

# arm NEOVERSE

Arm Neoverse N2: Arm's 2<sup>nd</sup> generation  
high performance infrastructure  
CPUs and system IPs

*Andrea Pellegrini,  
Distinguished Engineer*

Arm Infrastructure Line of Business

Hotchips 33  
August 23<sup>rd</sup> 2021

# Outline

- Roadmap
- Technical objectives and market targets
- Core architecture & uarchitecture
- System architecture & uarchitecture
- Performance projections
- Conclusions

In development

# Arm Neoverse Platforms Roadmap

IP Date

Production

Now

2022+

Platform features

7nm  
PCIe Gen4, DDR4, HBM2  
CCIX 1.0

7/5nm  
PCIe Gen5, DDR5, HBM2e  
CCIX 1.1


5nm  
PCIe Gen5, DDR5, HBM3  
CCIX 2.0, CXL 2.0

5/3nm  
PCIe Gen5/6, DDR5, HBM3  
CCIX next, CXL next

## V-series

**V1 Platform (Zeus)**  
N1+50% ST perf. uplift  
SVE 2x256b, bFloat16


**Poseidon Generation Platforms**  
  
+30% infra WL perf.  
ML/Vector uplift  
Greater core density



## N-series

**N1 Platform**  
A72+60% ST perf. uplift

**N2 Platform (Perseus)**  
N1+40% ST perf. uplift  
SVE2 2x128b, bFloat16

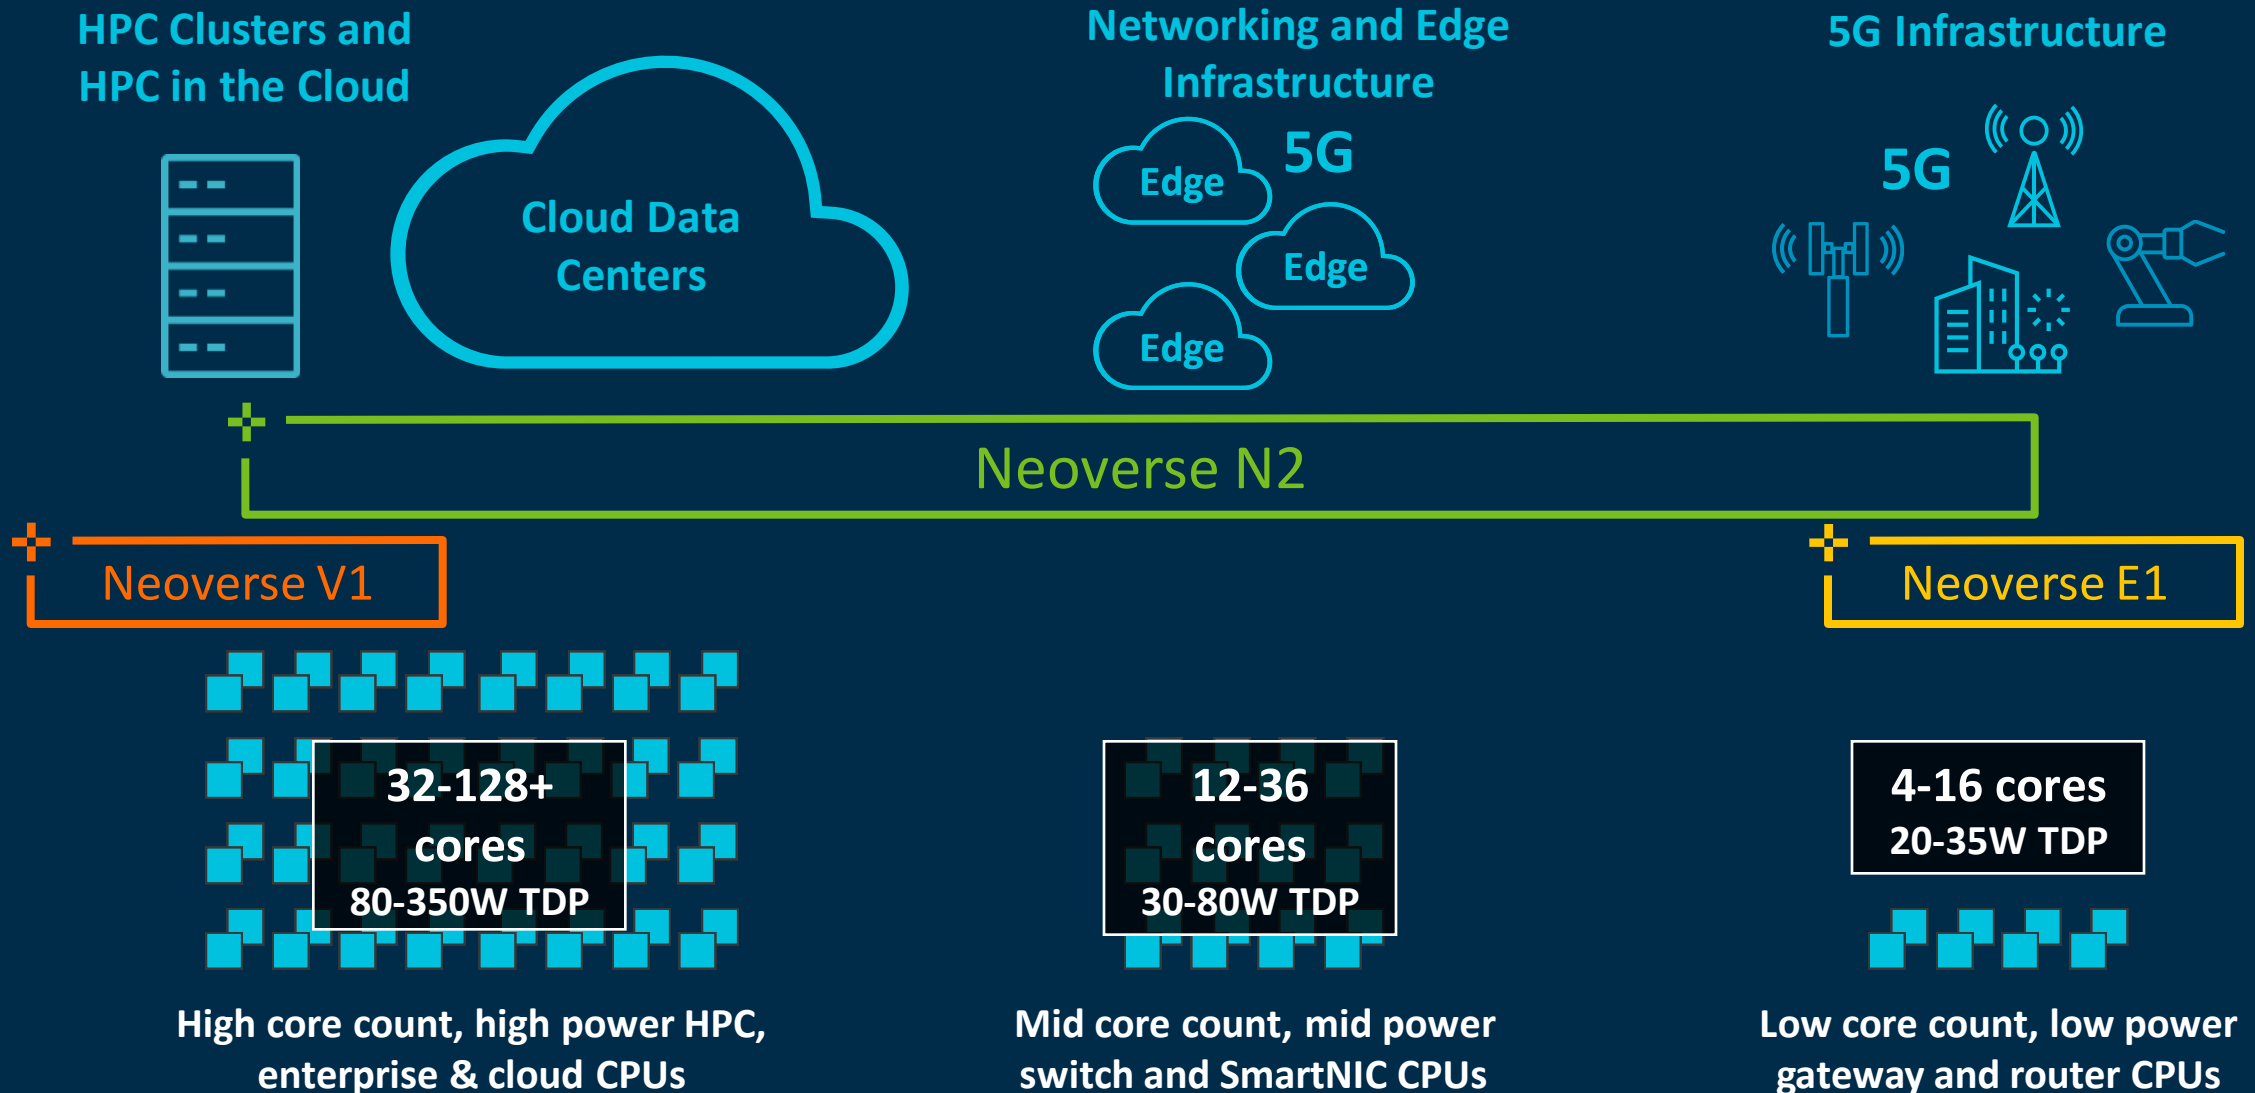


## E-series

**E1 Platform**  
A53+2.7x throughput

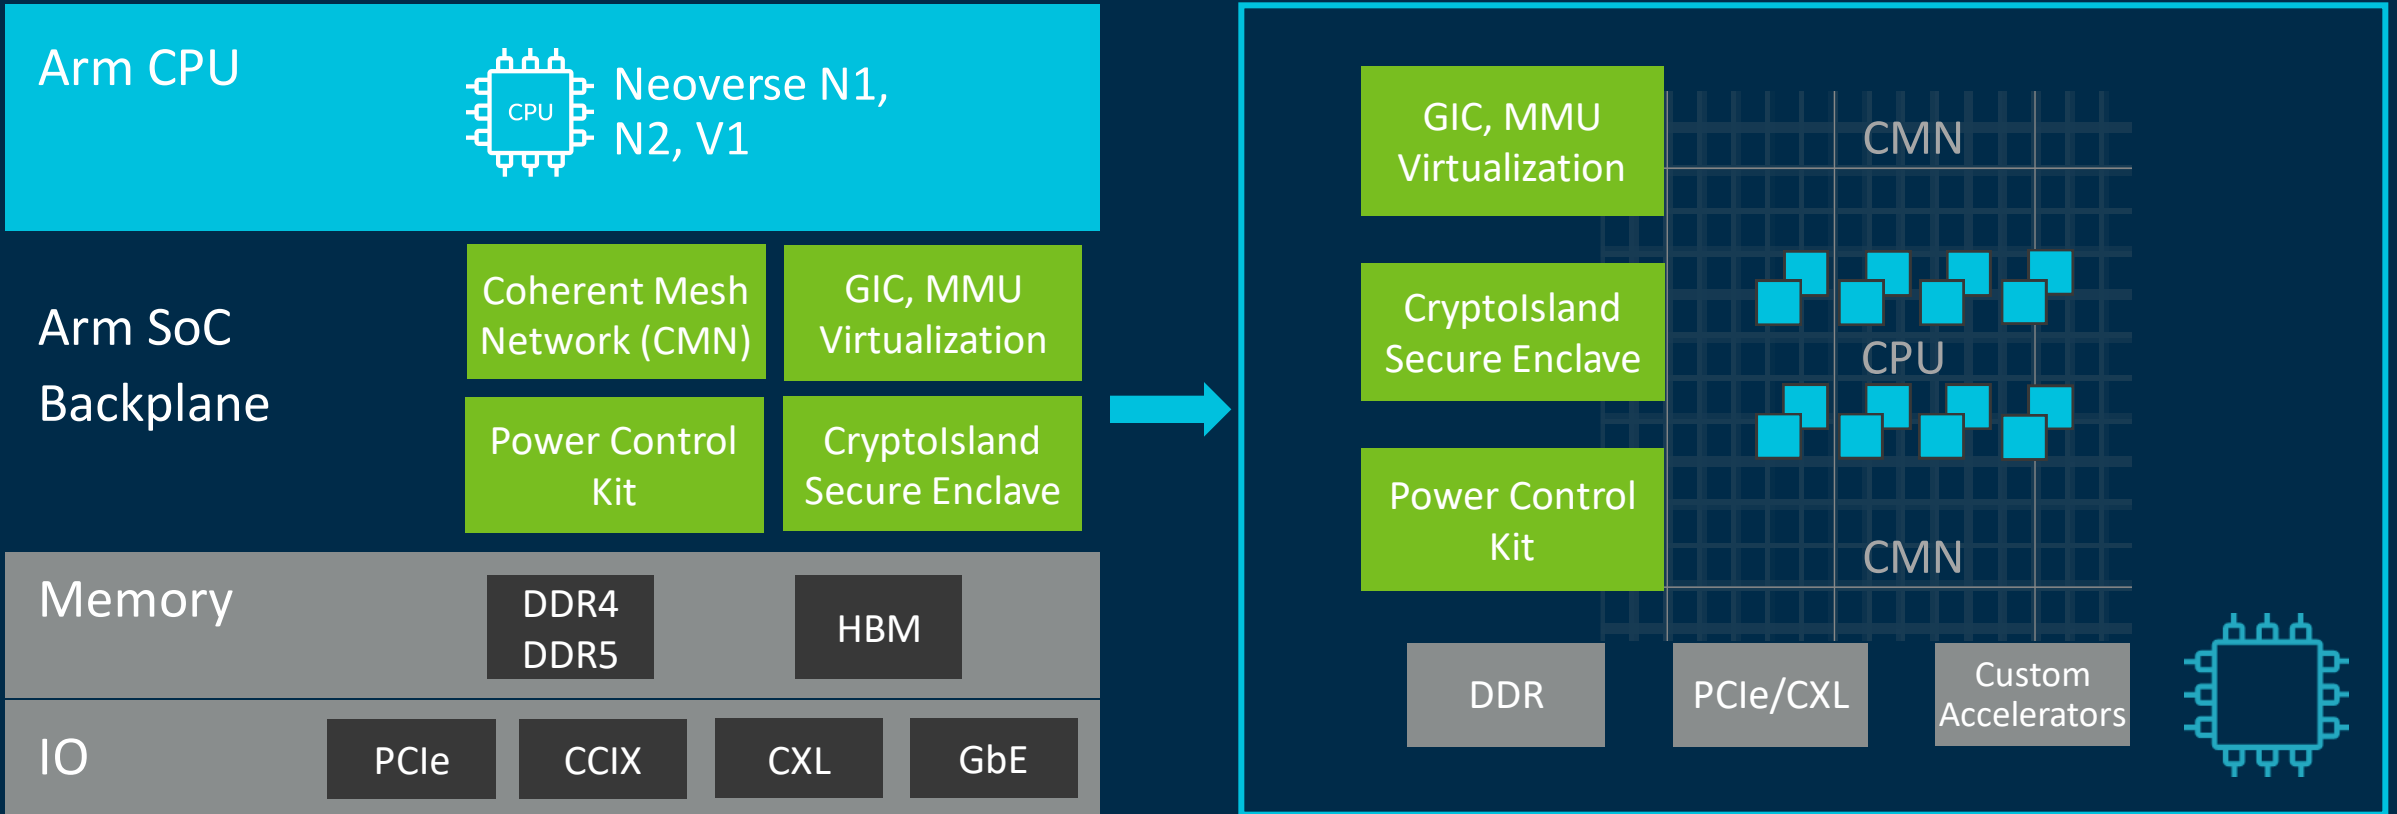


# Arm Neoverse Cloud-to-Edge Workload Positioning



# Arm Neoverse Scalable Compute Platforms

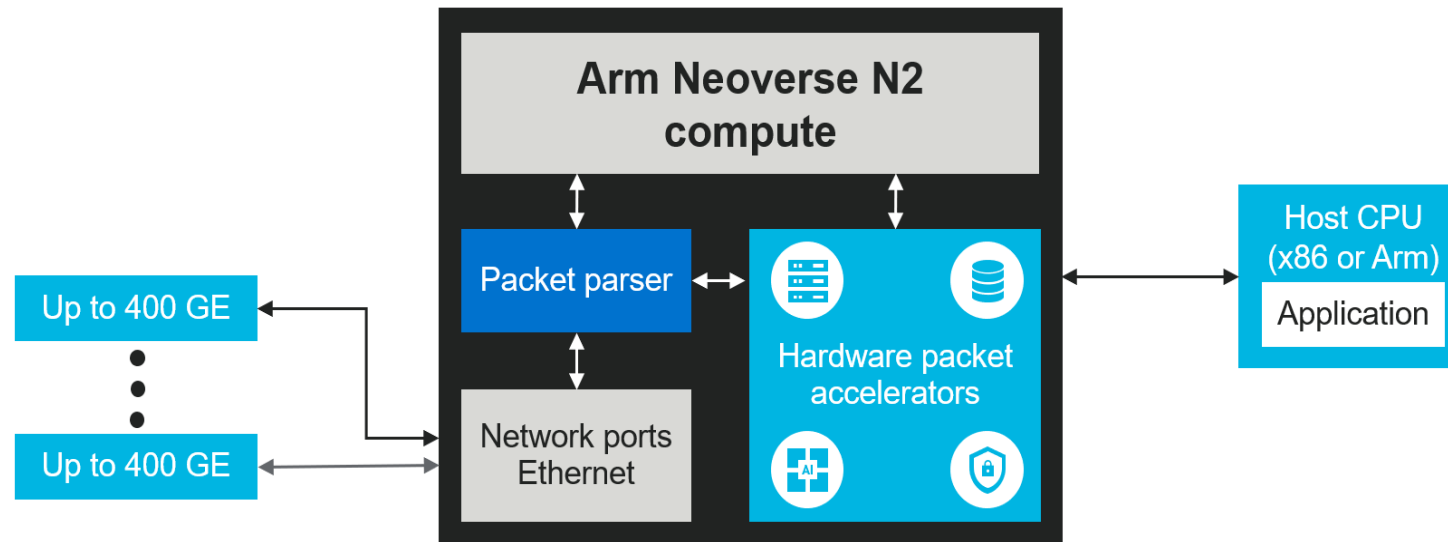
GIC: Generic Interrupt Controller  
MMU: Memory Management Unit



Common Software Platform, SBSA, SBBR, Arm SystemReady SR  
Arm Architecture v8.x-A, AMBA

# Arm Neoverse N2... Design Implementation

## Marvell OCTEON 10 DPU family



### OCTEON 10 performance specifications

- Compute: Up to 36 N2 cores, 1000+ SPECint
- TSMC 5nm
- Datapath: 400G+
- Ethernet ports: Up to 400GE
- AI/ML: 100's TOPS
- Security: 400G+ of IPSEC and SSL

© 2021 Marvell. All rights reserved.

