

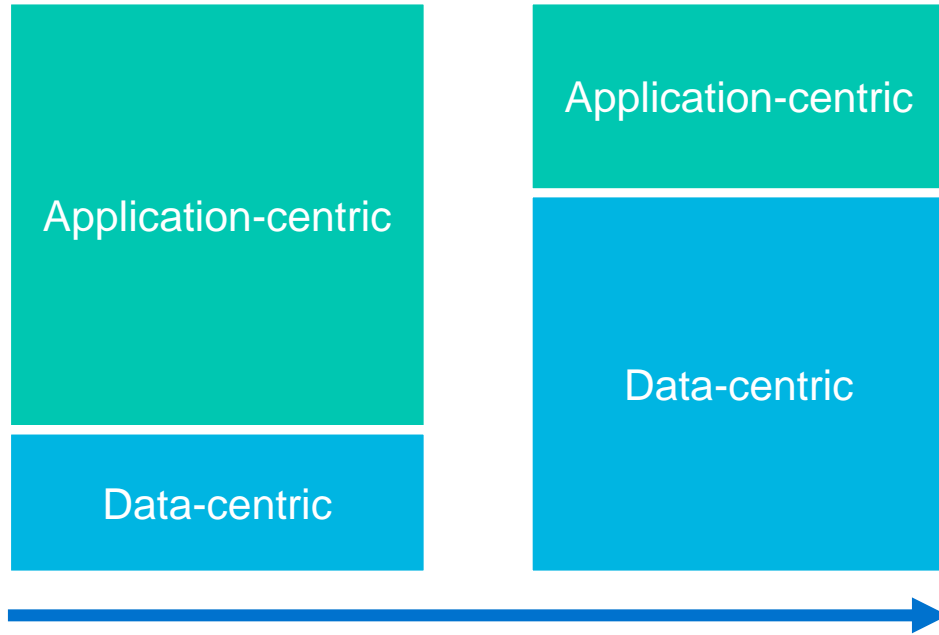


# OCTEON 10 DPU Family

Accelerating the data infrastructure transformation

June 2021

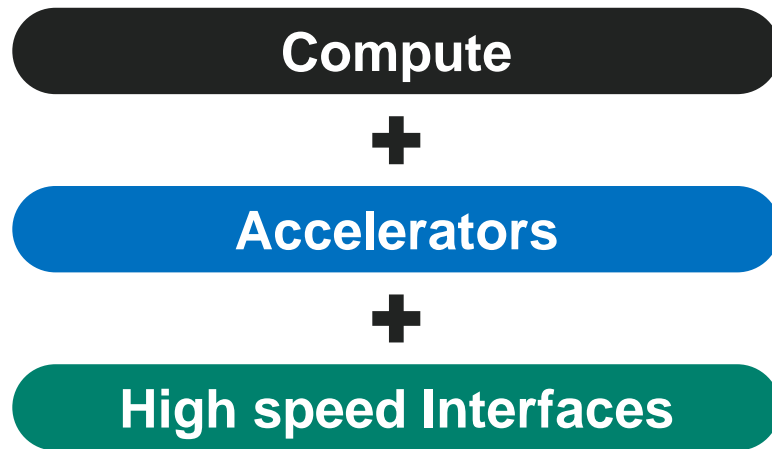
# Workloads are shifting to data-centric compute



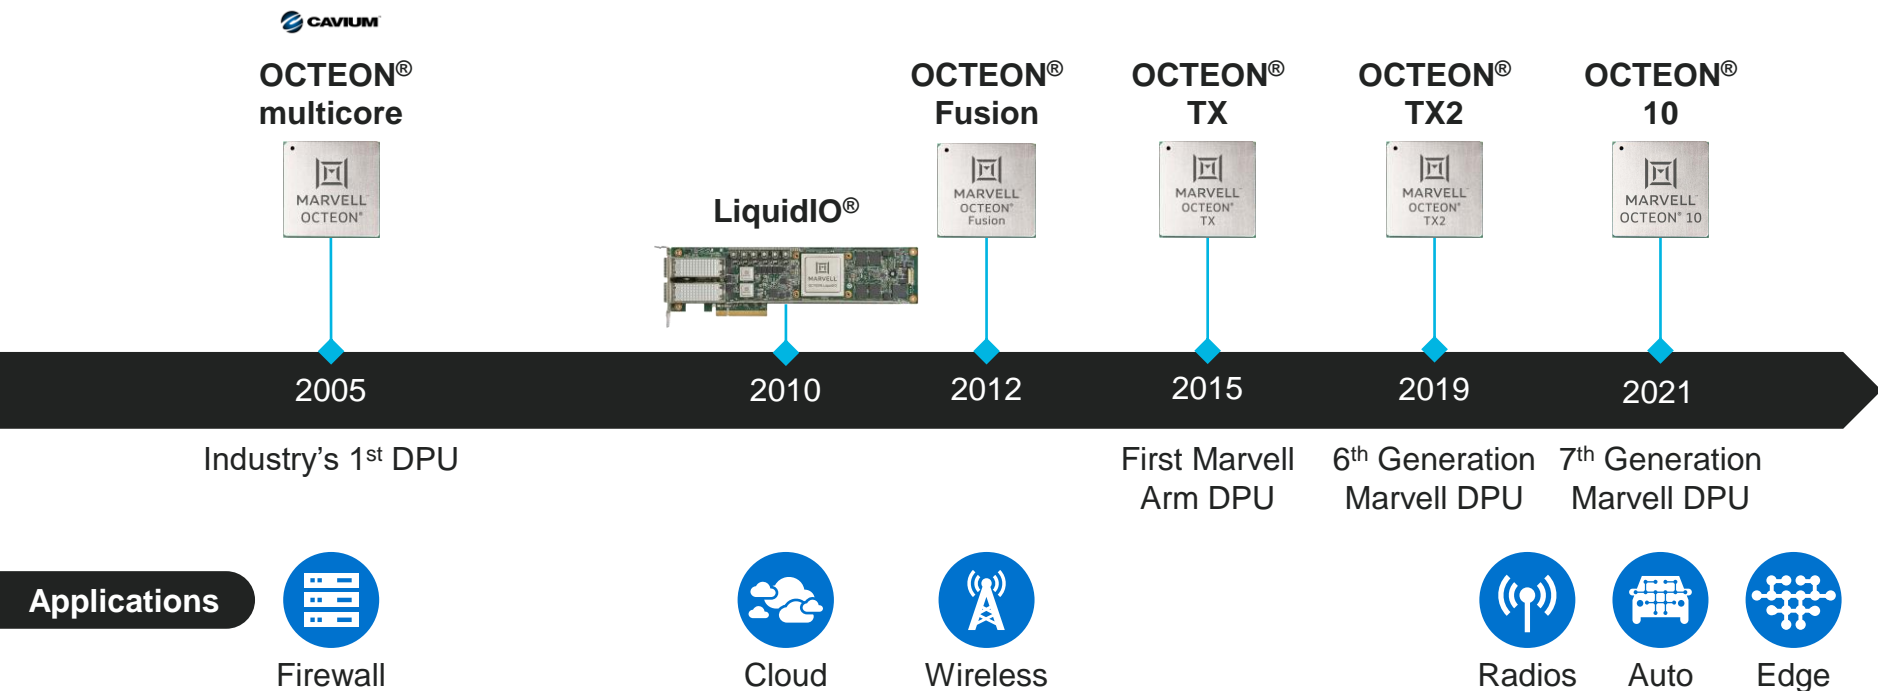
**AI, Networking, security, video and storage virtualization**

