Hardware Acceleration for High-density Datacenter Monitoring

Datacenter IaaS Workshop 2014

Denis Matoušek
matousek@invea.com

Company Introduction

- Czech university spin-off company
- Tight cooperation with CESNET, a Czech NREN operator
- Established in 2007
- 40+ employees
- Key focus
  - **Hardware acceleration and FPGA Solutions**
  - **Flow Monitoring and Network Behavior Analysis**
  - **Lawful Interception and Data Retention**

CESNET
Czech Education and Scientific NETwork

e-infrastructure for science, research and education

Technology Agency of the Czech Republic

Motivation I

- High-speed links in data centers
  - **40G Ethernet**
  - **100G Ethernet**:
    - Bleeding-edge technology
    - Very high load: new packet every 6.72ns
    - Transition from 10GE challenge (10× 10G high-density mode)
    - Many variants (IEEE standards, CFP MSA specifications, SFF specifications)
- Even newer technologies are emerging
  - **25G/50G Ethernet:** 25G Ethernet Consortium

→ **Dedicated and flexible hardware** needed!

Motivation II

- Network attacks
  - Which types of attacks are effective to be mitigated on high-speed links?

  → **(D)DoS attacks**: filtering out the traffic based on previous traffic analysis

  → **Black listing**: dropping traffic from known malicious sources
  - It is necessary to process the traffic at **wire-speed**!

Platforms I

- **Commodity hardware**
  - Cheap and flexible
  - Limited I/O performance

- **Dedicated hardware**
  - High I/O performance
  - Expensive, limited flexibility

- **Commodity hardware + Hardware acceleration**
  - **Multi-core CPUs + FPGA network interface card**
  - **High I/O performance**
  - **Reasonable price**
  - **Flexible**