# Arm Neoverse N2... Coming Soon in Silicon


## Marvell OCTEON 10 Family

- |                     |  |
|---------------------|--|
| <b>Compute</b>      | <ul style="list-style-type: none"><li>▪ Up to 36 Neoverse N2 cores</li><li>▪ <u>SPECint</u> &gt; 1000</li></ul>  |
| <b>Accelerators</b> | <ul style="list-style-type: none"><li>▪ inline ML engine</li><li>▪ inline <u>IPSec</u></li><li>▪ SSL/TLS</li><li>▪ Vector packet processing (VPP)</li></ul>                    |
| <b>SW</b>           | <ul style="list-style-type: none"><li>▪ DPDK networking suite for Control, Management and Fast path stacks</li><li>▪ SDK with Linux kernel and user plane extensions</li></ul> |
| <b>Memory</b>       | <ul style="list-style-type: none"><li>▪ DDR5-5200 + ECC on-board memory</li></ul>  |
| <b>I/O</b>          | <ul style="list-style-type: none"><li>▪ Up to 400GE</li><li>▪ Integrated 1Terabit switch</li><li>▪ PCIe 5.0</li></ul>  |

© 2021 Marvell. All rights reserved.




## OCTEON 10 DPU Family



Available in Q4


# Arm Neoverse N2: Arm's First v9 Implementation



Neoverse N1

*Datapath*

● ———— Arm v8.3 through v9 & SVE ———— →



Neoverse N2

- Large system extensions (Atomics)
- Int8 dot product instructions
- Scalable vector extension (SVE, SVE2)
- RCpc memory ordering
- FP round to constrained int.
- Javascript FP conversion
- Enh. cond. flag manipulation
- Complex number support
- Bfloat16
- Int8 matmul

*Security and Crypto*

- Privileged and Stage2 Access-Never
- Unprivileged access override
- Side-channel protection
- SHA512, SHA3, SM3, SM4
- Data independent instructions timing mode
- Branch target identification
- Speculation control
- RAS extensions
- Pointer authentication
- Secure EL2

*Telemetry and system management*