# DPU definition


A Data Processing Unit (DPU) is a compute entity that is used to move, process, secure and manage data, as it travels or while at rest, to make it available and optimized for application



# OCTEON: The original DPU platform



Compute leadership  
with Arm Neoverse  
N2 cores



Integrated  
hardware ML  
engine



# OCTEON 10

## Industry Firsts

Based on TSMC  
5nm process



Integrated 1terabit  
switch



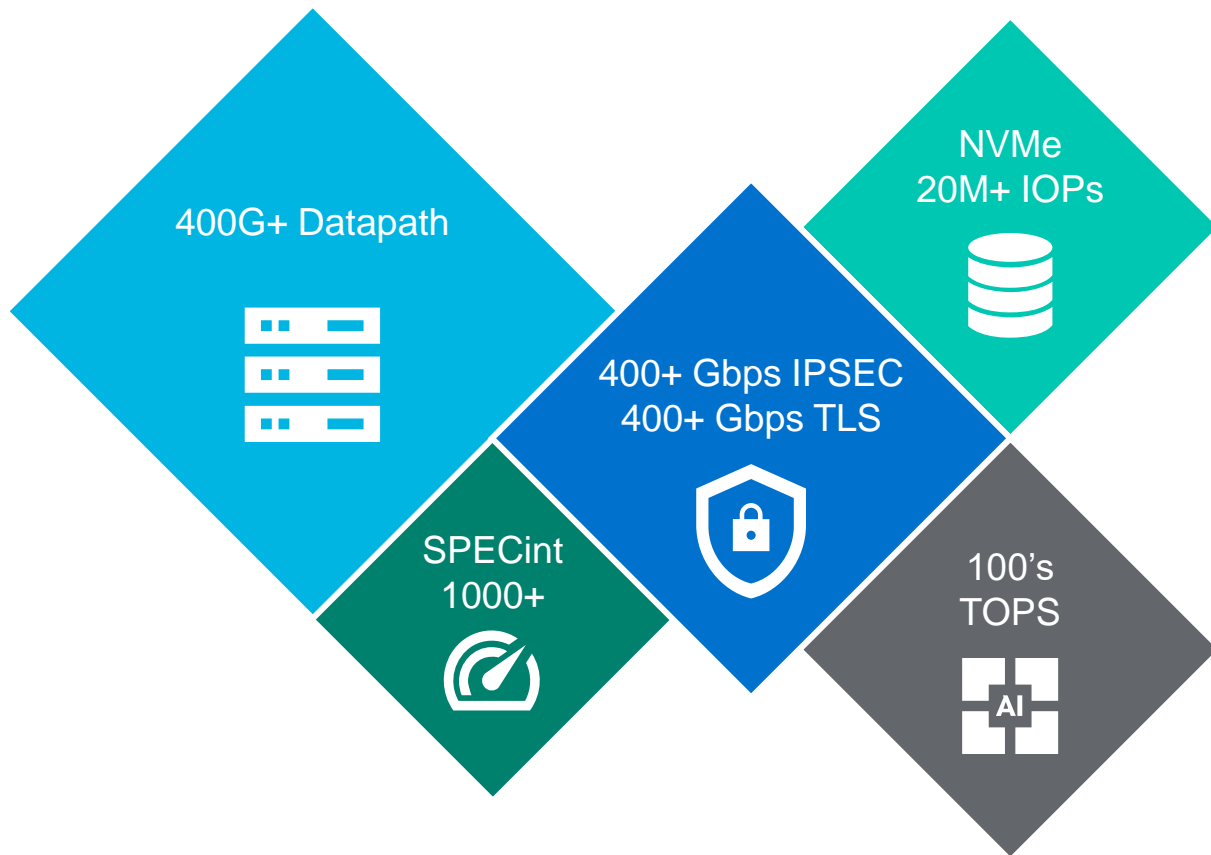
VPP hardware  
acceleration



Advanced inline  
crypto accelerators



**Compute leadership with industry-leading performance per Watt**



# OCTEON 10 Platform

Scalable system  
