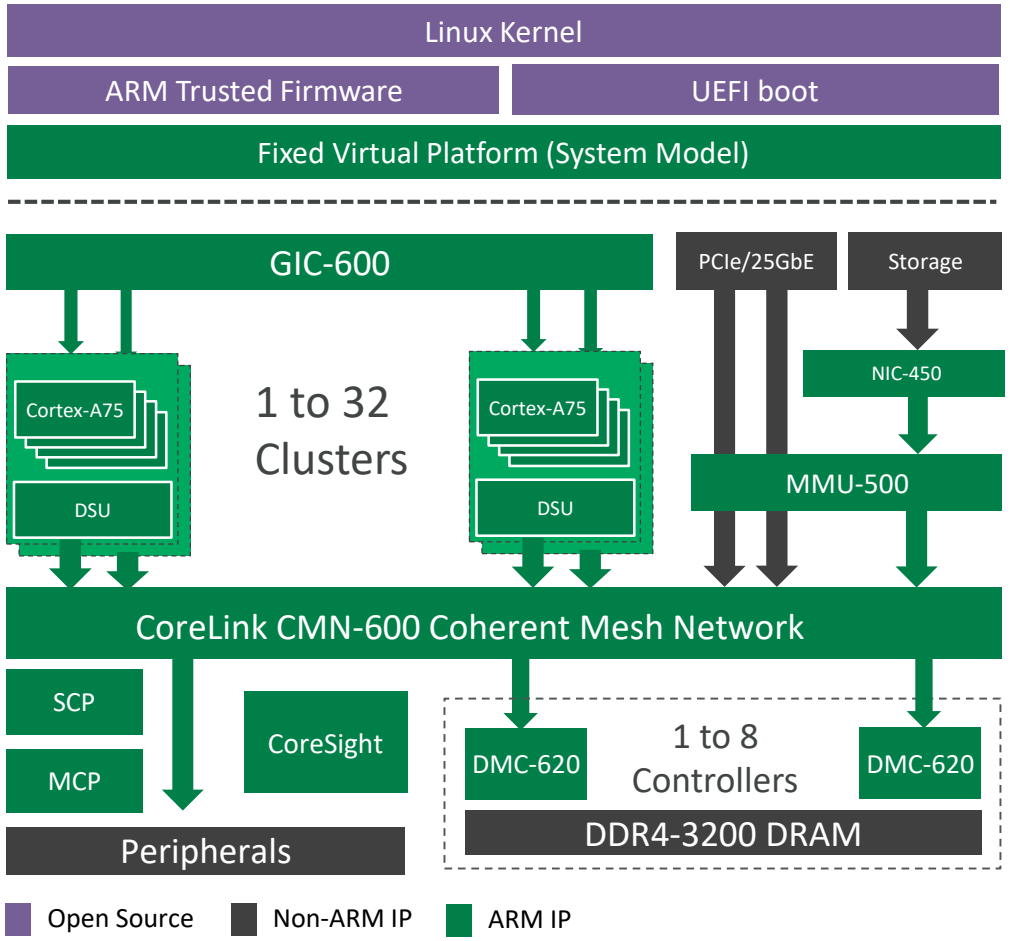

ARM

CMN-600: Delivering maximum compute density

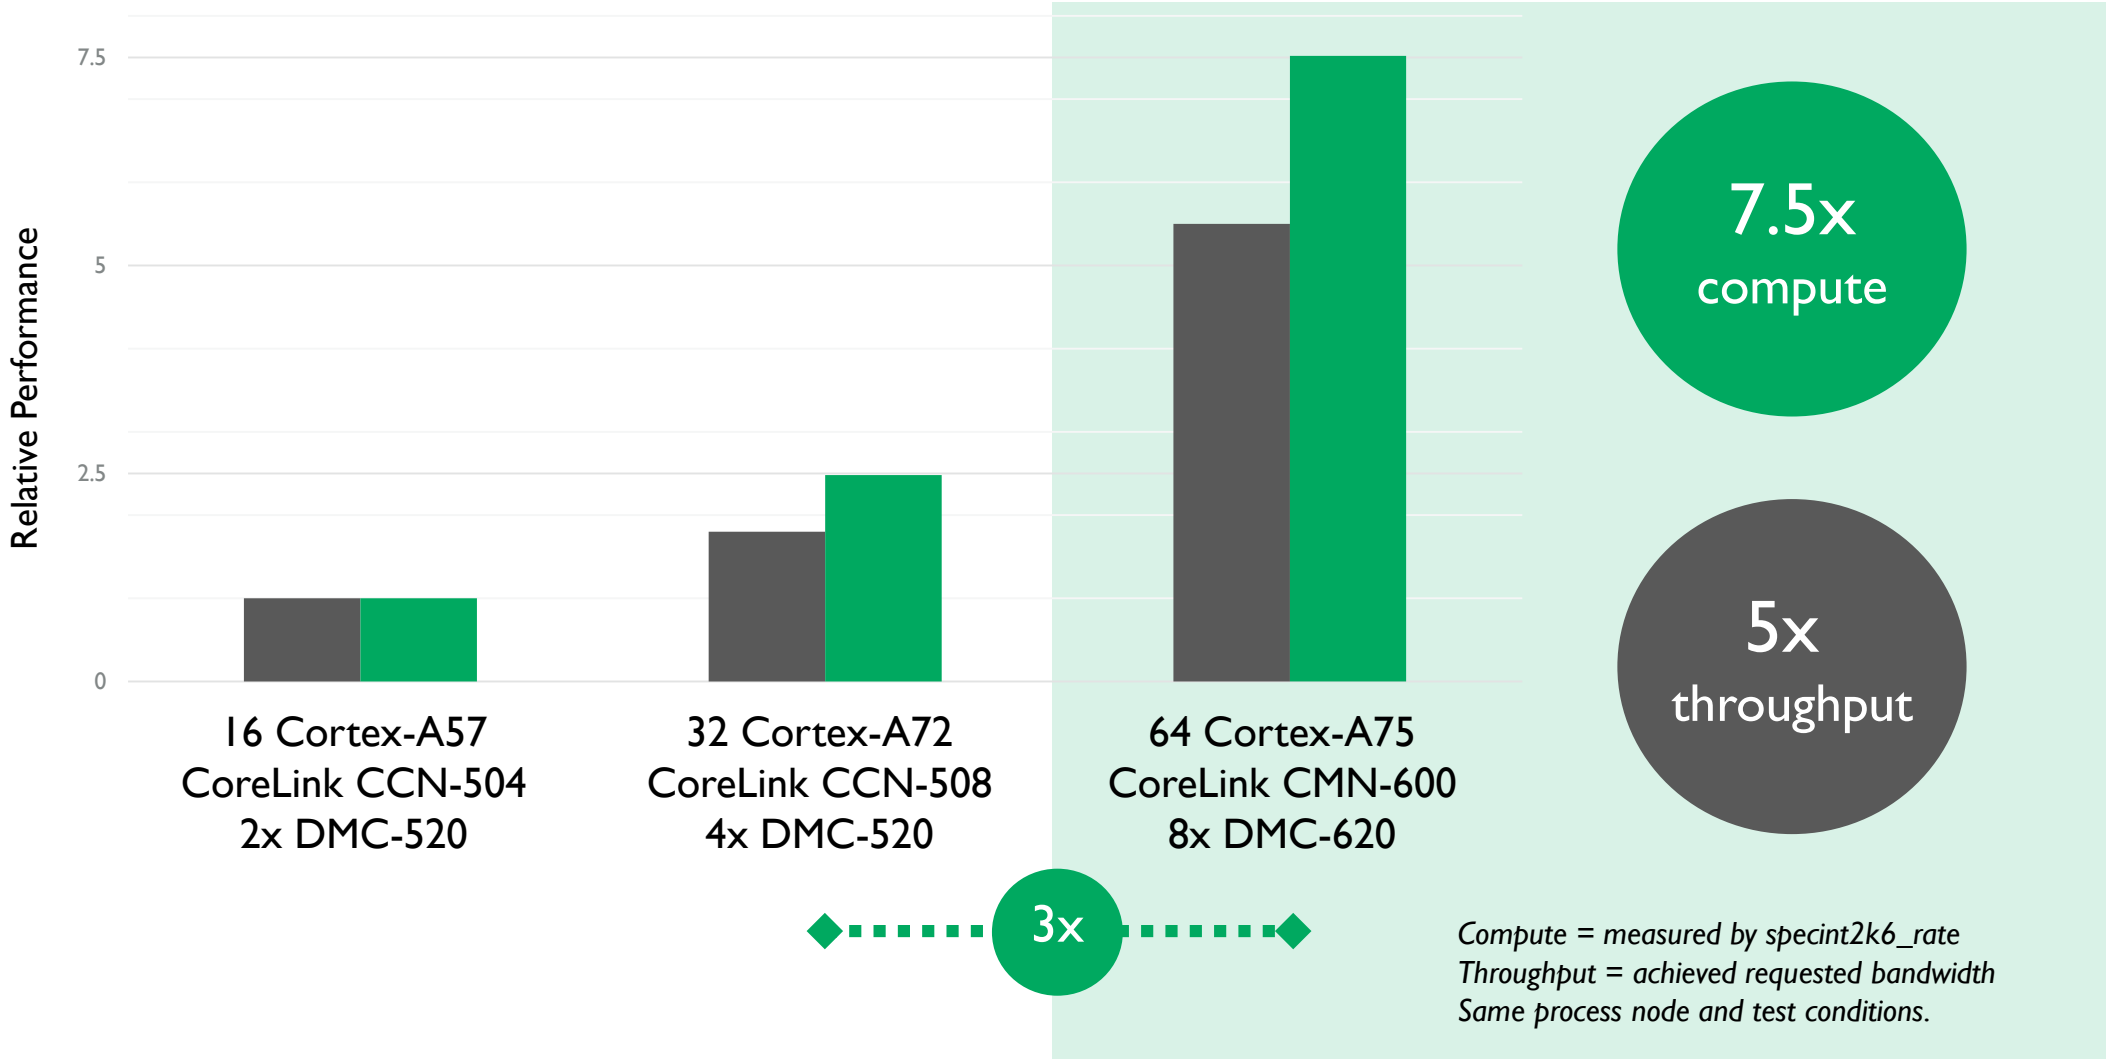

ARM

Scalable solutions from core to edge

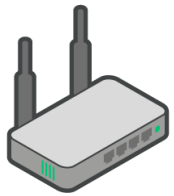

Access point

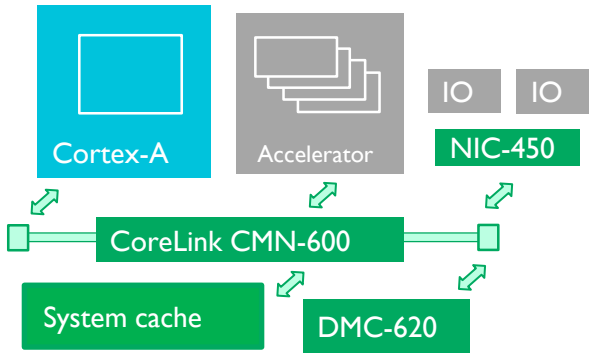

Scalable system configurations

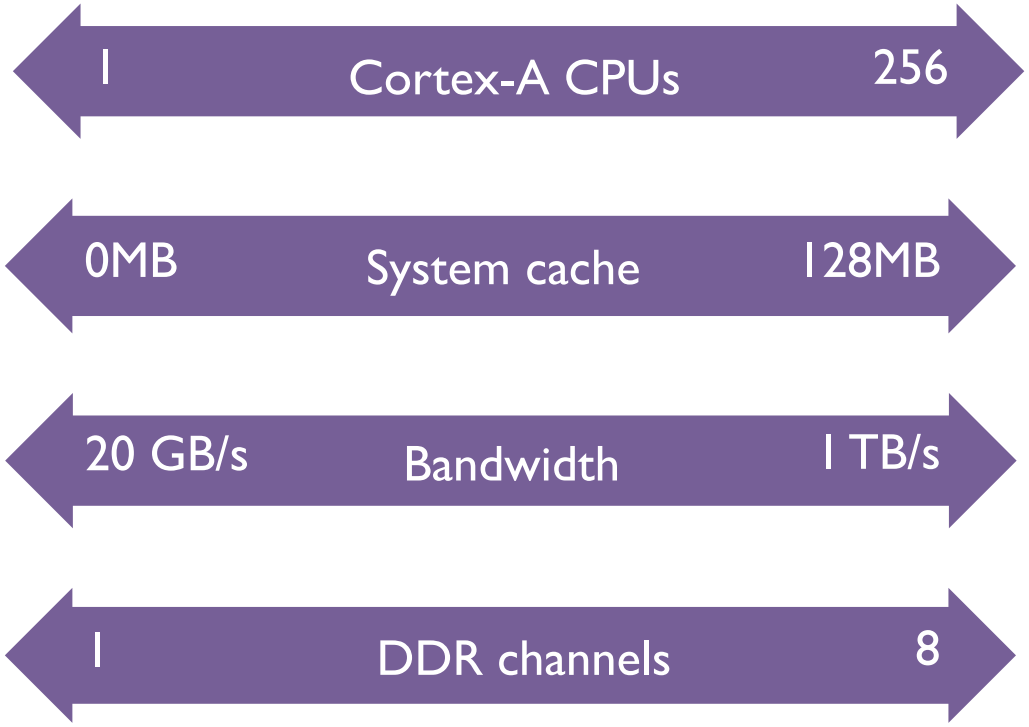

