

SmartNICs: Giving Rise To Smarter Offload at The Edge and In The Data Center

Jeff Defilippi

Senior Product Manager

Arm


#Arm Tech Symposia

Copyright © 2018 Arm Limited, All rights reserved.



arm NEOVERSE

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



SmartNIC and Smart Offload Definition

Smart

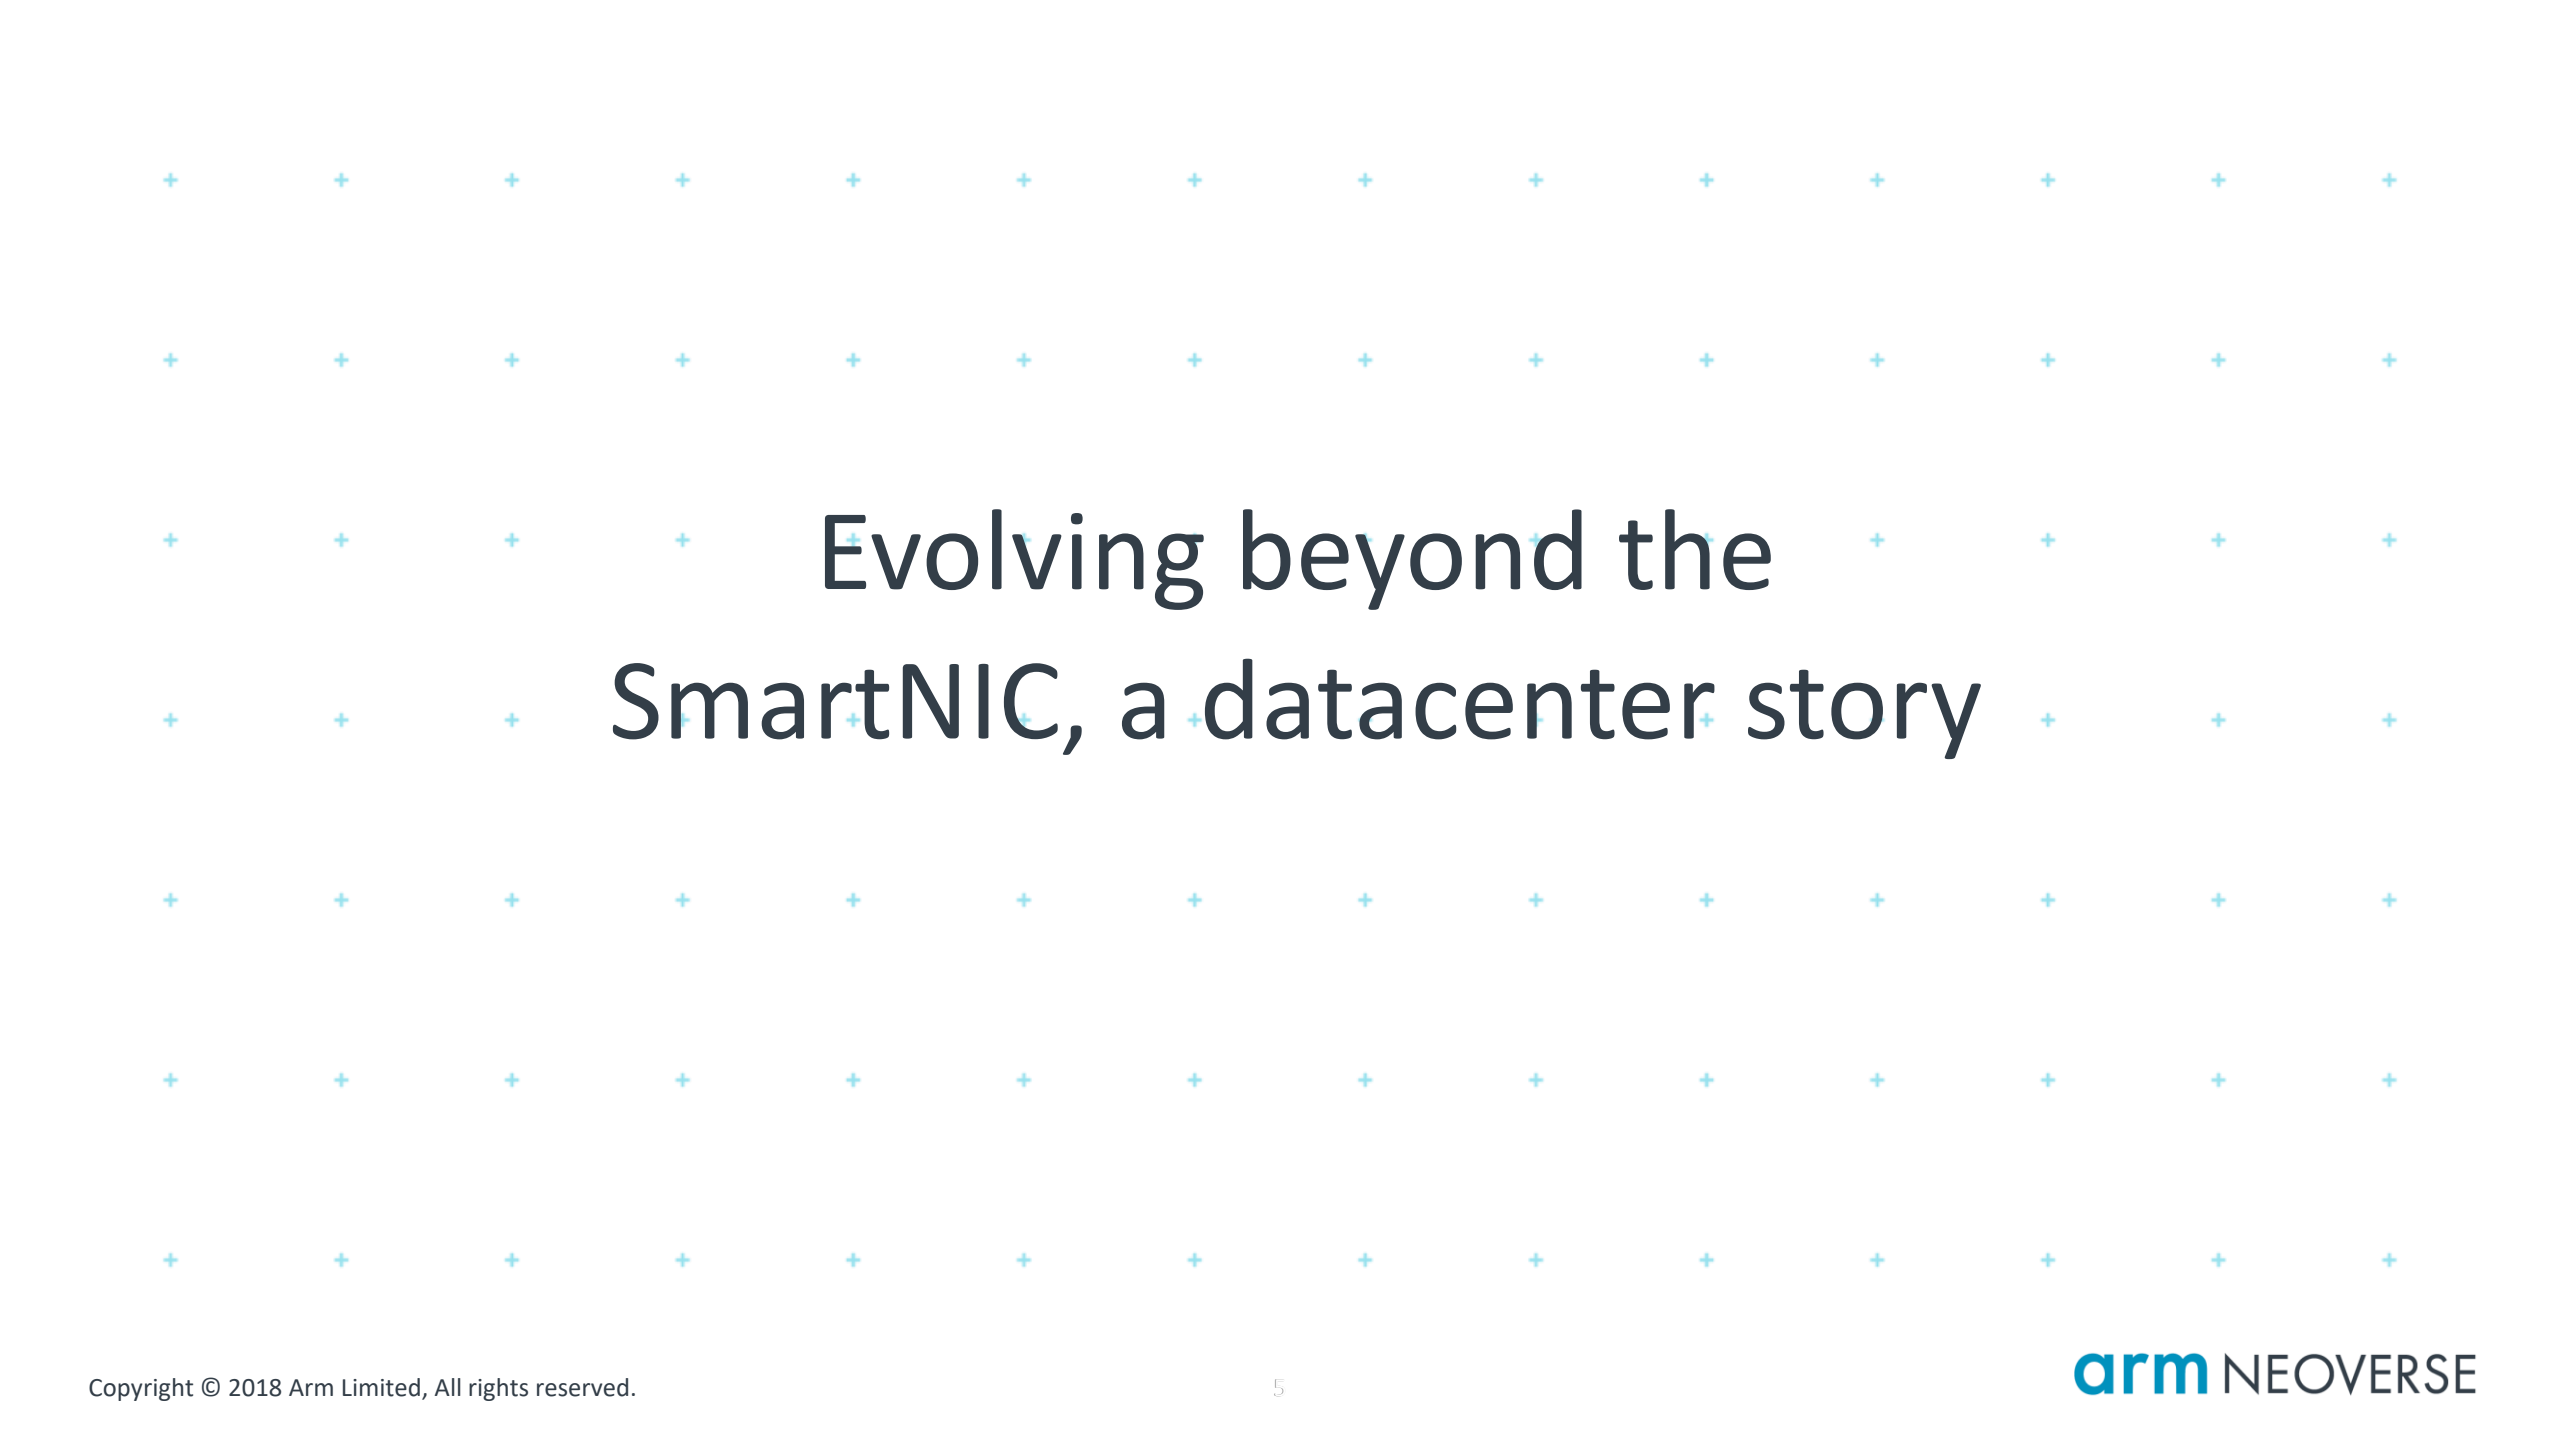
- Includes application processing capable of running rich operating system (Linux) and distributed cloud stacks.