performance

# OCTEON 10 platform scalability

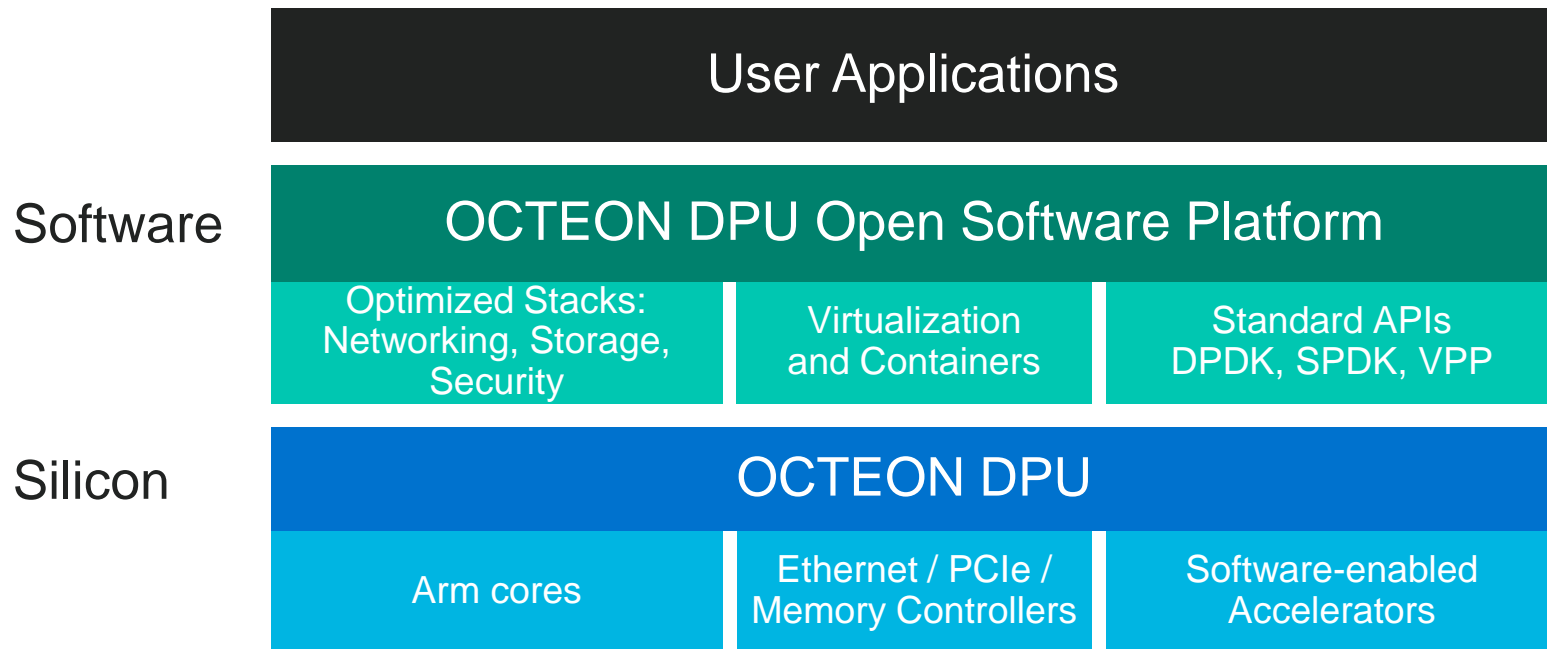


## Scalable features

<b>Ethernet port speed</b>	1G	→	400G
<b>Datapath</b>	50G	→	400G+
<b>Compute SPECint</b>	275	→	1000+
<b>Security IPSEC/SSL</b>	50G	→	400G+

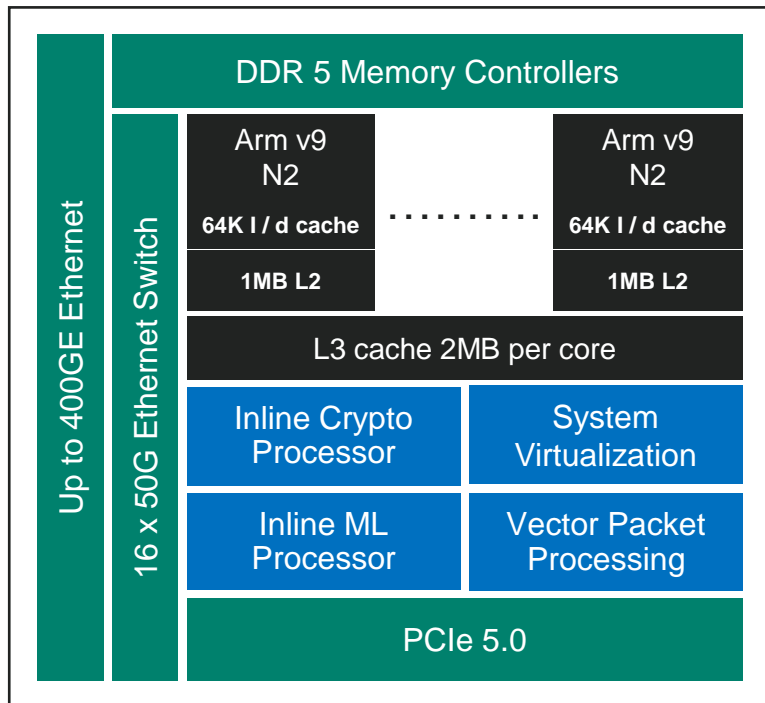
## Common software

# OCTEON DPU platform



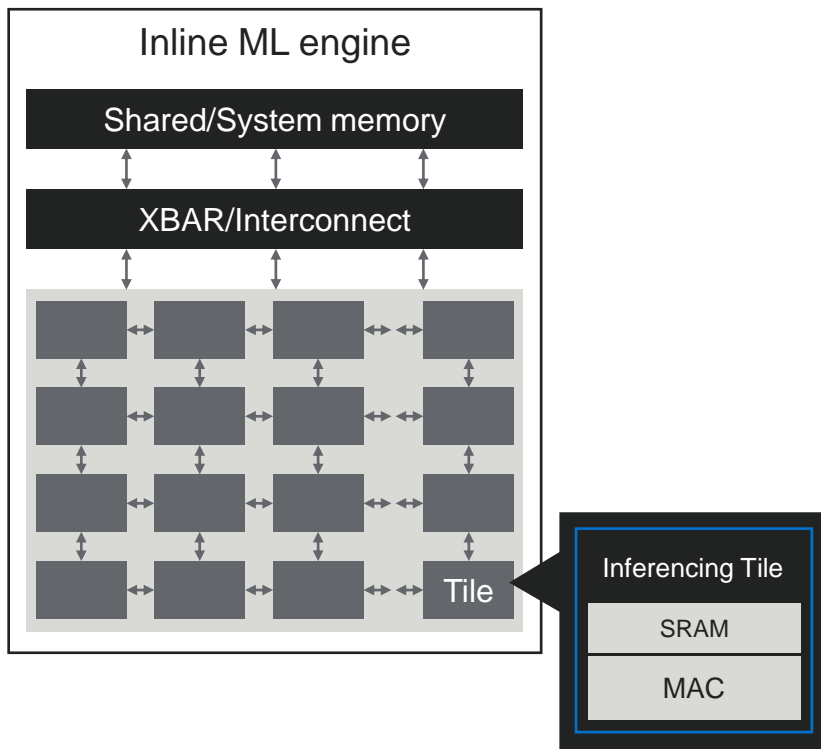


# OCTEON 10 innovations



- 5nm TSMC process
  - Enables fanless designs
- First inline DPU ML Engine
- Hardware VPP acceleration
- Inline crypto processor
- Arm Neoverse N2 cores
  - Highest SPECint in industry
- PCIe 5.0, DDR5 support
- Integrated with 16x 50GE switch
- 56G SerDes