Platforms II

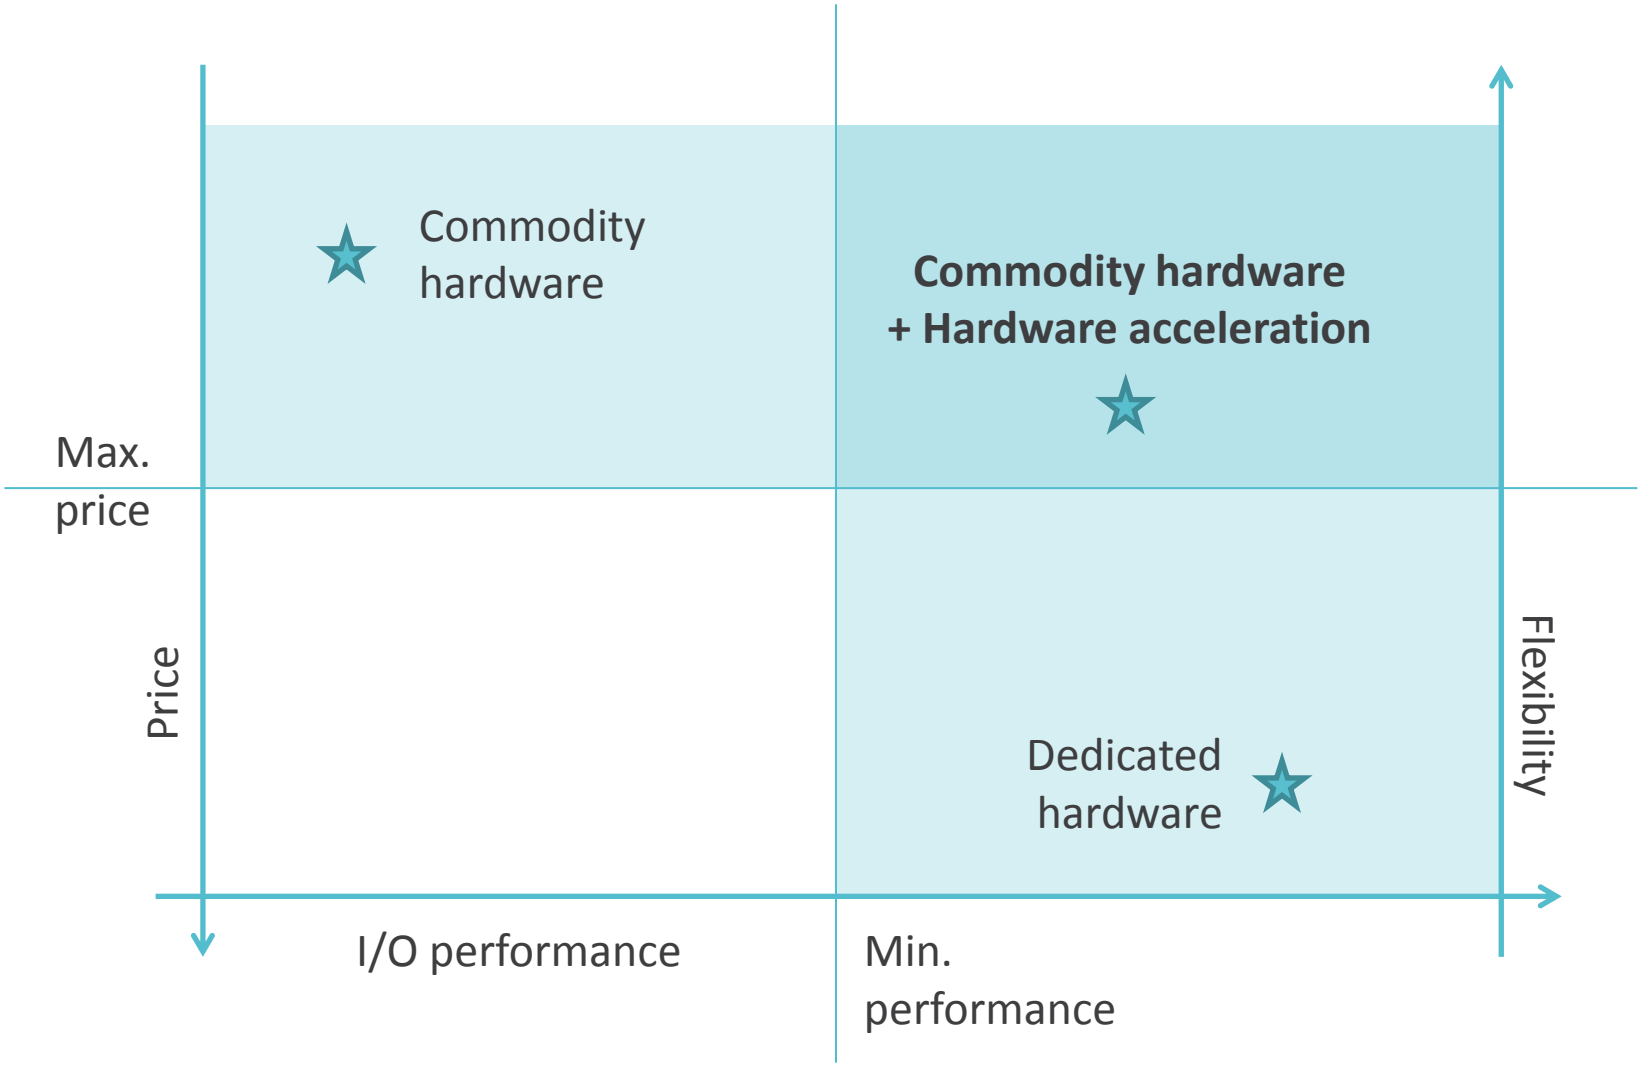

INVEA-TECH Solutions

- Hardware-acceleration in **FPGA** chip
  - **Flexibility** and **high performance**
  - One hardware for multiple applications (traffic recording, IDS/IPS, SDM…)
  - HW/SW codesign

