Sponsored By

Google intel STORDIS The Open Networking Expert

PENSANDO

Programmable Data Plane Architecture for the Network Edge

Mario Baldi, Diego Crupnicoff, and Silvano Gai
Distinguished Technologist
Pensando Systems, Inc.

Goals/Outline

- Distributed Services at the Network Edge: the Pensando Platform
- Distributed Services Card Architecture
- Representative use cases
  - How the card architecture supports them
- Performance evaluation of a sample use case

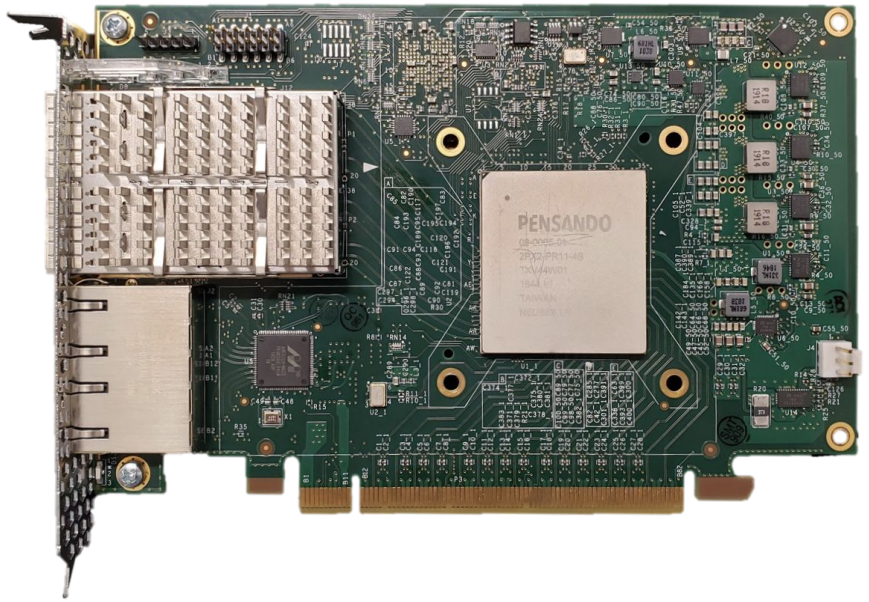

PENSANDO