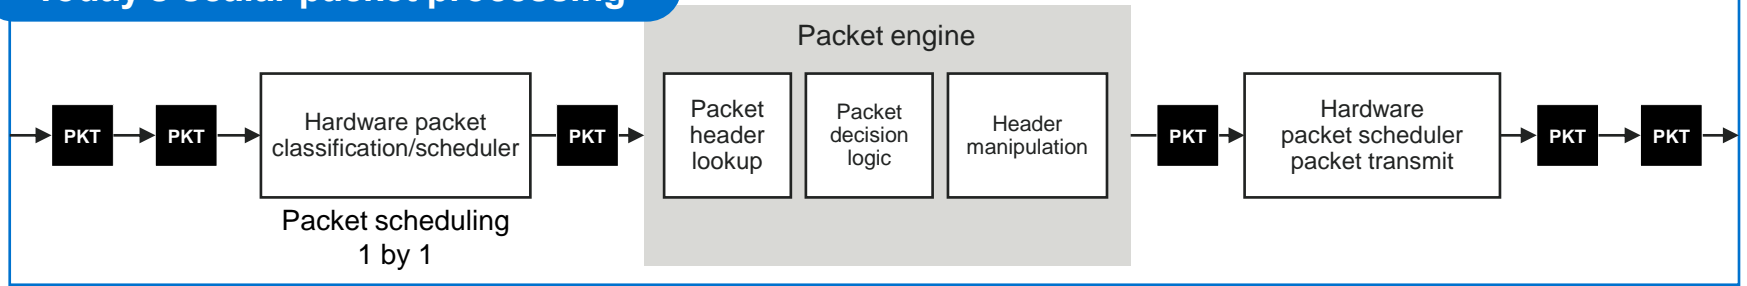
# Integrated ML engine



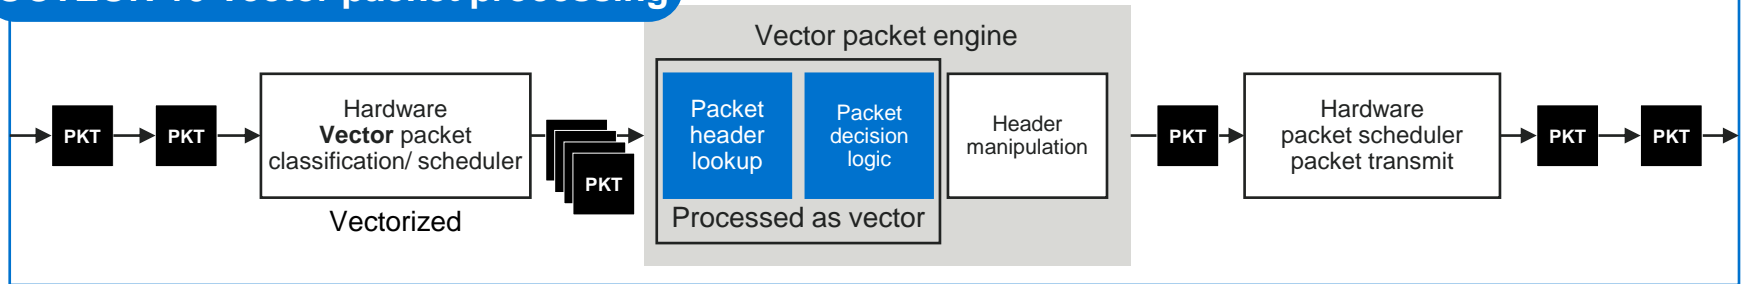
- **Best-in-class DPU inferencing**
  - Directly in the data pipeline
  - Each ML tile contains private SRAM
- **Up to 100x performance vs SW**
  - Supports Int8, FP16
- **Use cases**
  - Threat detection
  - Context-aware service delivery
  - QoS
  - Beamforming optimization
  - Predictive maintenance