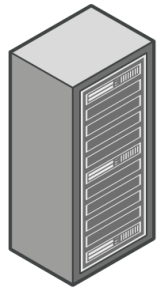

Data center compute

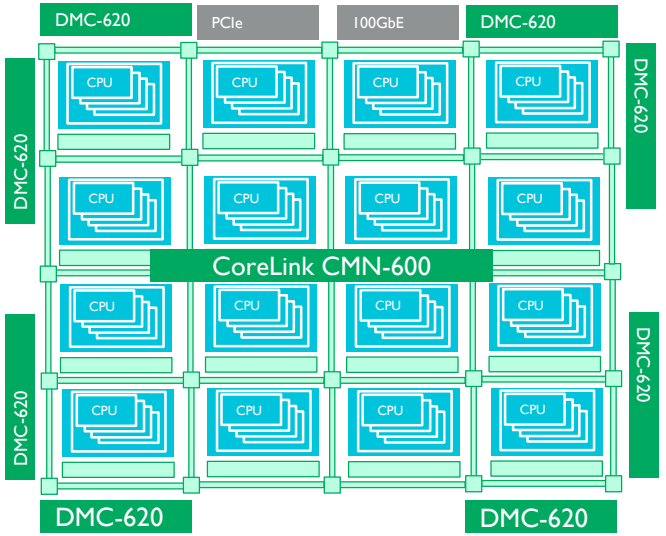

ARM

ARM AMBA 5 CHI coherency protocol

- Credited, non-blocking, coherence protocol
- Capable of achieving high frequencies (>2GHz)
- Layered architecture definition
  - Protocol/Routing/Link/Physical layers
- Flexible topologies for target applications

