arm

As the volume and variety of connected devices continues to soar, network and compute infrastructure needs are evolving fast. To manage this explosion of traffic requires transformation from the cloud all the way to the edge.

Real-time decision making at the network edge will require new levels of compute, as well as new types of compute with accelerators and offload—with even more constraints on power and space.

Arm Neoverse E1

The Arm Neoverse E1 delivers the efficiency, ecosystem support, and high scalability from edge to core that's required to maximize throughput compute efficiency.

Optimized Performance for Today's Workloads

✛ 2.7x throughput performance
✛ 2.4x throughput efficiency
✛ 2.1x compute performance[1]

Intelligent design for highly efficient throughput

Arm Neoverse E1 was architected from the ground up to balance compute and efficiency requirements and deliver maximum throughput.

✛ Proven Armv8.2-A AArch64 CPU is aligned with Neoverse N1 CPU
✛ Advanced CPU cluster technology supports multiple performance domains, streamlines traffic across bridges, and offers scalable interfaces for edge to cloud applications
✛ Advanced power management features enhance efficiency
✛ Simultaneous multithreading supports two threads concurrently running at different exception levels, OSes, and more

Builds on the diverse Arm software ecosystem

Arm is focused on delivering optimized solutions built on a common software ecosystem for cloud to edge applications.

✛ Enables OSS, virtualized software stacks on traditional blackbox networking equipment versus fixed-function chip/offload
✛ Scalable for a wide range of devices utilizing a common AArch64 architecture and robust reference design
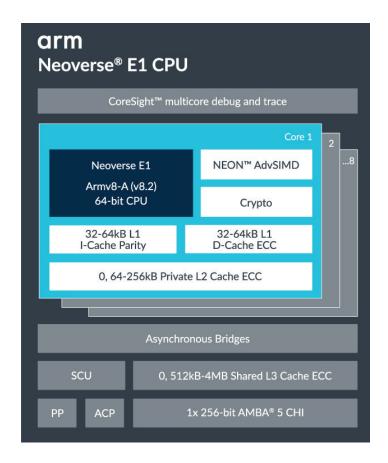
[1]Relative to Cortex A53

Highly scalable throughput for edge to core data transport

Arm Neoverse E1 is ideally suited for data plane compute workloads such as 4G/5G transport, software defined networking, software defined storage, and SD-WAN. This innovative platform features a scalable architecture suitable for 10Gb wireless/wireline devices to high-performance 100G+ Dataplane Processing Unit (DPU).

✤ Flexible architecture enables fixed-function accelerator integration via low-latency Accelerator Coherency Port

✤ High efficiency design supporting 25Gbps+ throughput in less than 4W power budget

✤ Common Armv8.2-A architecture enables integration with Neoverse N1 platform for high-performance, multi-port 100Gbps networking devices

Specifications and Features

✤ Simultaneous Multithreading (SMT) supporting two threads concurrently

✤ Up to 8 cores (16 threads) per cluster

✤ Superscalar, out-of-order pipeline

✤ Configurable private L2 cache

✤ Configurable L3 cache

✤ Low-latency Accelerator Coherency Port (ACP) for closely coupled accelerator integration

✤ Support cache stashing into L2/L3 cache

arm
Neoverse® E1 CPU

CoreSight™ multicore debug and trace

Core 1

Neoverse E1
Armv8-A (v8.2)
64-bit CPU

NEON™ AdvSIMD

Crypto

32-64kB L1
I-Cache Parity

32-64kB L1
D-Cache ECC

0, 64-256kB Private L2 Cache ECC

2

...8

Asynchronous Bridges

SCU

0, 512kB-4MB Shared L3 Cache ECC

PP   ACP

1x 256-bit AMBA® 5 CHI