# VPP hardware acceleration

## Today's scalar packet processing

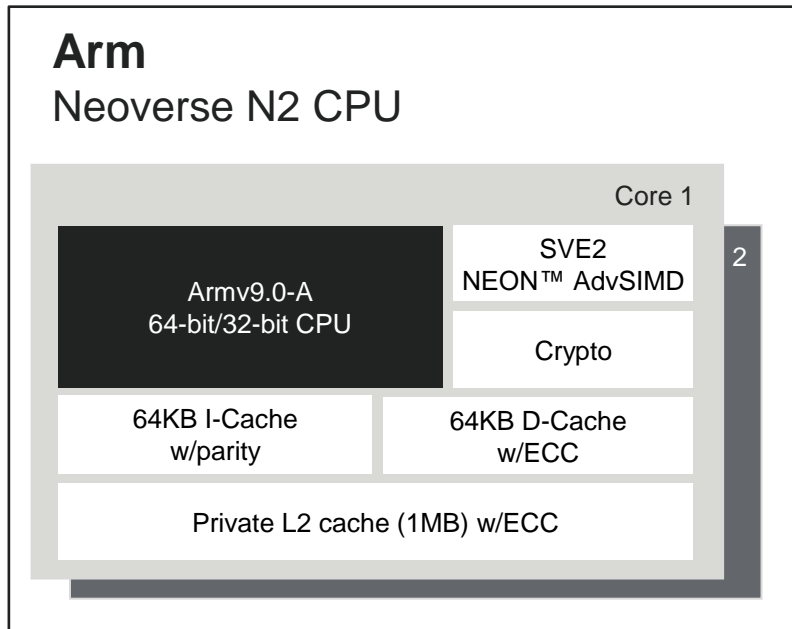


## OCTEON 10 Vector packet processing



**Up to 5X system level performance gains**

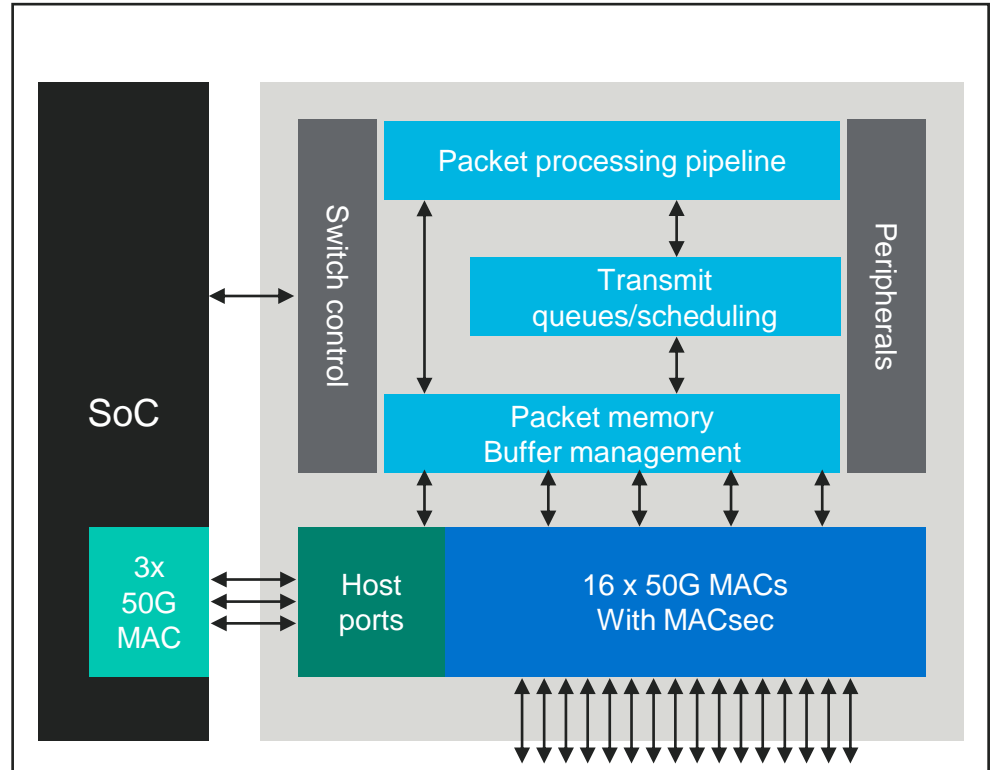
# Arm Neoverse N2 advantages



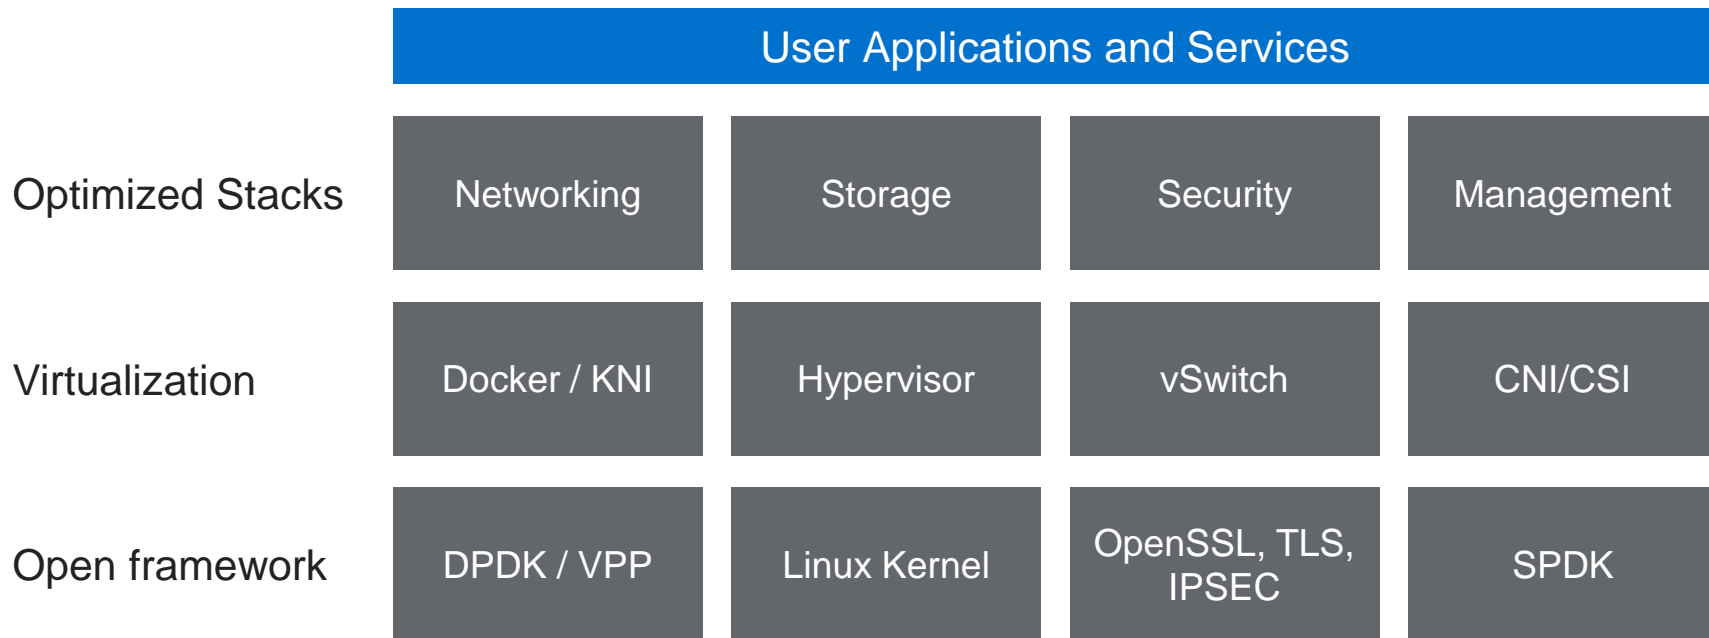
- **Maximizes performance per watt**
- **3x single-threaded performance**
  - Same software runs faster
- **Lower application latency with 1M L2 cache**
- **Higher performance and scalability**
  - SVE2 for hyperscan/DPI and ML support
  - Enhanced cryptography instructions
- **3x latency reduction to HW accelerators**
  - Enabled by hardware scheduling attached cores