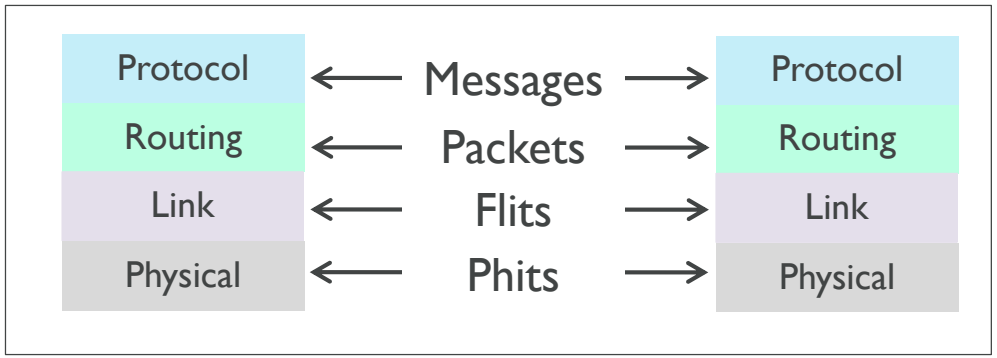

Topology agnostic

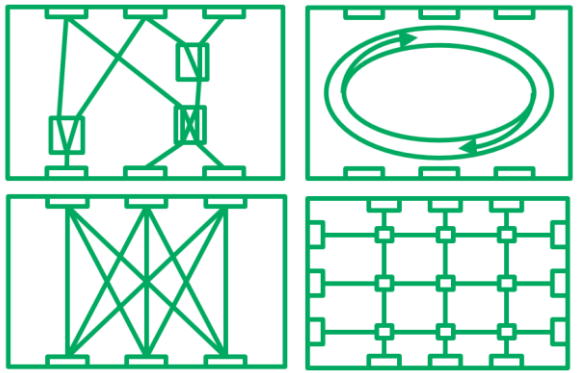

ARM AMBA
Interconnect standards

Innovations to increase throughput

- Up to 3x more throughput with cache stashing
  - Used for acceleration, network, storage use-cases
- Critical data stays on-chip
  - Accelerator or IO selects data placement
  - Low latency access by a CPU or group of CPUs
- Stash to any cache level
  - CPU L2 cache – private to a single CPU thread
  - DynamIQ L3 Cache – shared by CPUs in a cluster
  - Agile System Cache – shared by all DynamIQ clusters

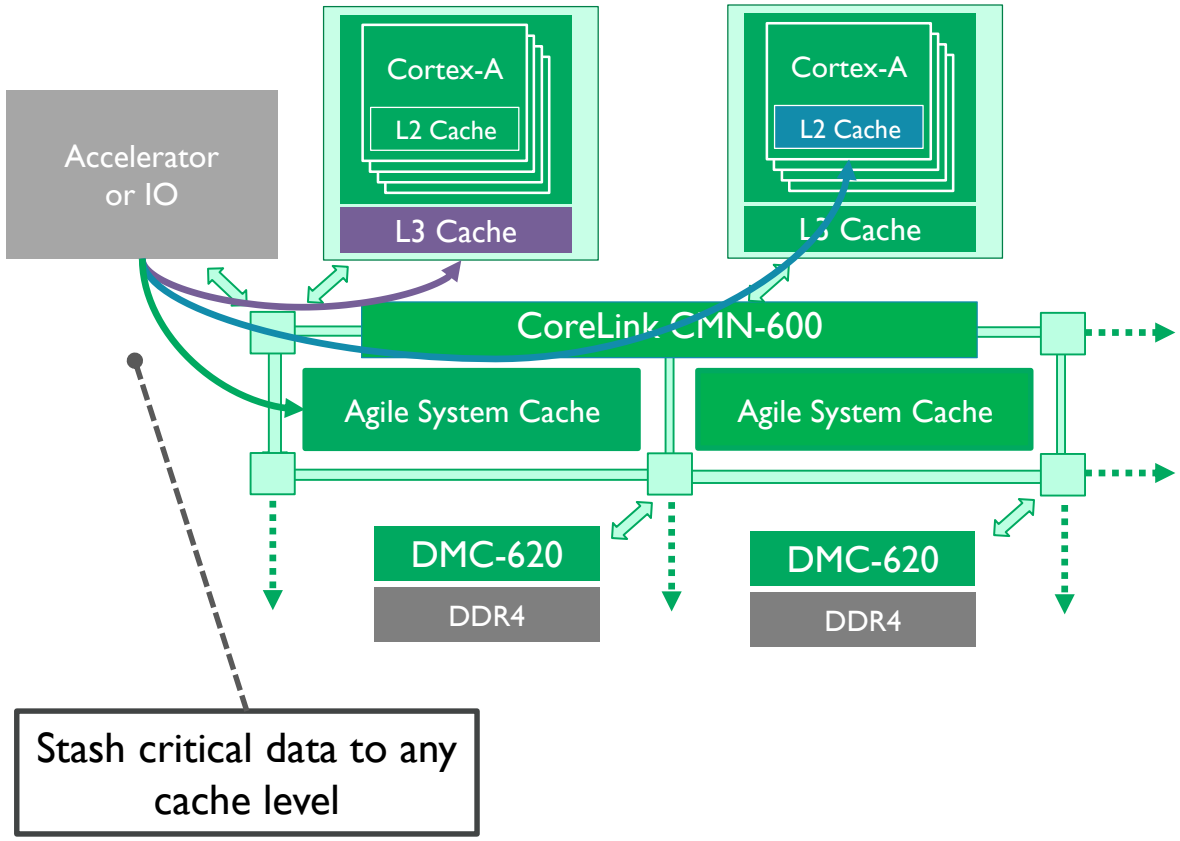

ARM

Customized, tiered acceleration platform

Closely coupled cluster accelerator
- Low latency control and data exchange

Global SoC accelerator or IO
- Direct, high bandwidth path to memory

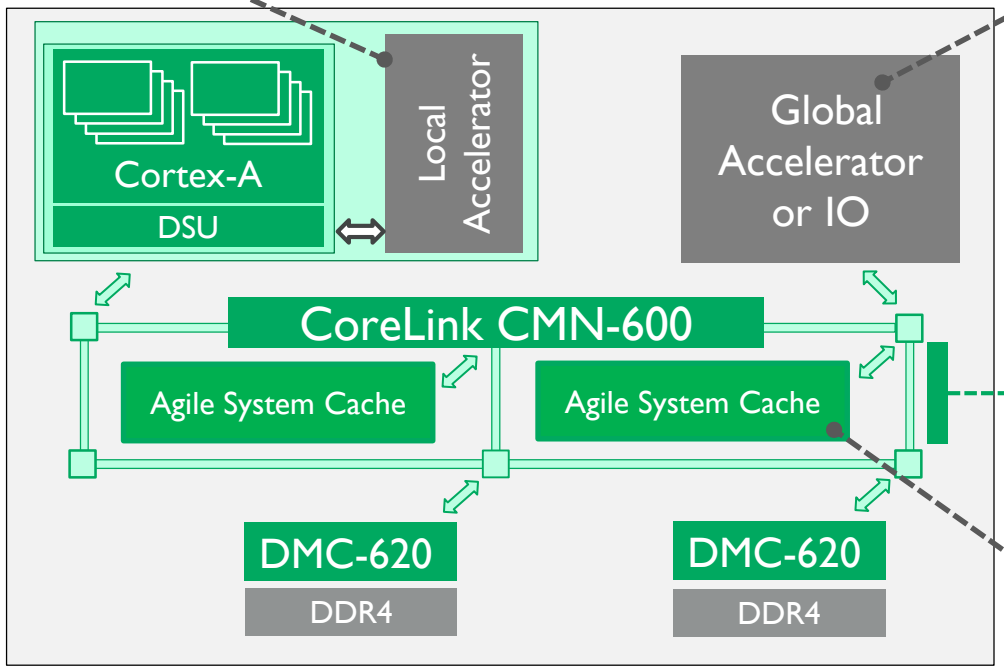

Off-chip coherent accelerator
- High performance CCIX connectivity

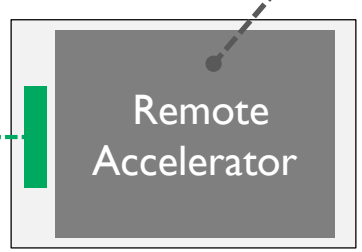

Agile System Cache
- Shared heterogenous system cache

ARM

Coherent multichip

Interconnects needs at different scale

SoC interconnect

Connectivity for on-chip processors, accelerator, IO and memory elements.