- Support for persistent memory
- HW update of access/dirty bits
- Enhanced PMU
- MPAM
- Deep persistence support
- Self-hosted trace
- Embedded trace extension

*Virtualization*

- 16-bits VMID
- Virtual host extension (type2 hyp)
- Nested virtualization
- TLB maintenance ops
- Small page tables
- Enhanced virtualization traps

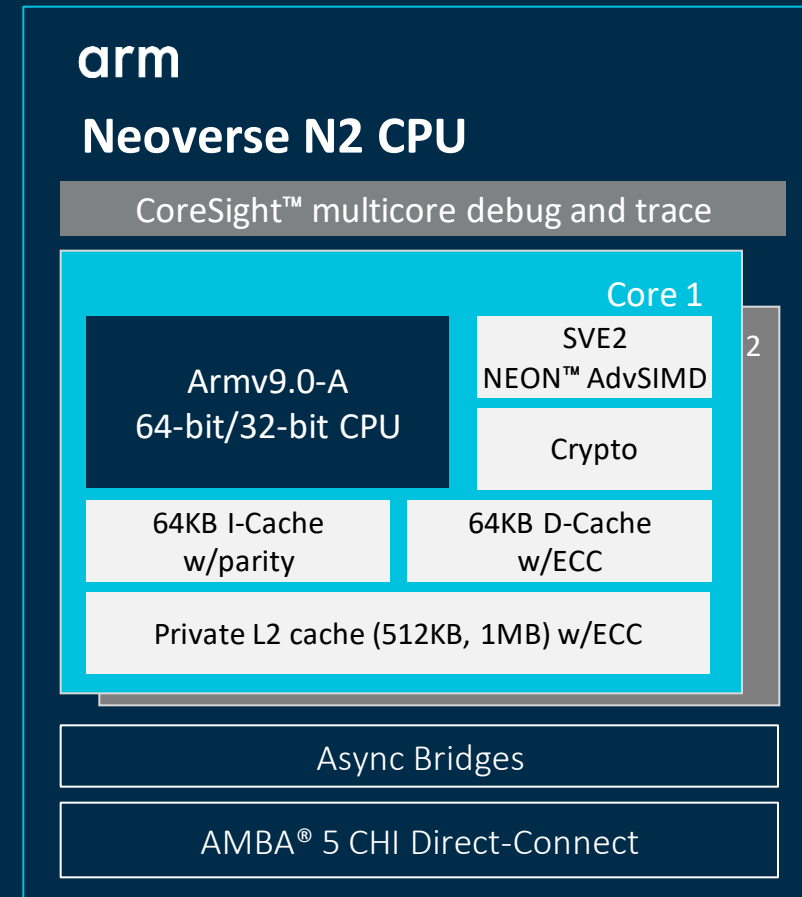


# Arm Neoverse N2: Core Implementation

Major push in efficient performance over Neoverse N1

+40% IPC performance uplift


- Micro-architectural improvements across the board
- Improves benchmarks and server workloads
- Maintains similar power- and area-efficiency as N1, maximizes perf/Watt



# Arm Neoverse N2: Core Implementation

Neoverse N2 maintains the same perf efficiency trajectory as Neoverse N1

Neoverse N1




On 7nm node:

1.0-1.8W / core+L2

1.15-1.4 mm<sup>2</sup> core + L2 (512K/1M L2)

Neoverse N2



On 5nm node:

1.0-2.0W / core+L2

1.1-1.3 mm<sup>2</sup> core + L2 (512K/1M L2)

*Max core freq 3.6GHz*


# Neoverse N2 Core Pipeline Upgrades: Front-End

Branch Prediction	Neoverse N1	Neoverse N2	Improvement
Branch Prediction Width	8 instrs	2 x 8 instrs (up to 2 taken per cycle)	<b>2x</b>
Nano BTB (0 cyc taken-branch bubble)	16 entry	64 entry	<b>4x</b>
Conditional branch direction state	1x	1.5x	<b>1.5x</b>
Main BTB	6K entry	8K entry	<b>1.33x</b>
Alt-Path Branch Prediction	No	Yes	<b>New</b>
Fetch	Neoverse N1	Neoverse N2	Improvement
L1 Instruction cache	64KB	64KB	-
Mop Cache	N/A	1.5K entry	<b>New</b>
Fetch Queue	12-entry	16-entry	<b>1.33x</b>
Fetch Width	4 instr	4 instr from I\$, 5 instr from MOP\$	<b>Up to 1.5x</b>
Early branch redirect	Yes (unconditional)	Yes (uncond + cond)	<b>Improved feature</b>
Decode width	4 (I-cache)	4 (I-cache) or 5 (Mop cache)	<b>Up to 1.25x</b>