Pensando Distributed Services Platform

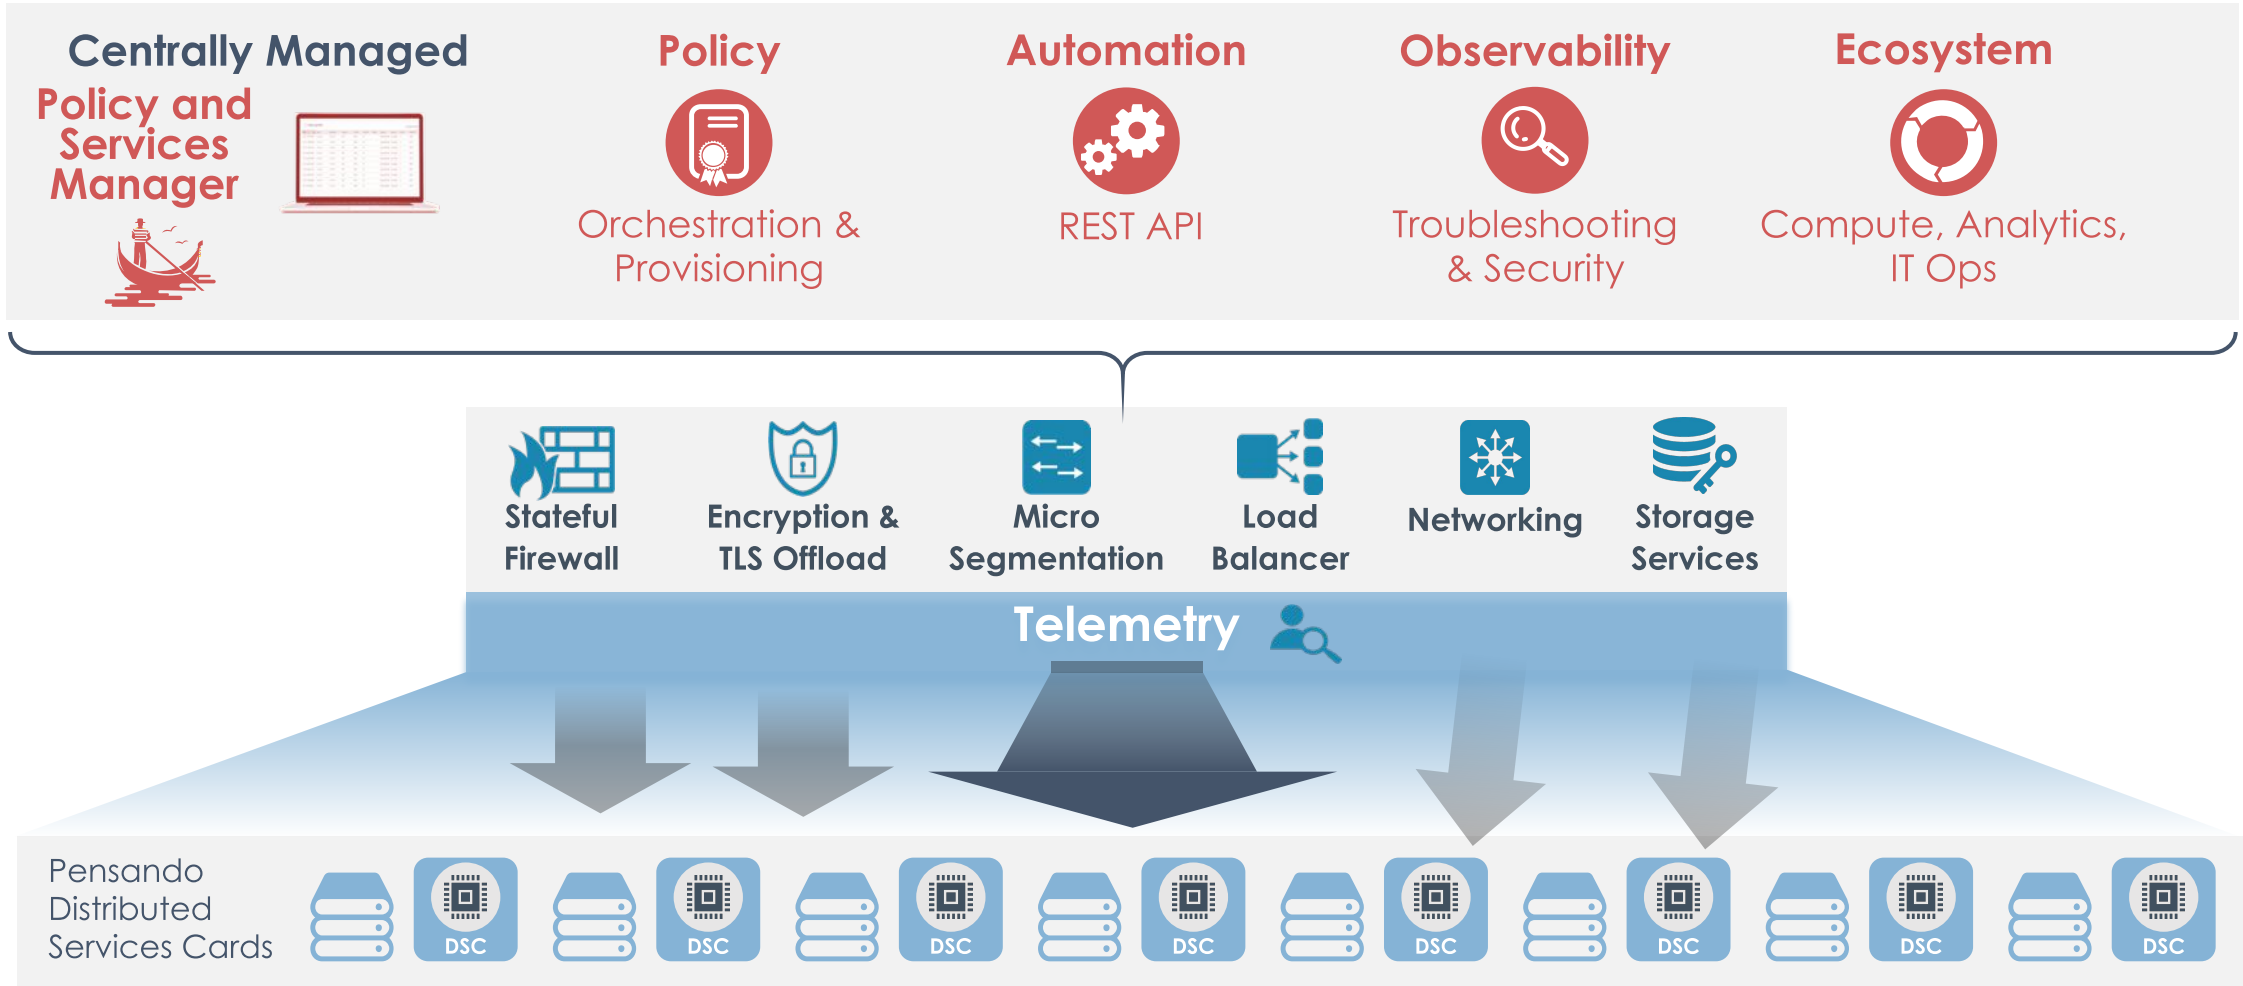

PENSANDO

DSC Architecture

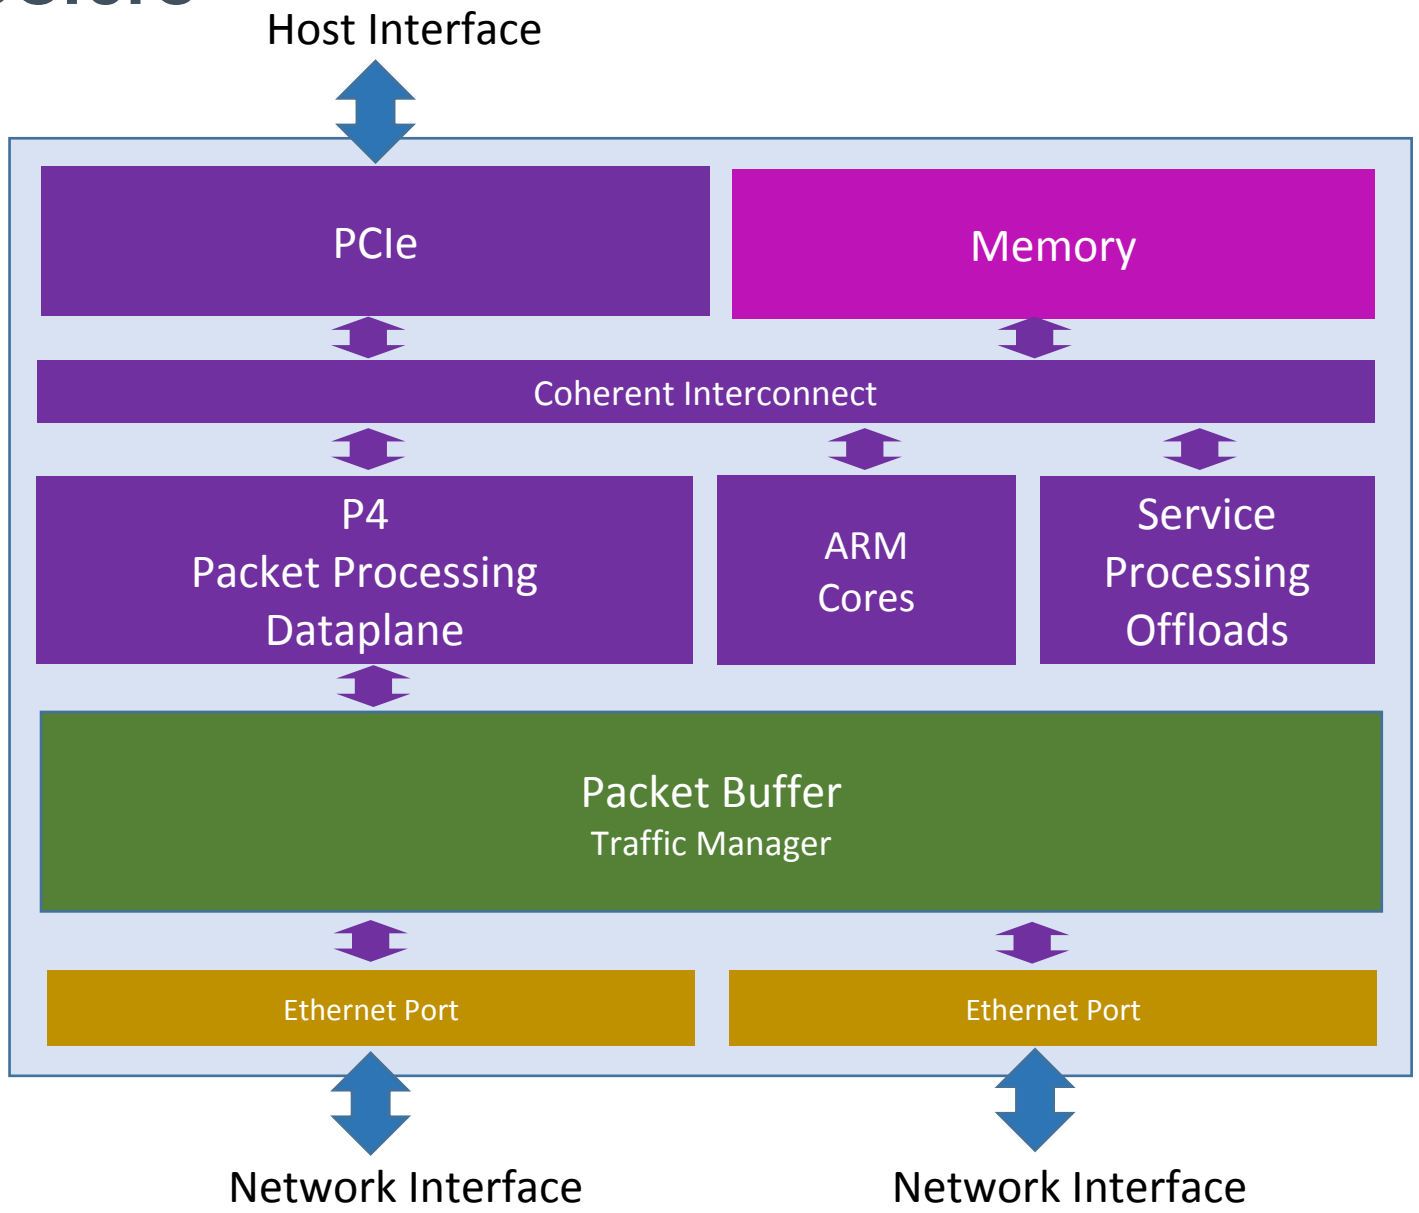

PENSANDO

Sample Use Cases

PENSANDO

TLS Offload

- Protocol processing offload is a typical SmartNIC application
  - The DSC is not a SmartNIC, but can be used as such
- Especially the ones requiring significant resources
  - Memory
  - CPU
- TLS Offloading is a great example
  - TCP connection management
  - TCP state handling
  - TLS session management
  - Data encryption and decryption