Offload

- Offload special purpose hw that 'offloads' a specific task (ex: network, storage, security) to maximize efficiency and performance.
- NIC - class of offload for ethernet controllers.

Agenda

- Evolving beyond the SmartNIC, a datacenter story
- Arm platform for smart offload and accelerated edge compute



Evolving beyond the SmartNIC, a datacenter story

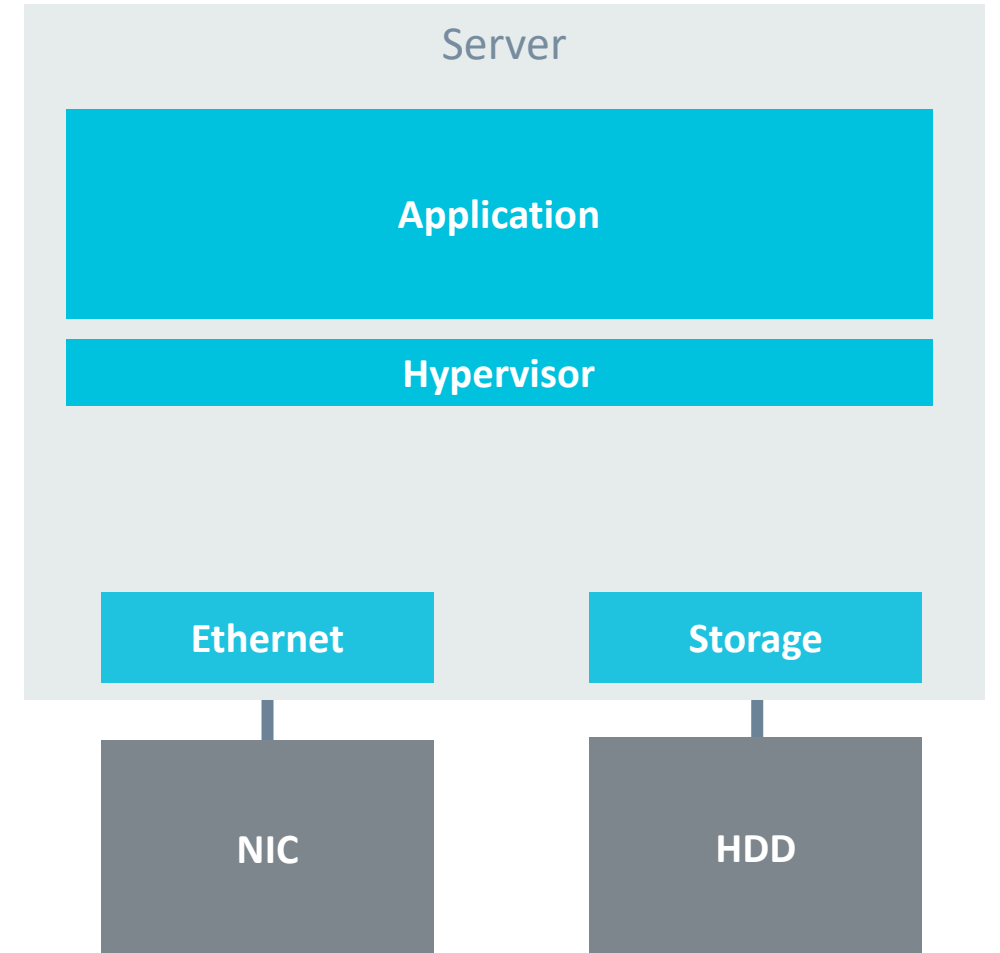
Traditional server deployment

Server

- Runs all OS, virtualization and application stacks
- IO stacks for fixed function off-load

Fixed function off-load

- NIC – Network Interface Controller
- Storage – HDDs, SSDs, etc



Web and mobile traffic drove fundamental shift to IO

Global IP traffic hit 1 Zettabyte in 2016

- Will increase to >3 Zettabytes by 2021

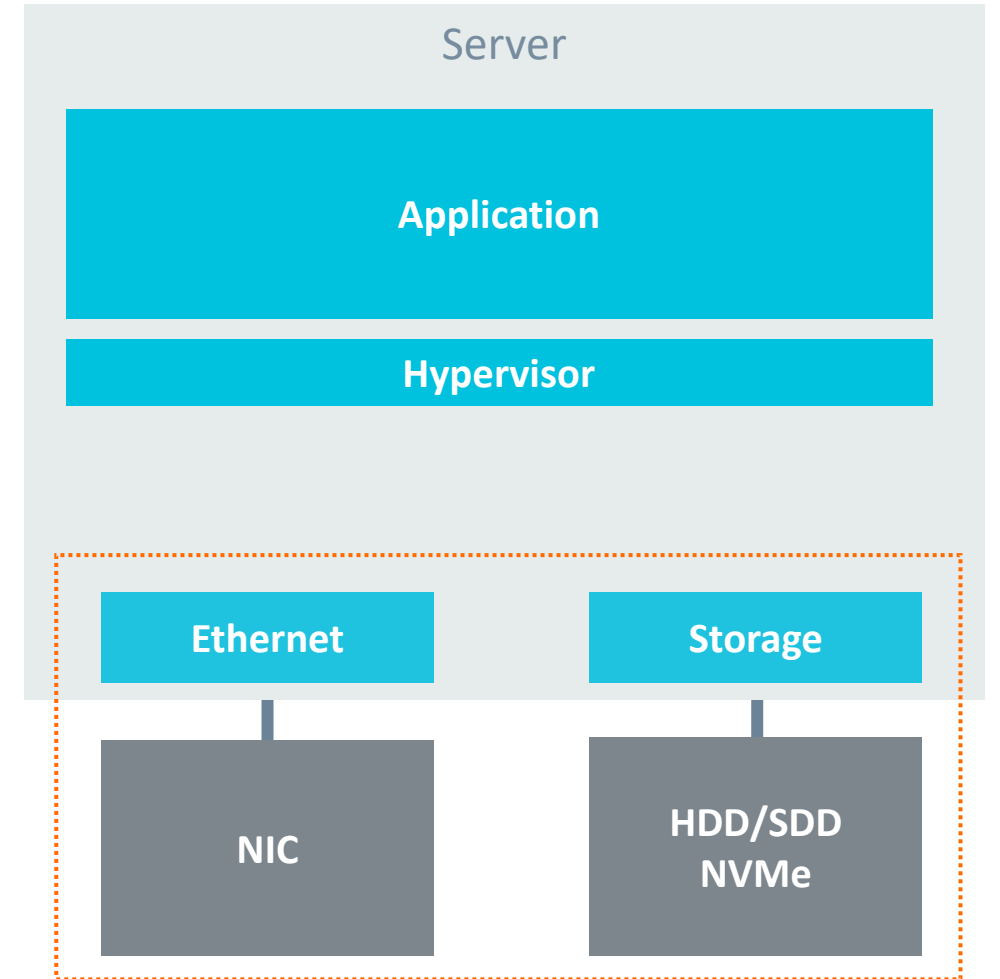
Direct impact on server network bandwidth

- 10Gb Ethernet ->100Gb Ethernet++

Orders of magnitude more storage requests

- NVMe/SSDs - millions of operations per second

Impact on Server: Increased CPU utilization for network/storage IO



Storage and memory disaggregation

Pull data from anywhere, quickly

- 'Over Fabric' such as NVMe-oF

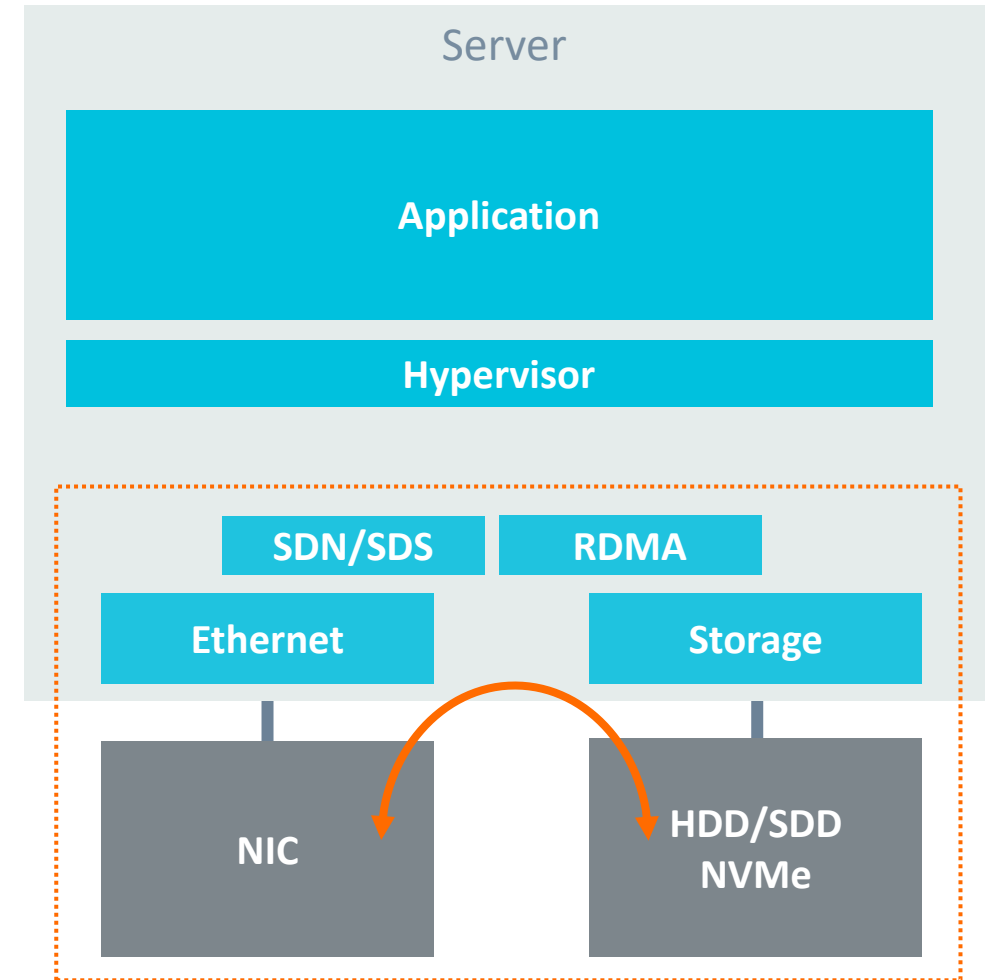
Software defined network/storage (SDN/SDS)