HANIC-80G

HANIC-100G

Network Technologies

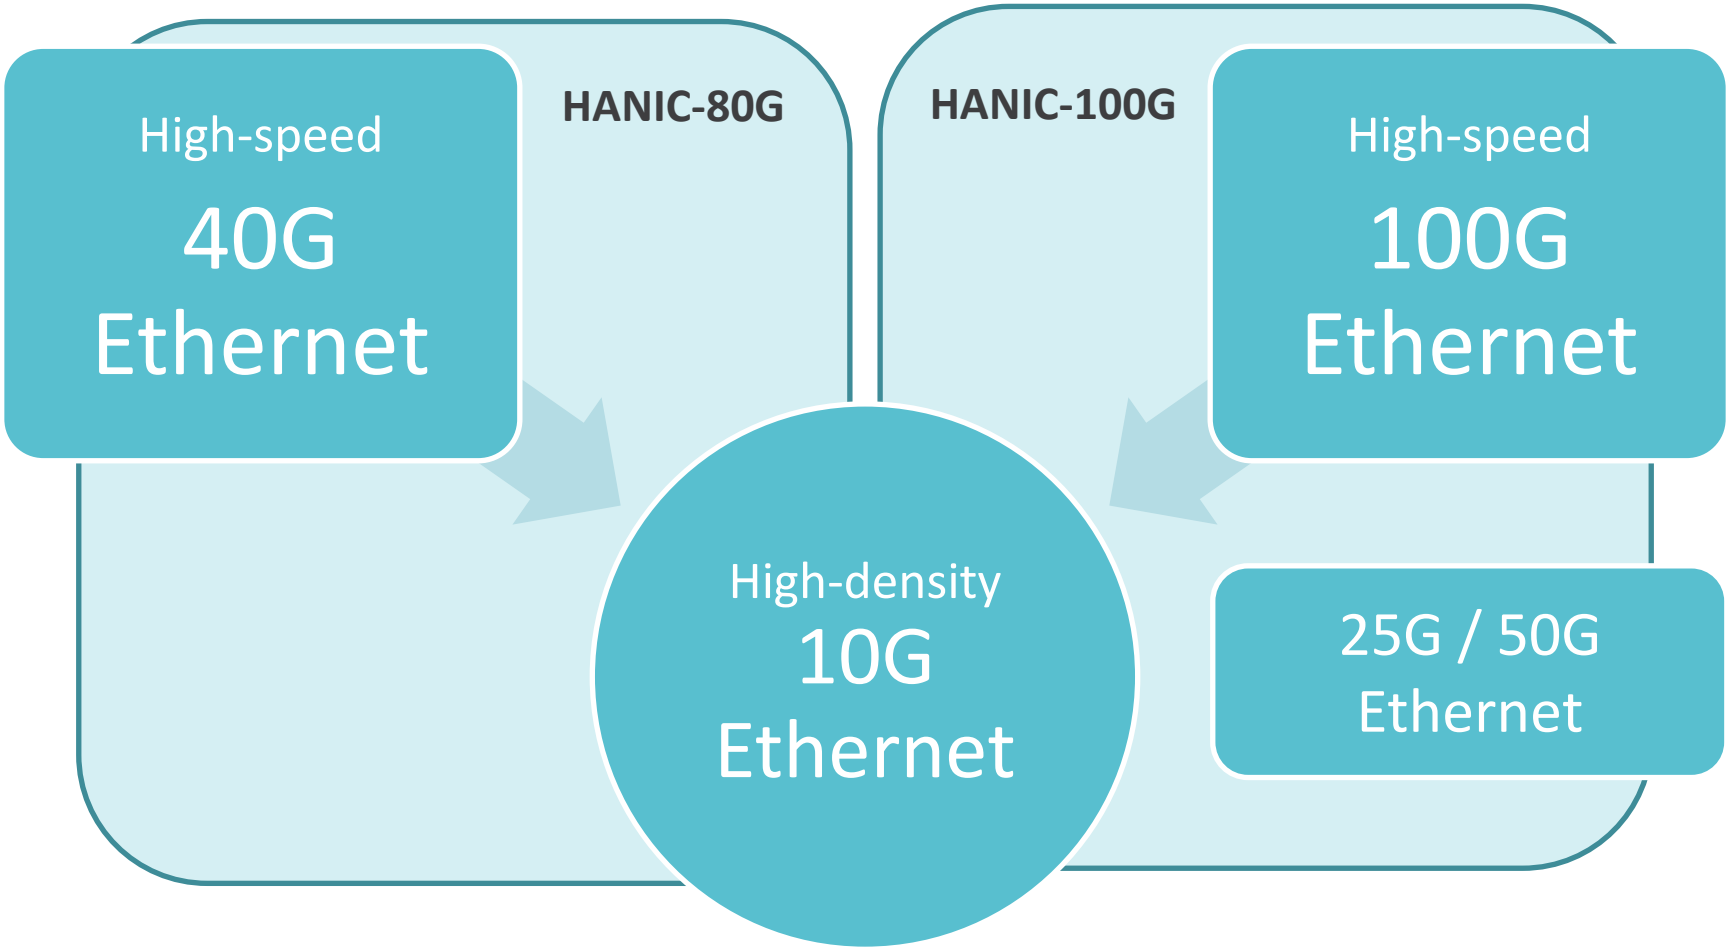

→ **No change in hardware is needed!**

Deployment of 100GE

- Evolution of optical modules:
  - CFP → **CFP2** → CFP4 / QSFP28
  - Smaller & cheaper
  - Allowing even multiple 100G interfaces on one adapter

- **HANIC-100G** is the world's first 100G FPGA adapter available
  - 100GBASE-xR4, -SR10
  - PCI Express Gen 3 x16

High Density

- The nature of 40G and 100G Ethernet technologies allows to use one link as multiple independent 10G links
- SM (single mode) vs. PSM (parallel SM)
- FPGA provides flexibility that allows to support many Ethernet standards
  - 40G + 10G + even 1G!