# OCTEON 10 Switch integration

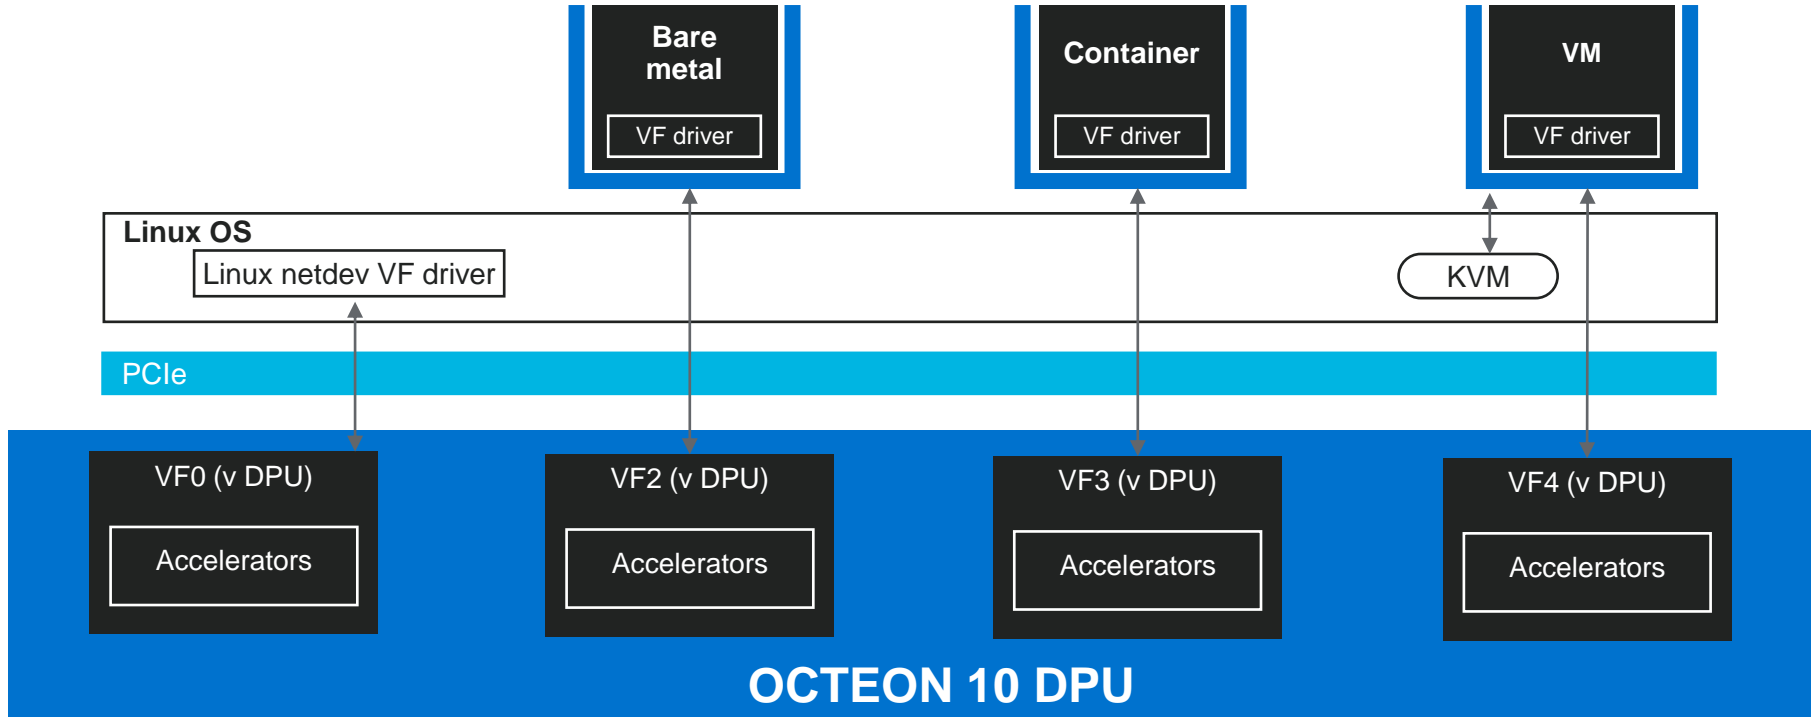
- 1T switch with 16 x 50GE ports
  - Support from 1GE to 100GE
- Feature support
  - 256bit MACsec
  - Network overlay for VxLAN/GRE/MPLS
  - Network analytics: sFlow, IPFIX
  - Flow aware-processing
  - Line rate telemetry
  - TSN timing
- Example use cases
  - 5G: front haul, back haul, side haul
  - Edge switching
  - Enterprise ethernet port fan out