PENSANDO

Possible Implementation: Proxy

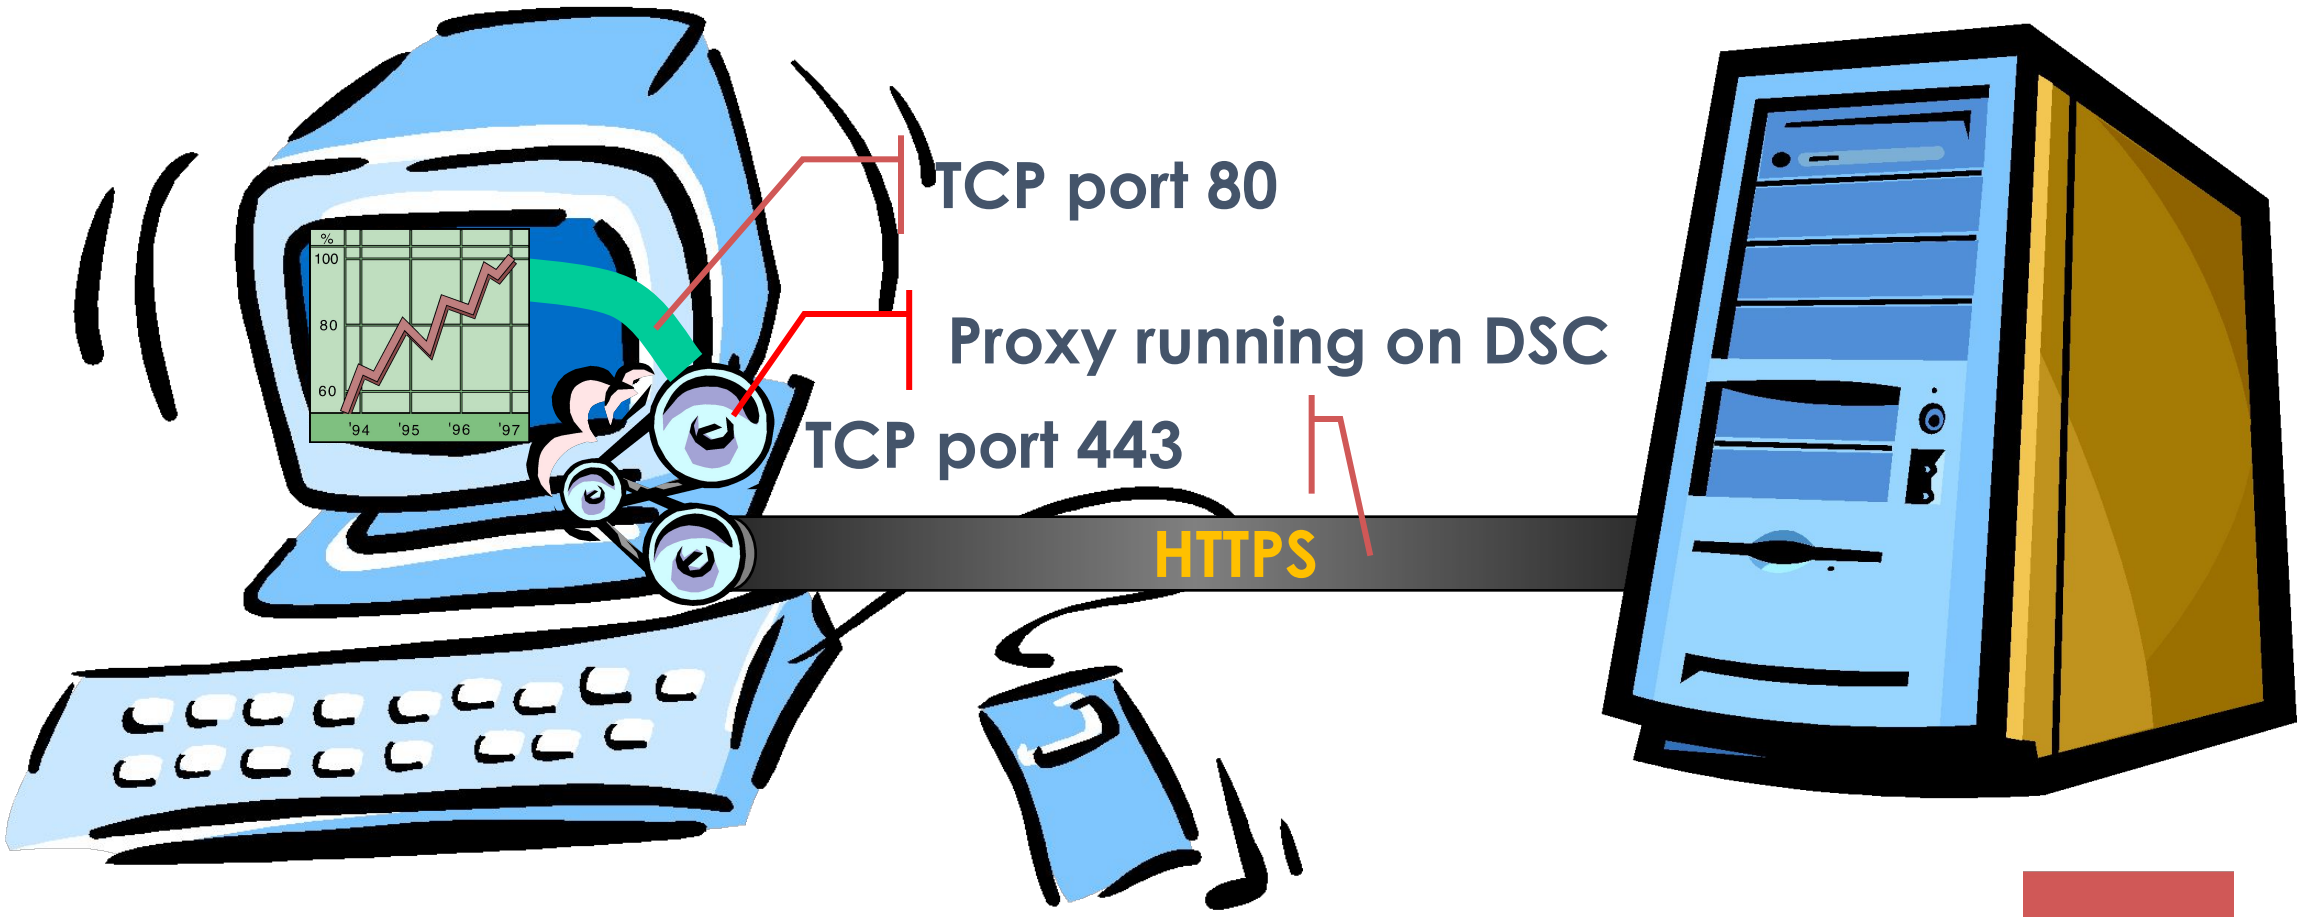

PENSANDO

TLS Offloading Support

- ARM involved only in connection/ session setup/tear down
- Pipeline ensures
  - Wire speed throughout
  - Minimal delay
  - Minimal jitter

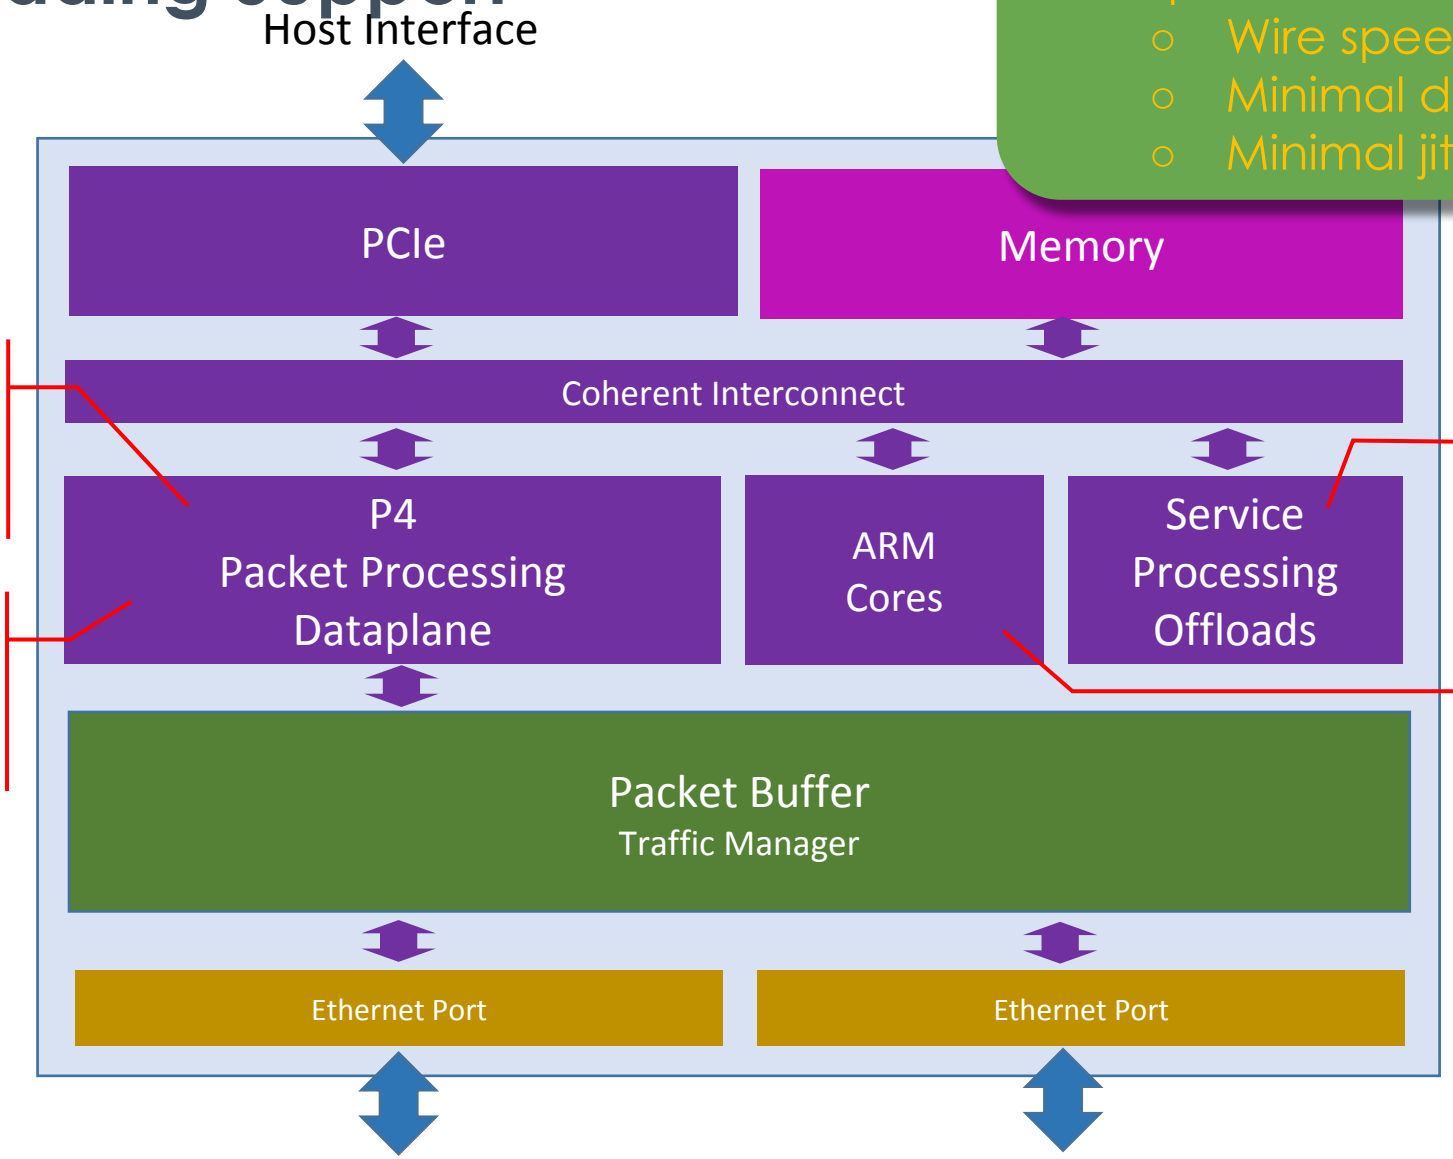

(1) packets are generally processed by the pipeline

(3) subsequent packets are processed in the pipeline

(3.1) encryption/ decryption performed by service processing offload

(2) TCP connection and TLS session initiation packets are forwarded to the ARM cores for software processing

