AMD

Heterogeneous, programmable silicon for high-performance network processing

25/5/2023

Kimon Karras

Agenda

Overview	• Motivating factors behind SmartNICs & DPUs • AMD Networking Solution Portfolio
Solution – AMD 400G Adaptive SmartNIC SoC	• Offload Acceleration Blocks for Host, Network, and Embedded Processing • Pervasive Security and Confidential Computing • Use Cases & Examples
Solution – AMD Infrastructure Accelerators	• Infrastructure Accelerator Architecture Overview & Roadmap • Programmable Silicon through P4 • Use Cases & Examples
Summary	
Q&A	

Motivations for Adaptable, Intelligent Infrastructure

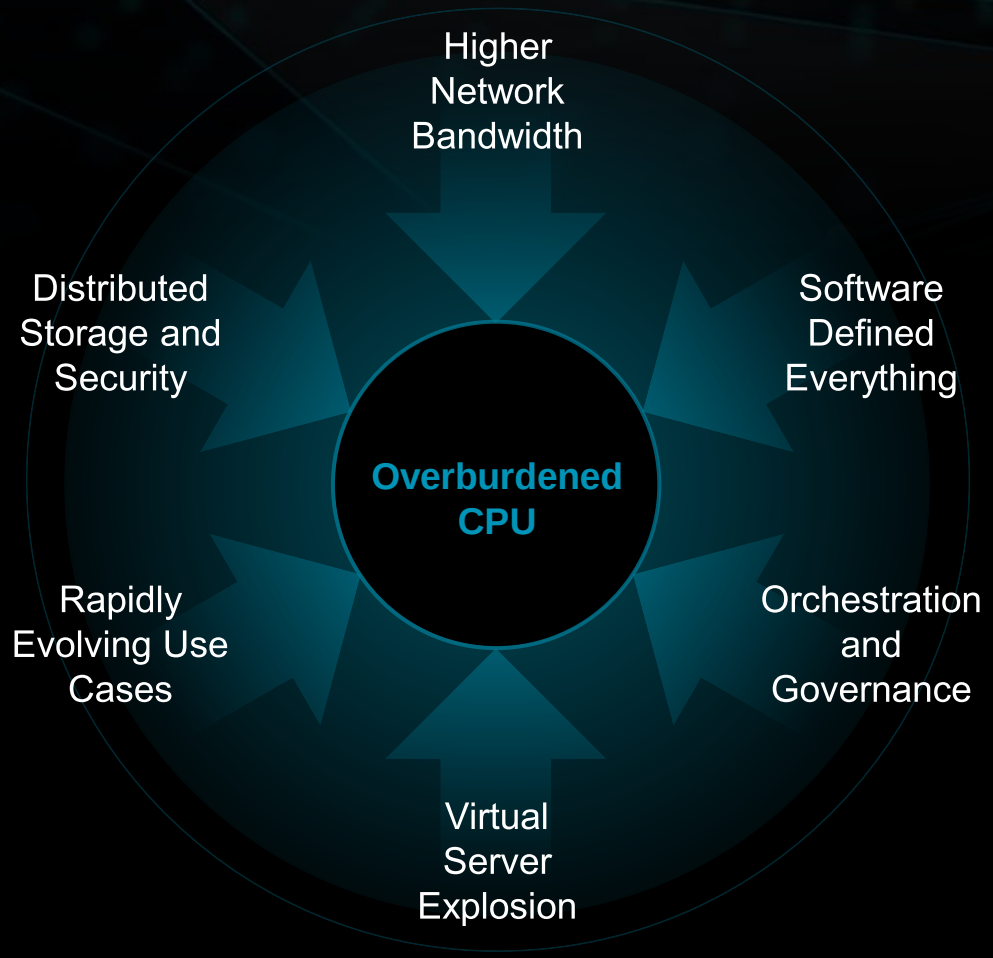

New Trends Expose Current Infrastructure Challenges

Intelligent Infrastructure for Next Gen Applications

Next Generation Bandwidth
400Gb Data Processing

Infrastructure Acceleration
Maximize CPU revenue potential

High Security
Multi-tenant Isolation and Protection

We are at an Architectural Inflection Point

Tug-of-War between

ASIC-like Fixed Function Logic
Frequent vs. occasional functions

Embedded Processor Cores
Workload flexibility vs. processing offload

FPGAs or Programmable Logic
Adapting new and existing functions

AMD Networking Portfolio

- The AMD networking portfolio spans a wide range of offerings
- Broad network hardware and software offering spanning DPU, SmartSwitch, FGPA and FinTech
- Allowing for the creating of powerful heterogeneous architectures that combine programmability & performance
- Continuing to innovate with partners around standards, line rates, and heterogenous integration

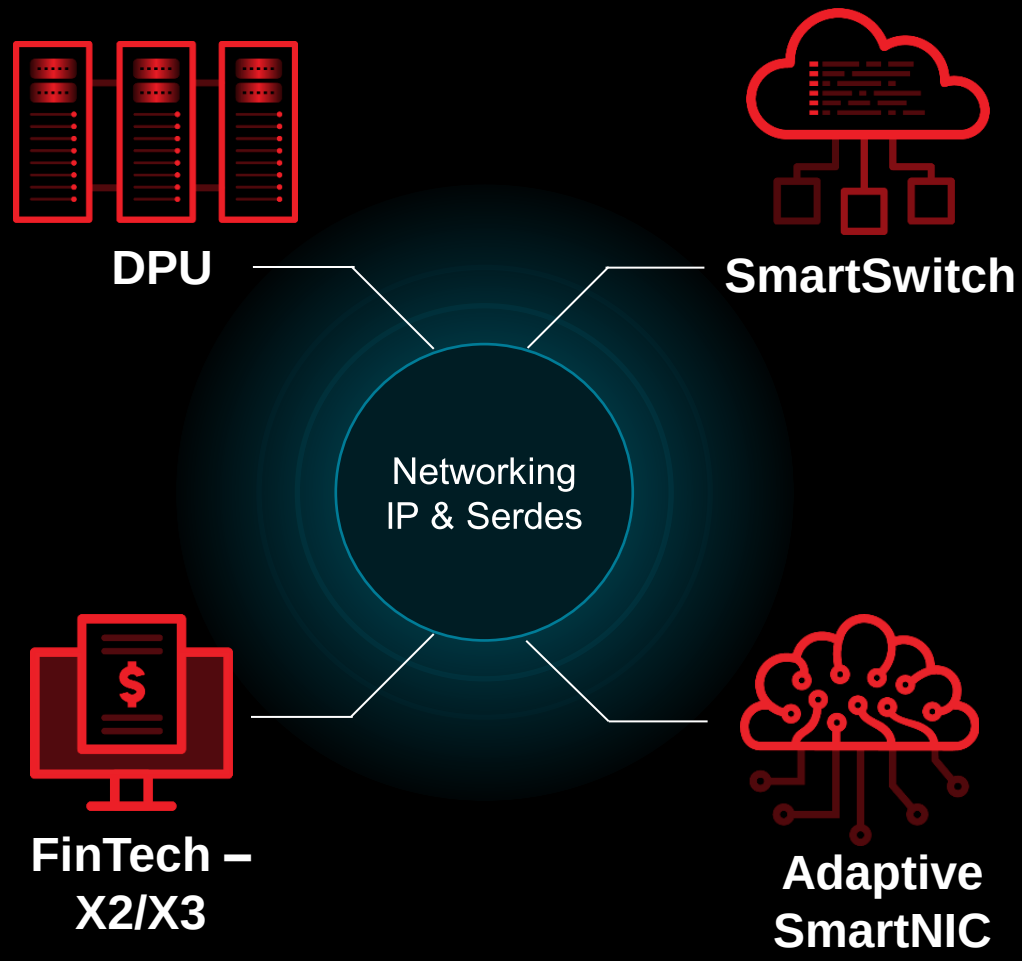

AMD

Adaptive, Programmable SmartNIC Overview

AMD 400G Adaptive SmartNIC SoC Objectives

Balanced Architecture

- Architect an adaptive SoC for the SmartNIC domain
- Strike the right balance between the heterogeneous elements
- Present a unified software view of the heterogeneous element SoC

Performance and Adaptability

- Latest Interfaces/Speeds: 112G Ethernet, PCIe Gen5, CXL 2.0, LPDDR5/DDR5
- Adaptive Interface APIs for evolving workloads and customization

Security

- Latest security standards: PCIe, OCP and NIST
- Symmetric and asymmetric crypto offload for data encryption and key exchange
- Physical and logical isolation for confidential compute