- Reconfigurable virtual networks

RDMA – fast path to remote memory requests

- RDMA - Remote Direct Memory Access

Impact on Server: SDN, SDS, RDMA CPU utilization and increased east/west traffic



Rise of cloud computing

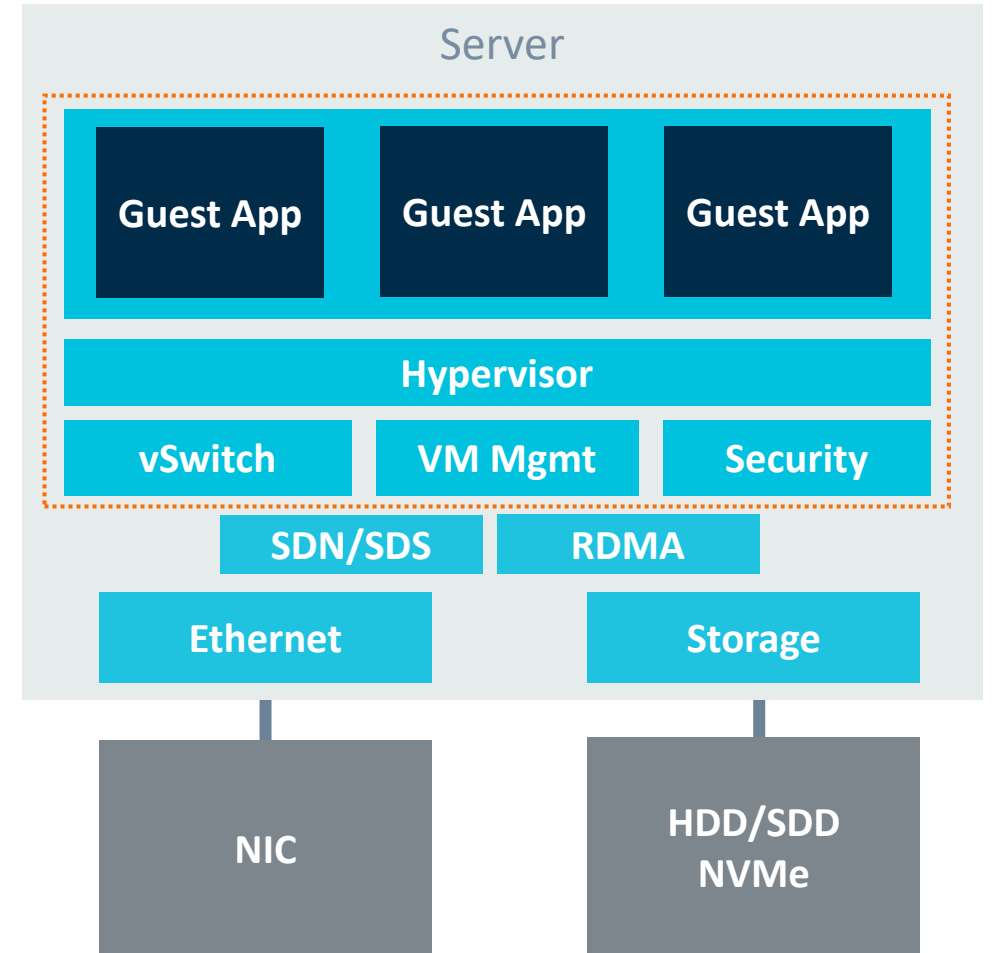
New business models built on 'as a service'

- SaaS, IaaS, PaaS, FaaS

Required new management and security

- Virtual machine orchestration, management
- Virtual machine switching
- Protecting shared resources from guest applications

Impact on Server: More CPU cycles devoted to VM/container management and security



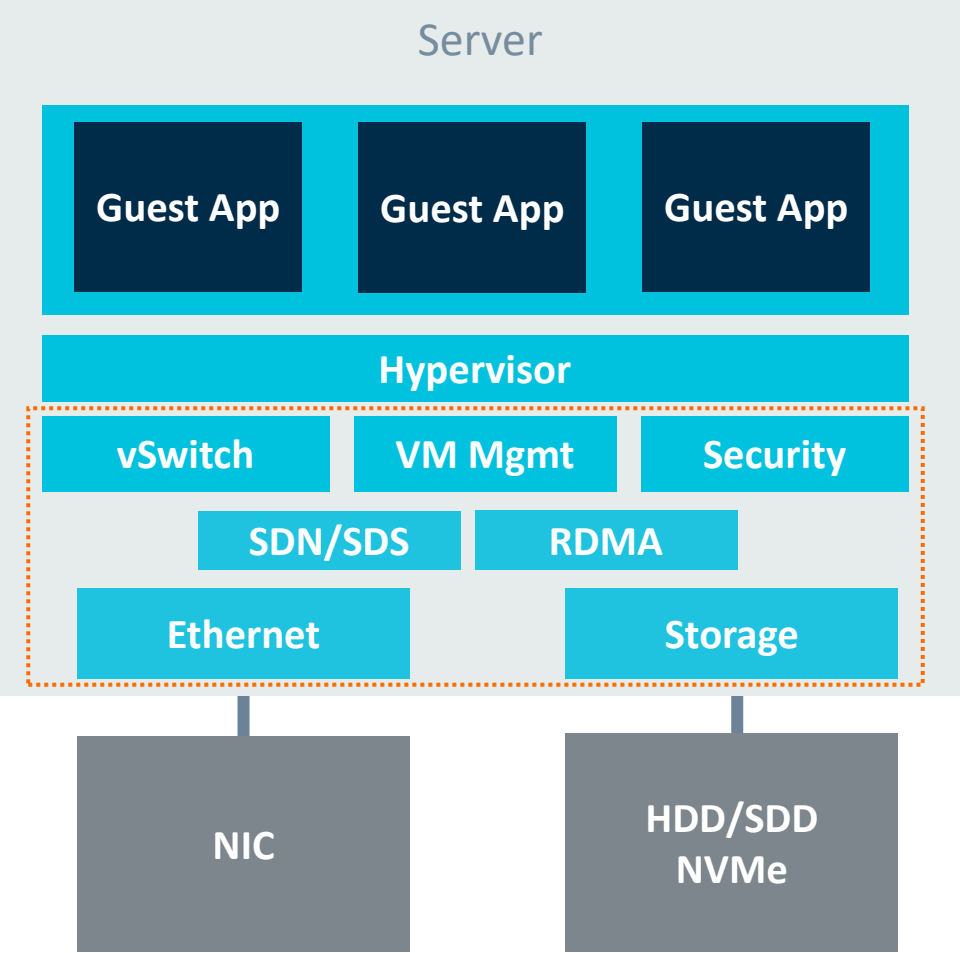
Significant overhead cost for offload and management

\$3.5B lost revenue opportunity per 1M servers/yr

- 16 vCPUs⁽¹⁾ consumed for offload tasks
- Cloud CPU rate = \$0.025⁽²⁾ per vCPU/hr
- \$220/vCPU/yr
- 1M servers * 16 vCPU * \$220/vCPU/year = \$3.5B

⁽¹⁾ Estimated based on 50GbE NIC with 96 vCPU instance

⁽²⁾ Based on average Amazon AWS M5 and Microsoft Azure Dv3 reserved pricing per core



SmartNIC Offload Solution – a better approach

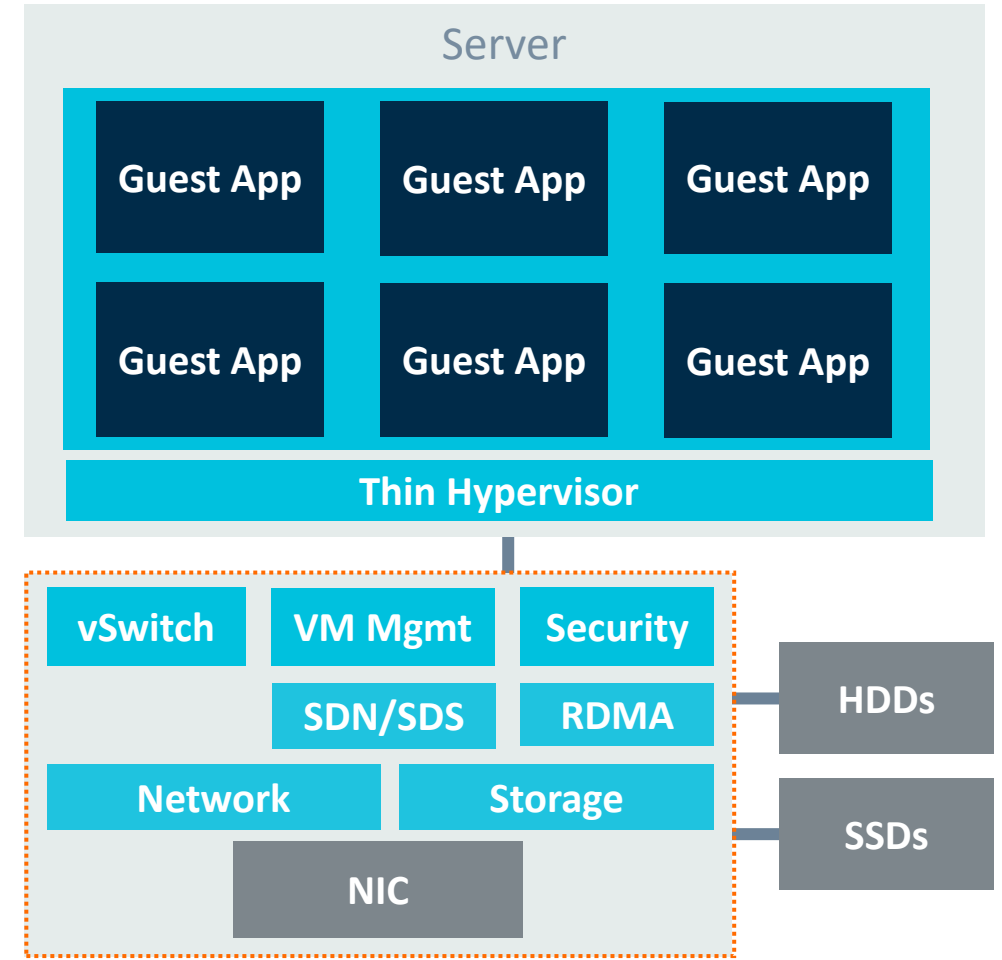
Add the 'Smart' in NIC

- Add compute capable of running Linux OS & offload SW stacks

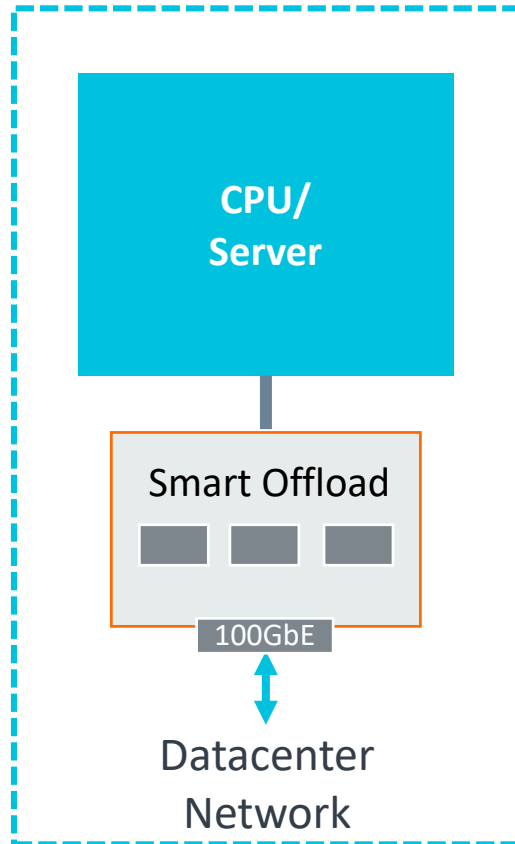
But it's more than a NIC

- Orchestration, VM management, security, storage
- Run vSwitch, SDN/SDS software stacks