# Arm Neoverse N2: Front-End

Re-design front-end microarchitecture to achieve higher performance and security

- Higher fetch bandwidth and lower latency
  - Branch predict up to 16-inst/cycle, 2-taken/cycle
  - New Macro-op (MOP) cache with 1.5k entries
- More accurate branch predictor:
  - 50% larger branch direction predictor



- More efficient instruction handling:
  - More instruction fusion via MOP cache
- Special effort to accelerate application with large code footprints
  - Improved Branch Predictor Directed Prefetch by larger fetch queue for more outstanding requests
  - 33% larger BTB with shorter average latency
  - Early re-steering for conditional branches that miss the BTB
- Enhanced security:
  - Branch target is tagged by software context number (SCXTNUM, Arm v8.5)

Lower branch misprediction and lower I-cache miss penalty for server workloads

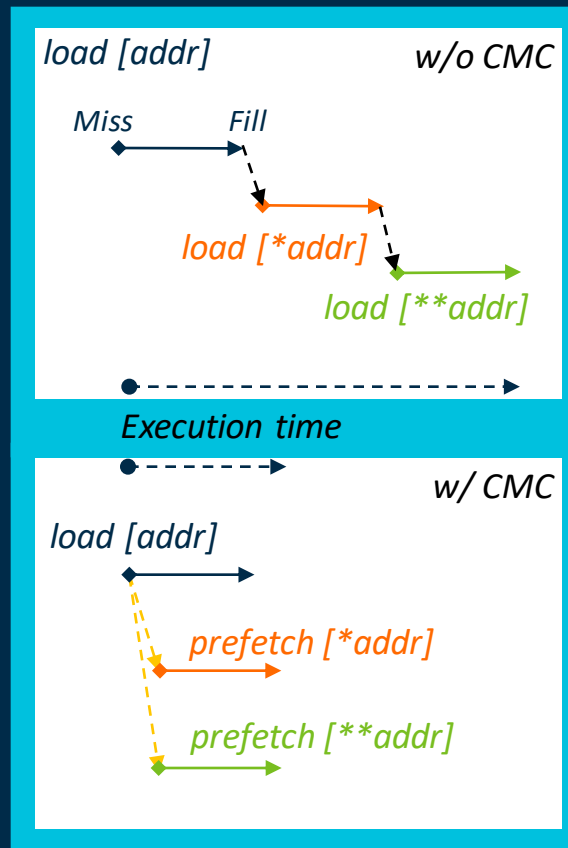
# Neoverse N2 Core Pipeline Upgrades: Mid-Core & Back-End

Mid-Core	Neoverse N1	Neoverse N2	Improvement
Rename width	4 instrs	5 instrs	<b>1.2x</b>
Rename Checkpointing	N	Y	<b>New</b>
ROB size	128	160+	<b>1.25x</b>
ALUs	3	4	<b>1.33x</b>
Branch resolution	1 per cycle	2 per cycle	<b>2x</b>
Overall Pipeline Depth	11 cycles	10 cycles	<b>1.1x</b>
Back-End	Neoverse N1	Neoverse N2	Improvement
L1 Data cache	64KB	64KB	-
L2 Cache	Private 512KB / 1MB	Private 512KB / 1MB	-
AGUs	2-LD/ST	2-LD/ST + 1 LD	<b>1.5x</b>
L1 LD Hit bandwidth	2x 16B/cycle	3x 16B/cycle	<b>1.5x</b>
Store data B/W	16B / cycle	32B / cycle	<b>2x</b>
L2 bandwidth	32B read + 32B write	64B read + 64B write	<b>2x</b>
L2 Transactions	48	64	<b>1.3x</b>
Data Prefetch Engines	Stride, spatial/region and stream	N1 + temporal, stride, and tablewalk improvements	<b>New and improved</b>

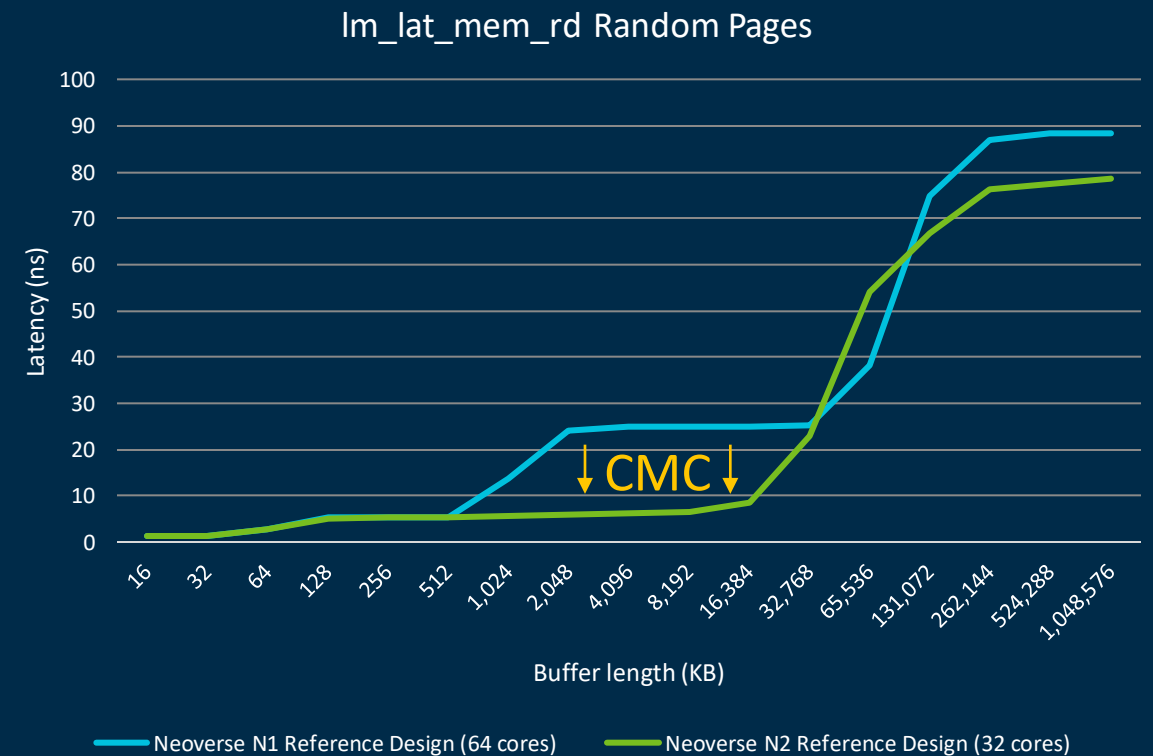
# Arm Neoverse N2: Prefetchers and Caches

## Correlated Miss Caching (CMC) prefetching

- Temporal prefetcher can prefetch arbitrary, repeating access patterns to accelerate workloads with temporal streams, such as pointer chasing
- Records hard to predict memory patterns that the prefetcher can leverage when needed