Solution

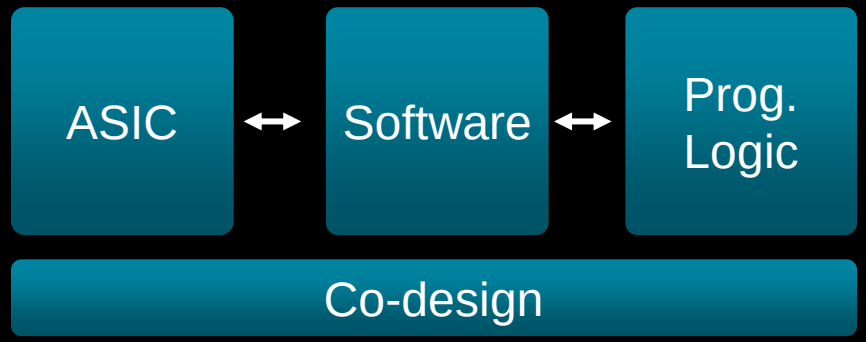

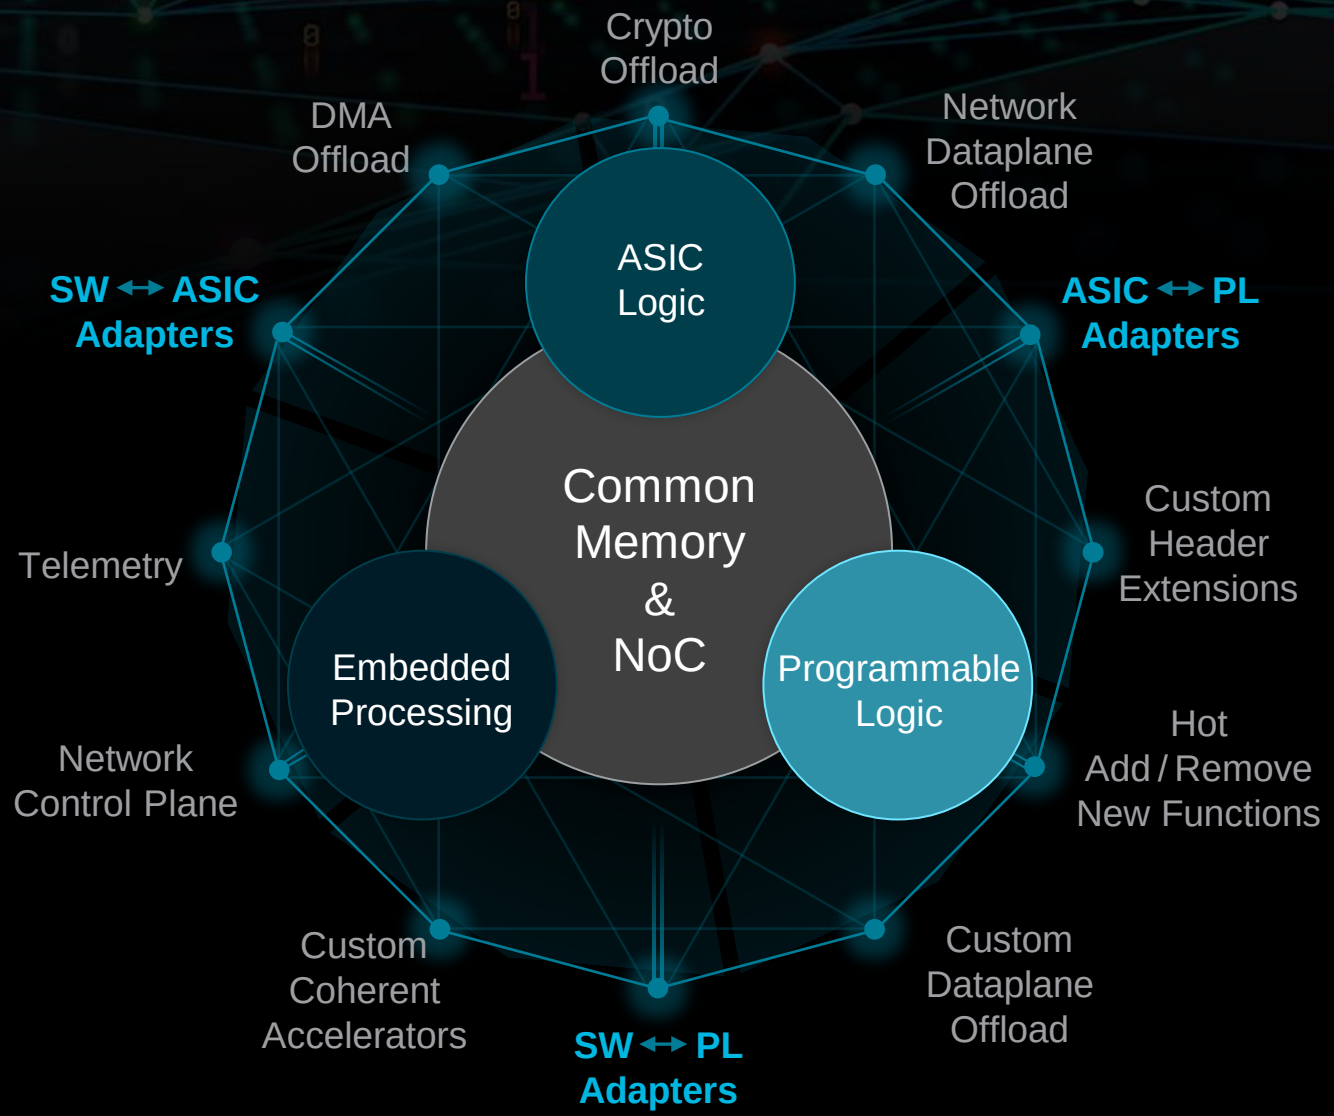

AMD

AMD 400G Adaptive SmartNIC Architecture & Performance

Tightly Coupled SmartNIC Blocks

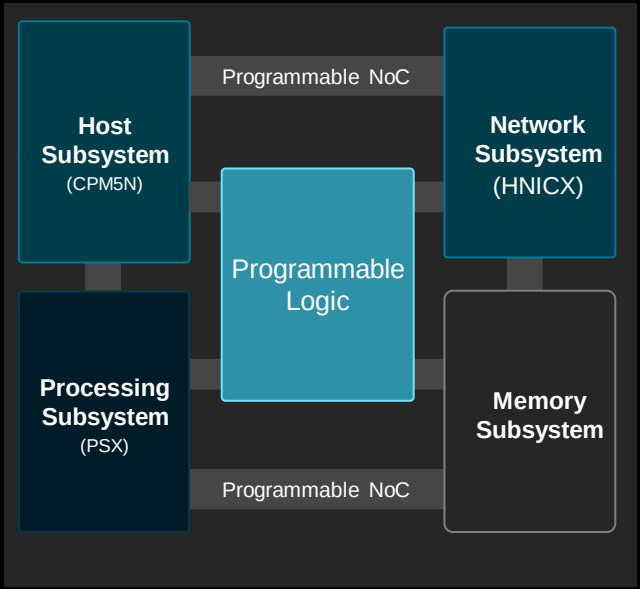

Architectural Resources

Host subsystem	PCIe Gen5 x16 PCIe CMA/DOE/IDE CXL v2.0
Network subsystem	2x 200G Ethernet Virtual switch offload Stateless offloads
Application Processing Subsystem	16x A78-AE
Real-time Processing Subsystem	4x ARM R52
PKI (TLS1.3 Offload Engine)	Yes
DDR subsystem	8x 32b LPDDR5/DDR5 ECC Inline encryption 1.638Tb/s LPDDR5 BW

SmartNIC SoC Performance†

Programmable logic packet rate[1]	400Mpps Ingress + 400Mpps Egress
Host ⇔ Network packet rate[2]	400Mpps RX + 400Mpps TX
Full Virtio.NET offload BW[3]	400Gbps RX + 400Gbps TX
AES-XTS offload BW[4]	800Gb/s
AES-GCM offload BW[5]	800Gb/s

† See Endnotes

AMD

Pervasive Security

Layers of Security

CHECK

Secure Boot, Key Exchange and Attestation

AMD keys, SmartNIC owner keys, and SmartNIC tenant keys

PROTECT

Memory and Peripheral Protection Units

SoC-wide hardware firewalls for tenants and accelerators

SHIELD

Secure Data in Flight, at Rest, in Use

Encryption for data, host/network interfaces, and DRAM
Secure monitors for external, internal, physical attacks
Crypto engines with DPA countermeasure protection