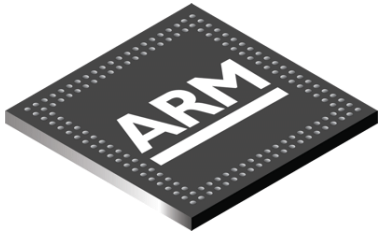

Server node interconnect - ‘scale-up’

Simple multichip interconnect (typically PCIe) topology on a PCB motherboard with simple switches and expansion connectors

Rack interconnect - ‘scale-out’

Scale-out capabilities with complex topologies connecting 1000’s of server nodes and storage elements.

Scale-up server node compute and acceleration

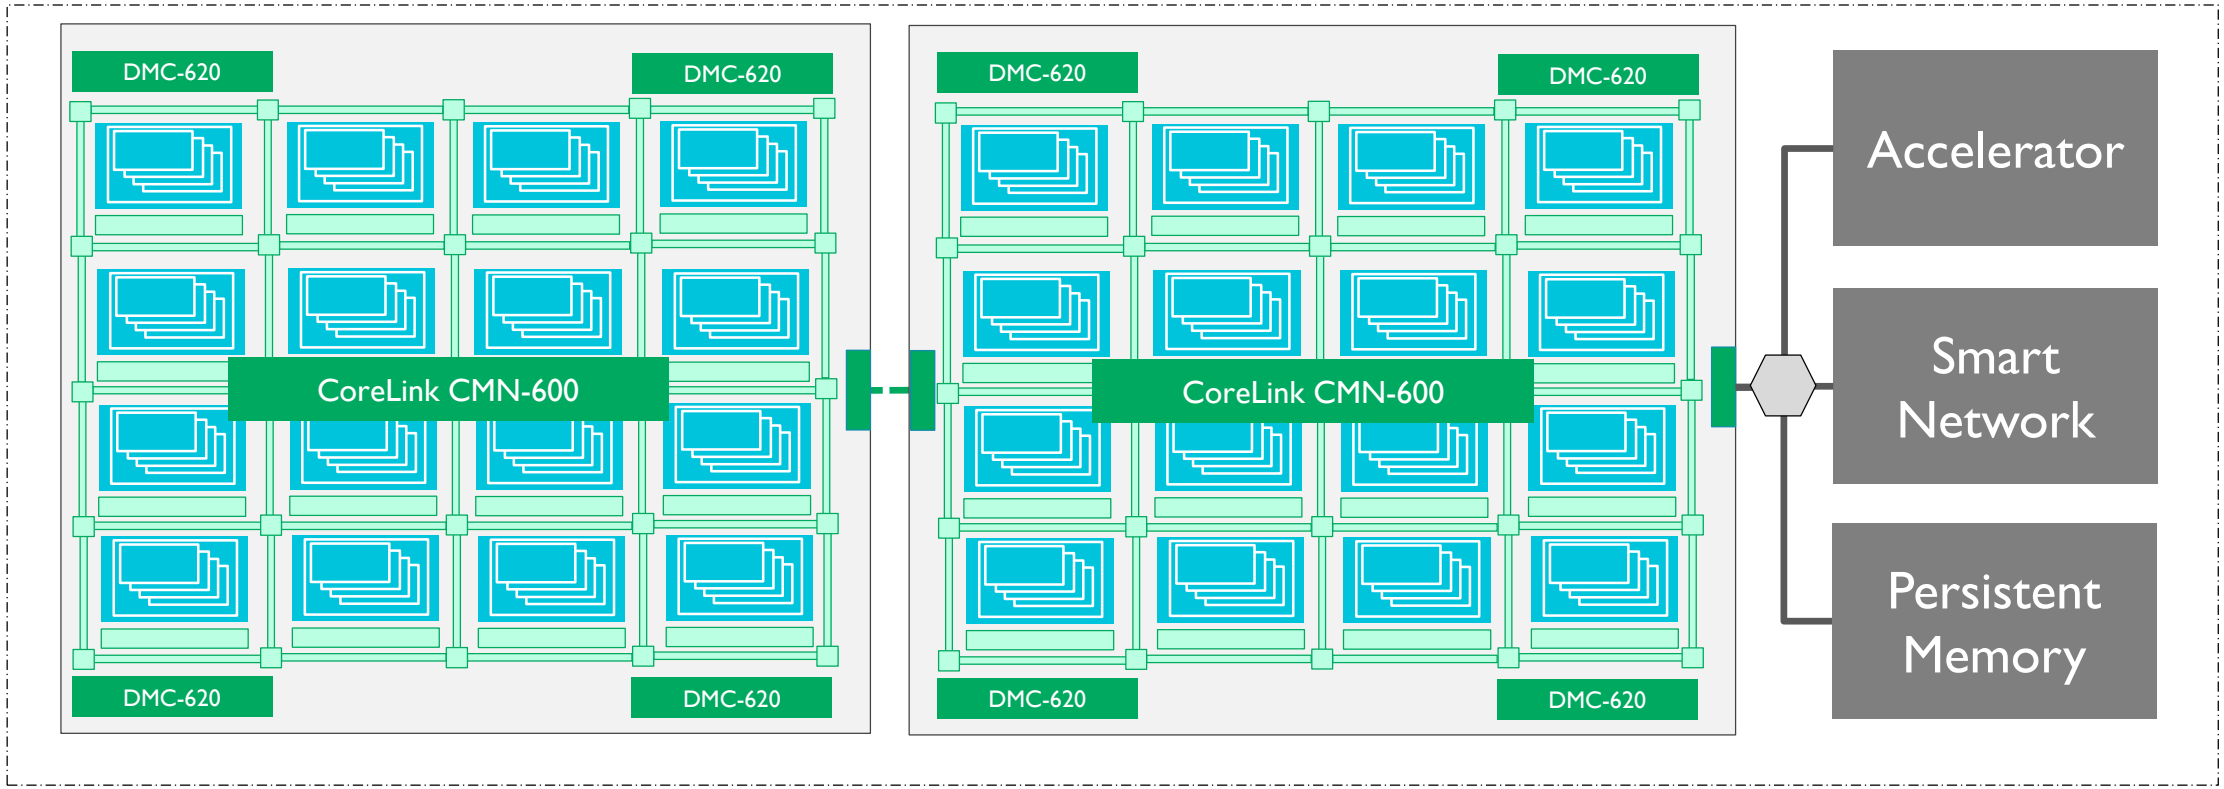

Shared virtual memory system

Cache coherent interconnect for accelerators

- Extends cache coherency to the multi-chip use cases
  - Intelligent network/storage, machine learning, image/video analytics, search and 4G/5G wireless
- Open standard to foster innovation, collaboration and choice to solve emerging workload challenges
- Hardware specification available for designs start for member companies