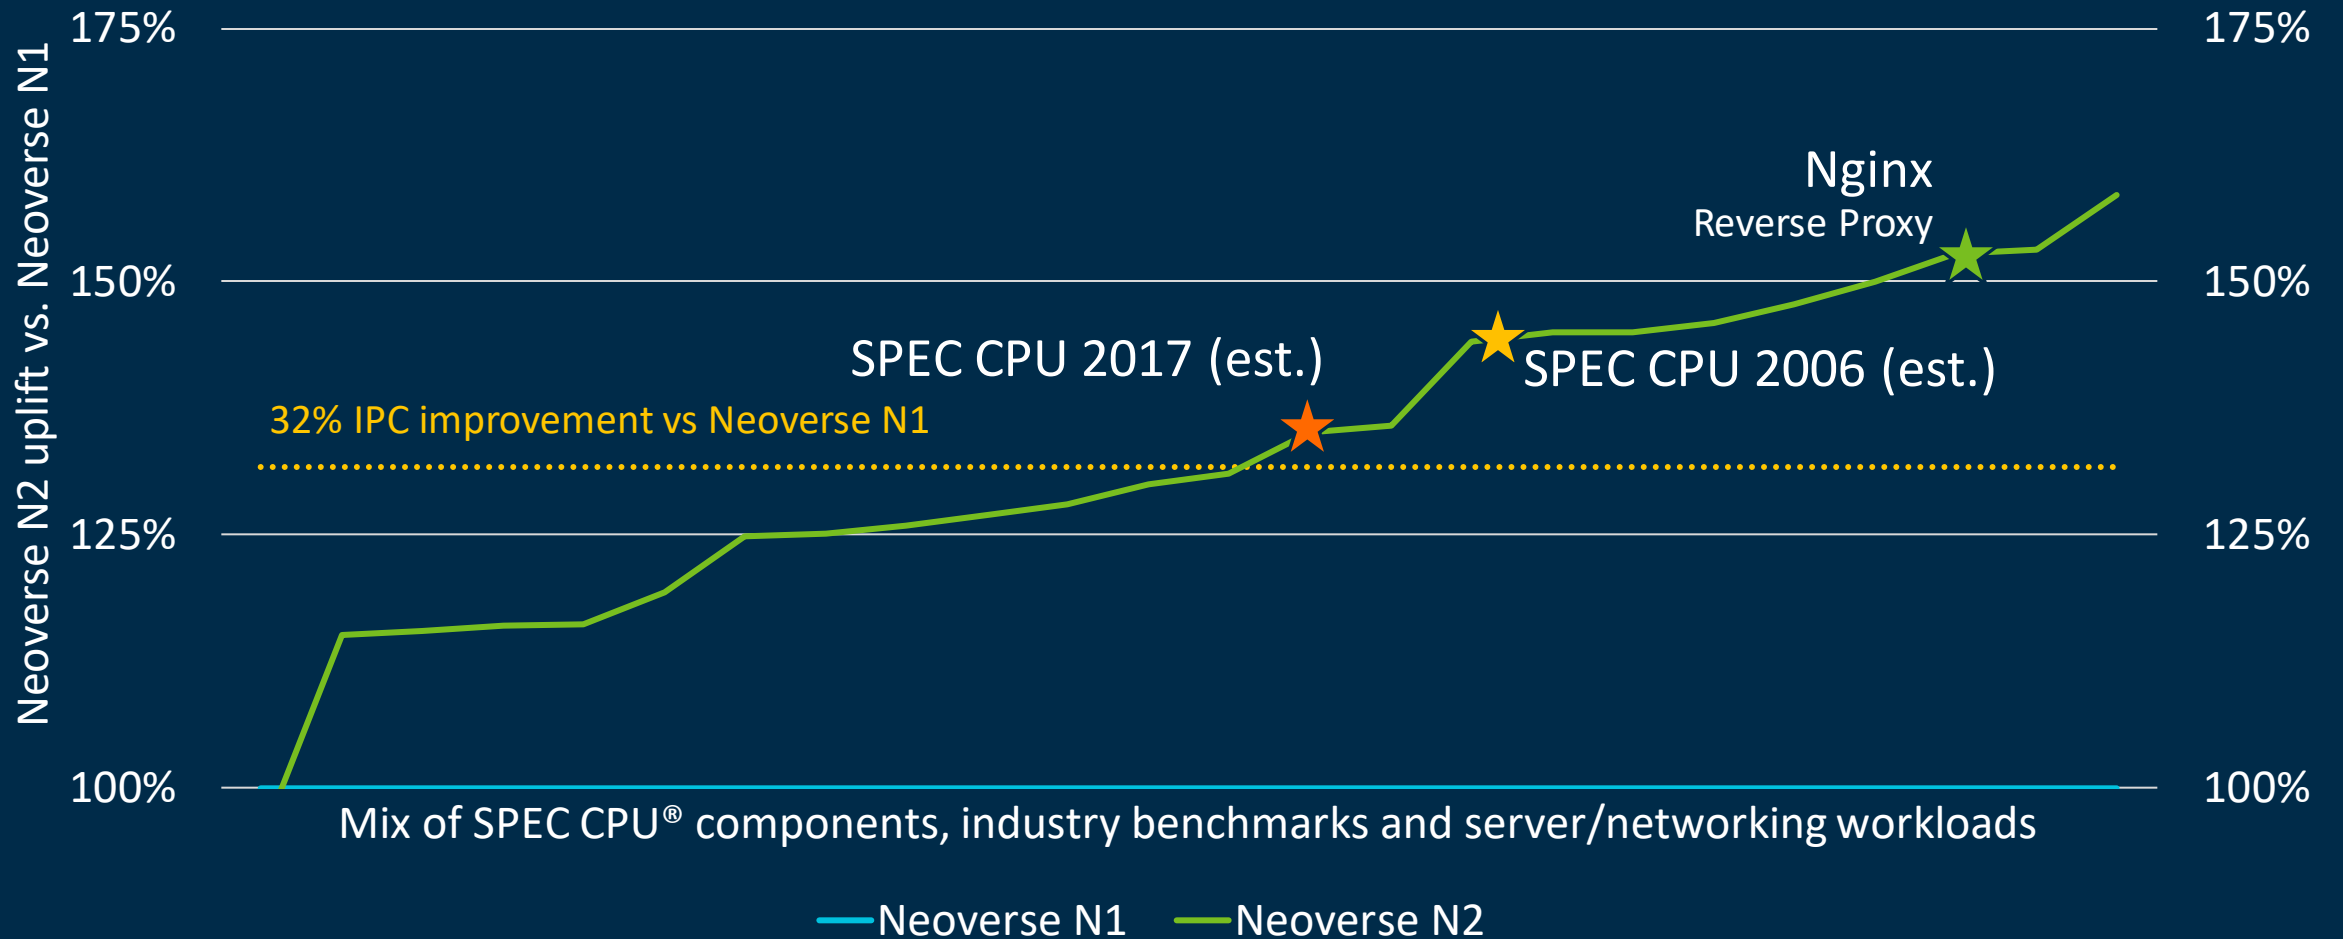


- This technology can significantly reduce load-to-use latency and increase the leverage on the L2 cache