Use-cases

Full OVS Offload

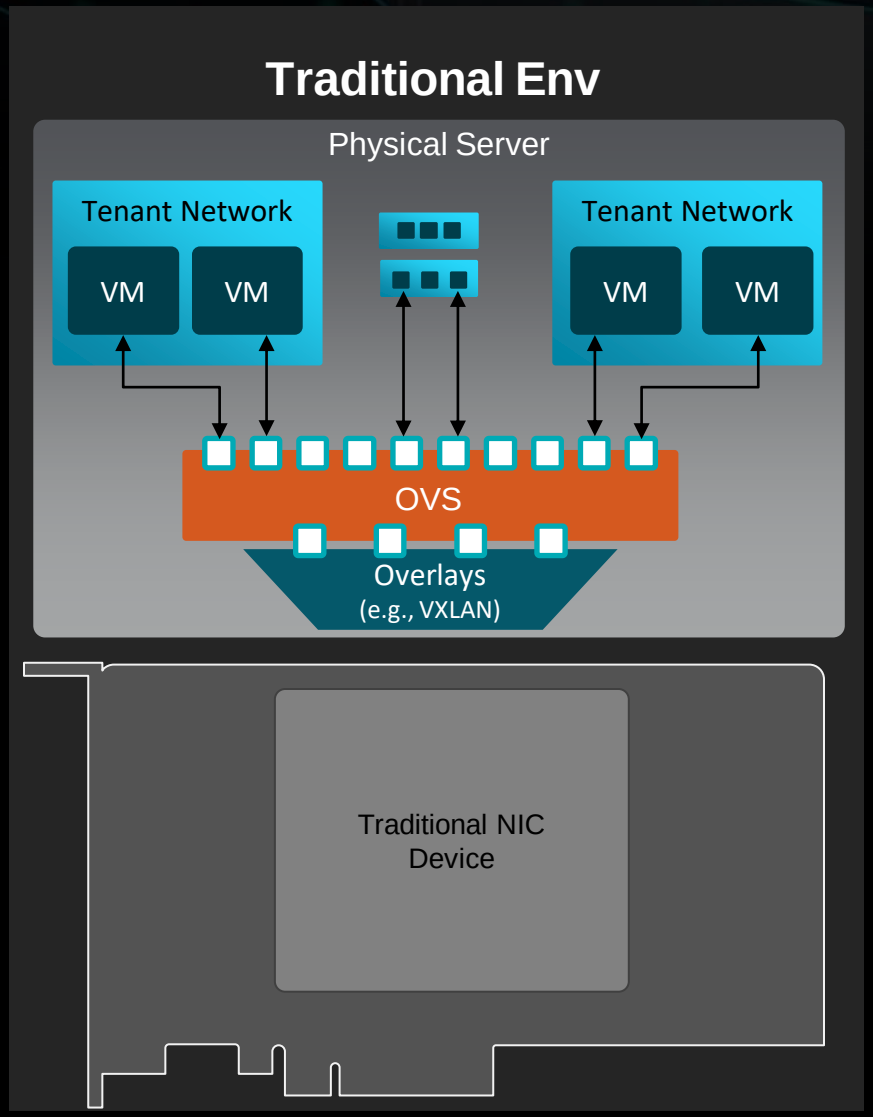

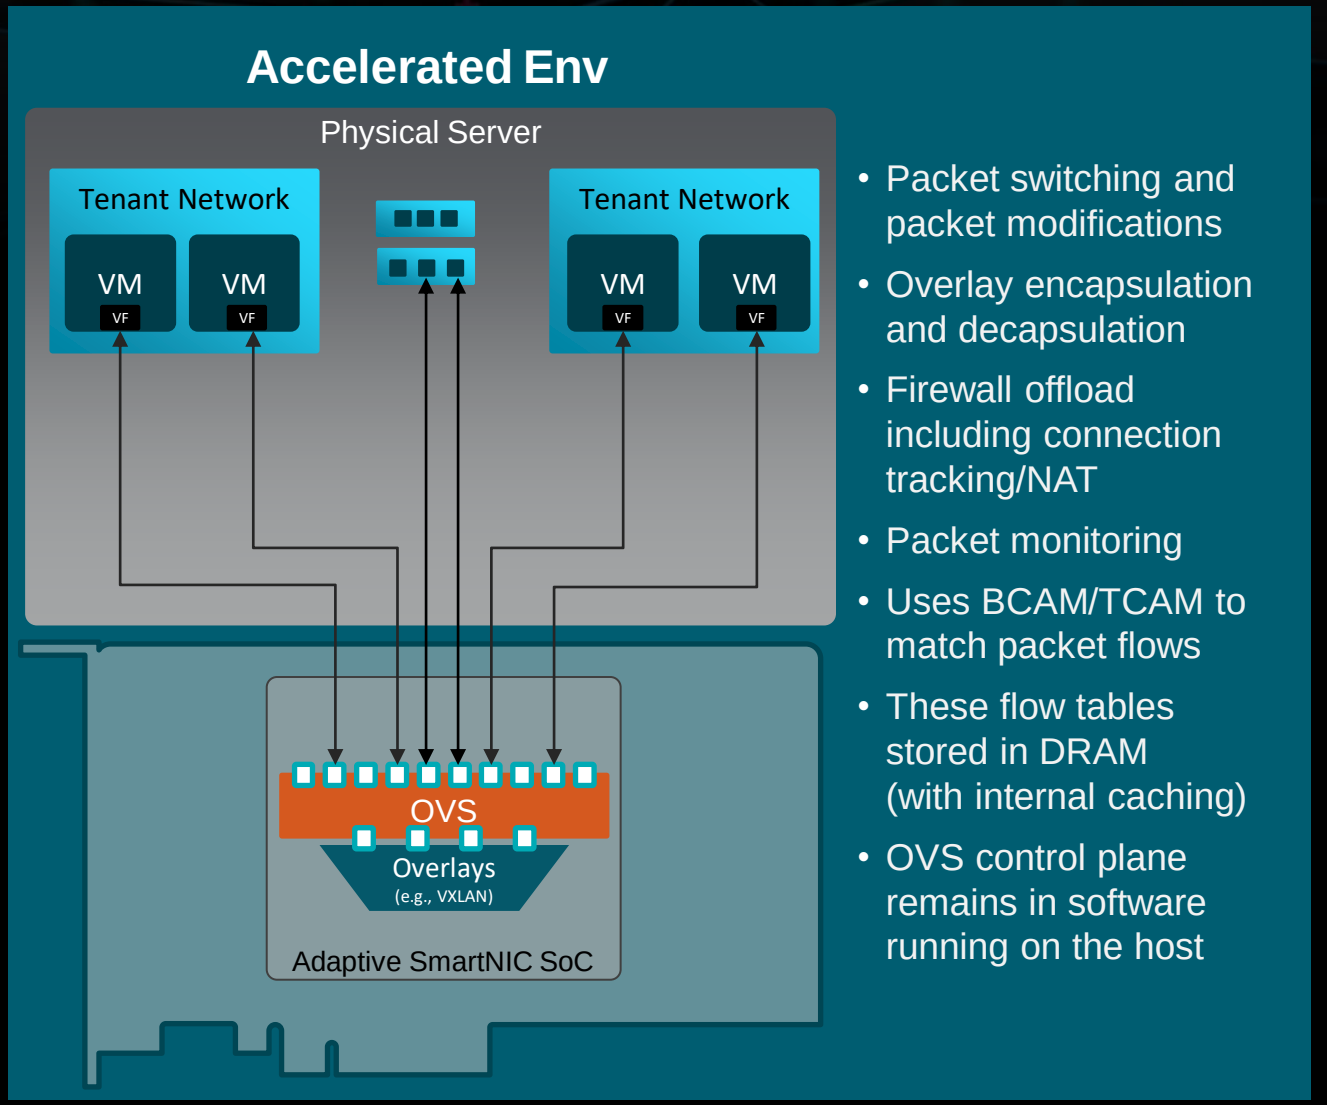

- Packet switching and packet modifications
- Overlay encapsulation and decapsulation
- Firewall offload including connection tracking/NAT
- Packet monitoring
- Uses BCAM/TCAM to match packet flows
- These flow tables stored in DRAM (with internal caching)
- OVS control plane remains in software running on the host

PCIe Endpoint 400G SmartNIC w/ with Host Storage and Network Offload

- Standard Ethernet Packet Flow
- Custom Header, Encapsulation, or Encryption Flow
- Host Connection with Advanced Switch Hot Plug