Impact on Server: Frees all CPUs for applications



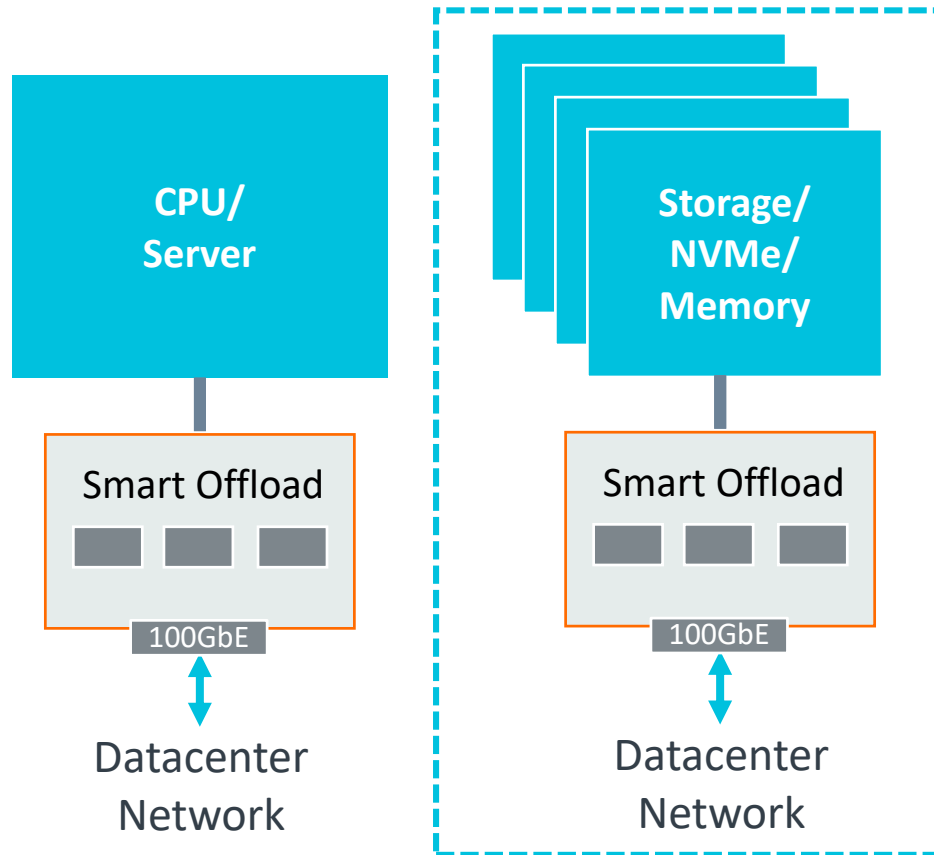
Smart offload becomes the datacenter orchestration point



Datacenter orchestration:

- Decide where to deploy applications/VMs/containers
- Map/re-map storage to applications
- Map/re-map datacenter network topologies

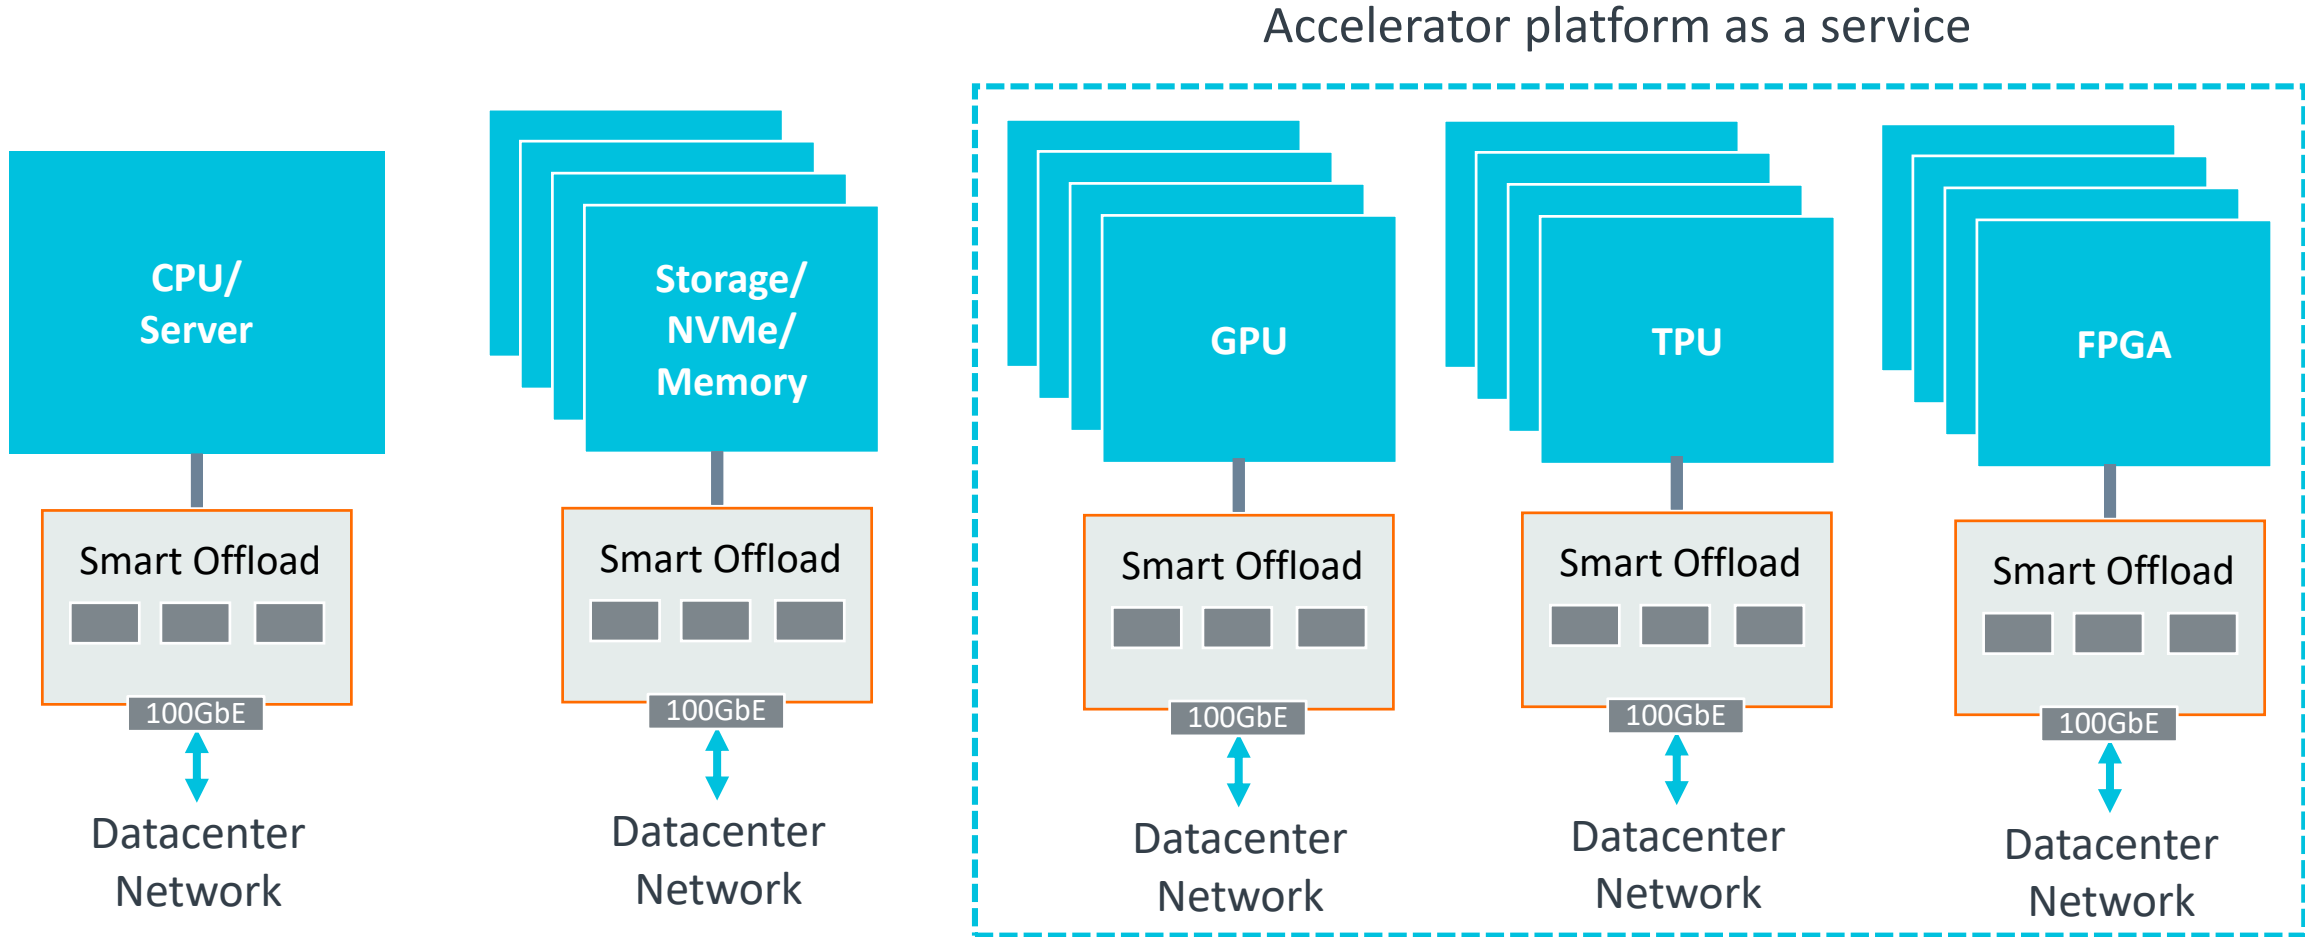
Smart offload as a host for cloud storage



Host for disaggregated storage and memory:

- Pools of HDDs, SDDs, NVMe and DRAM memory
- Operates cloud object, file and block storage

Smart offload as a host for cloud acceleration



Challenges of moving compute to smart offload

10x efficiency gain required


Servers:

- 100W~200W processors focus on compute perf
- Large footprint, large heat sink, sophisticated cooling

Offload devices:

- ~20W-35W ASIC budget for mainstream
- Attached via PCIe slots or Mezz card
- Smaller footprint, smaller heat sink, much less sophisticated cooling

OCP 3.0 Mezz Card – 20W ASIC budget



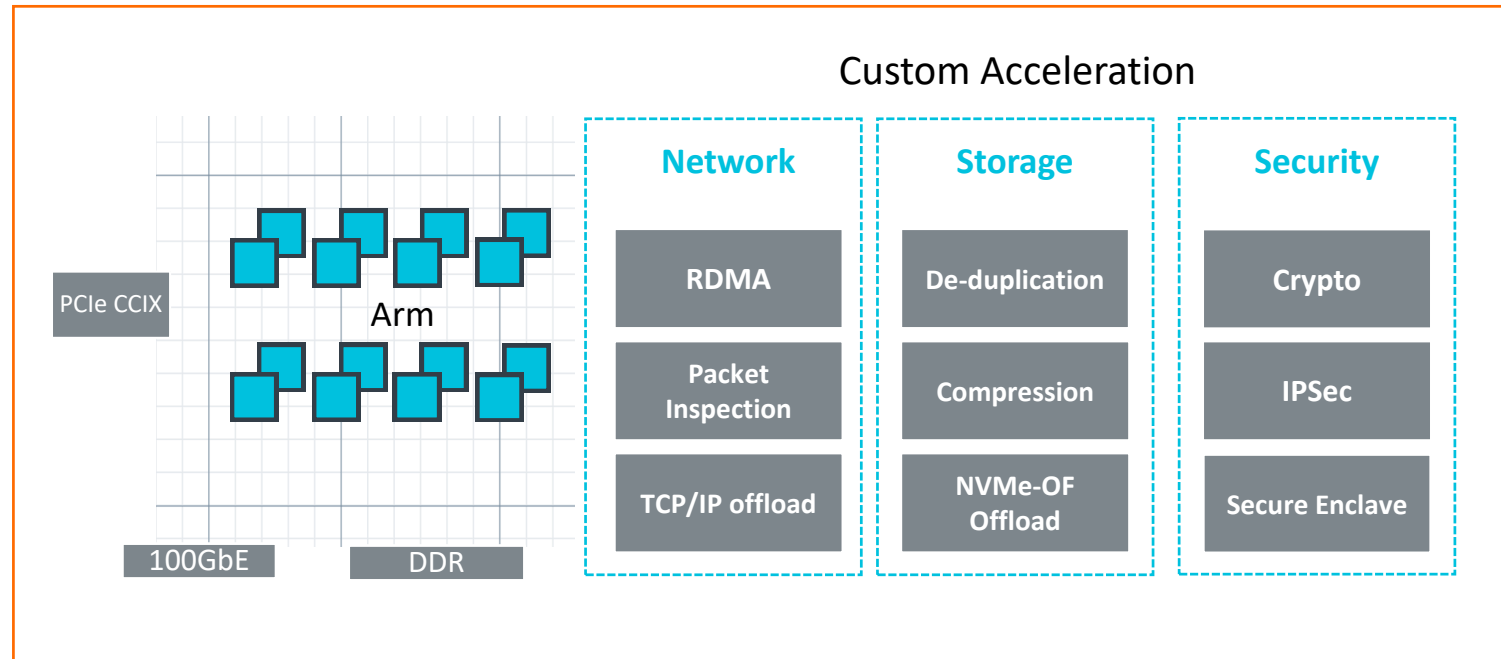
<https://www.opencompute.org/files/OCP18-EngWorkShop-OCP-NIC4.0-v05-20180305-add-note.pdf>

The path to 10x efficiency gain

Arm = better Perf/W


Move functions from server CPUs to more efficient accelerators

Customize with 'secret sauce'