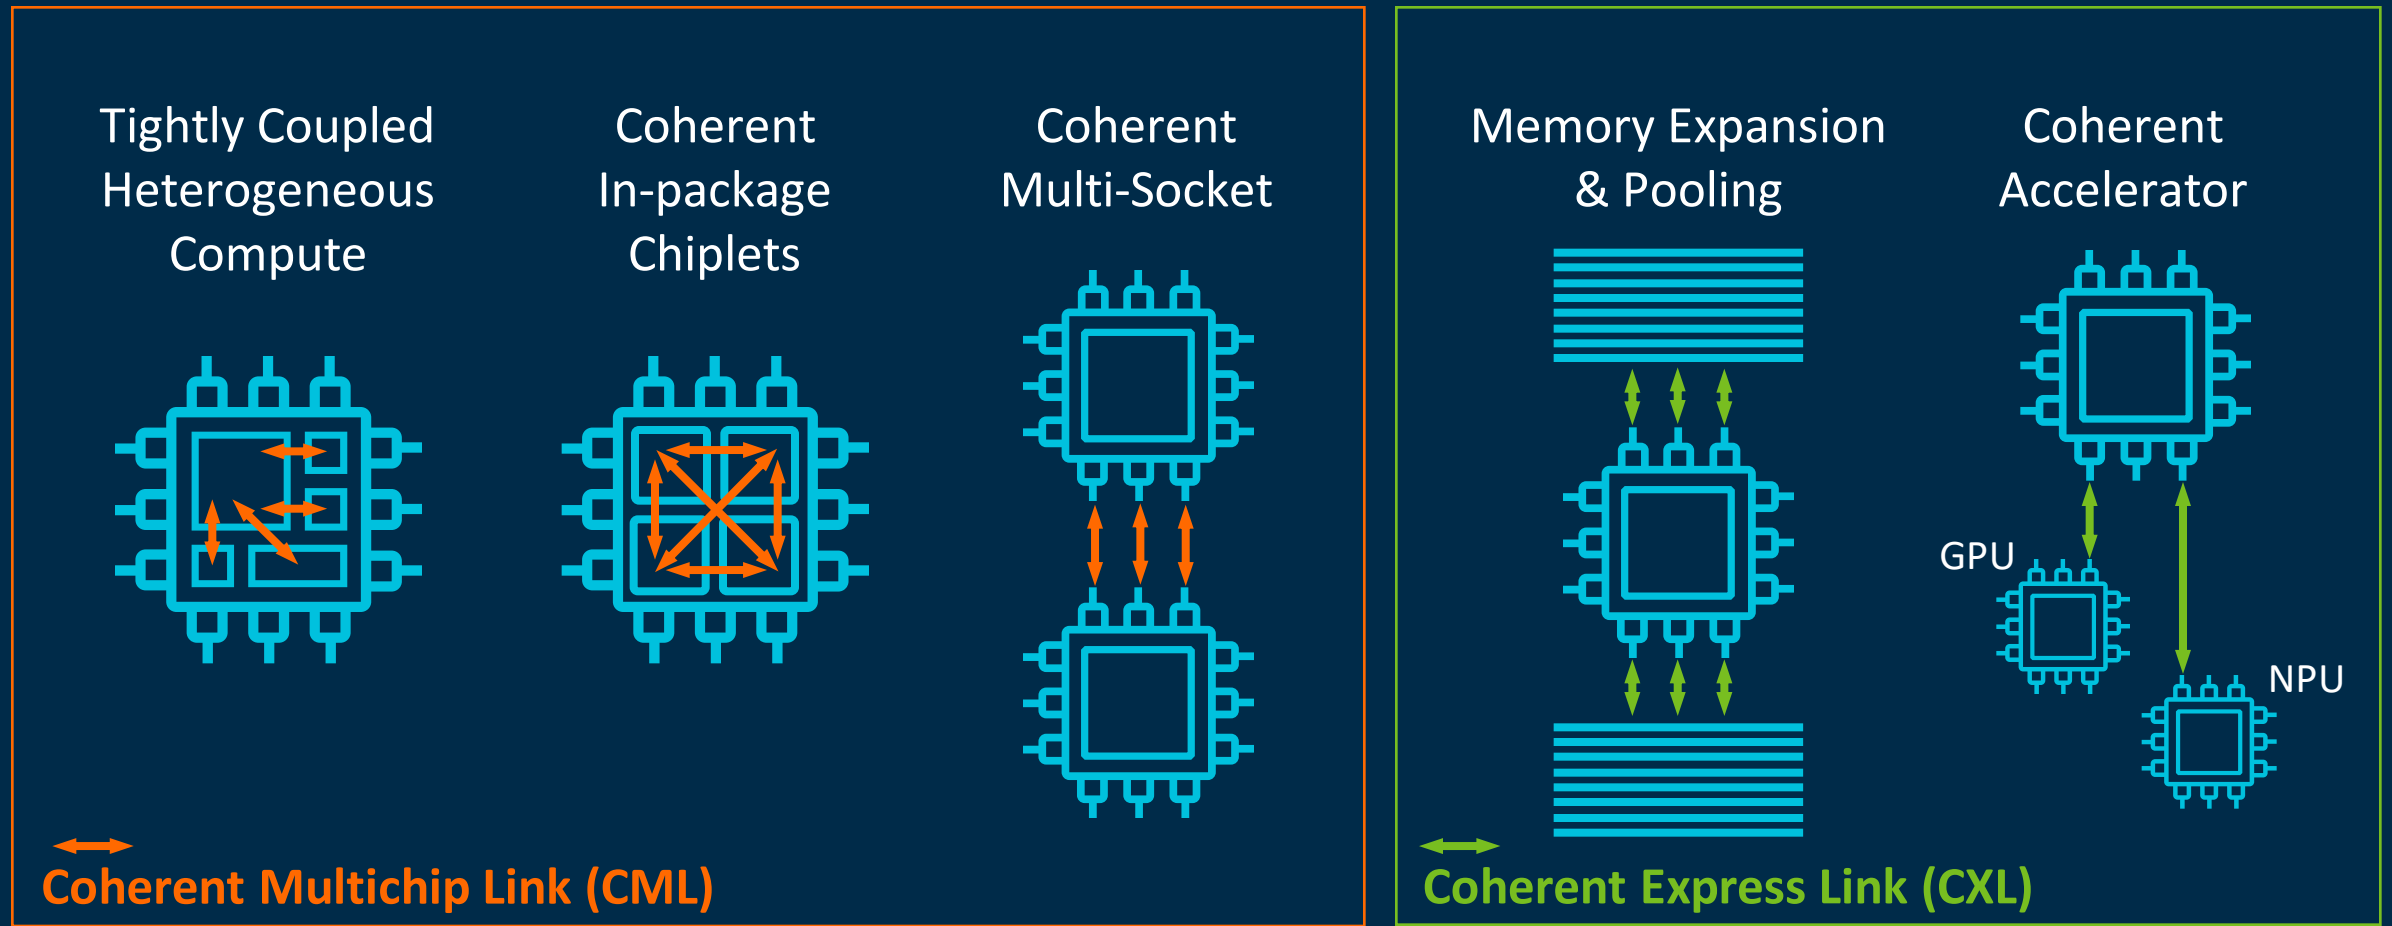
# Arm Neoverse N2 IPC Performance Uplift vs. Neoverse N1

ISO-frequency, unconstrained memory subsystem, single core



# Multichip and CXL Leadership for the Intelligent Era

CMN-700 enables next-gen use cases for multi-chip, memory expansion and accelerators





# Arm Neoverse CMN-700

## + Step Function in Performance on Every Vector

- Larger mesh, double wide channels
- **3x** increased cross-sectional BW
- Hot spot reroute capability
- CHI enhancements for optimized data movement
- Upgraded IO bandwidth (PCIe Gen 5)

## + Composability for a Customized Datacenter

- CXL host, device and memory expansion support
- Enables homogeneous and heterogeneous topologies
- Optimized multi-protocol gateway (CXL and CML)

Platform Capabilities	CMN-600	CMN-700	Improvement
# cores supported per die / system	64 <sup>1</sup> / 128	256 / 512	<b>4x</b>
System Level Cache (SLC) size per die	128MB	512MB	<b>4x</b>
Nodes (cross points) per die	64 (8x8)	144 (12x12)	<b>2.25x</b>
# memory device ports (ex, DRAM, HBM) per die	16	40	<b>2.5x</b>
Multichip ports per die	4	32	<b>8x</b>

# Composable Datacenter SoCs

## Example multi-die SoC with Super Home Node


### Super Home Node Functions and Composability

- Home node for local memory
  - System level cache (SLC)
- Cluster cache for remote homed memory
  - Local coherency node and cache (LCC)
- Can be configured as Unified (SLC+LCC) or Decoupled (LCC and SLC)

### Composability (Multichip and CXL attached)

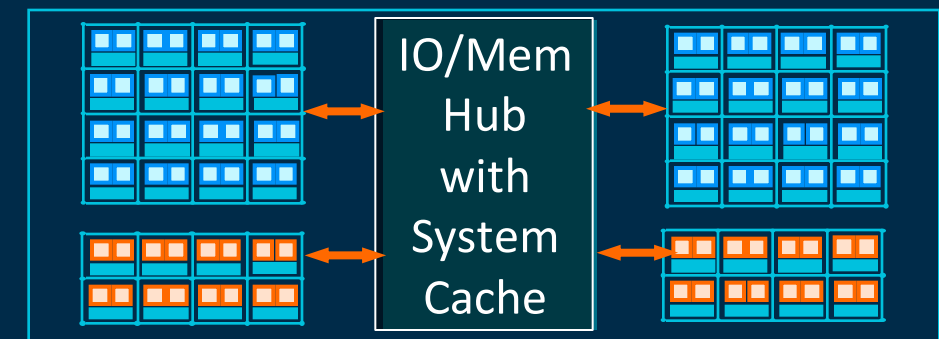
- Unified SLC+LCC for **Homogeneous compute**
- Decoupled LCC for compute chiplets and HNF/SLC in **Heterogeneous Hub & Spoke**
- Home Node for CXL Host attached memory and LCC for Host managed Device memory (HDM)
- CXL Device Coherency Node (DCOH) for CXL Device