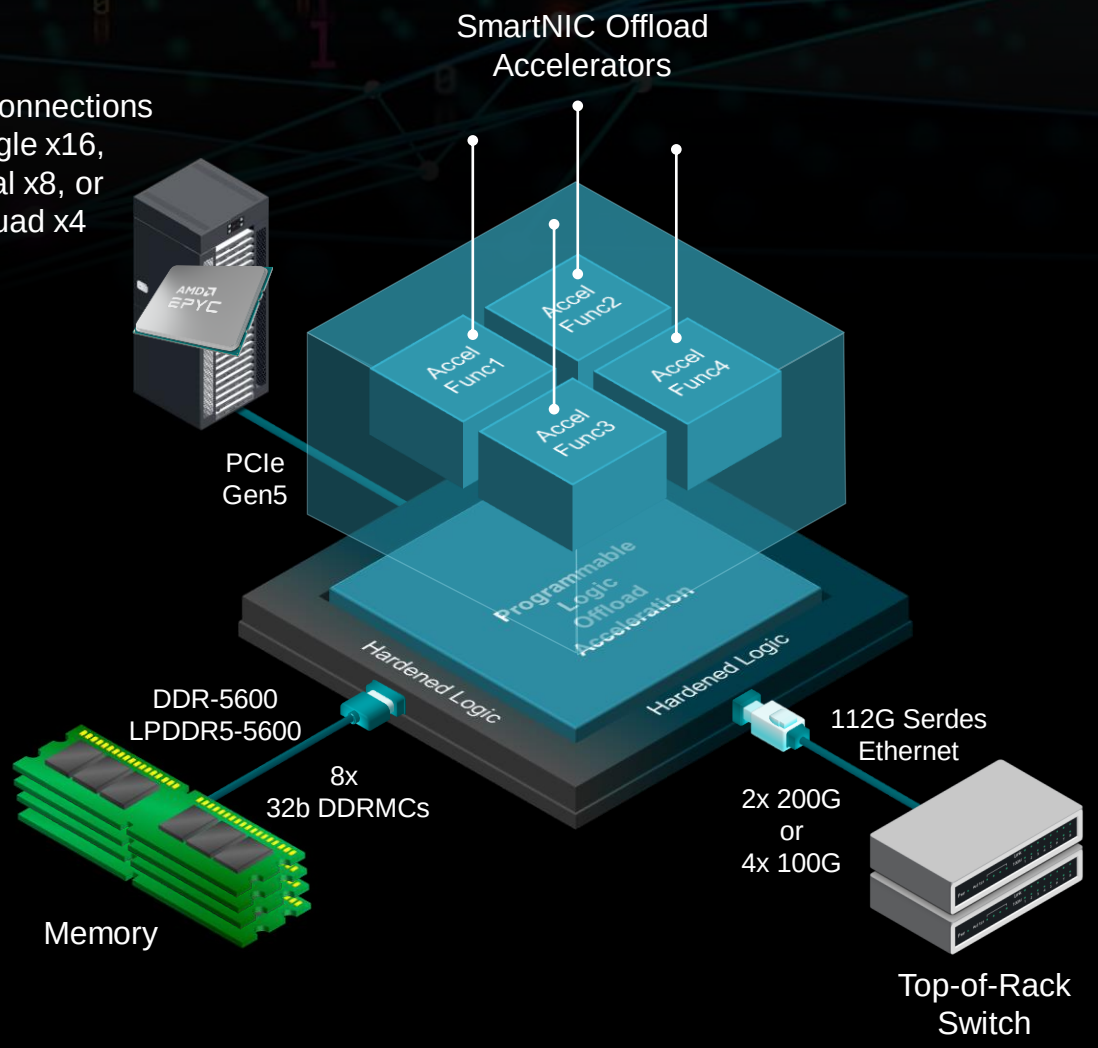

AMD

PCIe Endpoint 200G SmartNIC w/ 2 Gen5 x4 NVMe SSDs

- Host endpoint connection NVMe virtualization
- SmartNIC Root-Port connections for dual-NVMe SSDs w/ AES-XTS offload acceleration
- 2x 100G or 4x 50G network connection

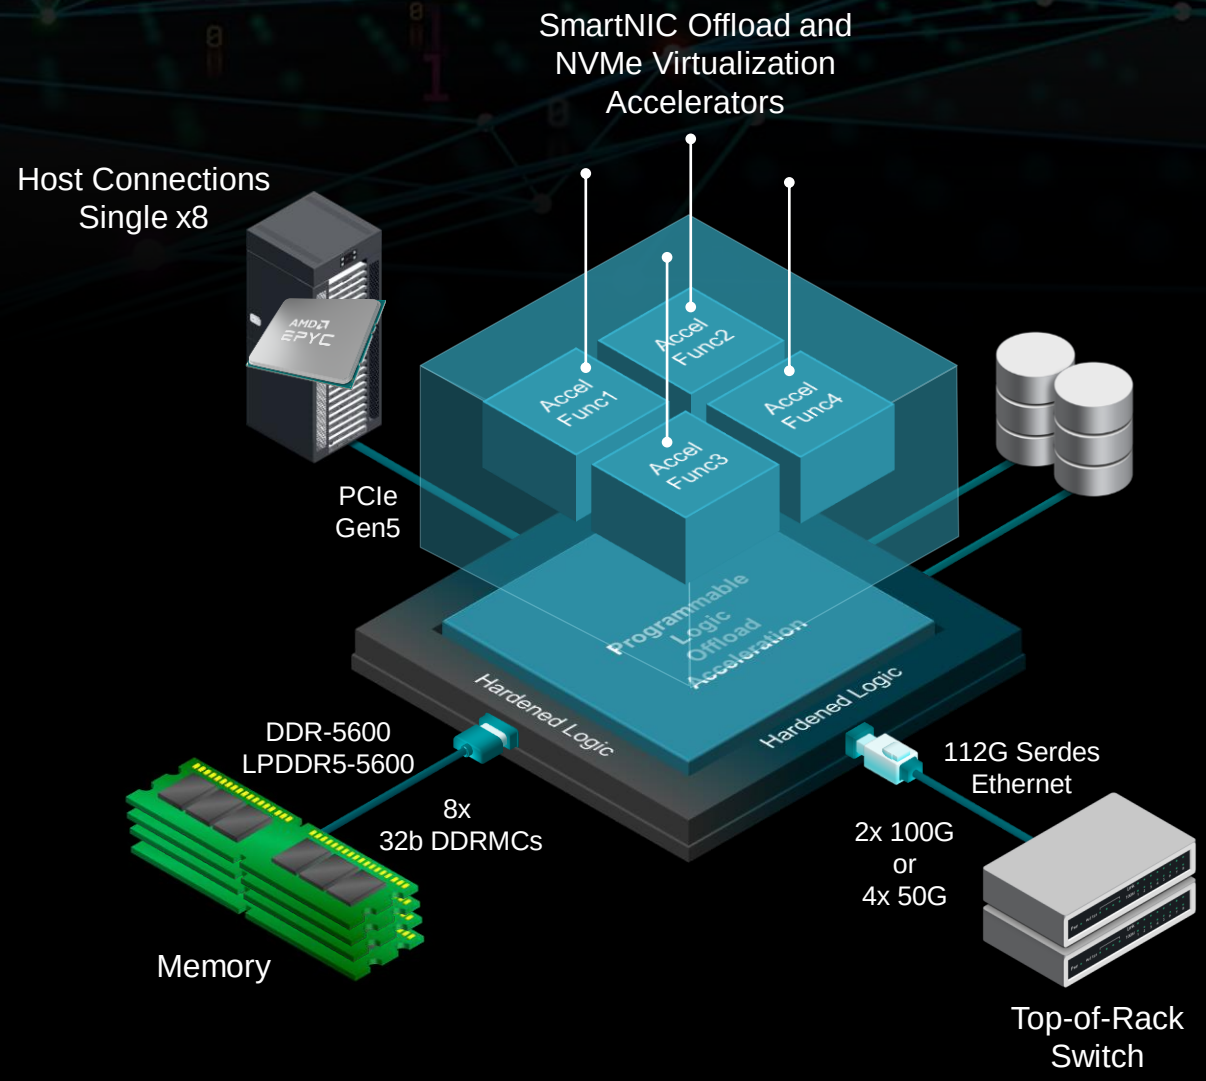

AMD

PCIe Root Complex 400G NVMe Storage Node

- Network-attached storage appliance
- SmartNIC Root-Port connections for quad-NVMe SSDs
  w/ AES-XTS offload acceleration
- 2x 200G or 4x 100G remote storage connections