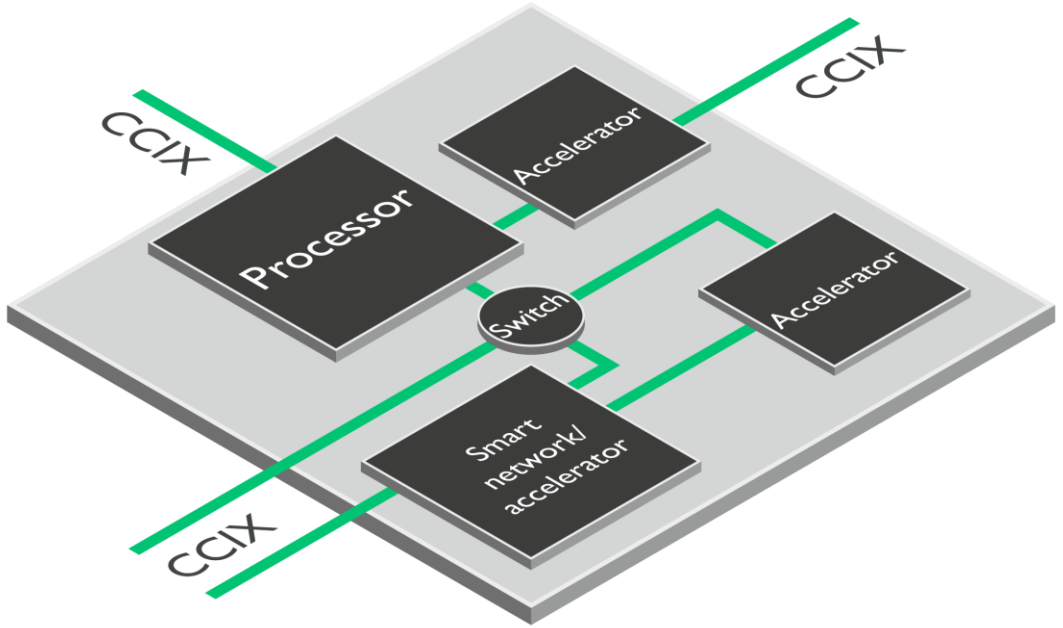

CCIX™ Cache Coherent Interconnect for Accelerators

http://www.ccixconsortium.com

CCIX multichip connectivity and topologies

- New class of interconnect providing high performance, low latency for new accelerators use cases
  - CCIX defines up to 25GT/s (3x performance*)
  - Examining 56GT/s (7x performance*) and beyond
  - Enabling low latency via light transaction layer
- Flexible, scalable interconnect topologies
  - Flexible point-to-point, daisy chained and switched topologies
- Simplified deployment by leveraging existing PCIe hardware and software infrastructure
  - Runs on existing PCIe transport layer and management stack
  - Coexist with legacy PCIe designs

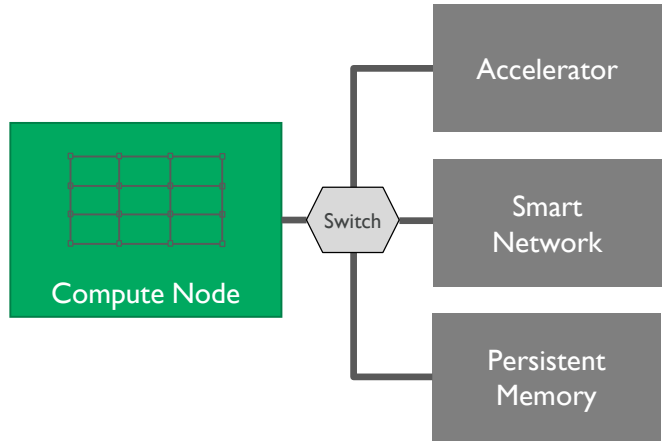

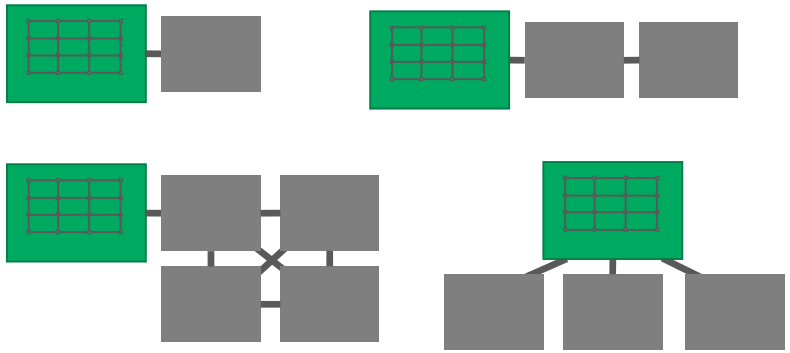

* Note: Based on PCIe Gen3 Performance

ARM

Extending the benefits of cache coherency

- CCIX provides true peer processing with shared memory
  - Accelerators may cache data from processor memory
- Coherency eliminates the software and DMA overhead of transferring data between devices
- Coherency protocol easily adaptable to other transport in the future
- Supports all major instruction set architectures (ISA) enabling innovation and flexibility in accelerator system design