- Handle connection establishment
- Install state in pipeline tables

PENSANDO

NVMEoF/TCP

Non-Volatile Memory Express Over Fabric over TCP Transport

- Access a remote disk as if it were local
  - Through a regular NVMe driver
- Multiple transports including
  - RDMA - Remote Direct Memory Access
  - TCP

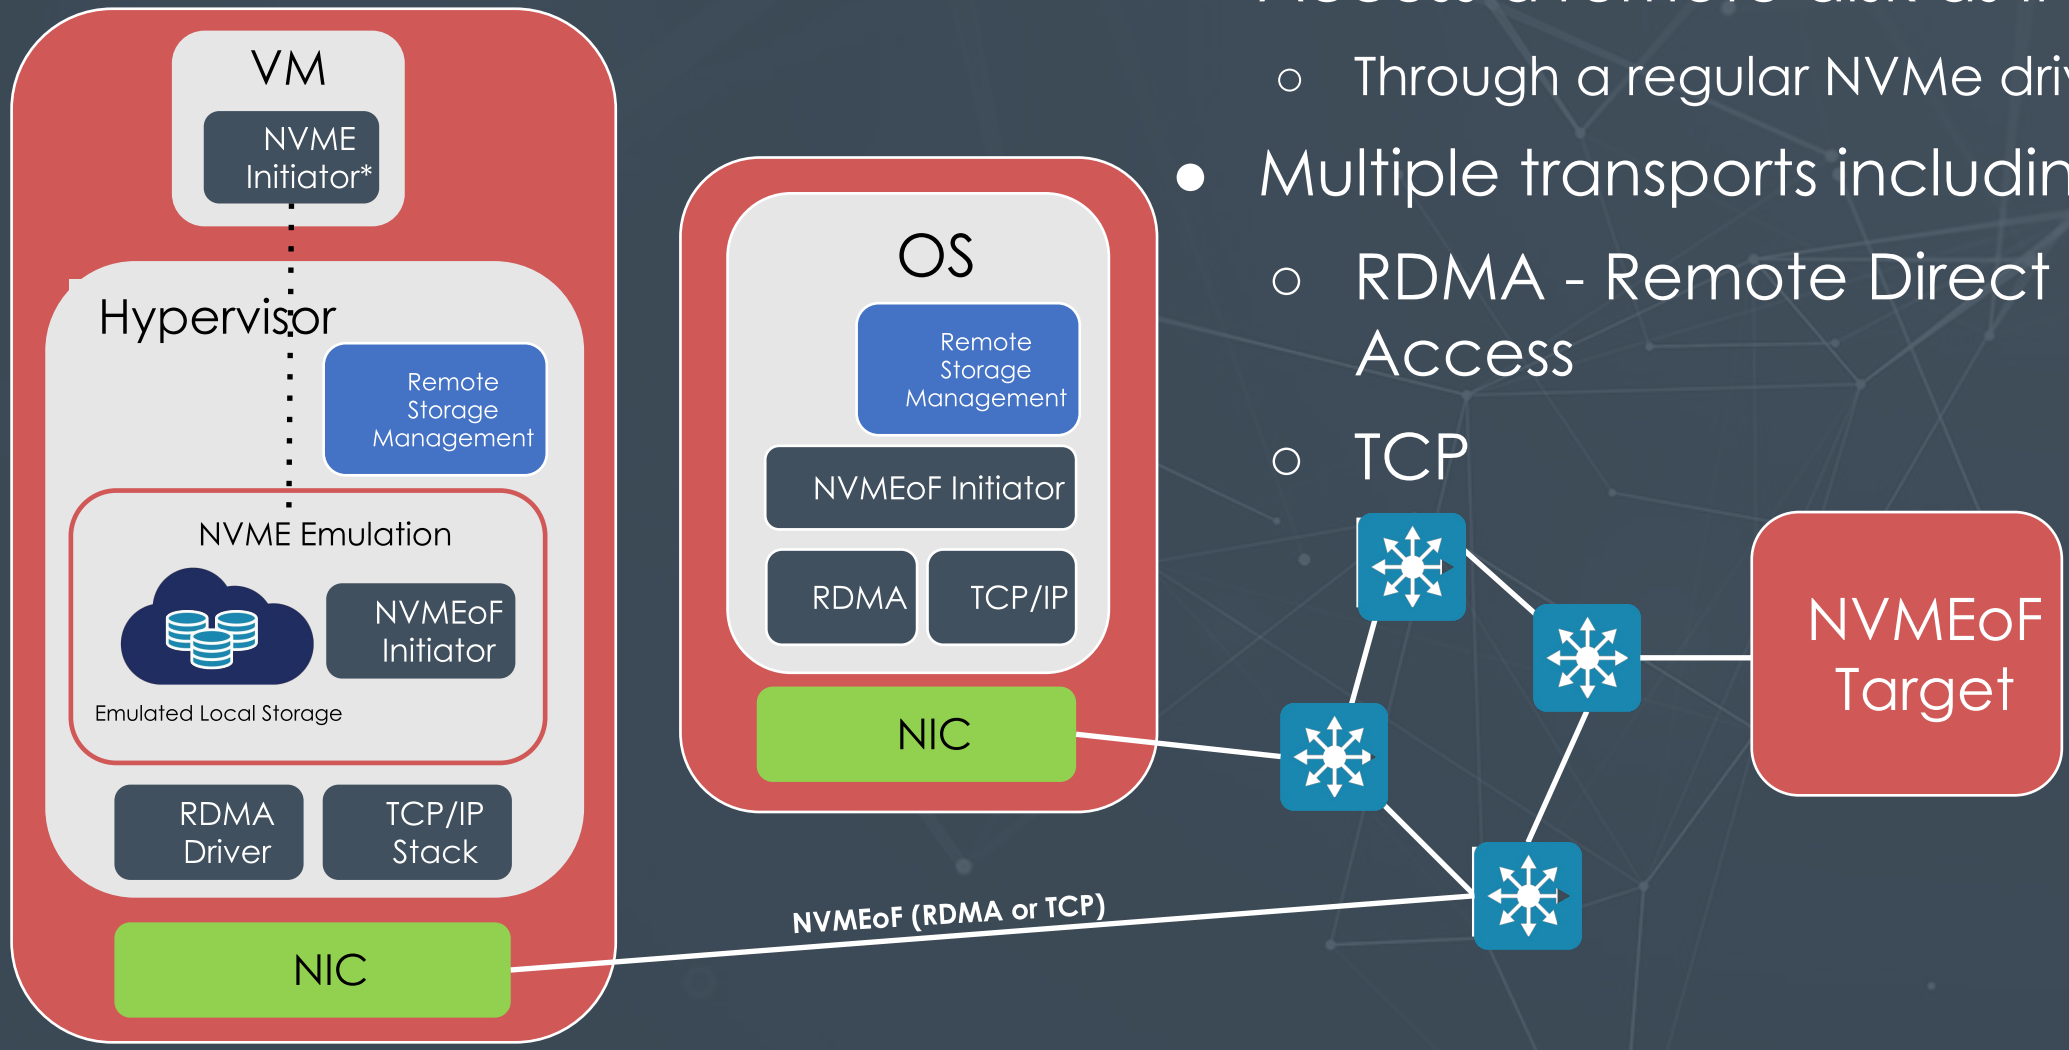

PENSANDO

NVMEoF/TCP Offload

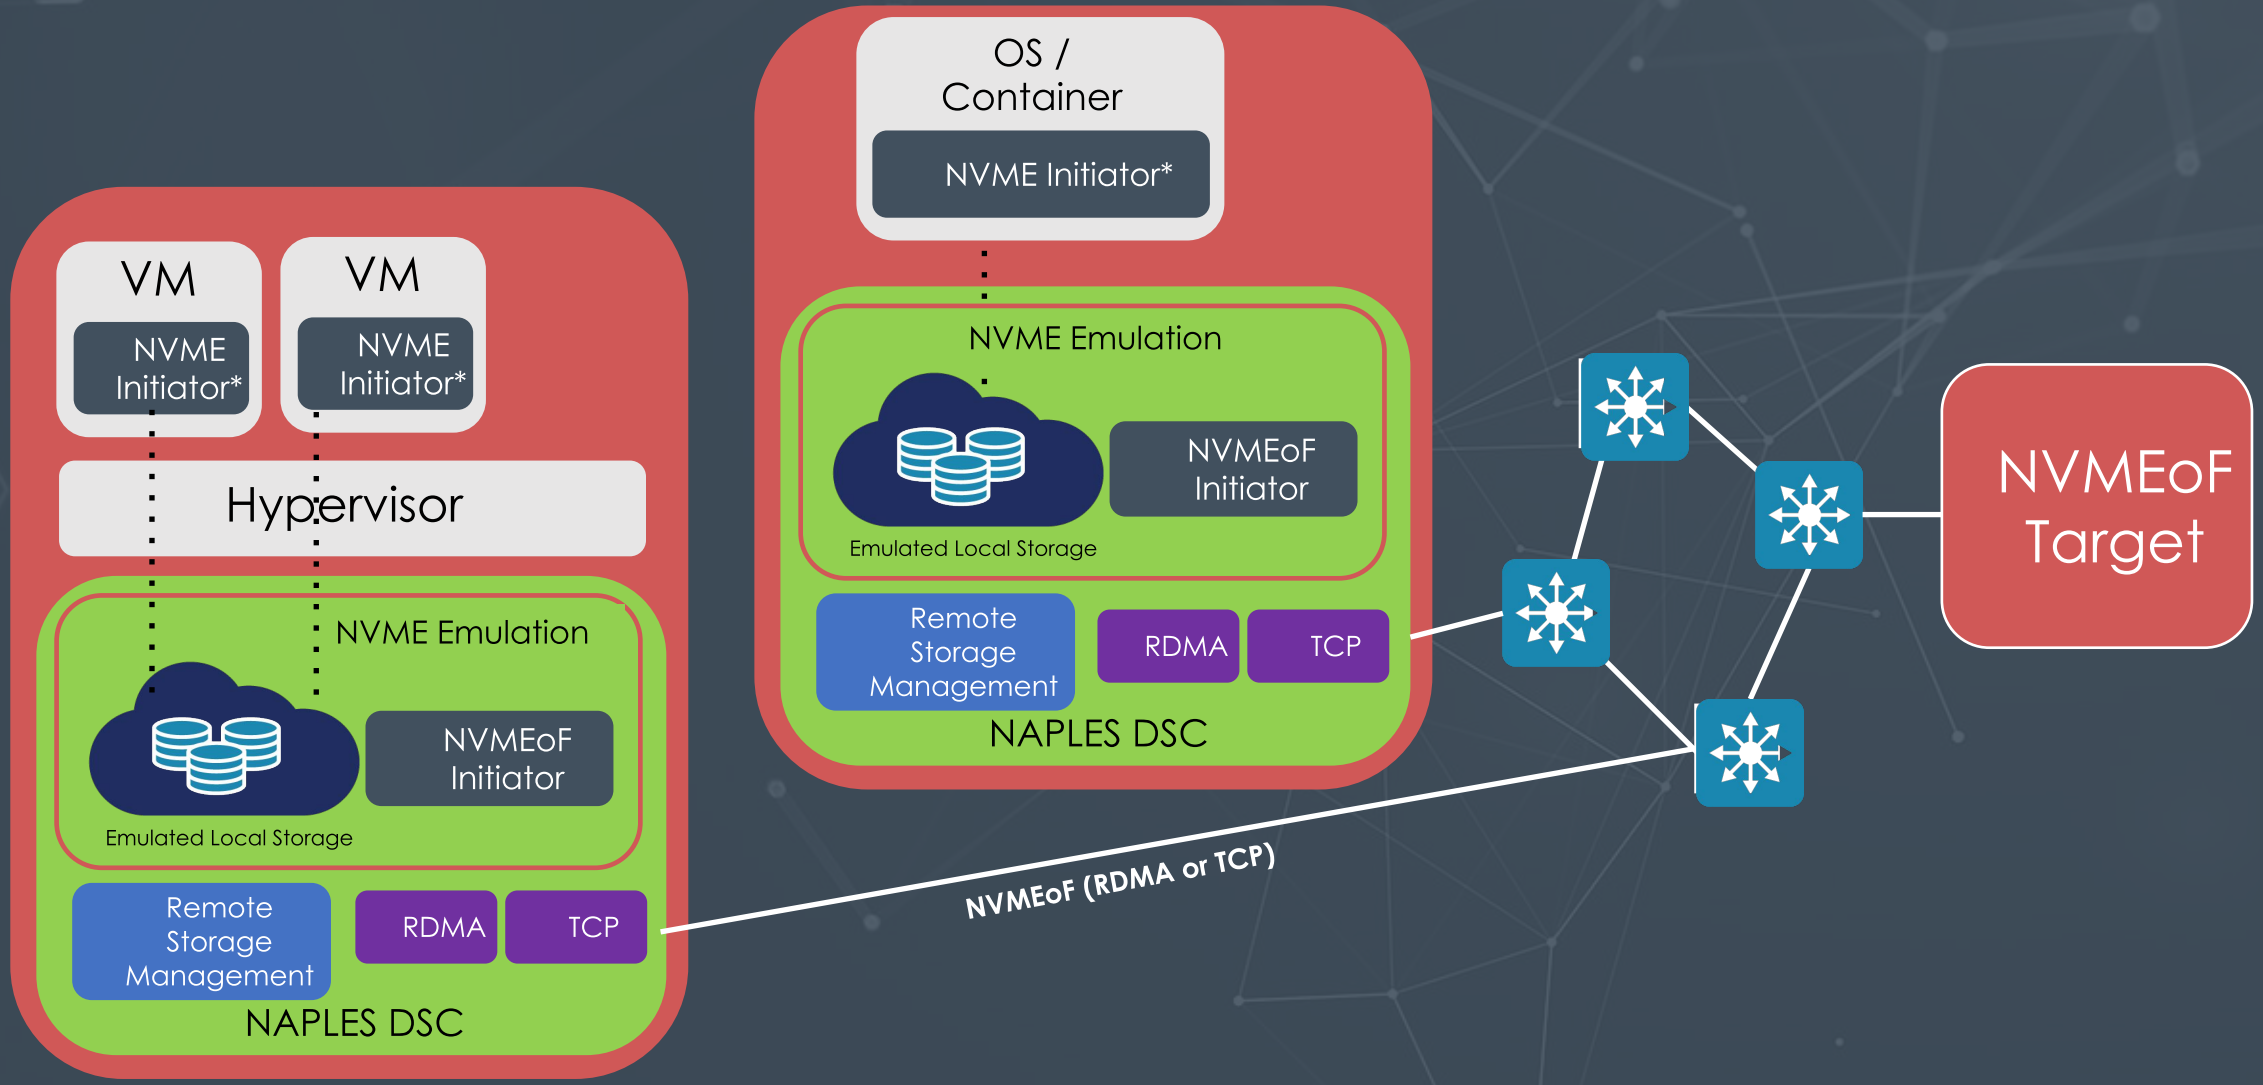