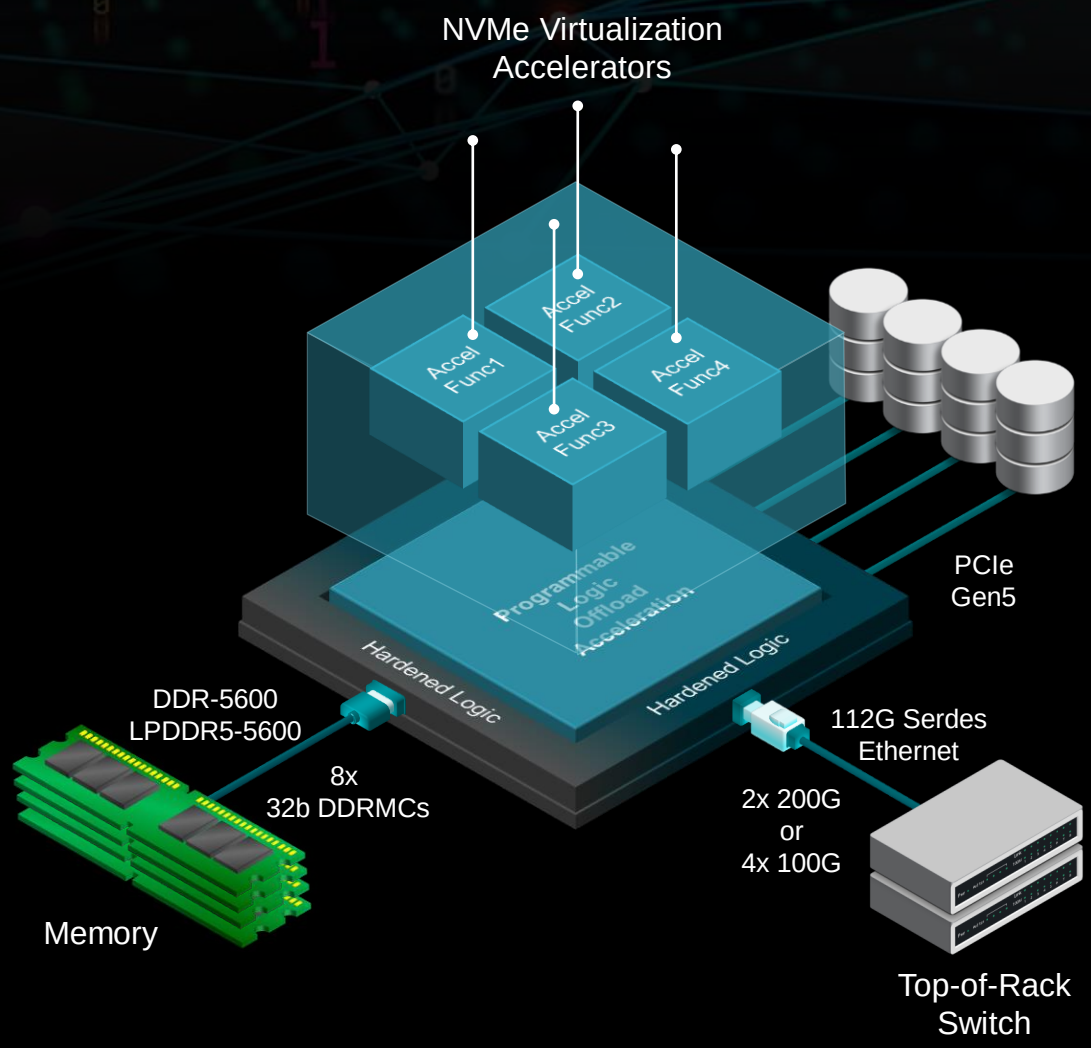

AMD Infrastructure Accelerators

AMD Pensando Cloud Solution with Turnkey Software Stack

Cloud Provider Controller

AMD Pensando
Policy and Services Manager(PSM)

Cloud Service gRPC API

AMD Pensando Cloud Software Stack

Network

Cloud VPC/VNet
VPC/VNet peering
VXLAN/NVGRE/Geneve
MPLSoUDP
NAT
Shared services (private endpoint)

Security

Secure Boot
Stateful security
Microsegmentation
Network ACL
IPsec/VPN
TLS/D-TLS
DDoS protection

Storage

NVMe and NVMe virtualization
NVMe/TCP initiator
Data-at-rest and Data-in-motion encryption

Telemetry and Infrastructure

SDK for customization

Capri/Elba

AMD P4 DPU Roadmap

- **Leadership Performance**
- **Architected to Scale with supporting multi-services Software Backward and Forward Compatibility**

- **2x100GbE**
- **16x PCIe Gen3**

- **2x200GbE/4x100GbE**
- **32x PCIe Gen3/4**
- **Inline AES-GCM**

- **Higher Performance**
- **Better memory throughput & density**
- **Lower Power**

2019 → 2024

Elba (7 nm)