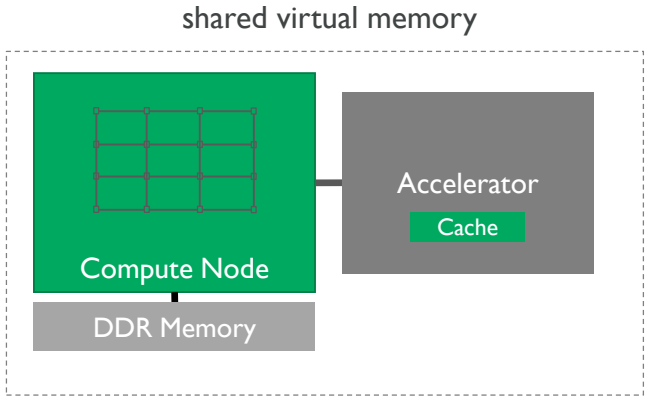

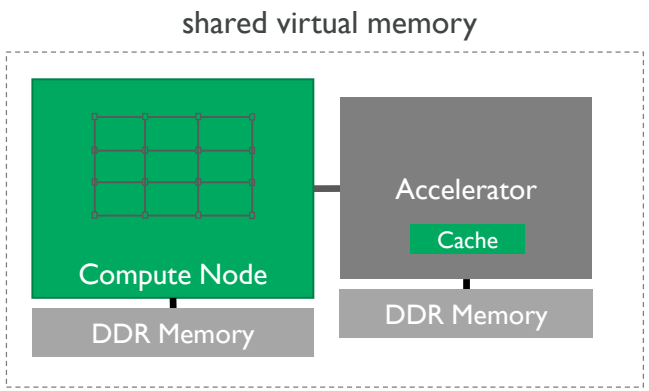

ARM

CCIX example request to home data flows

Accelerator shares processor memory

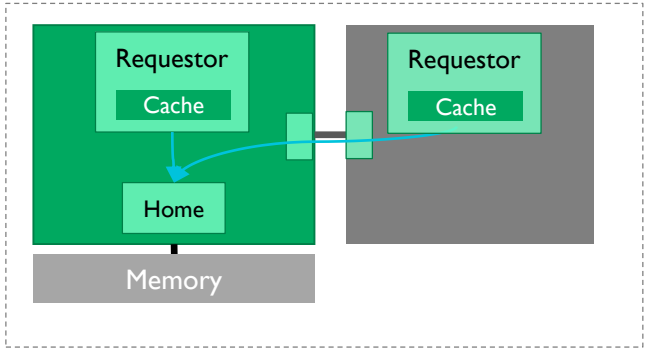

Shared processor and accelerator memory

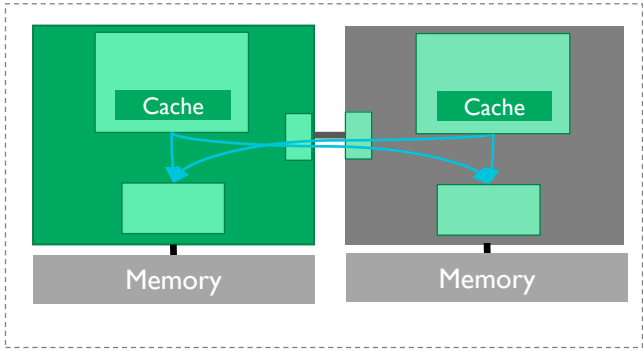

Daisy chain to shared processor memory

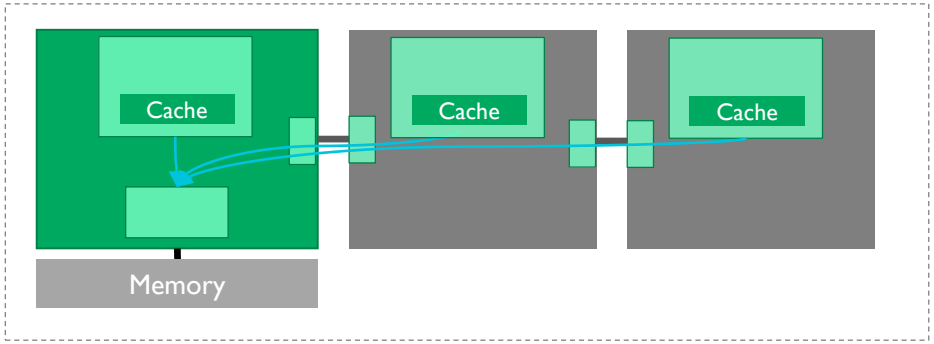

Shared memory with aggregation

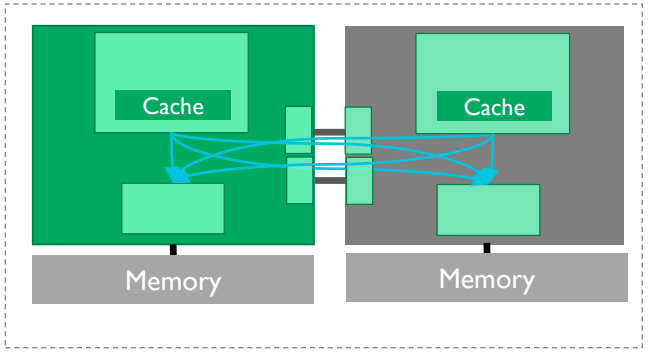

ARM

Today's existing PCIe integration

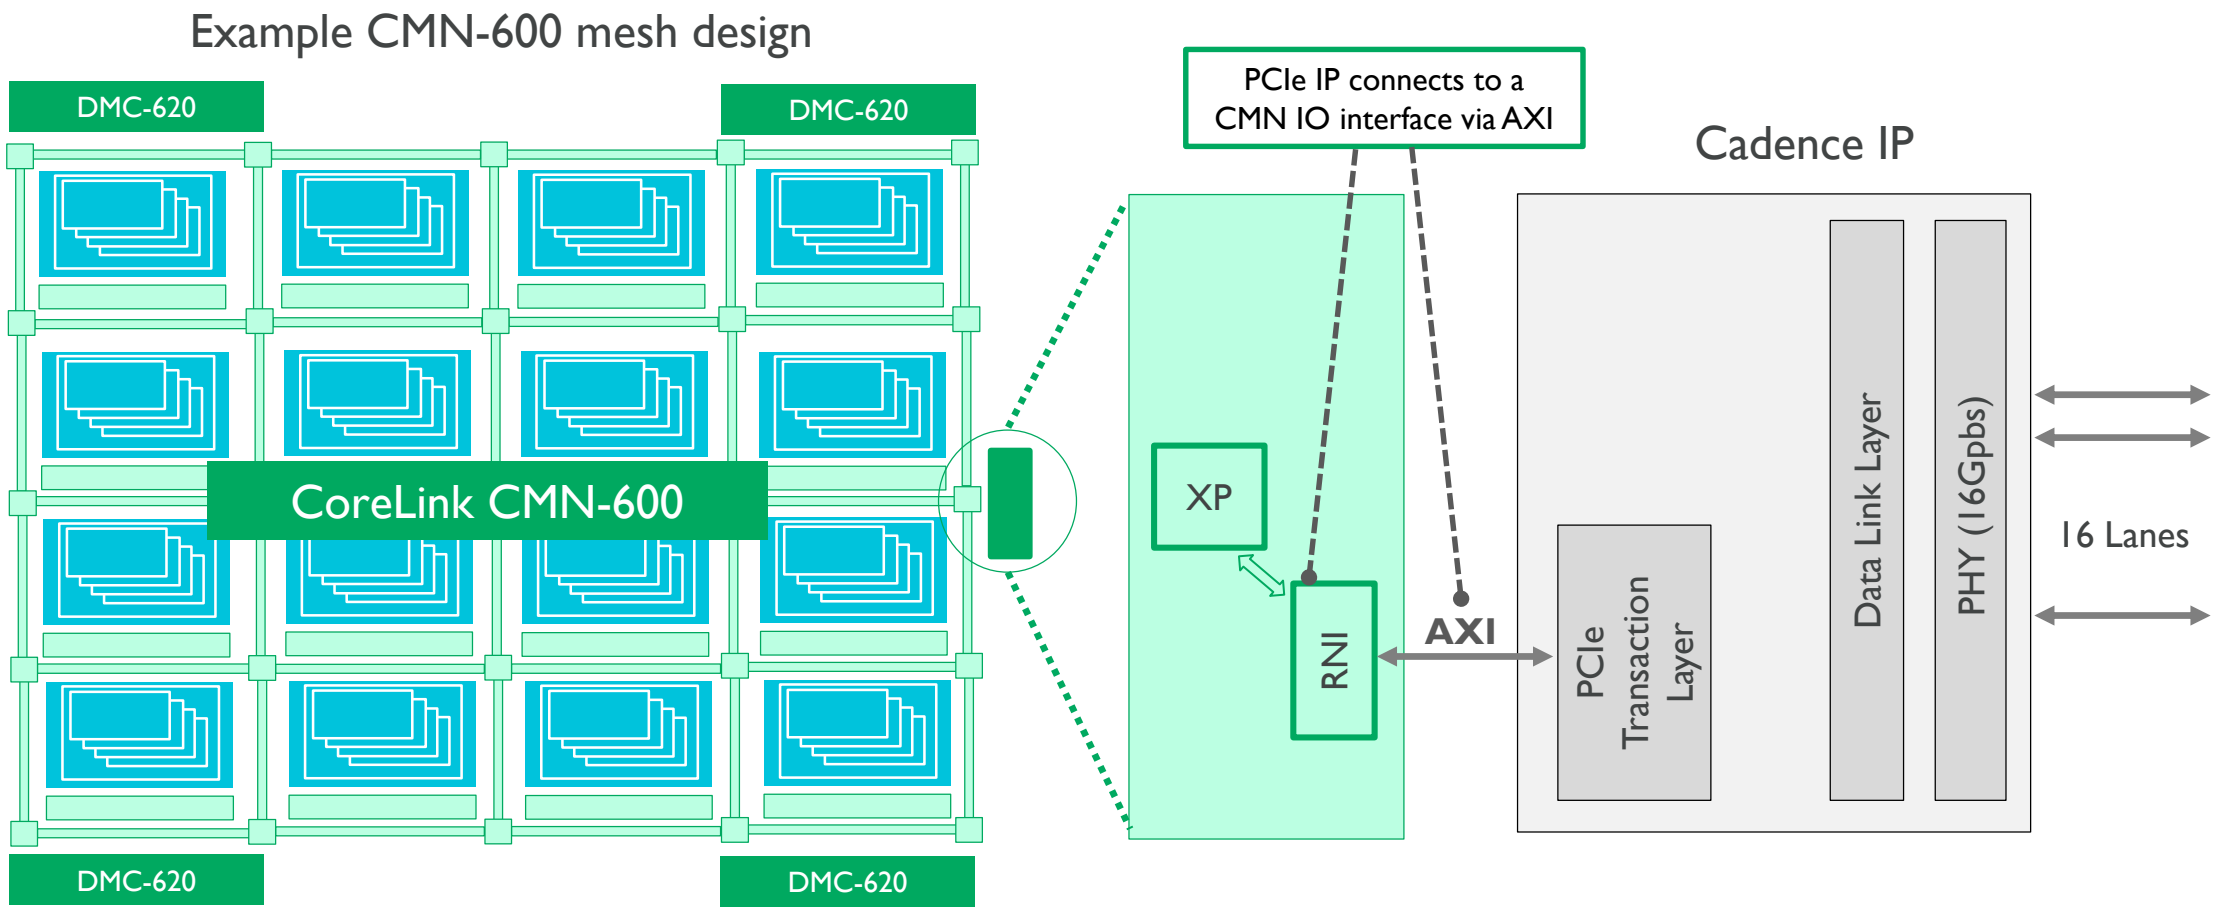

ARM

Coherent Multi-chip Link (CML)