### Homogeneous Compute



Performance optimized

### Heterogeneous Hub & Spoke

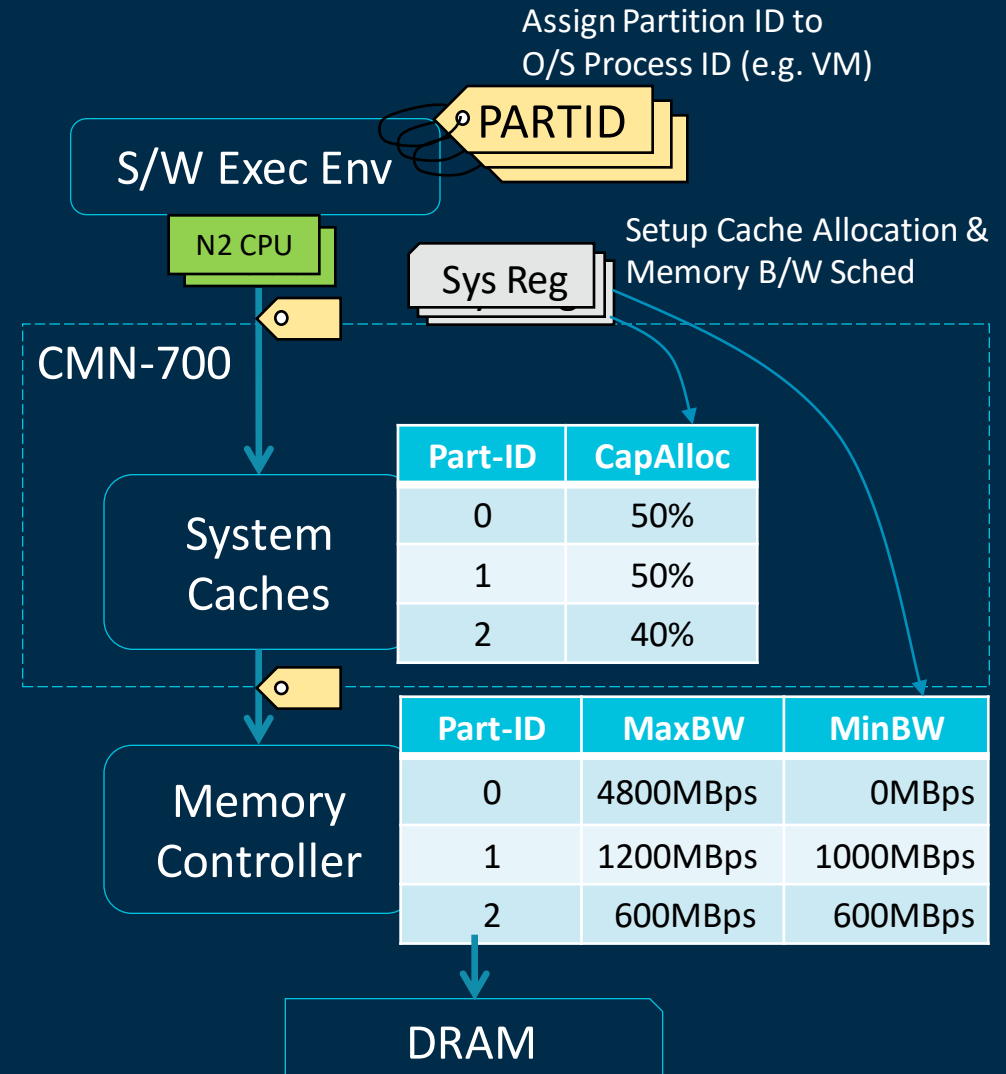


Flexible and re-use optimized

# MPAM: Memory Partitioning and Monitoring

MPAM bounds and monitors process interaction/interference in shared resources

- All memory requests issued by Processing Elements in the system are decorated with MPAM information:
  - Neoverse Cores
  - IO Devices can use the SMMU v3.2 or default MPAM values
- SW can measure resource usage matching two fields:
  - PARTID & PMG – Partition IDs and Perf Monitor Group
    - PMG can be used to probe whether a certain application within a PARTID should have its own PARTID
  - For Memory BW, read and write bandwidth are assessed independently
  - Can independently monitor memory requests for code and data – can have separate PARTIDs and PMGs
- Memory System Components provide controls for capacity or bandwidth

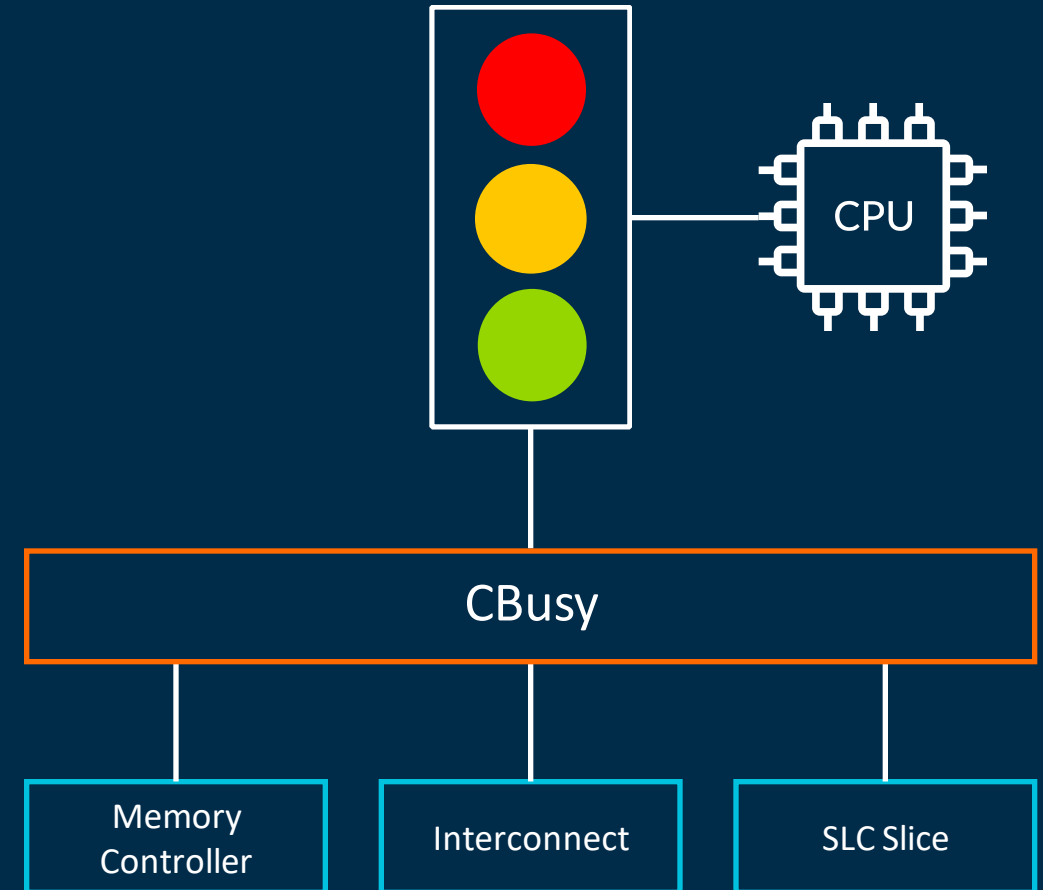


# CBusy: Scalable and Sustainable Performance

Memory bandwidth management is paramount to extract optimal and predictable performance on high core count, high performance systems

## Completer Busy (CBusy)

- Automatic regulation of CPU prefetch requests based on system congestion
- CPU aggregates and filters feedback from system
- Feedback affects hardware prefetcher aggressiveness
- CPU can also throttle its maximum outstanding transactions
- MPAM aware CBusy for VM based bandwidth control




# CBusy: Dynamic Prefetcher Throttling

Neoverse N2 cores support four levels of prefetcher aggressiveness:

- The most aggressive mode can generate prefetches with as low as 20% accuracy
- The most conservative mode will target 100% accuracy

When surplus bandwidth is available, more conservative prefetching results in significant performance loss (*shown in orange in the graph*)