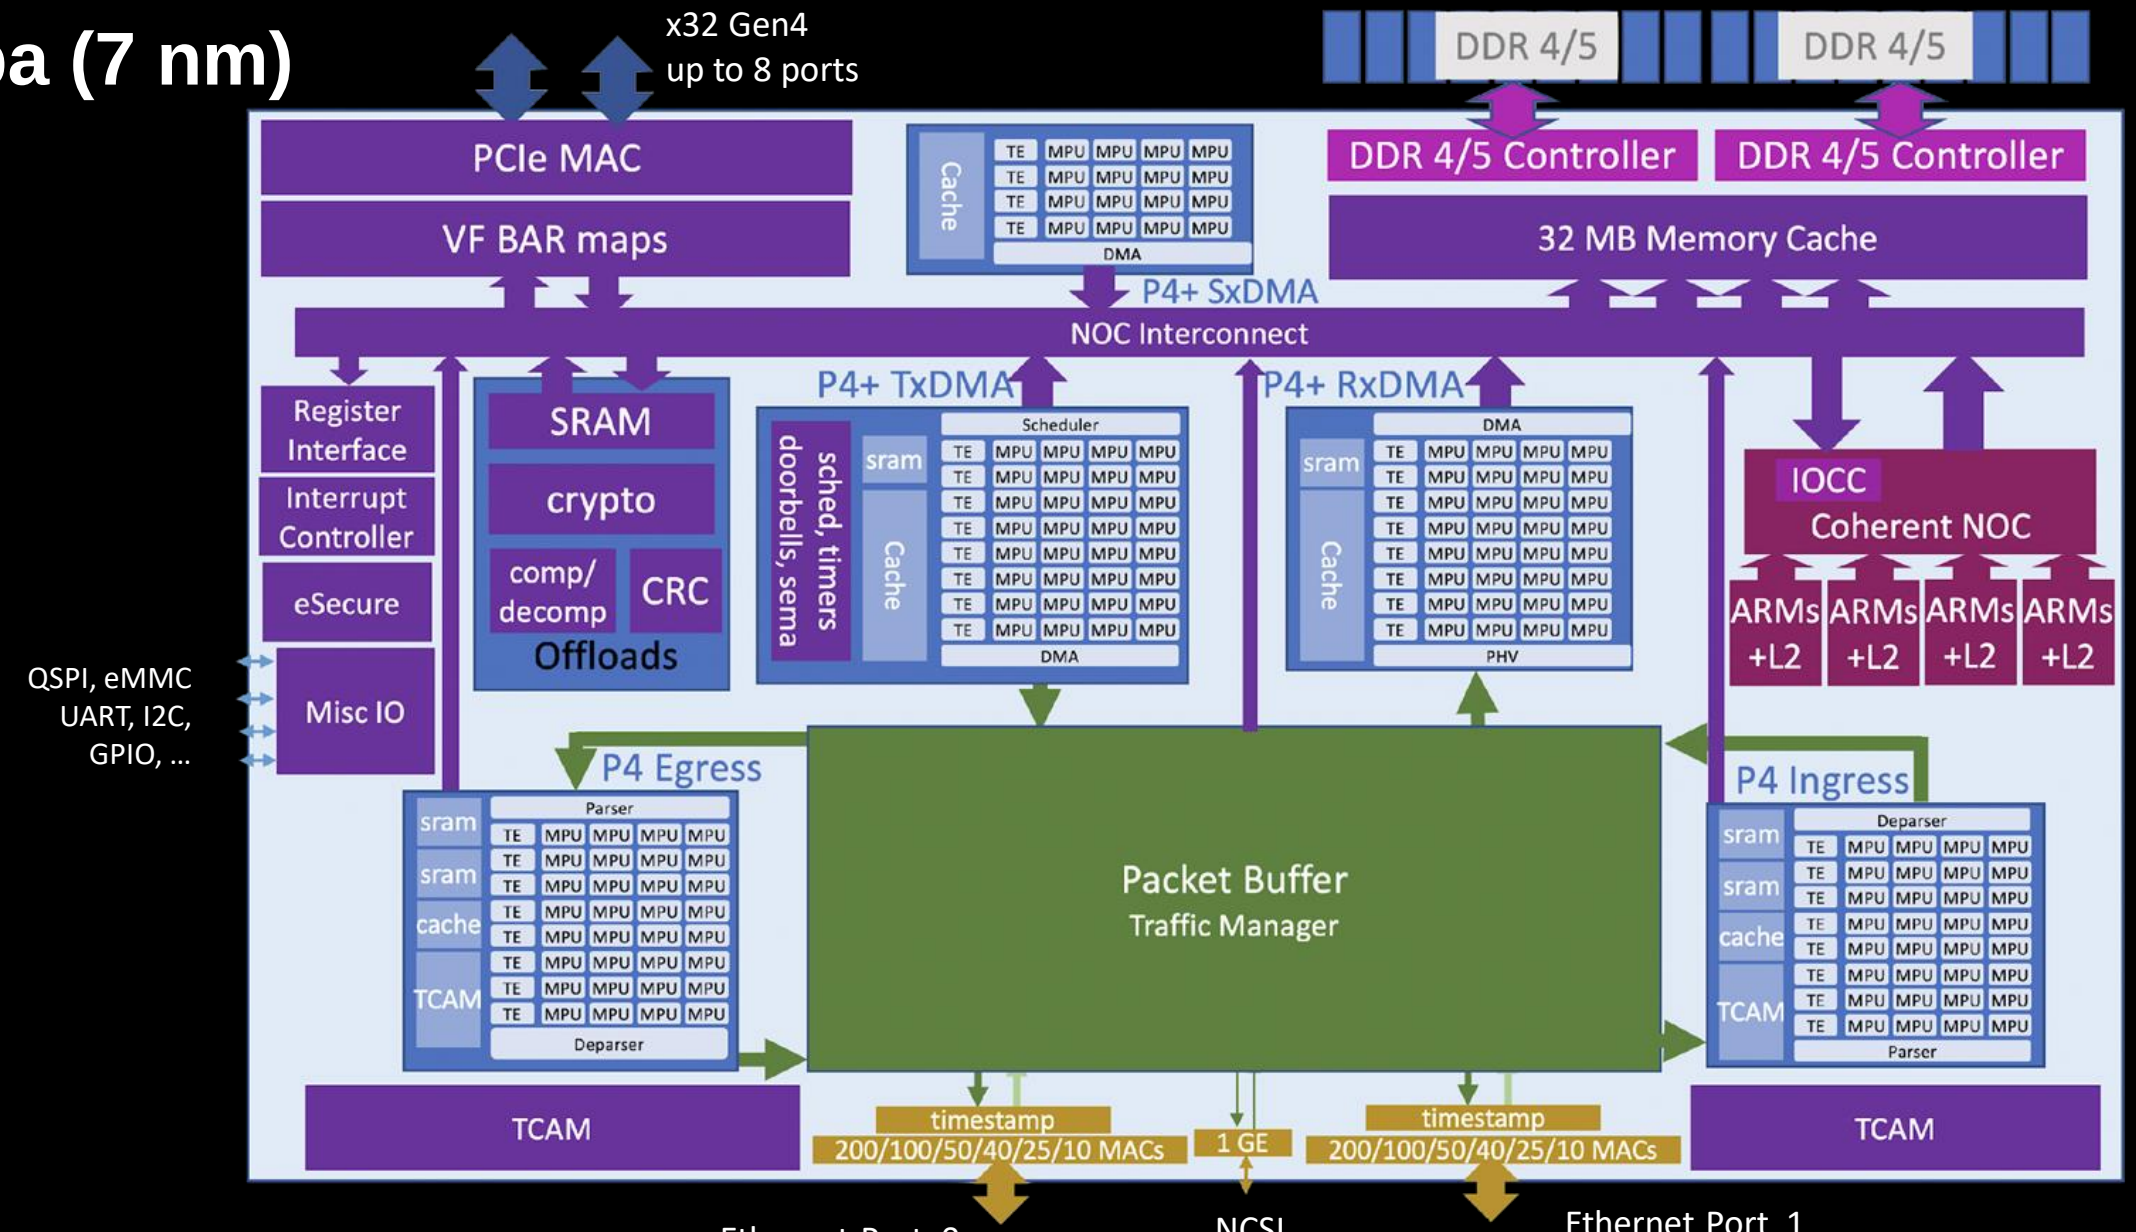

Accommodate Different Consumption Models

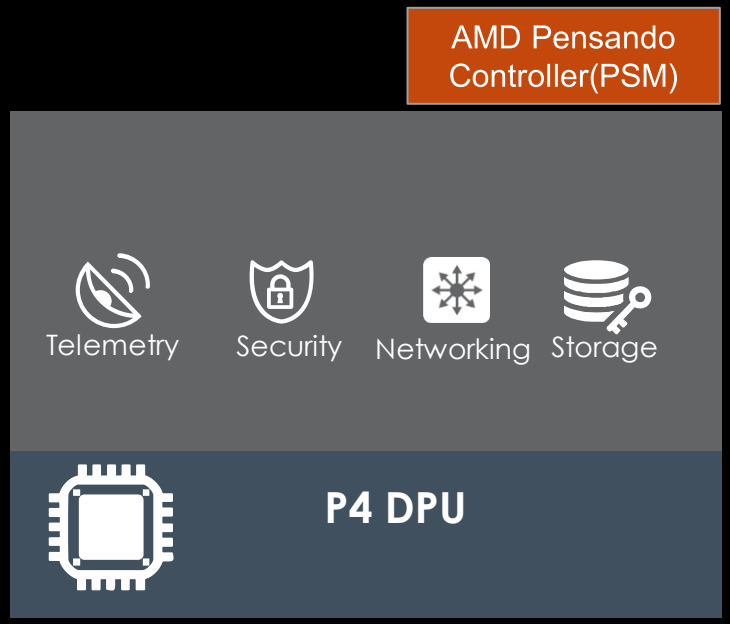

Turnkey solution

- Full software stack for networking, security, storage and telemetry
- Easy Integration with gRPC API
- Pensando PSM(Policy and Services Manager) for fabric wide policy provisioning and DPU mgmt

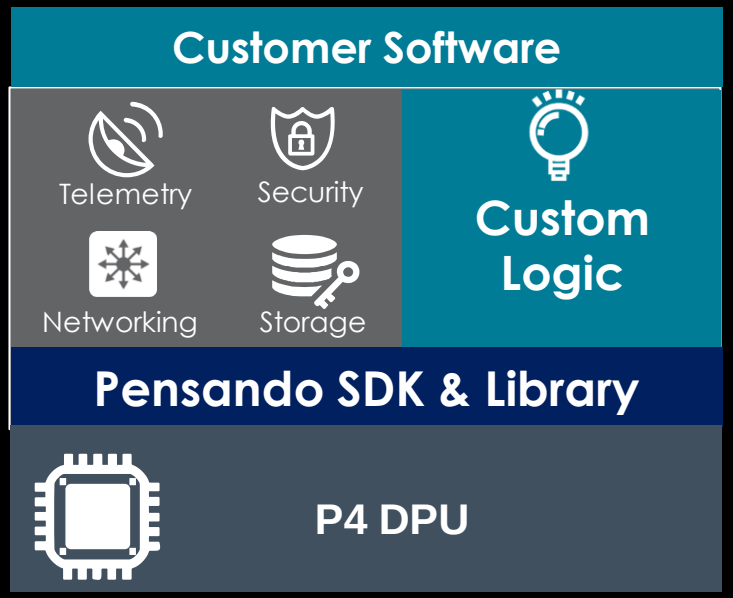

Co-Development Mode

- Open system to allow customers/partners to add/develop software defined services on top of Pensando provided software stack
- Support control plane, data plane or management plane customization

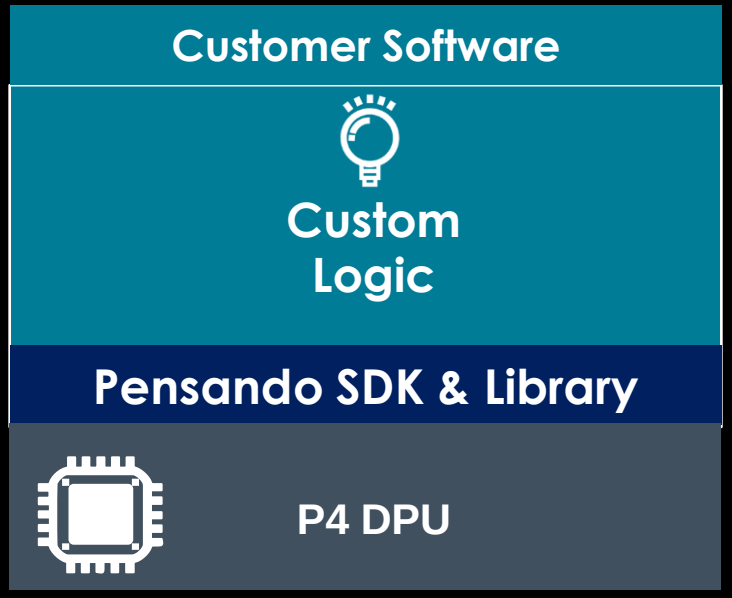

Customer's Business Logic

- Full Programmability at ALL Layers
- Support ANY Packet Format & Transformation/Processing Flows
- Leverage Pensando SDK and libraries, reference design