PENSANDO

NVMEoF Offloading Support

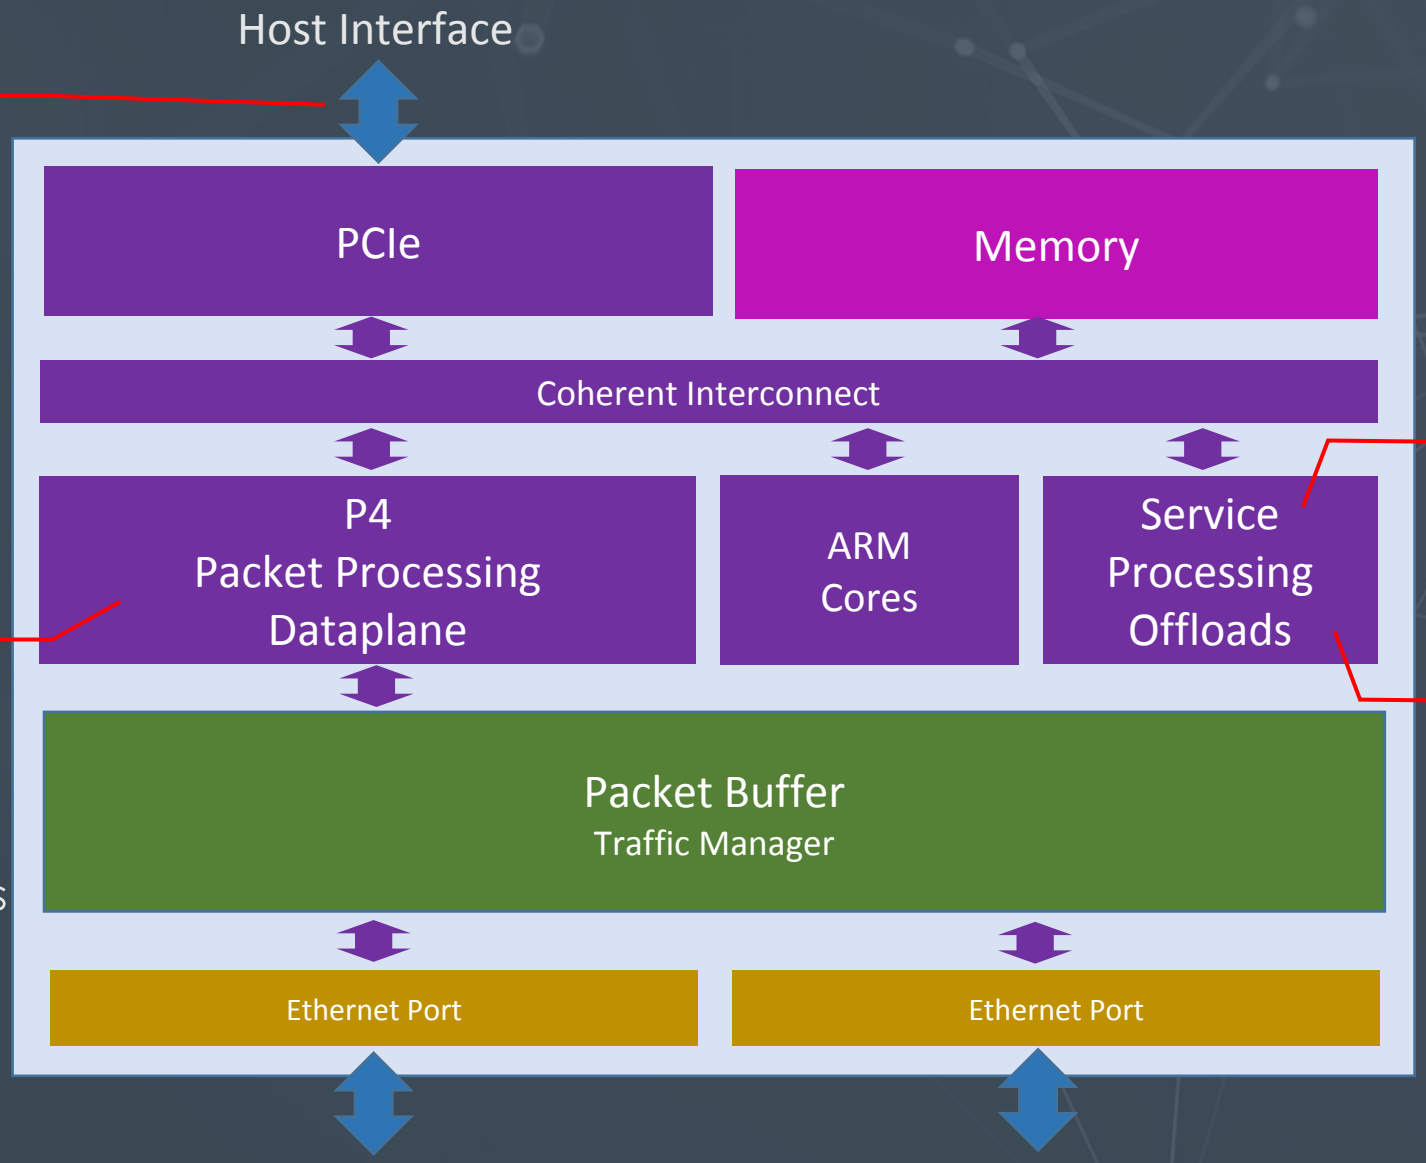

(1) NVMe commands

(2) NVMe commands translated into NVMEoF capsules

- Load balance across remote controllers
- Encapsulation
- TCP segmentation

(2.1) encryption/ decryption; data digest generation/verification

[(3) encryption/ decryption; of data at rest]

PENSANDO

Distributed Stateful E-W Firewall

Firewalling E-W traffic is particularly challenging