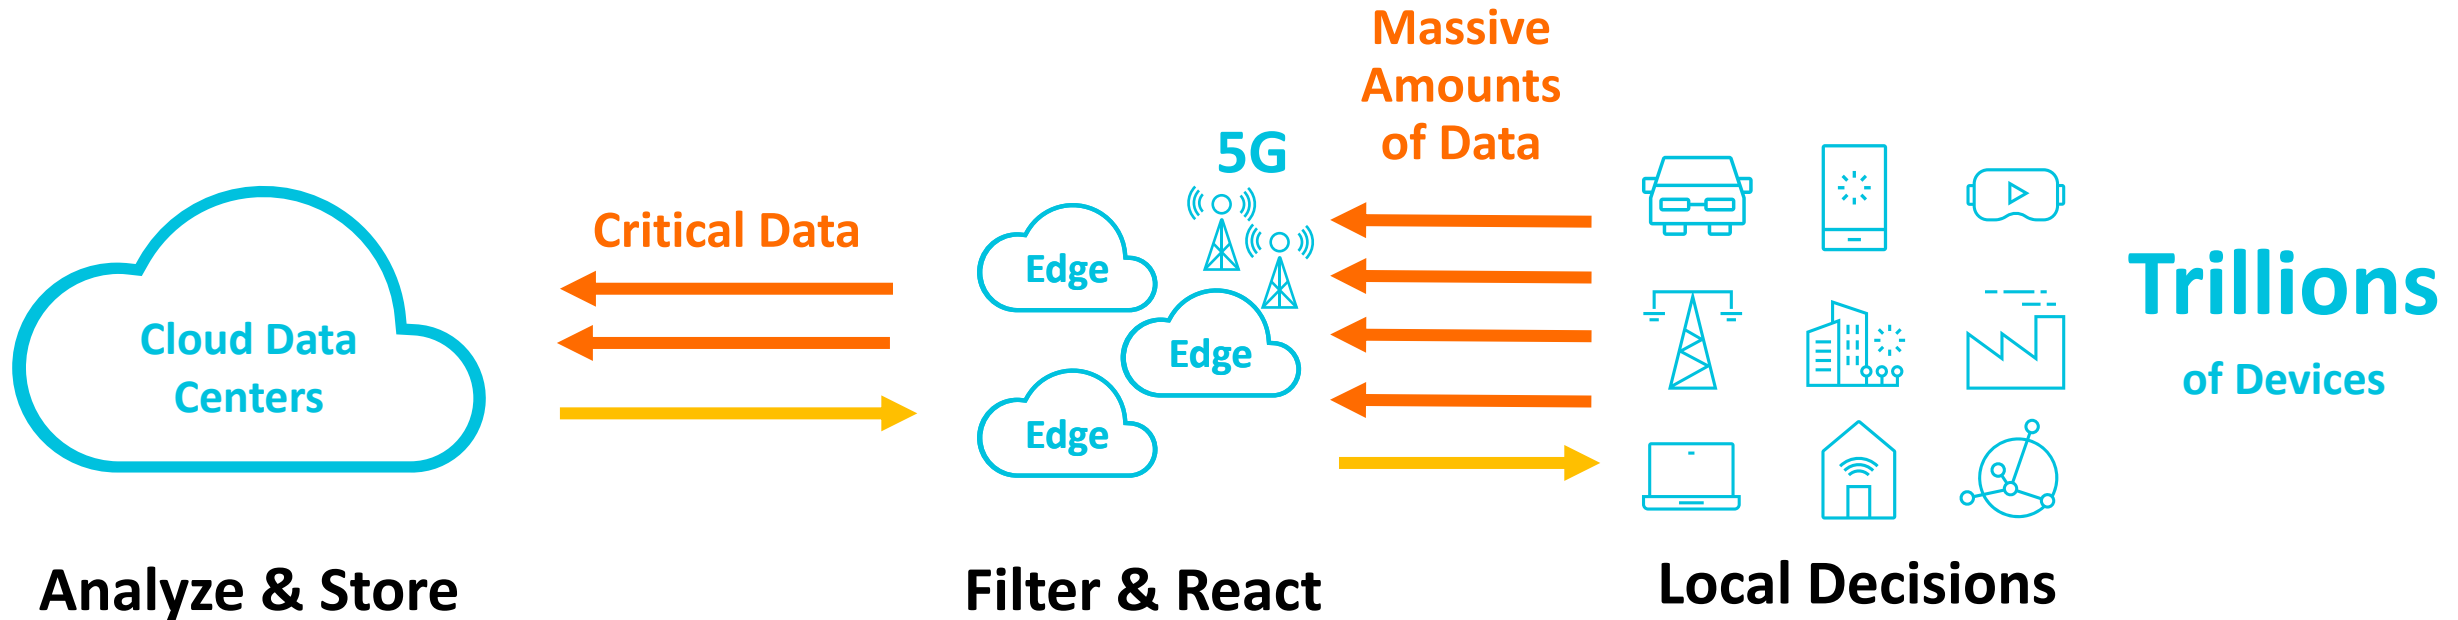


From smart offload in the datacenter
to accelerated edge compute

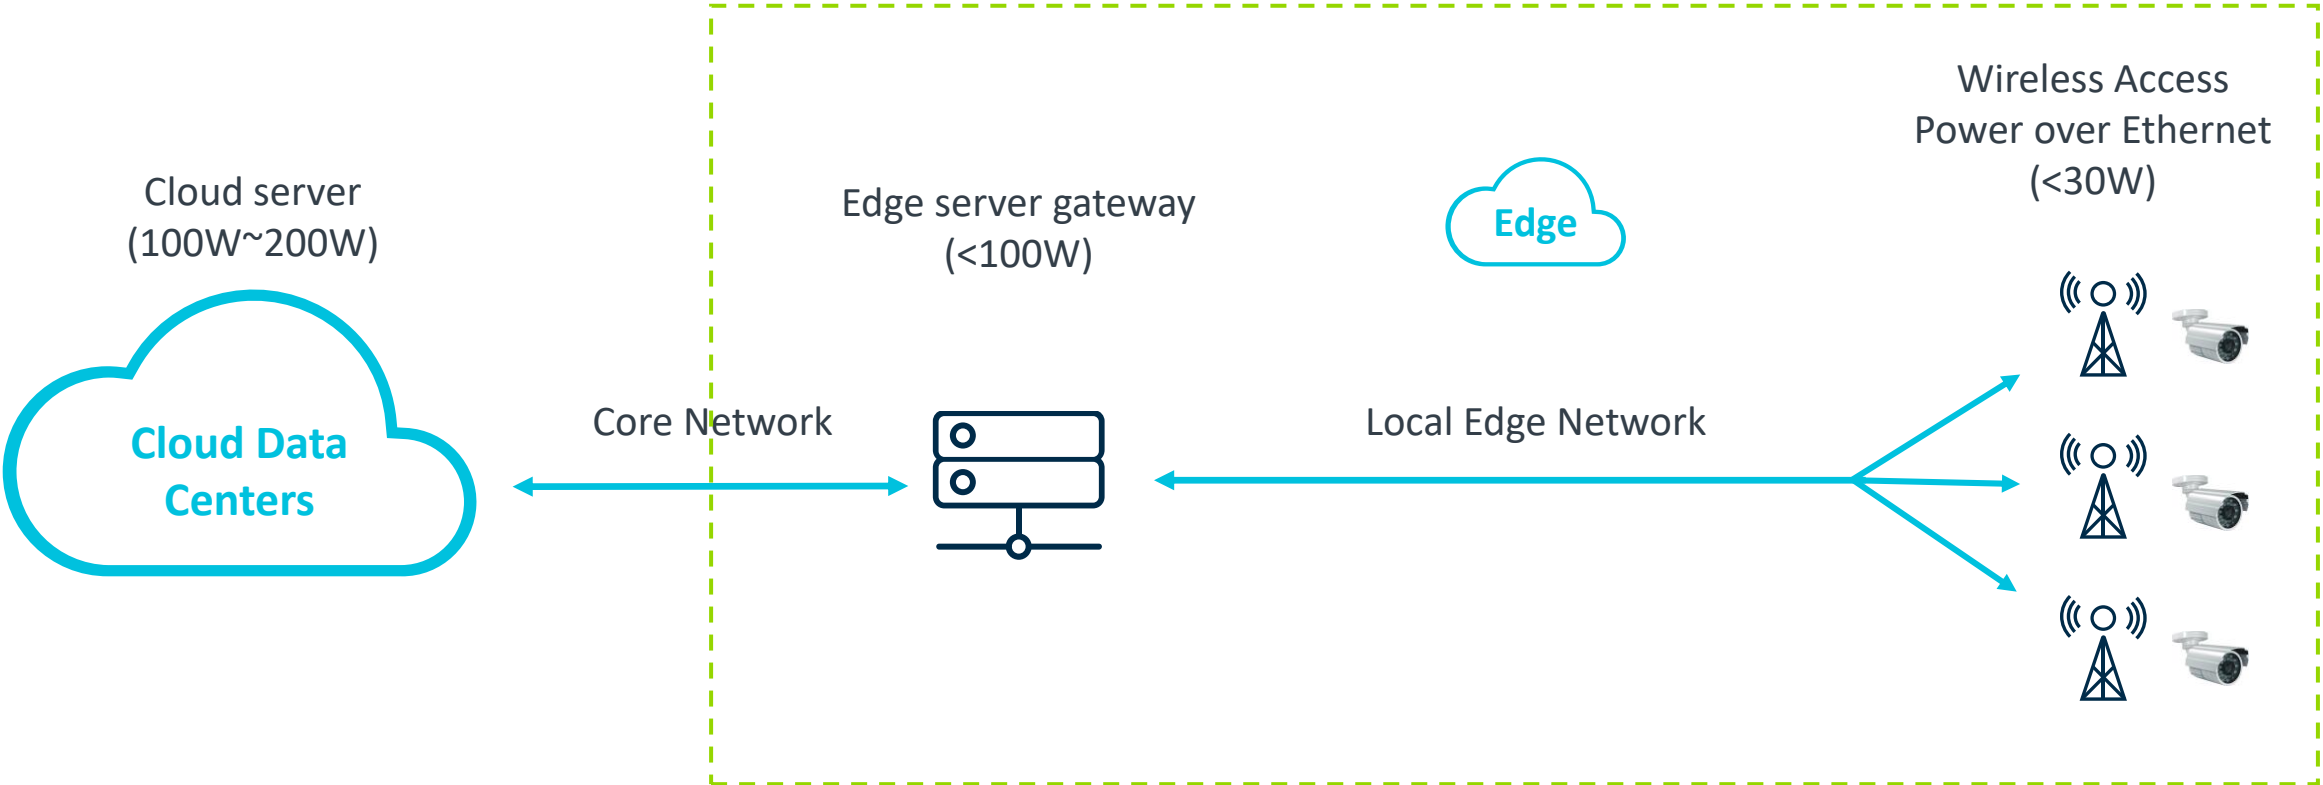
4G infrastructure latency is ok for today's mobile broadband



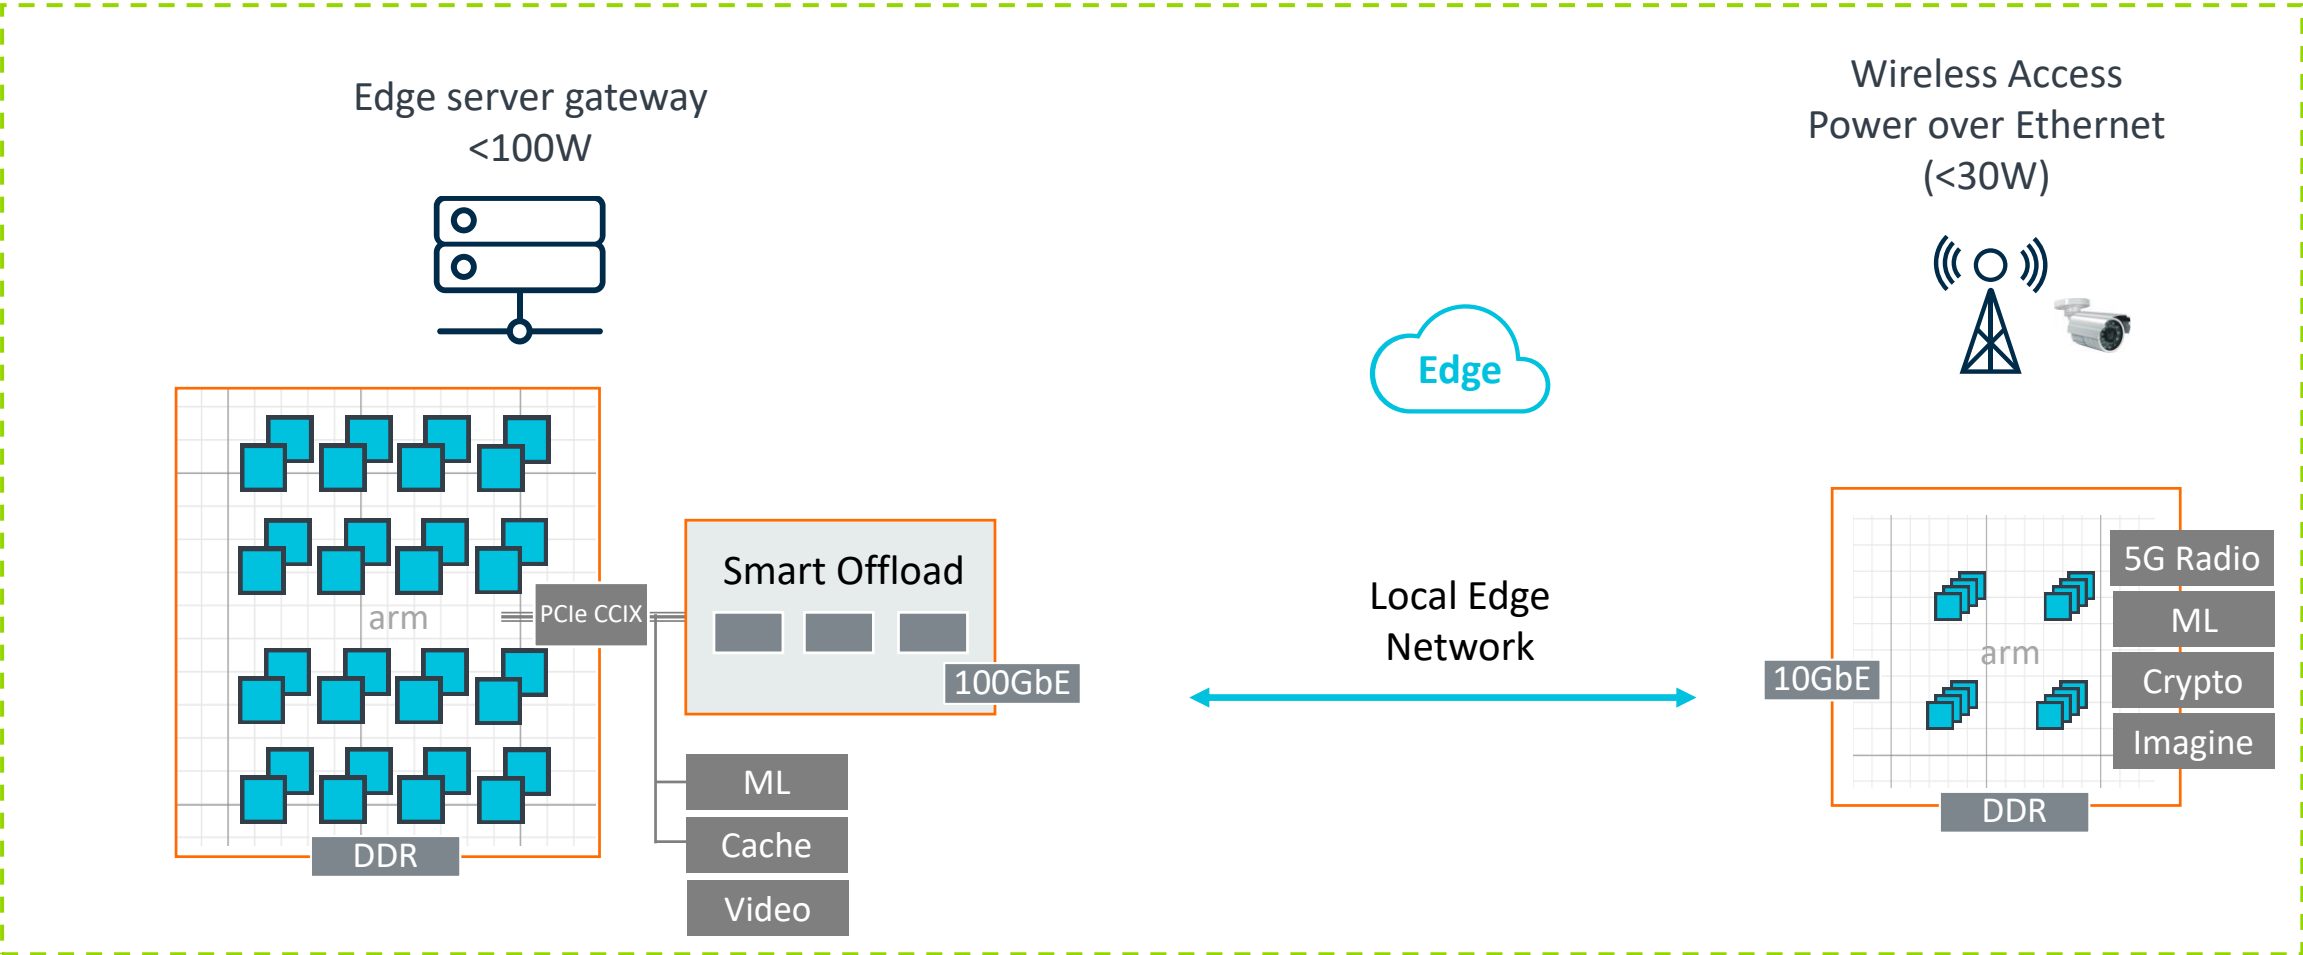
Data consumption and latency driving future designs