A Fully Customizable System

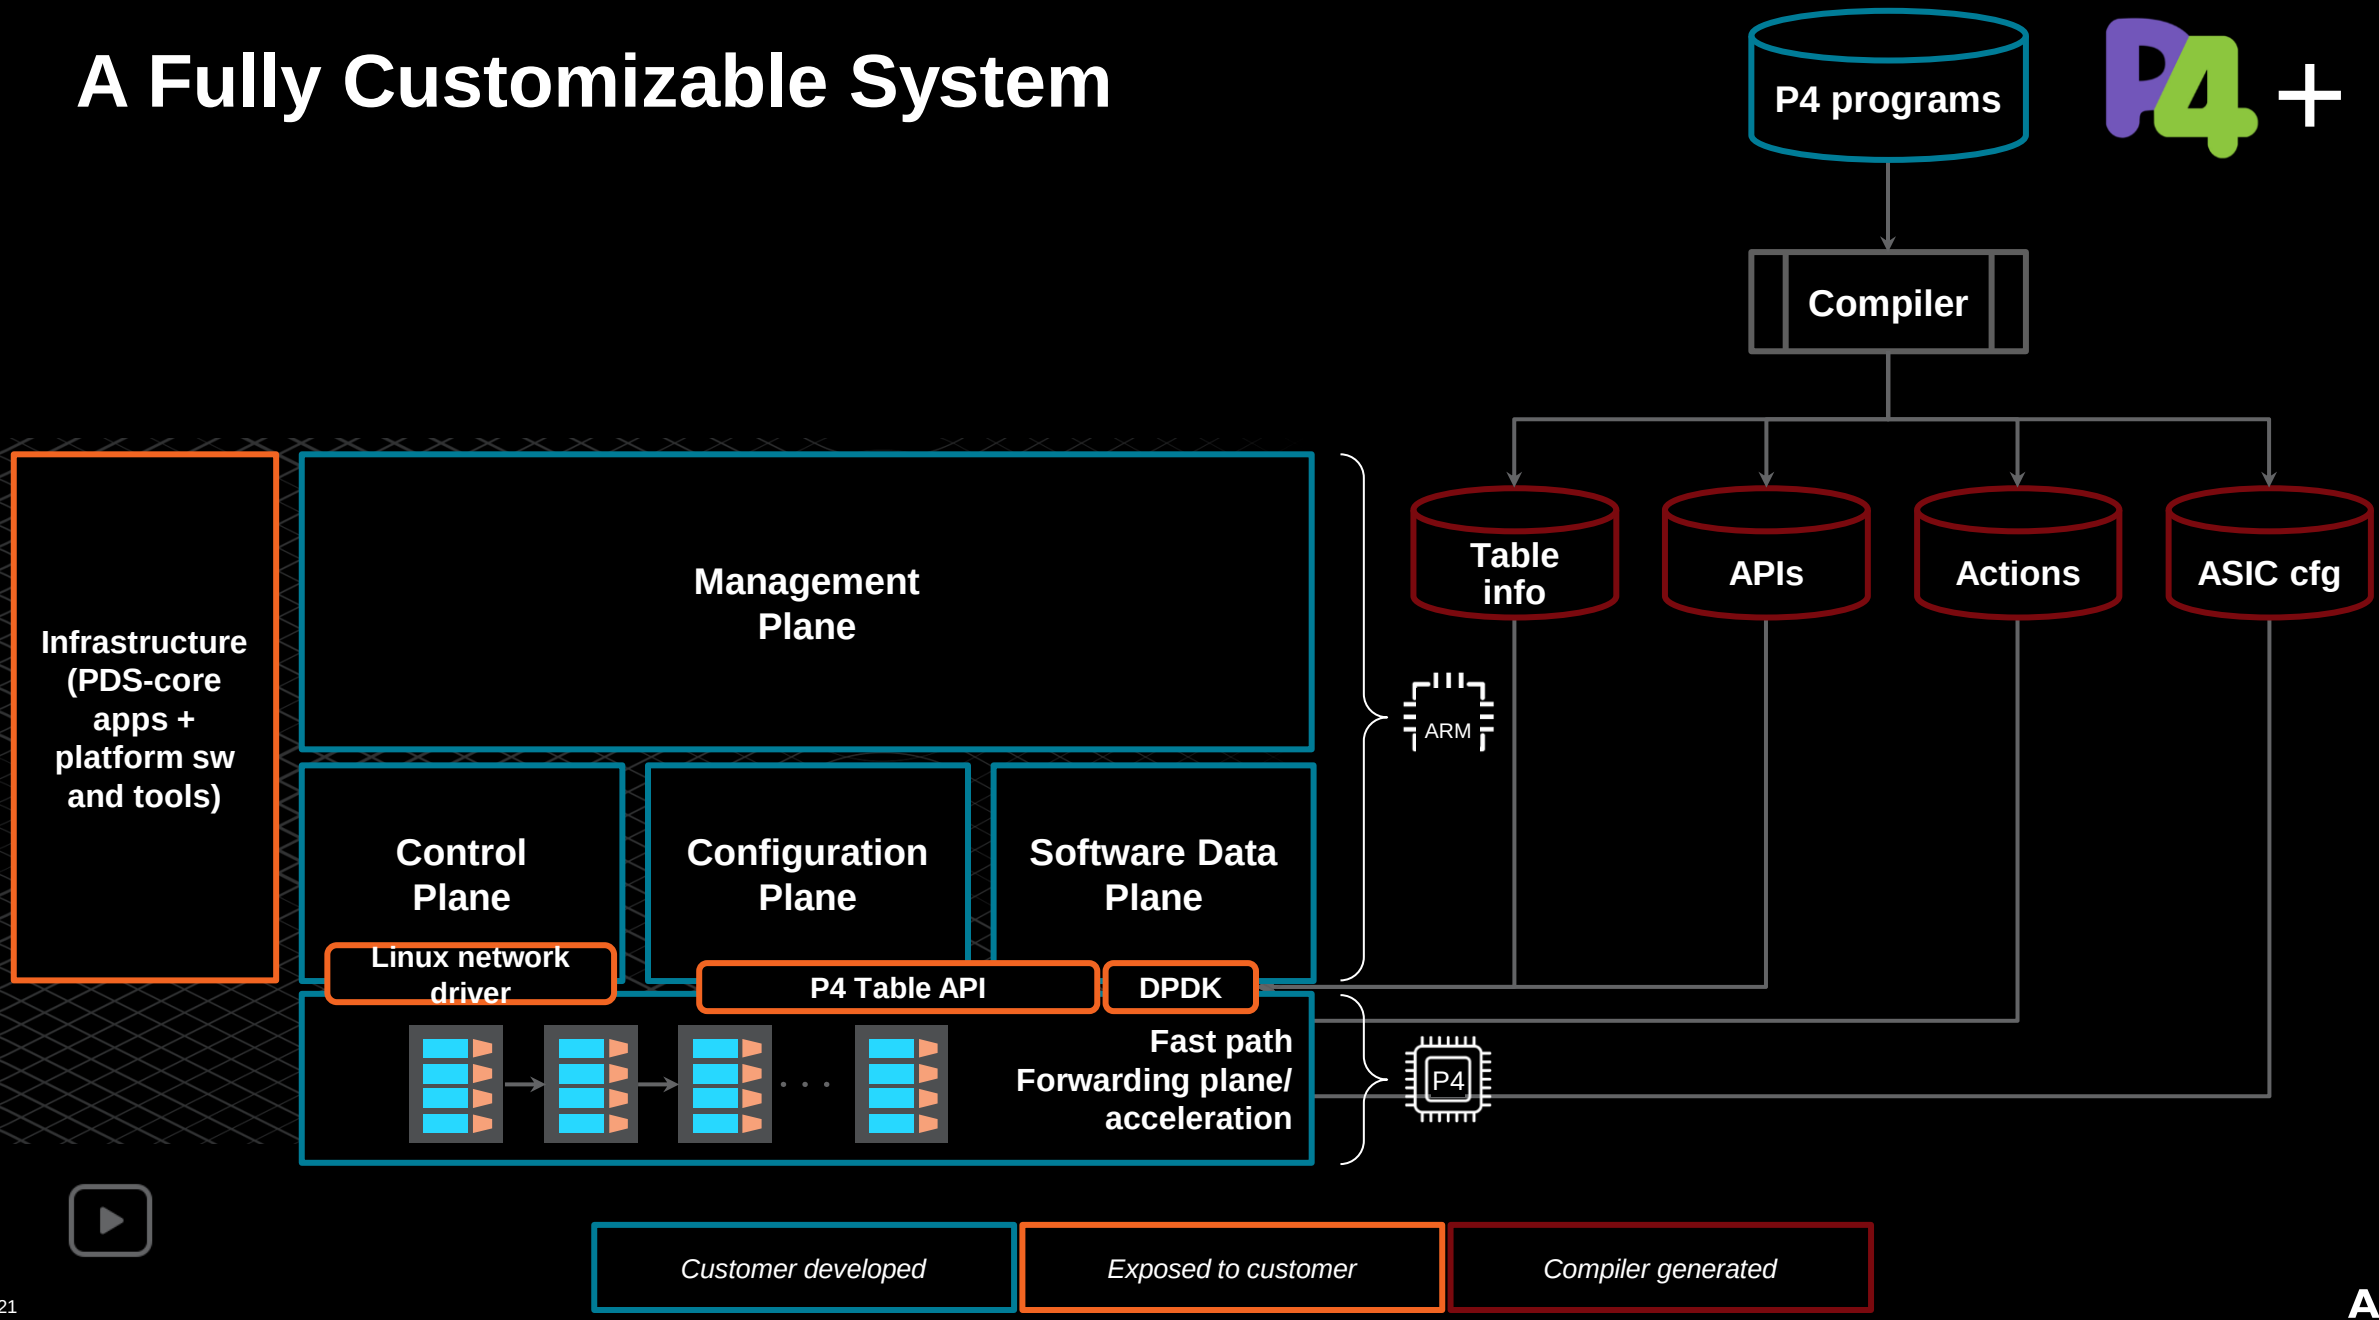

NVMe Virtualization and NVMe/TCP

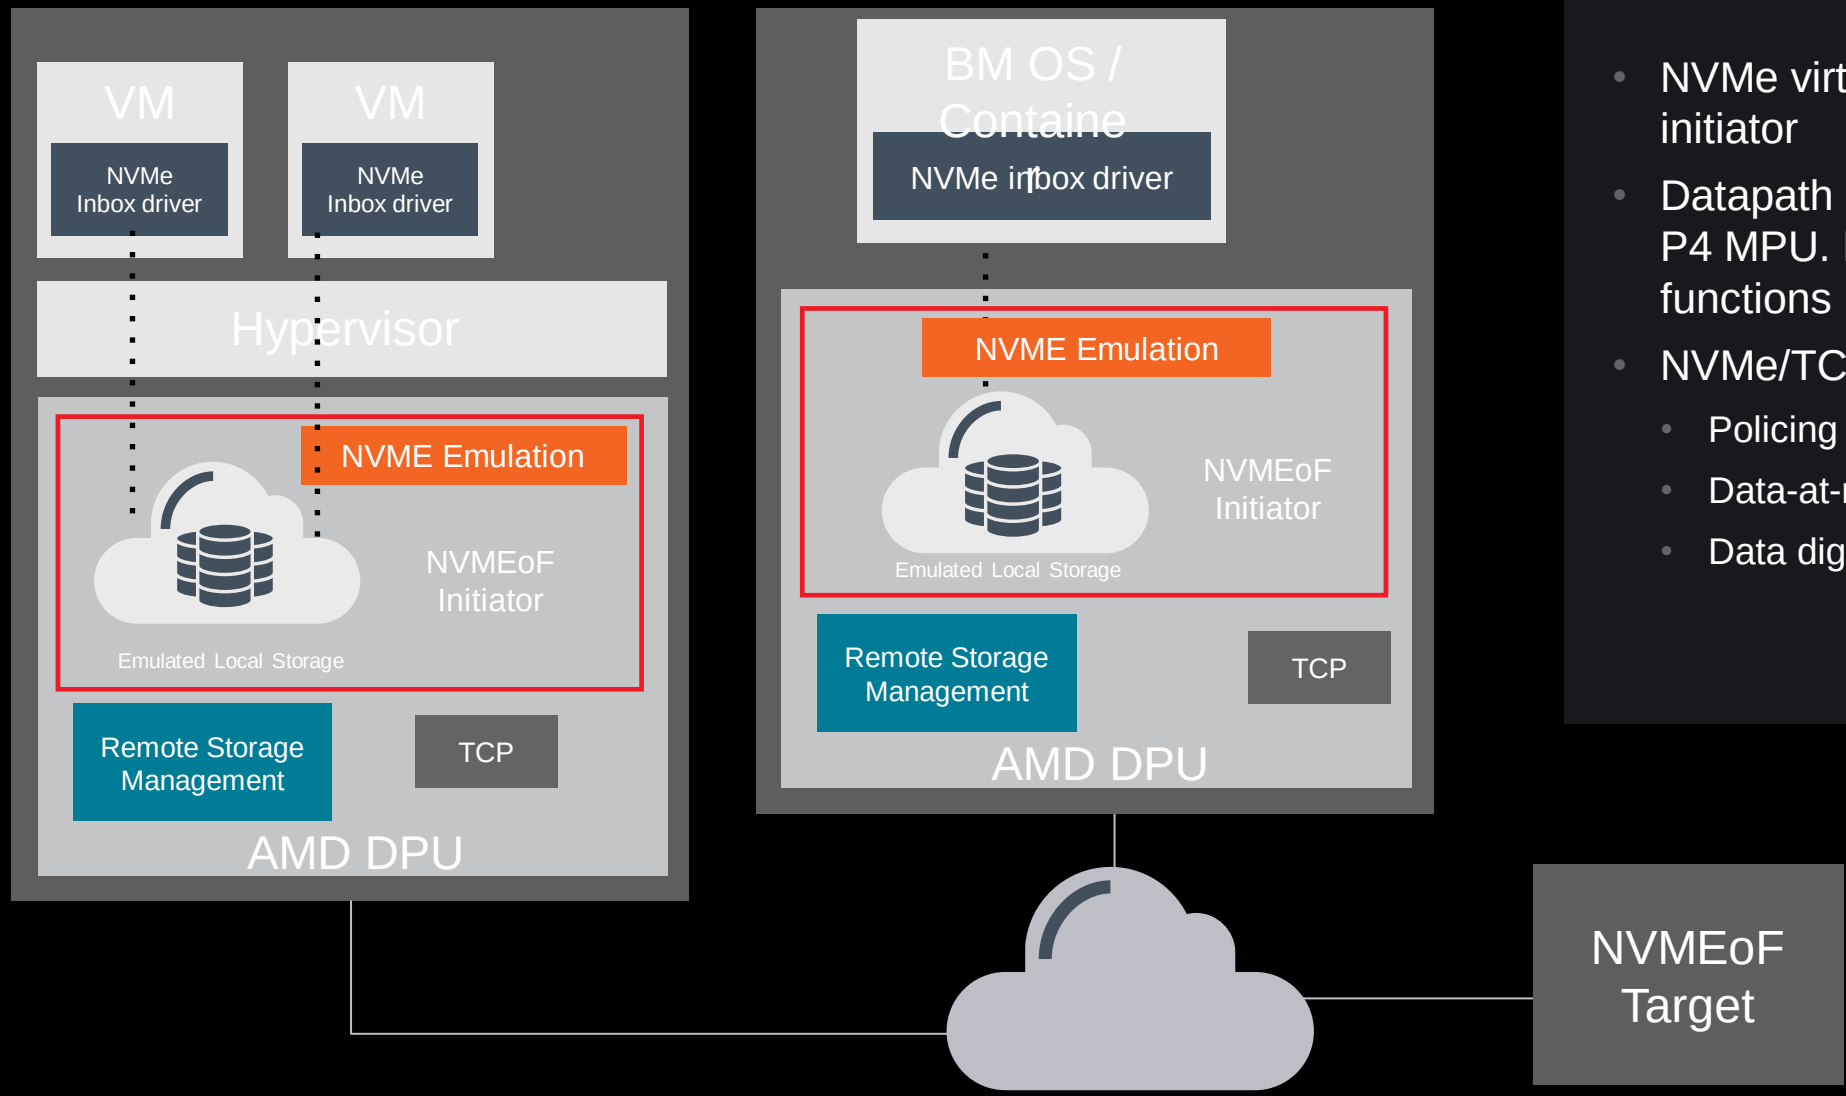

- NVMe virtualization and NVMe/TCP for initiator
- Datapath including TCP stack running on P4 MPU. Preserve Arm cores for other functions
- NVMe/TCP chained with other services
  - Policing and QoS
  - Data-at-rest and Data-in-motion encryption
  - Data digest

Cloud Solution Form Factor

Form Factors

DPU in Server

DPU in Appliance

DPU in SmartSwitch

DPU Offloads Examples:

SDN
Security Groups
NVMe over TCP
Telemetry/Statistics
Load Balancing
QoS (Policers, Rate-Limiting)
Encapsulation
Encryption
Address Translation
Compression
And many more….

Common Building Blocks and Benefits

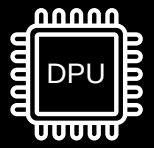

DPU

Cloud Software Stack

Performance/Scale

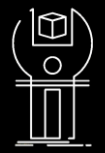

SDK

Summary

- AMD offers a wide range of flexible, high-performance NIC solutions
  - AMD SmartNICs
  - AMD Infrastructure Accelerators
  - AMD Low-latency NICs

- These combine a HW/SW approach to network & storage processing allowing the user to adapt the NIC's functionality:
  - Through P4 for the Infrastructure Accelerators
  - Through custom HLS or RTL blocks for the SmartNIC & Low-latency NICs

- This enables services from traditional OVS offload to NVMe/TCP acceleration.

Thank You for Participating

Endnotes

† SmartNIC SoC performance claims are based on simulation results of pre-silicon models with AMD traffic generators used to emulate external port traffic behavior as of August 23, 2022. Actual performance of production silicon may vary.

1. Interface packet rate between Programmable logic and the Network subsystem
2. Sustained TX and RX packet rate for 64B packets
3. Sustained TX and RX BW with 1500B payload per packet
4. Sustained BW for 1KB packets; Encrypt-only, Decrypt-only, or 400Gb/s Encrypt + 400Gb/s Decrypt with Network Interface configured to 2x200G Ethernet ports
5. Sustained BW for 1KB packets; 400Gb/s Encrypt + 400Gb/s Decrypt with Network Interface configured to 2x200G Ethernet ports
6. #4 and #5 were also measured concurrently enabled

Disclaimer and Attribution

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. 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.

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 MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY 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.

ATTRIBUTION