Performance Differential vs. Static Aggressive Prefetch

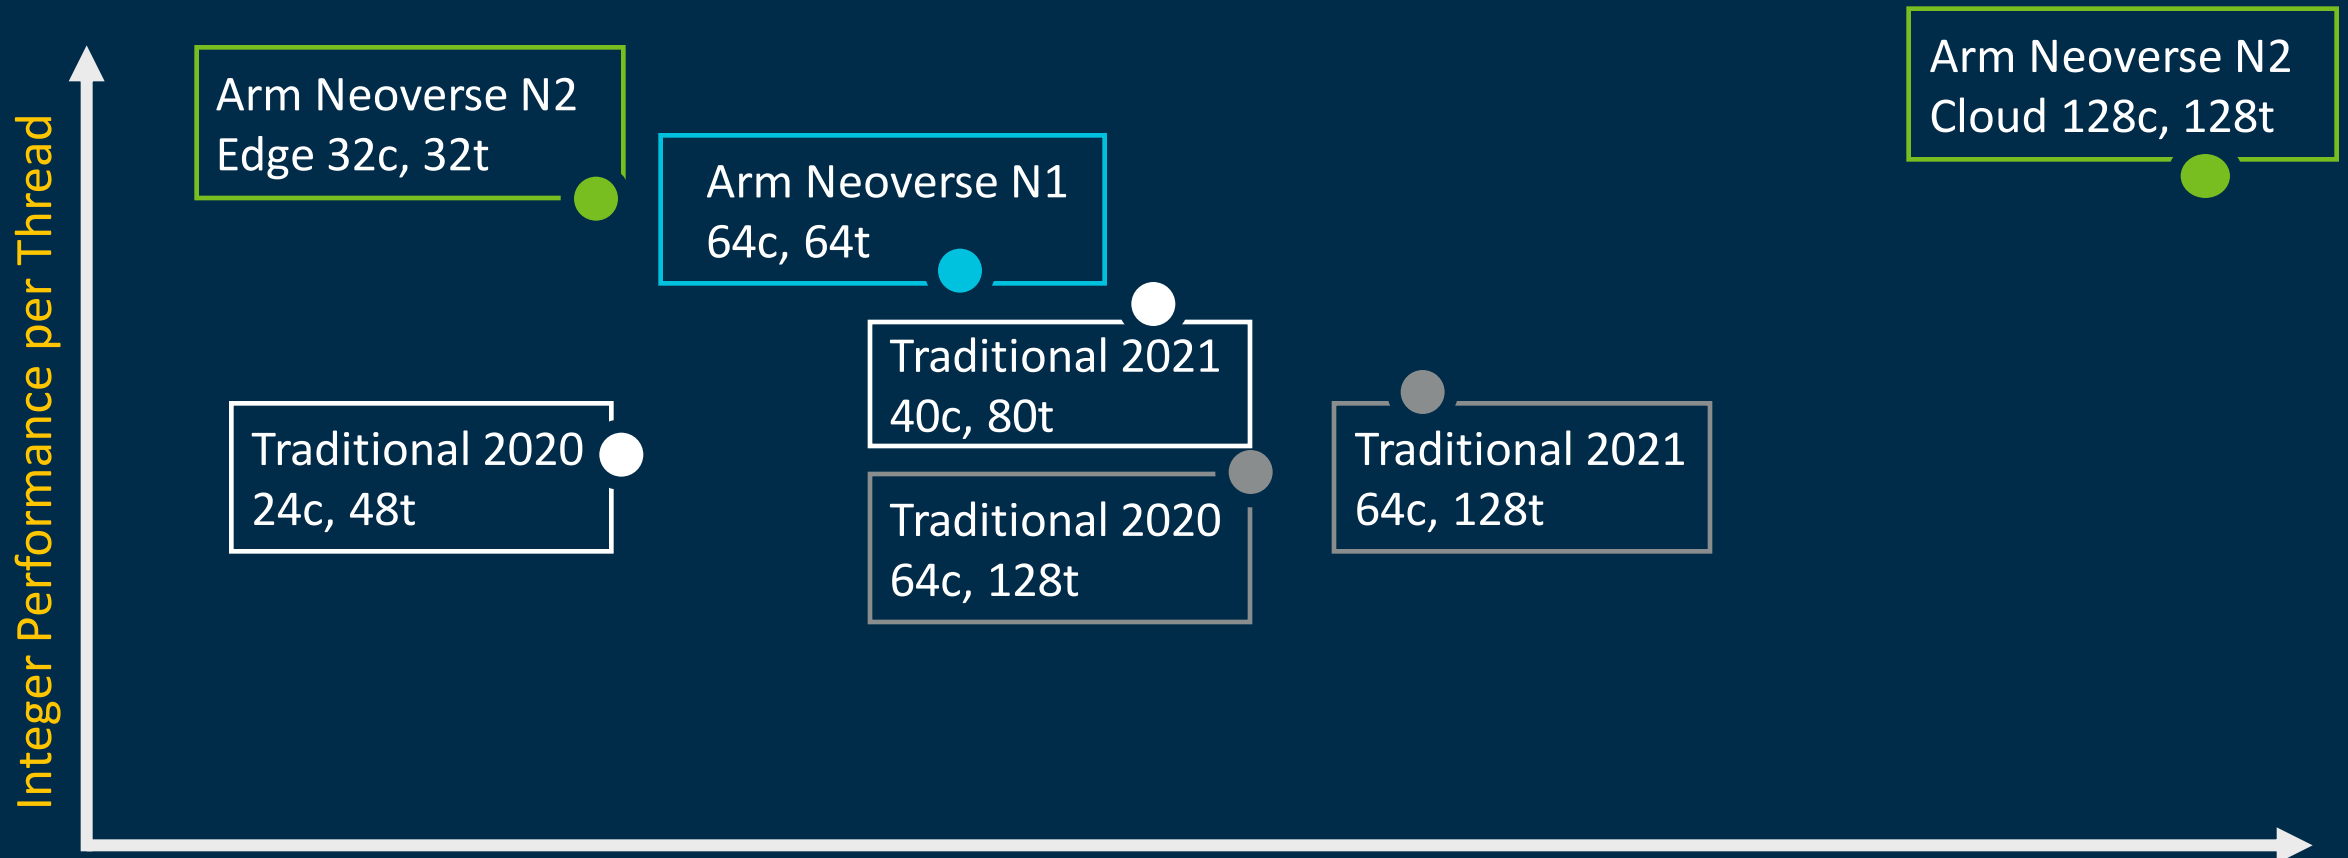


When memory BW is highly contended, higher performance can be achieved by avoiding more aggressive, less accurate prefetches

CBusy dynamically tunes prefetcher aggressiveness based on BW available, often resulting in best performance (*shown in blue in the graph*)

# Arm Neoverse N2: Performance Without Compromise

Hyper-threading won the datacenter, per-thread performance will win the cloud



Neoverse N2 Cloud: 128 cores, 3.0GHz, 8xDDR5-4800  
 Neoverse N2 Edge: 32 cores, 2.7GHz, 4xDDR4-3200


Integer Performance per Socket

Traditional 2020 is measured by Arm  
 Traditional 2021 data is projected by Arm  
 Arm Neoverse performance data is estimated by Arm


# Arm Neoverse N2: Relentless Improvements on Real Workloads

ISO process/frequency/core count (8 cores), Neoverse N2 significantly outperforms N1


## MemcachedD



## MySQL



## NGINX



*Approximate projections on reduced workload configurations on pre-silicon models. Server configuration under test limited to 8 cores for both configurations. Performance on silicon systems is subjected to change due to different SW configurations and partners' choices on HW implementations.*

# Arm Neoverse N2: Cloud-to-Edge Perf-per-Watt Leadership

- ✦ Comprehensively upgraded microarchitecture delivers +40% IPC uplift over Neoverse N1 at ISO process & frequency
- ✦ First Armv9 core supporting SVE2, Arm's state-of-the-art vector ISA
- ✦ Continues Arm industry leadership in performance / watt
- ✦ Up to 256-cores and 512MB of system-level cache per die
- ✦ Coming soon from Arm's partners who have the design freedom to add custom acceleration and tune system configurations for market leadership



# Performance and Benchmark Disclaimer

- This benchmark presentation made by Arm Ltd and its subsidiaries (Arm) contains forward-looking statements and information. The information contained herein is therefore provided by Arm on an "as-is" basis without warranty or liability of any kind. While Arm has made every attempt to ensure that the information contained in the benchmark presentation is accurate and reliable at the time of its publication, it cannot accept responsibility for any errors, omissions or inaccuracies or for the results obtained from the use of such information and should be used for guidance purposes only and is not intended to replace discussions with a duly appointed representative of Arm. Any results or comparisons shown are for general information purposes only and any particular data or analysis should not be interpreted as demonstrating a cause and effect relationship. Comparable performance on any performance indicator does not guarantee comparable performance on any other performance indicator.
- Any forward-looking statements involve known and unknown risks, uncertainties and other factors which may cause Arm's stated results and performance to be materially different from any future results or performance expressed or implied by the forward-looking statements.
- Arm does not undertake any obligation to revise or update any forward-looking statements to reflect any event or circumstance that may arise after the date of this benchmark presentation and Arm reserves the right to revise our product offerings at any time for any reason without notice.
- Any third-party statements included in the presentation are not made by Arm, but instead by such third parties themselves and Arm does not have any responsibility in connection therewith.

# End Notes

## Slide Title: Arm Neoverse N2: Performance Without Compromise

- Traditional 24c/48t CPU is Intel Xeon 8268
- Traditional 64c/128t CPU is AMD EPYC 7742
- Traditional next 40c/80t CPU is projected on a 40 core Intel Ice Lake system
- Traditional next 64c/128t CPU is projected on AMD EPYC 7763
- SMT and Turbo are enabled for measurements on both x86 SKUs
- GCC-snapshot-10. 2.1- fsf-10.100, 2MB Page size, unless noted otherwise
- Compiler Flags for Neoverse N1: `-march=armv8.2-a+crypto -O3 -ftree-vectorize -ffast-math`
- Compiler Flags for Neoverse N2: `-march=armv8.6-a+crypto+fp16+dotprod+sve -O3 -ftree-vectorize -ffast-math`
- Compiler Flags for Traditional systems: `-mcpu=native -O3 -ftree-vectorize -ffast-math`



# arm NEOVERSE

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

Thank You!