- Large volume compared to N-S
- Applications expect small latency
- Appliances are not suitable as they would create “traffic tromboning”

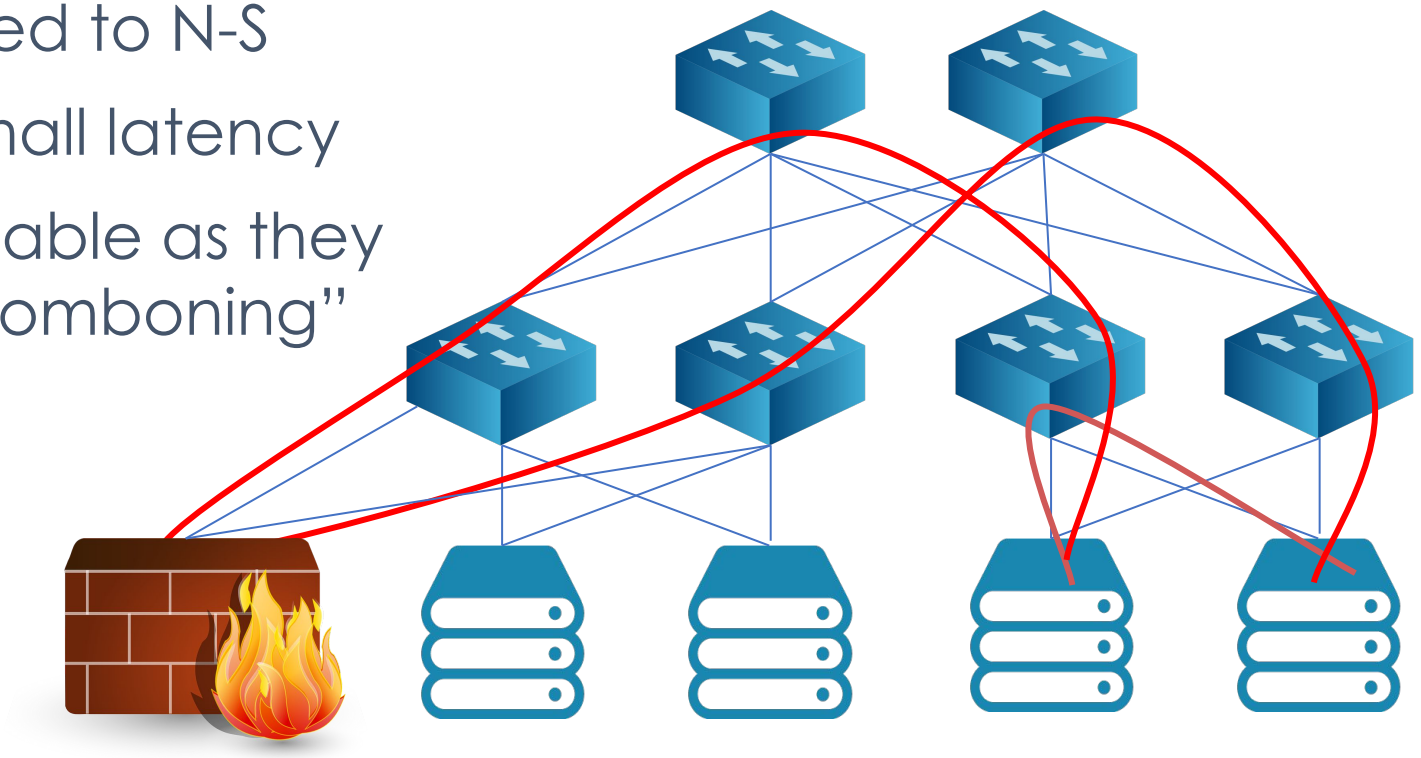

PENSANDO

The DSC is the perfect spot where to implement this

It is on the path of each packet

Flow caching to reduce latency

- Evaluate rules on first packet
- Install entry in flow cache table for handling following packets