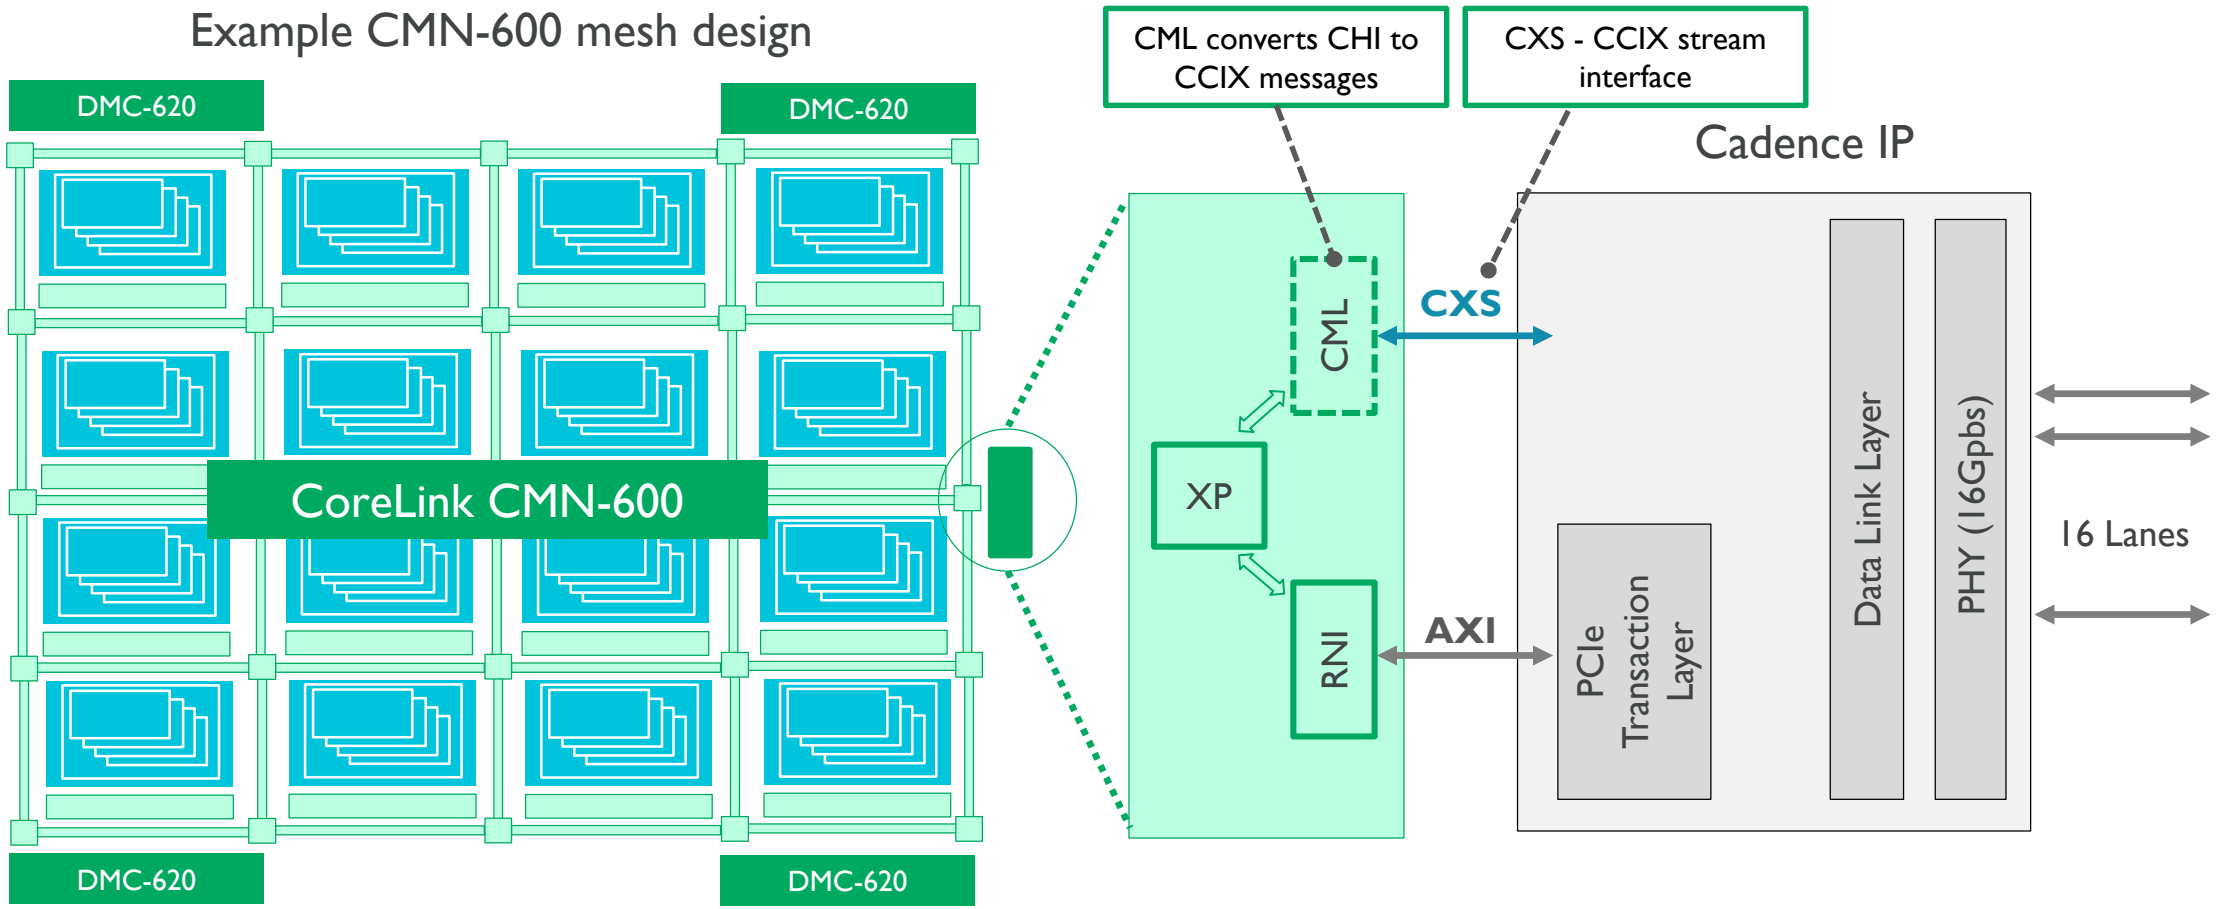

ARM

Cadence CCIX integration

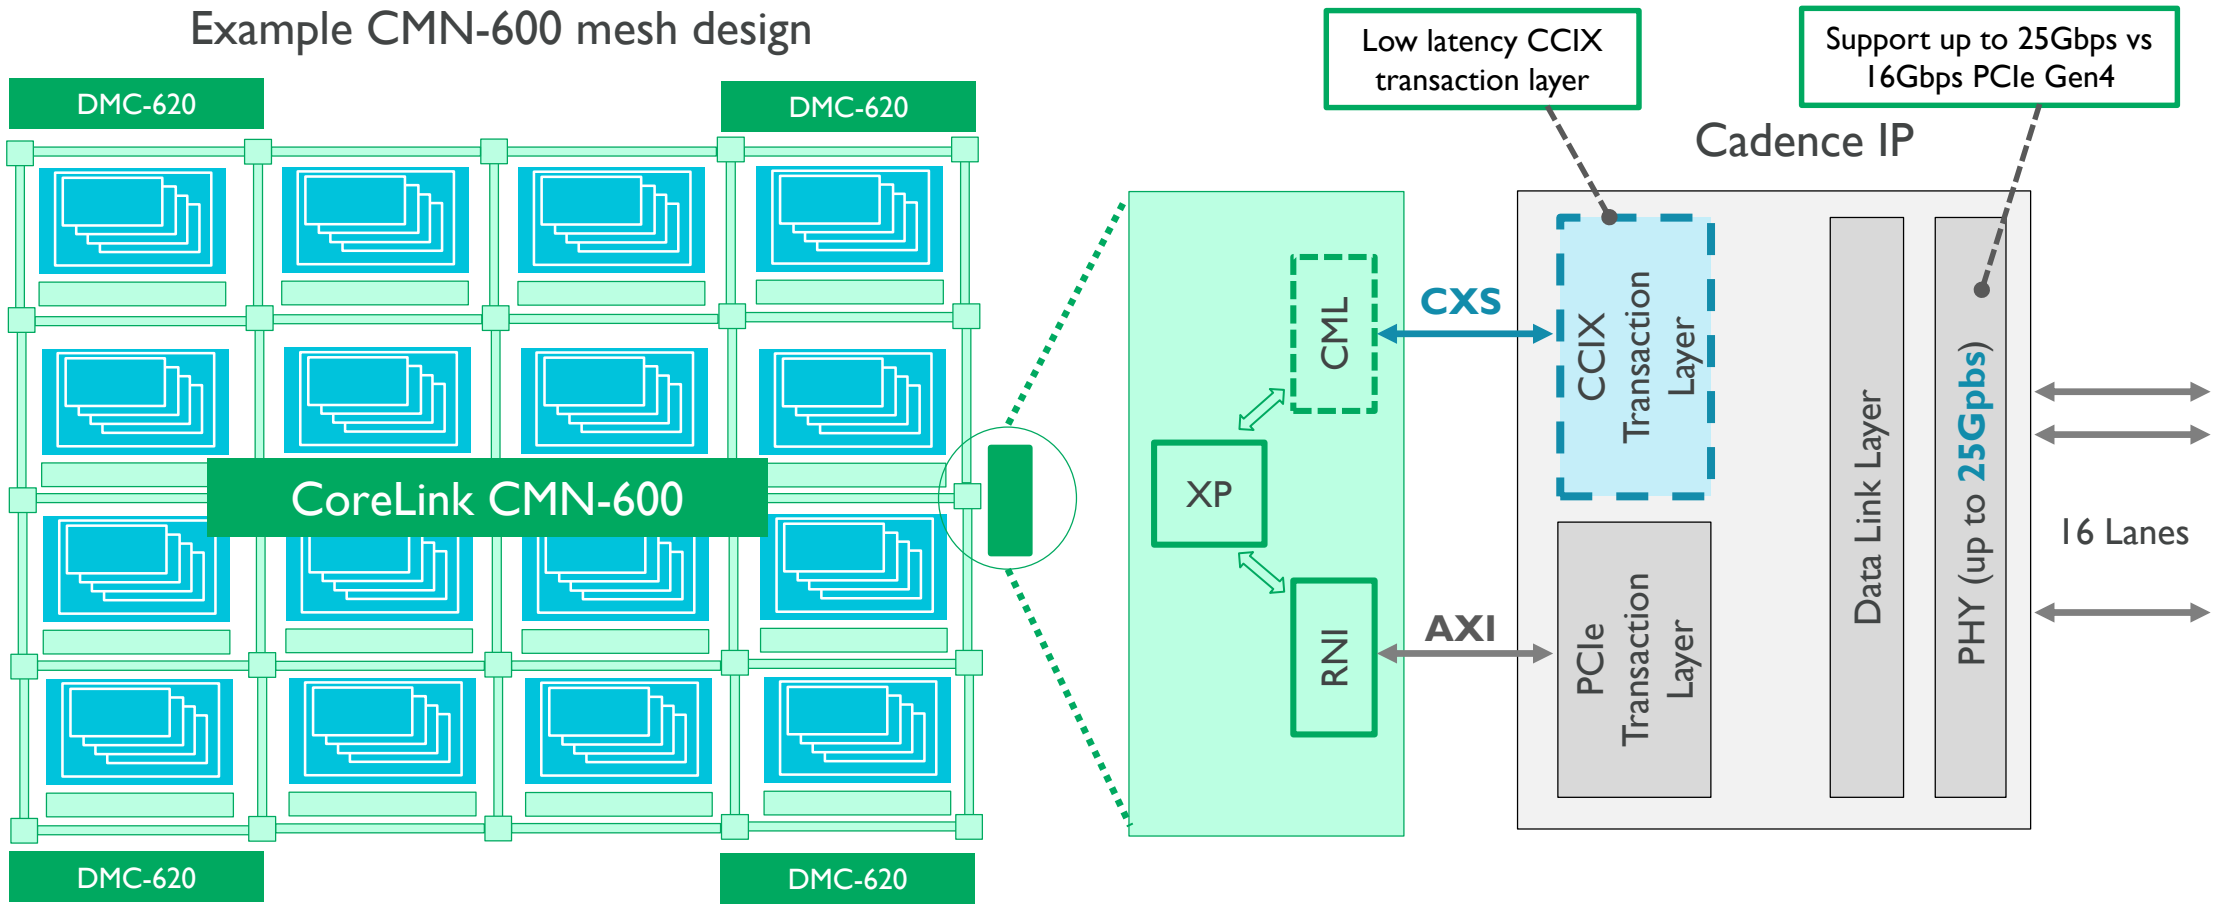

ARM

CCIX 25Gbps technology demonstration

- 3X faster transfer speed with CCIX vs existing PCIe Gen3 solutions
- Transferring of a data pattern at 25 Gbps between two FPGAs
- Channel comprised of an Amphenol/FCI PCI Express CEM connector and a trace card
- Transceivers are electrically compliant with CCIX
- Fastest data transfer between accelerators over PCI Express connections

Xilinx and Amphenol FCI first public CCIX technology demo