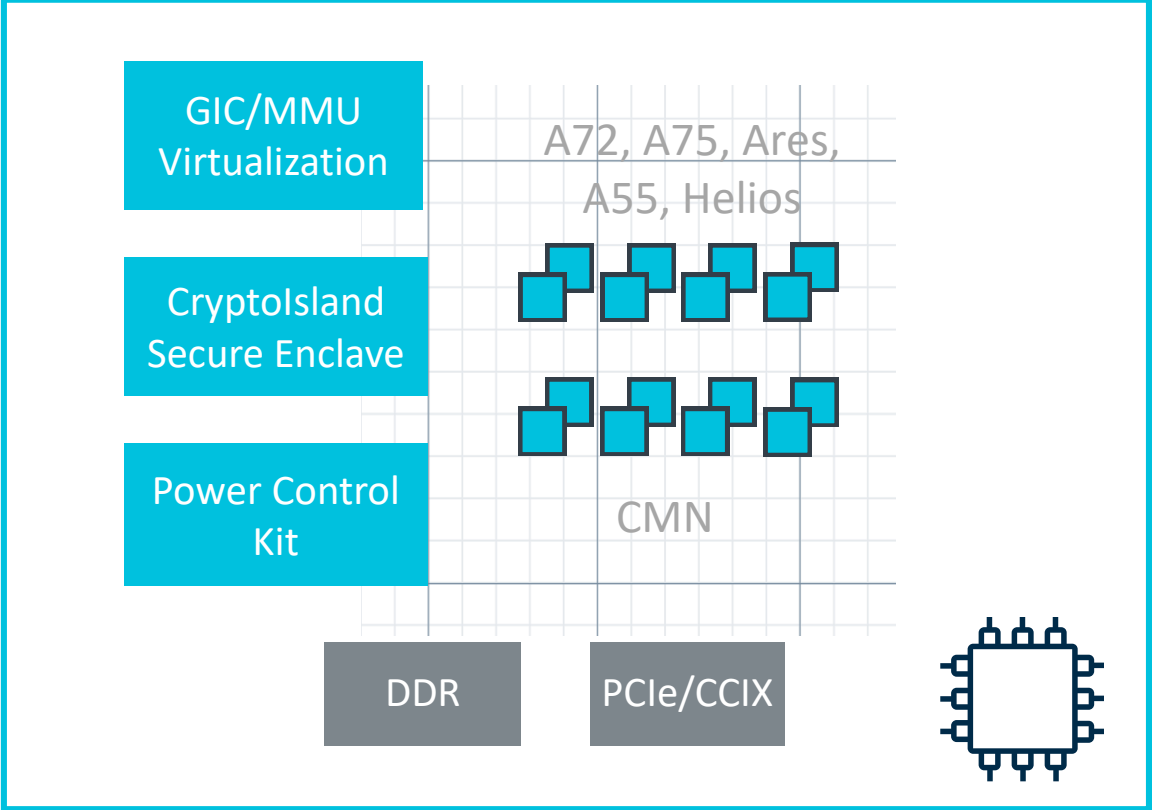
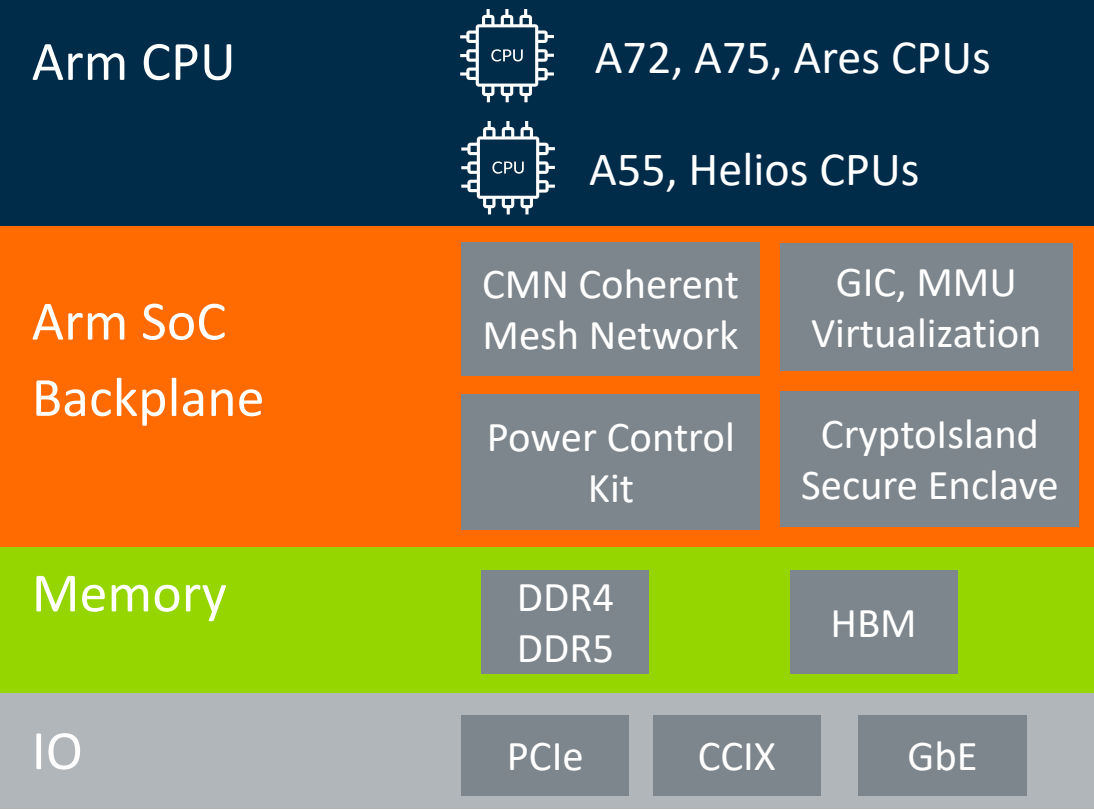
Accelerated edge compute example deployment



Accelerated edge compute example systems

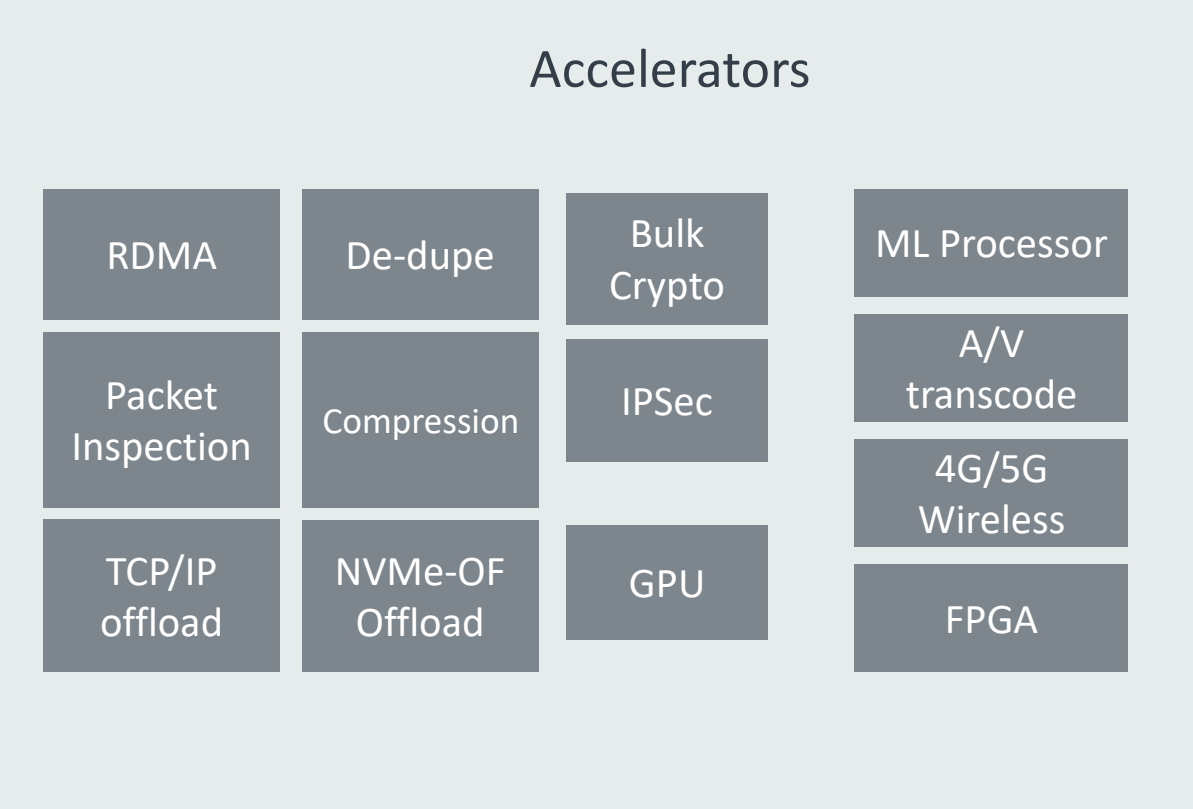
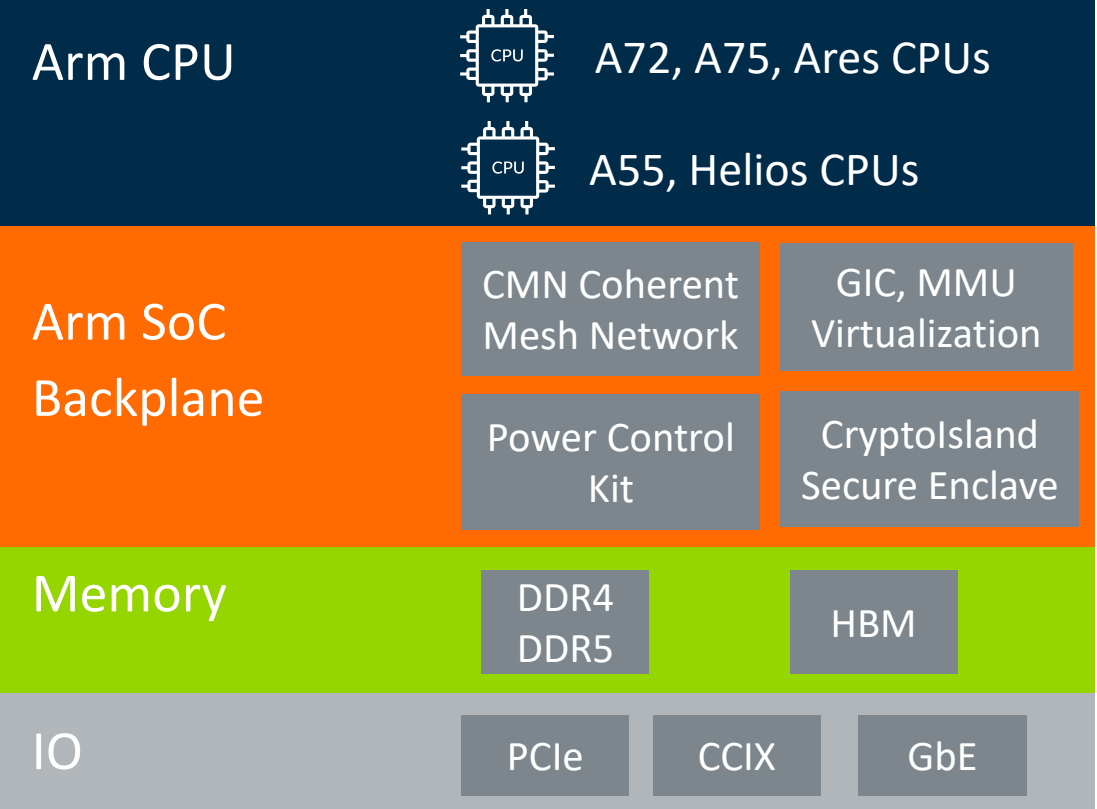