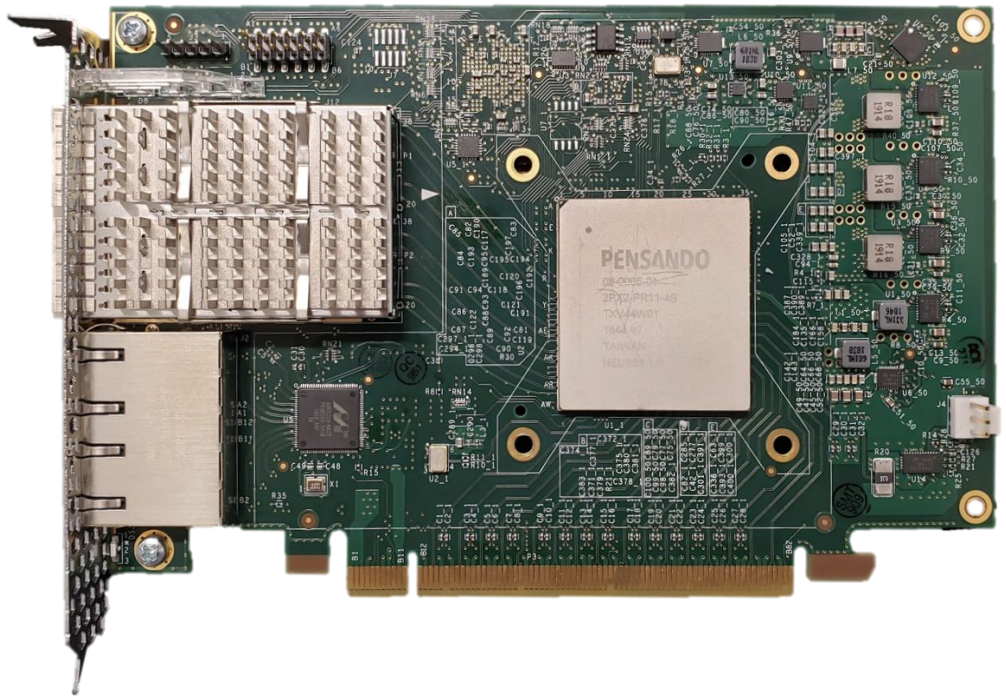

PENSANDO

Distributed Stateful Firewall Support

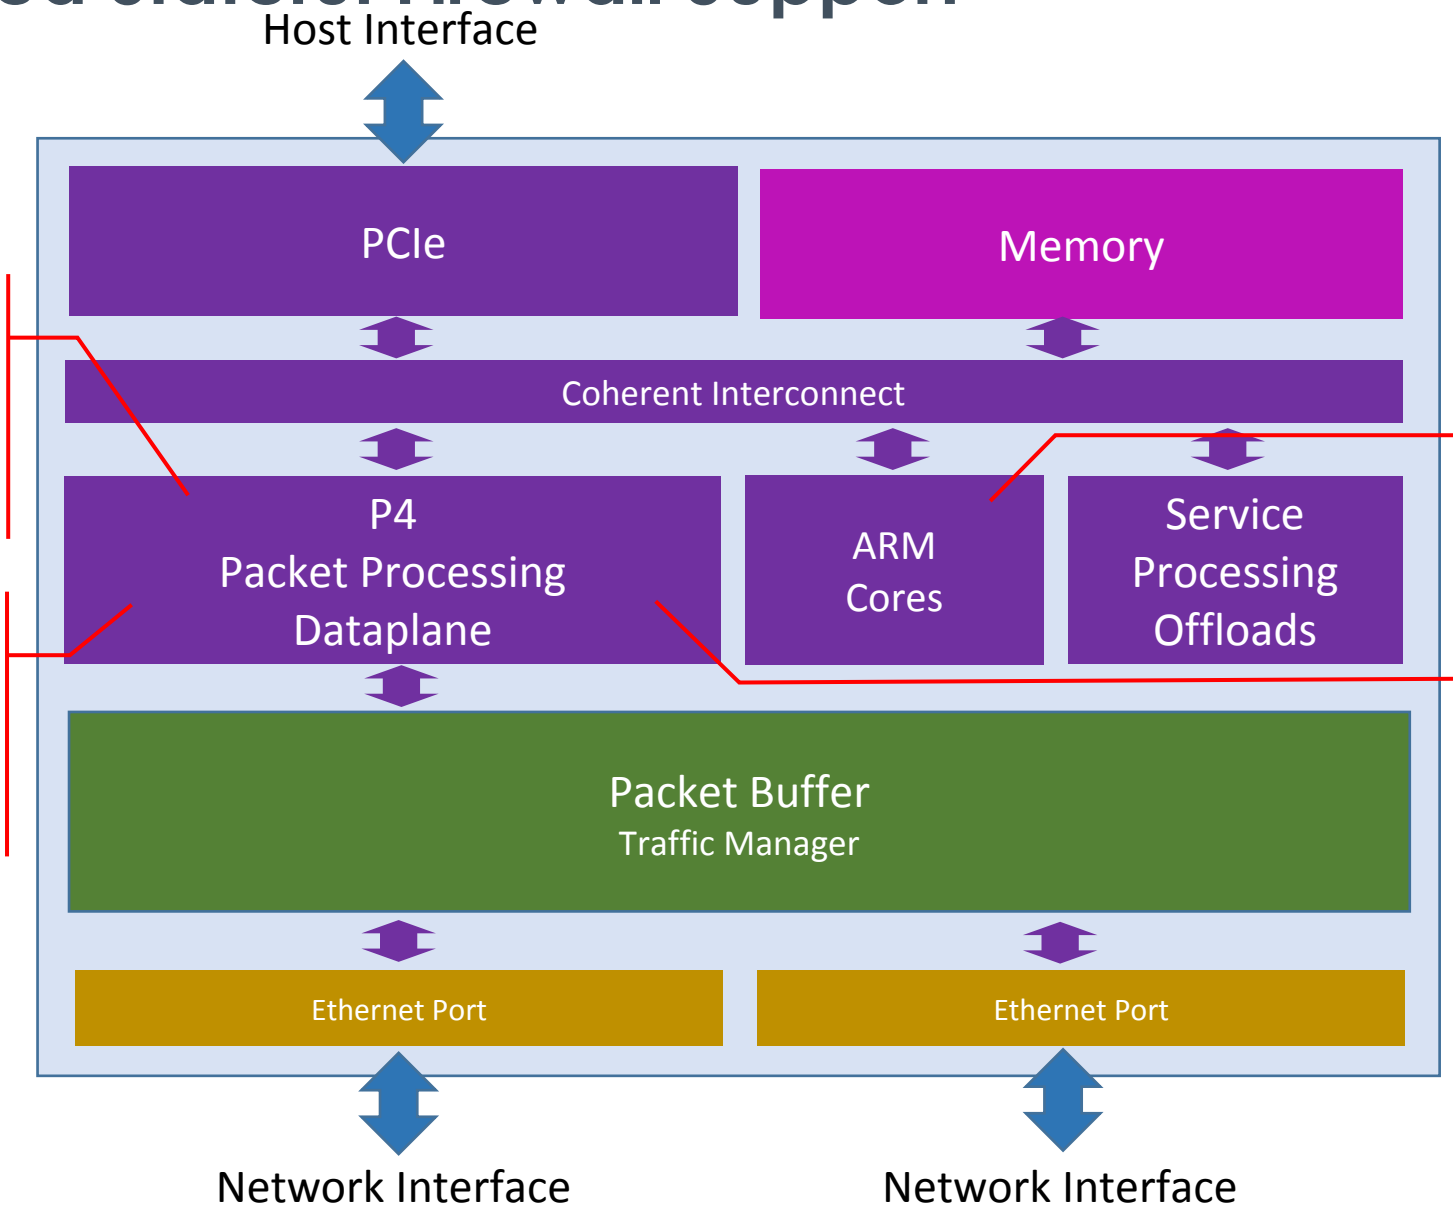

(1) packets belonging to a known flow are forwarded directly (flow cache table)

(2) packets of new flows are further processed in the pipeline to evaluate rules

Software creates forward and reverse flow entries in the flow cache table

(3) packet and corresponding action are passed to ARM cores

(4) packet is passed to pipeline for processing based on newly installed flow cache entry

Memory shared by ARM cores and pipeline ensures that entry is up-to-date

PENSANDO

New Flow Installation Options

Most critical task

ARM Heavy	ARM Light	Pipeline Heavy
• Packet and action passed to ARM • ARM software parses packet • ARM software creates forward and reverse flow entries	• Pipeline extracts relevant metadata and passes them to ARM • ARM software creates forward and reverse entries	• Pipeline extracts relevant metadata and passes them to ARM • Pipeline creates forward and reverse entries
1 M new flows per sec	3 M new flows per sec	**Work in progress**

Unique features of pipeline processing units + **Tight integration (hardware architecture)**