https://forums.xilinx.com/t5/Xcell-Daily-Blog/CCIX-Tech-Demo-Proves-25Gbps-Performance-over-PCIe/ba-p/767484

https://youtu.be/JpUSAcnn7VA

Scale-up server node performance with CCIX

- CCIX is a class of interconnect providing high performance, low latency for new accelerator use cases
- Easy adoption and simplified development by leveraging today's data centers infrastructure
- Optimize CCIX SoCs with ARM CoreLink CMN-600 and CCIX controller IP
  - Server, FPGAs, GPUs, network/storage adapters, intelligent networks and custom ASICs

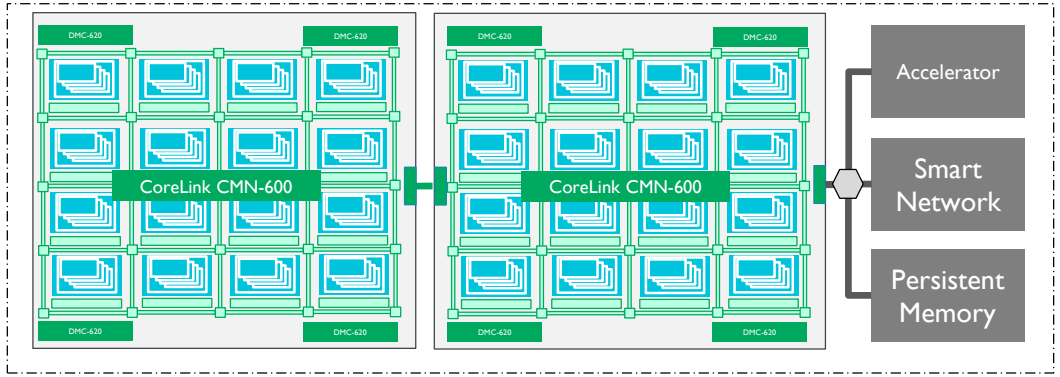

ARM

ARM