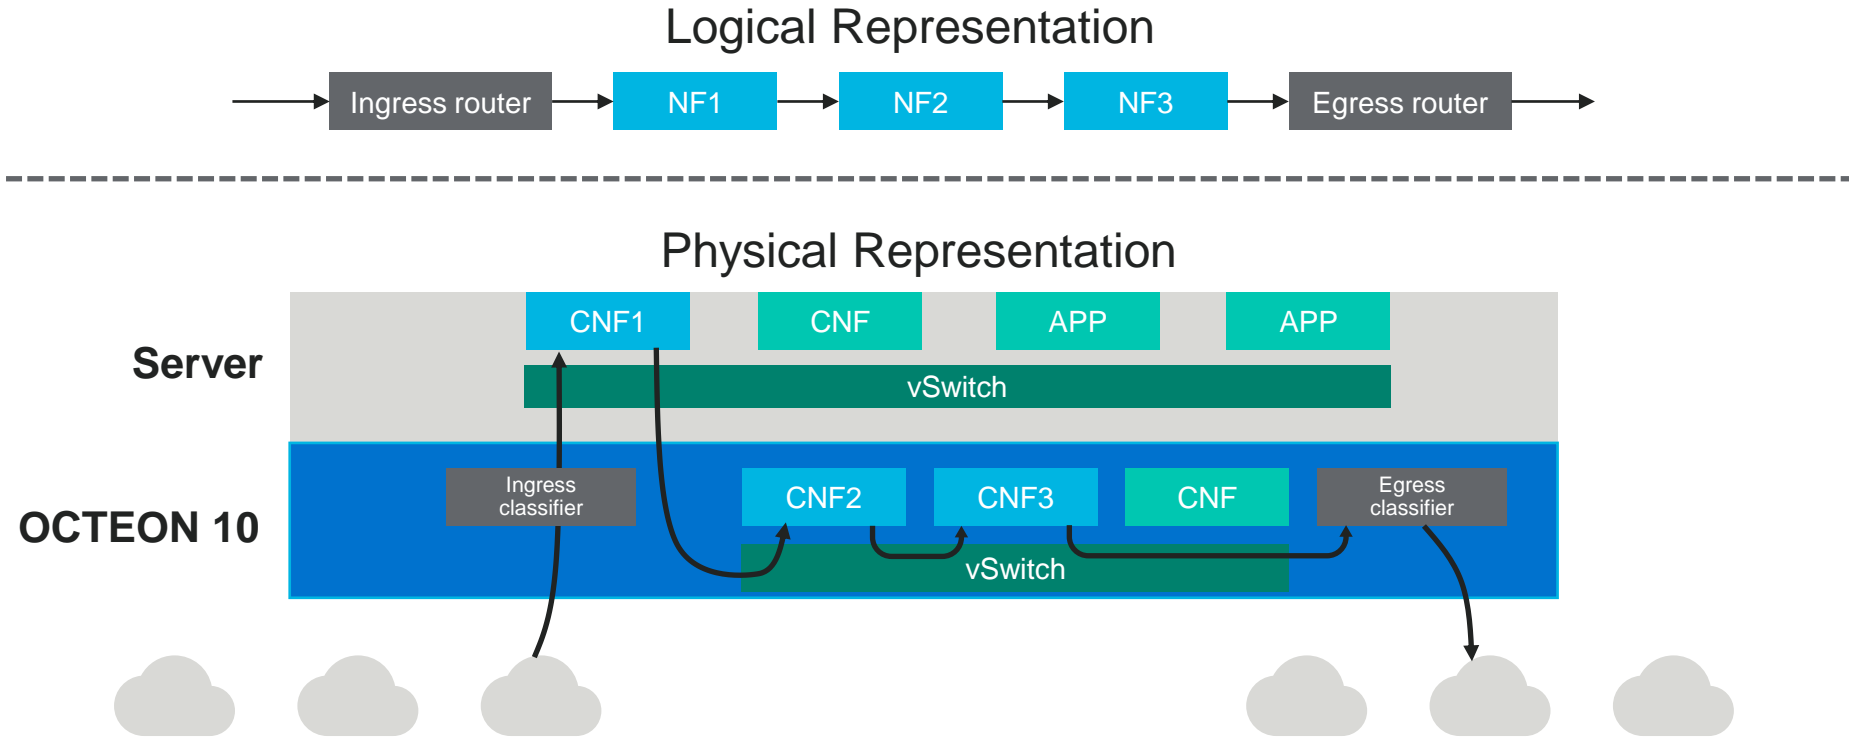
# OCTEON DPU software platform



# Virtualized Accelerators



# Service function chaining





# 4G/5G RAN Architectures

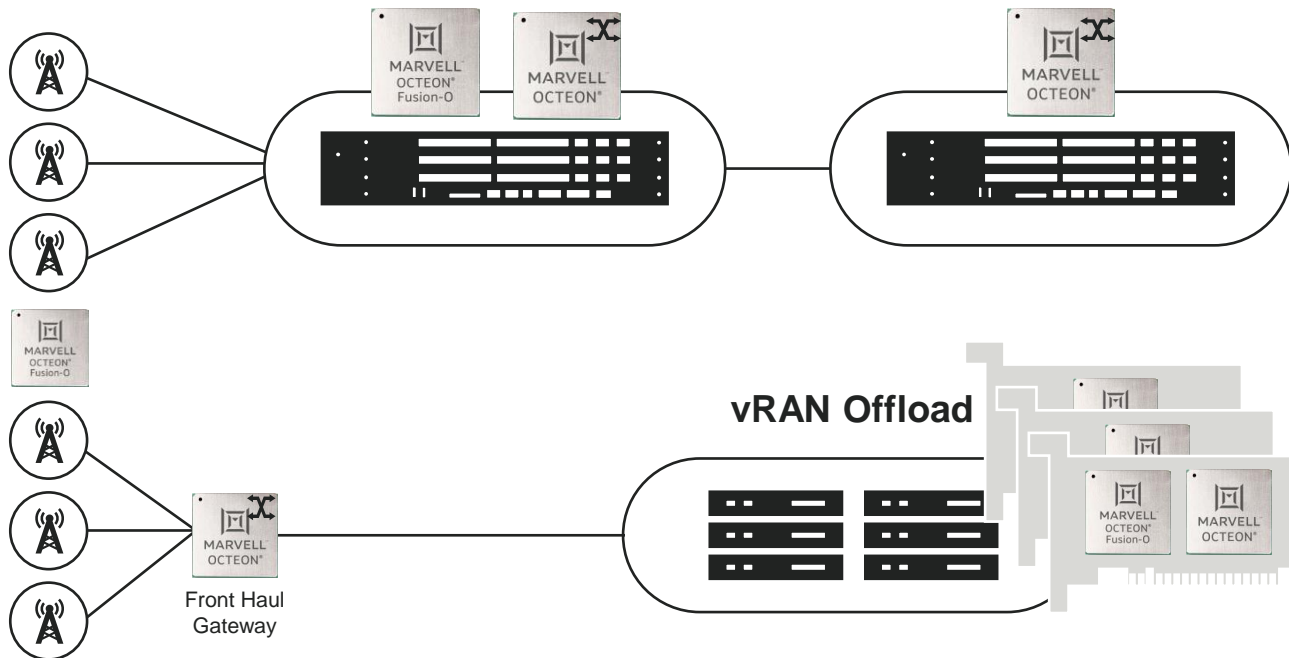
Use case

Disaggregated RAN

Radio Unit

Digital Unit

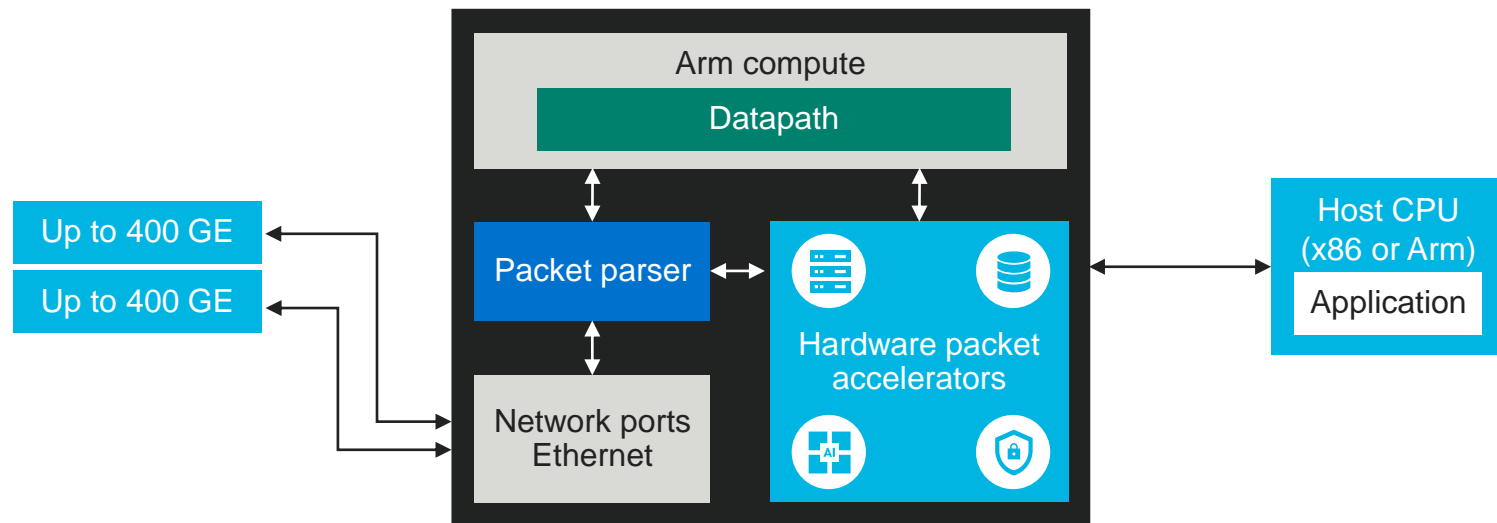
Central Unit



O-RAN/vRAN Offload

# Cloud and Datacenter DPU

Use case



- Compute: 1000+ SPECint
- Ethernet ports: Up to 400GE
- Datapath: 400G+
- Storage: 20M+ IOPS
- AI/ML: 100's TOPS
- Security: 400G+ of IPSEC and SSL

## Development Platform

---

### Compute

- OCTEON 10 DPU
  - 24 Neoverse N2 cores
  - SPECint > 800
- 

### Accelerators

- inline ML engine
  - inline IPSec
  - SSL/TLS
  - Vector packet processing (VPP)
- 

### SW

- DPDK networking suite for Control, Management and Fast path stacks
  - SDK with Linux kernel and user plane extensions
- 


### Memory

- 16GB DDR5-5200 + ECC on-board memory
- 

### I/O

- 2 x 100GbE QSFP56
  - PCIe 5.0
- 