PENSANDO

Performance

PENSANDO

Host Adaptor Mode

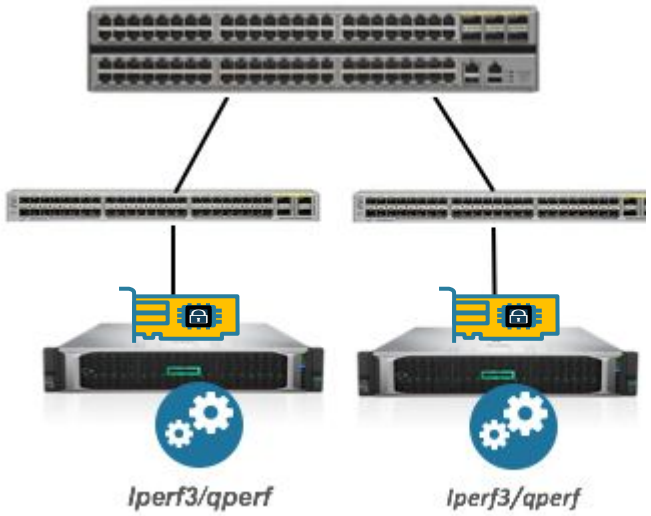

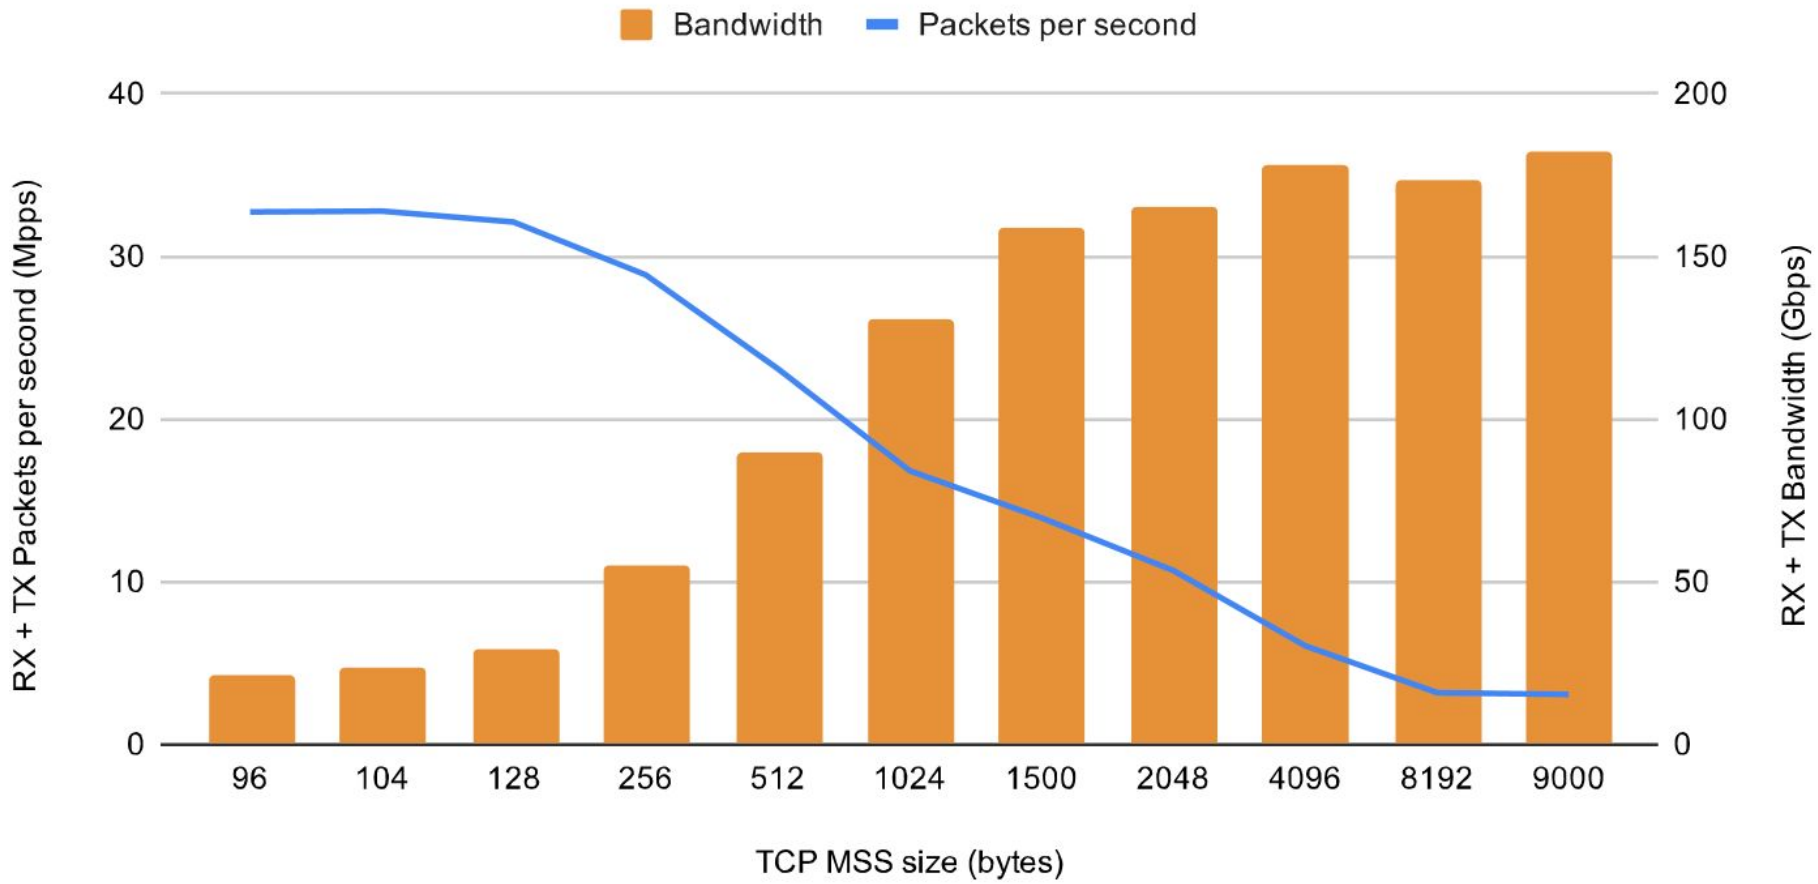

PENSANDO

Bump-in-the-Wire Mode

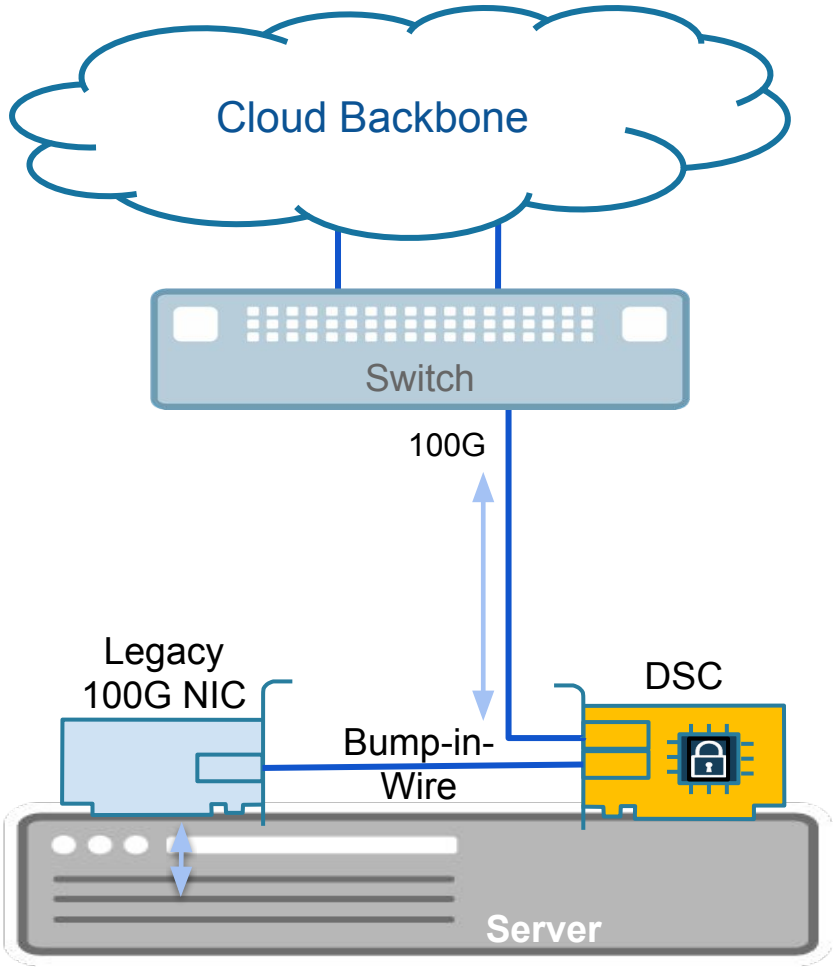

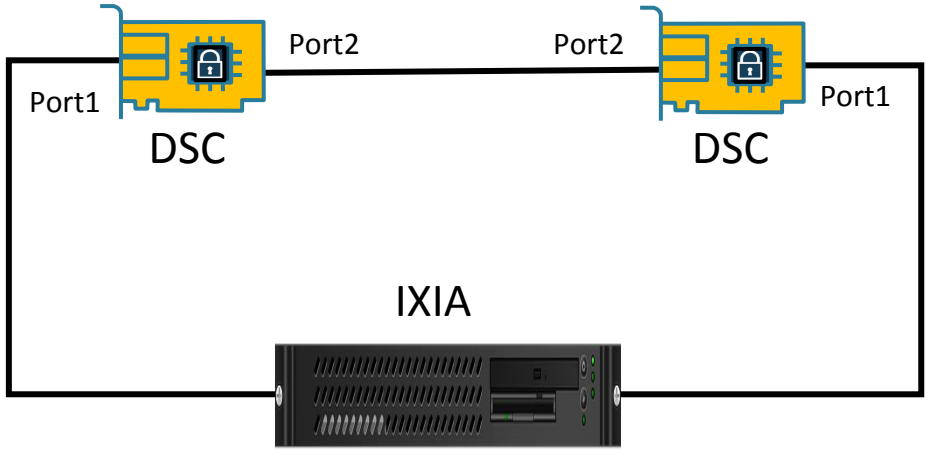

Throughput: 40 Mpps

Latency: 3 µs

Jitter: 35 ns

PENSANDO

In summary … to conclude

- Distributed Services Card Architecture
- How it can be leveraged to implement diverse services and applications
  - Possibly offloading the host
  - Moving them from somewhere else
- Achieve very high performance
  - High throughput
  - Low latency
  - Low jitter
- No performance hit on the host

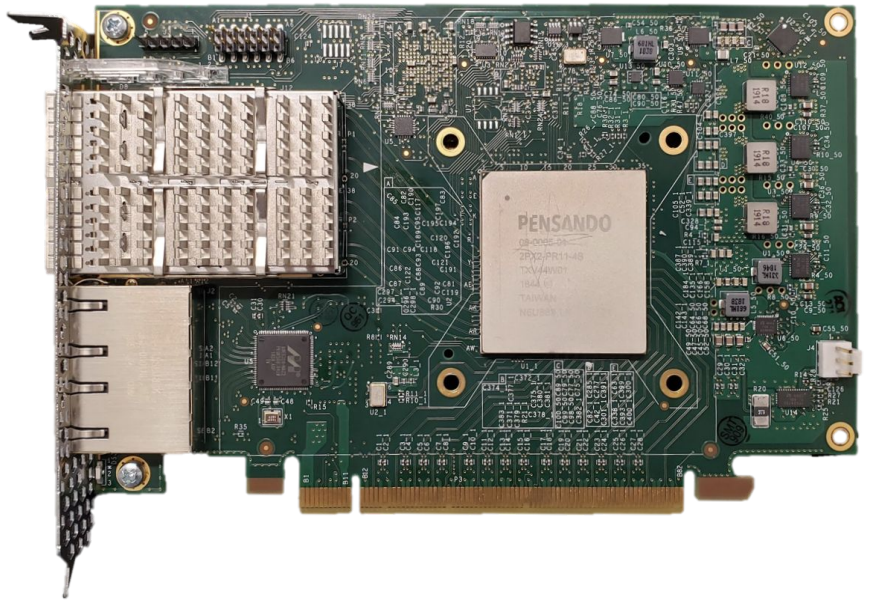

PENSANDO

Sponsored By

Google intel STORDIS The Open Networking Expert

PENSANDO

Thank You

baldi@pensando.io
www.pensando.io
blog.baldi.info
linkedin.baldi.info