Arm compute and IO subsystem



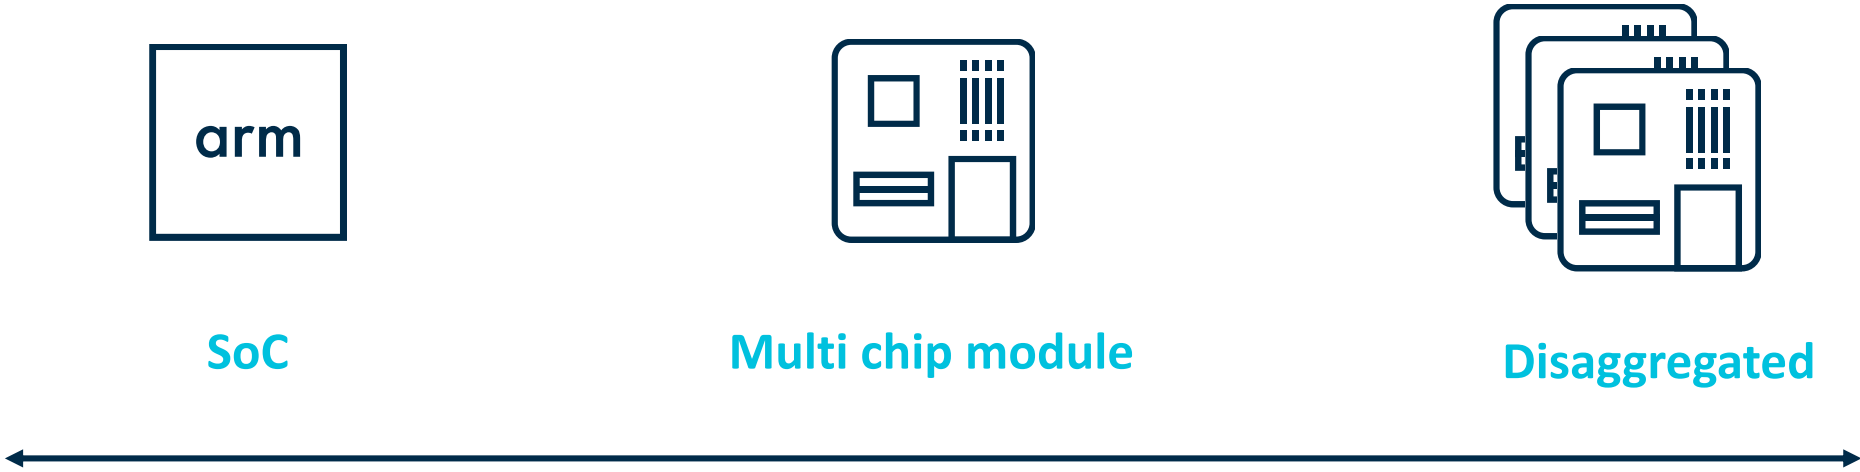
Common Software Platform, SBSA, SBBR, Arm ServerReady
 Arm Architecture v8.x-A, AMBA

Accelerator integration



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

Smart offload requires scalability



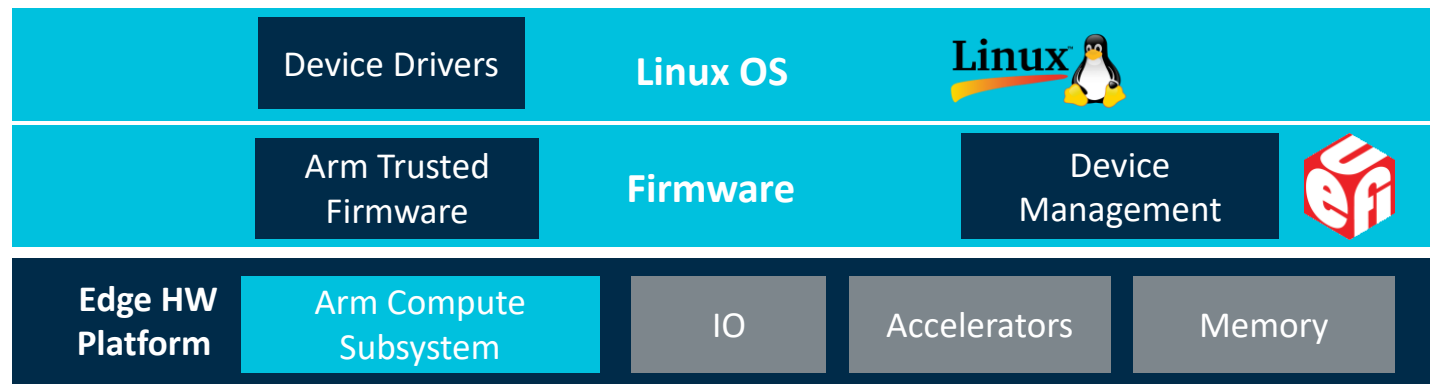
Platform Scalability from fully integrated SoC to multichip modules to fully disaggregated



Accelerated edge cloud stack

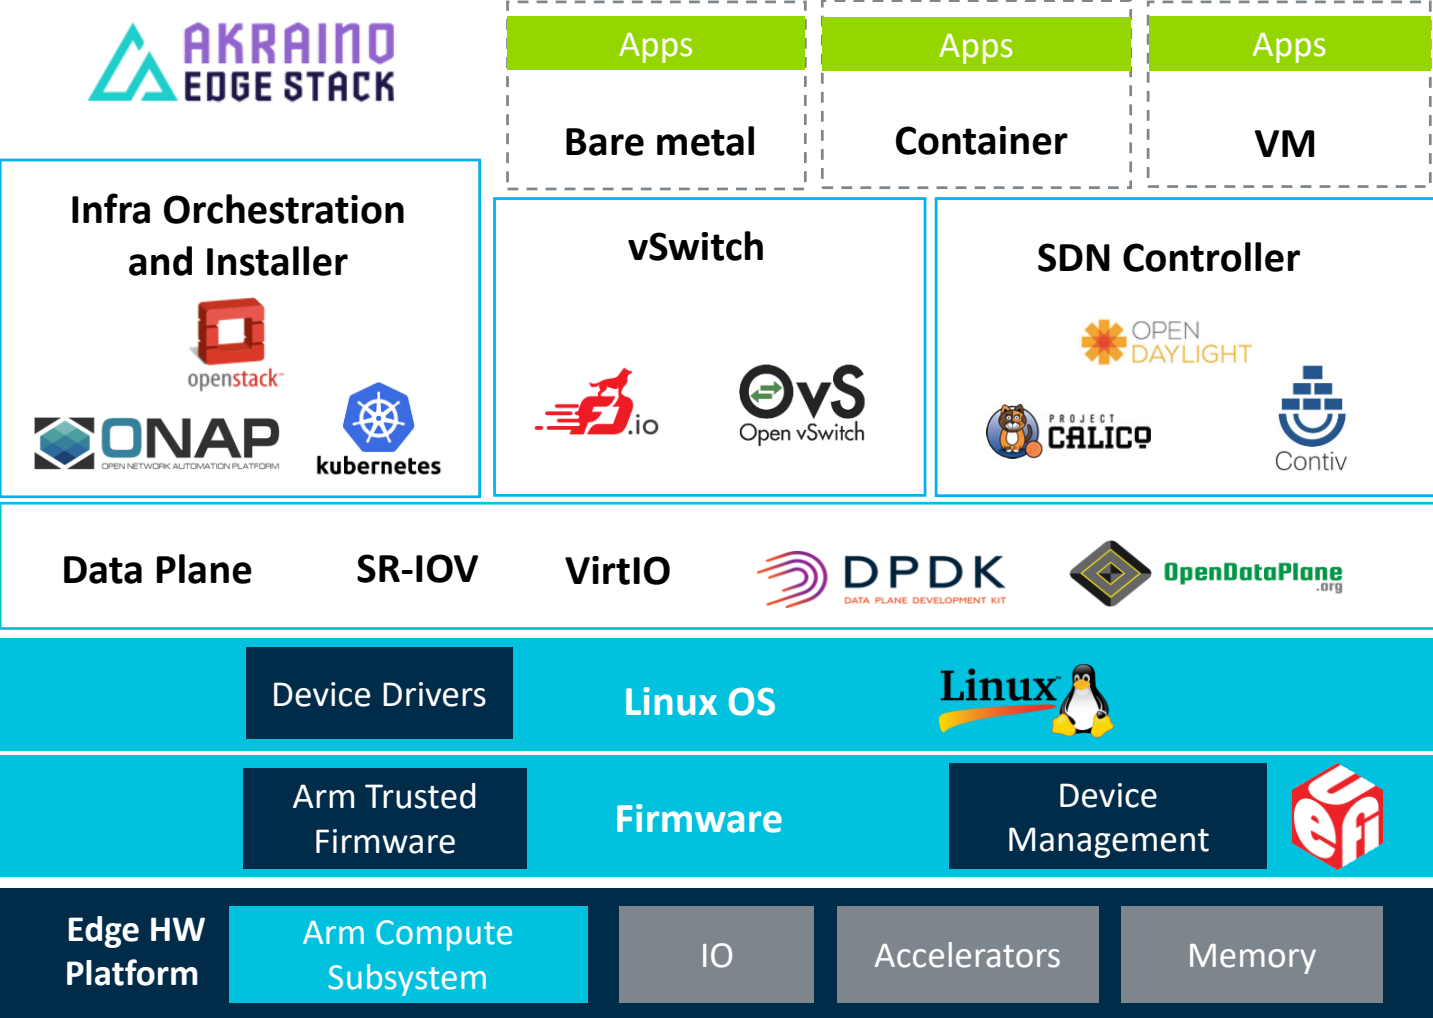
arm ServerReady

- Test suites to ensure major OS vendor compatibility
- SBSA: Server Base System Architecture
- SBRR: Server Base Boot Requirements



Foundation for cloud
deployments

Accelerated edge cloud stack



Orchestration and installation

- Automated deployment, discovery, management

Edge stack blueprints

- Network management, NFV
- Vertical deployments (ex: IIoT)
- Smart City and edge cloud

Arm powering smart offload and the accelerated edge



Scalable, secure, virtualized compute subsystems



Combine with custom acceleration for target offload functions



Ready for scale-out deployments in the datacenter and at the accelerated edge

Trademark and copyright statement

The trademarks featured in this presentation are registered and/or unregistered trademarks of Arm Limited in the EU and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

Copyright © 2018

Thank You!