# OCTEON 10 DPU Development Platform

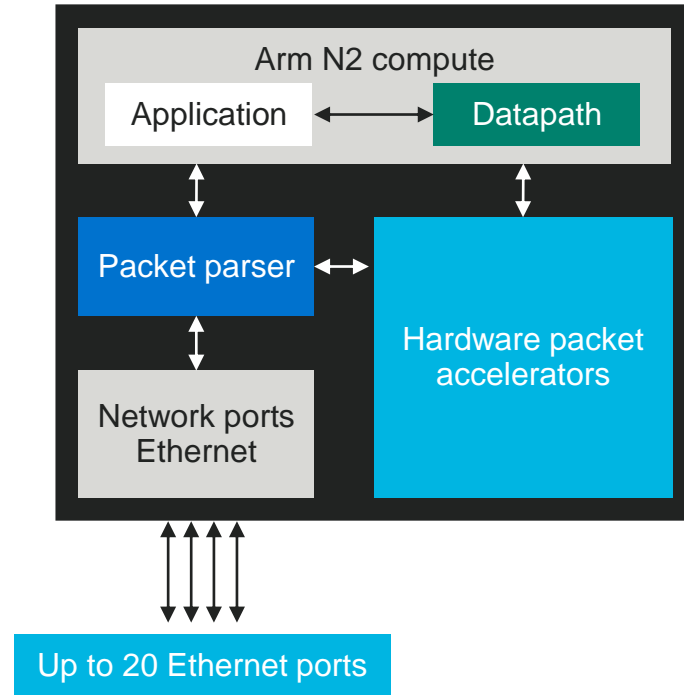


Available in Q4

# Enterprise router and firewall appliance

Use case

- Scalability from low to high end systems
- OCTEON 10 compute
  - Best single threaded performance
- Data plane acceleration supports:
  - NG Firewall, IPSec, TLS
  - L2/L3 forwarding
  - Advanced packet parsing
  - Inline hardware base AI/ML
- Support for up to 20 ethernet MACs



# OCTEON 10 platform initial family members

Metric	CN103XX	CN106XX	CN106XXS	DPU400
N2 Cores	Up to 8	Up to 24	Up to 24	Up to 36
Max Frequency	2.5GHz	2.5GHz	2.5GHz	2.5GHz
SPECint (2006)	>275	>800	>800	>1200
Cache (L2, L3)	8MB, 16MB	24MB, 48MB	24MB, 48MB	36MB, 72MB
DDR5 Controllers	2 at 4800MT/s	6 at 5200MT/s	6 at 5200MT/s	12 at 5200MT/s
Crypto	Supported	Supported	Supported	Supported
Ethernet	4x50G/25G/10G + 2x10G or 16x1G	4 x50G or 2x10/1G	16 x50G	Up to 400G
PCIe 5.0 controllers	Up to 6	Up to 6	Up to 4	Up to 8
Typical power	10-25W	40W	50W	60W

**Sampling in 2H21**



Essential technology, done right™



Thank You