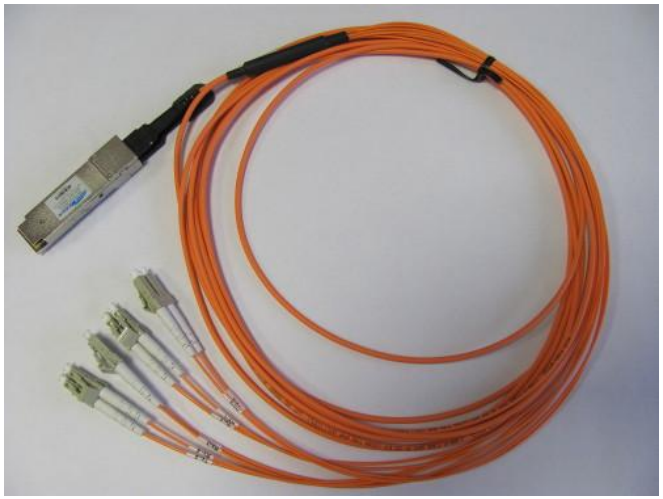

Optical break out cable
40G → 4× 10G

Use Case: 8× 10G with One Card

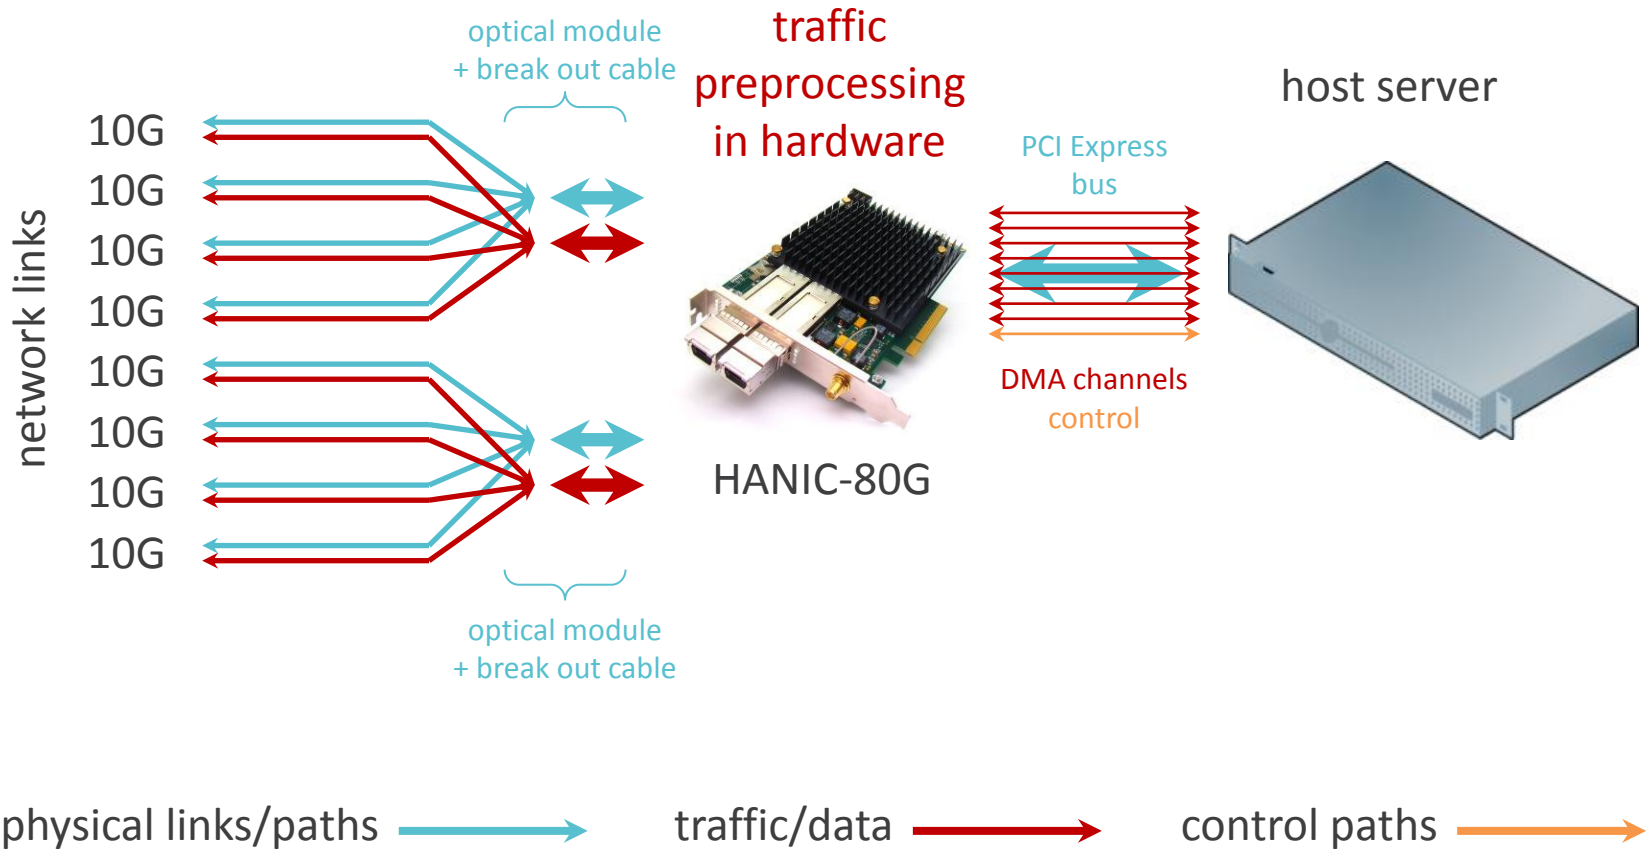

Use Case: 2× 40G with One Card

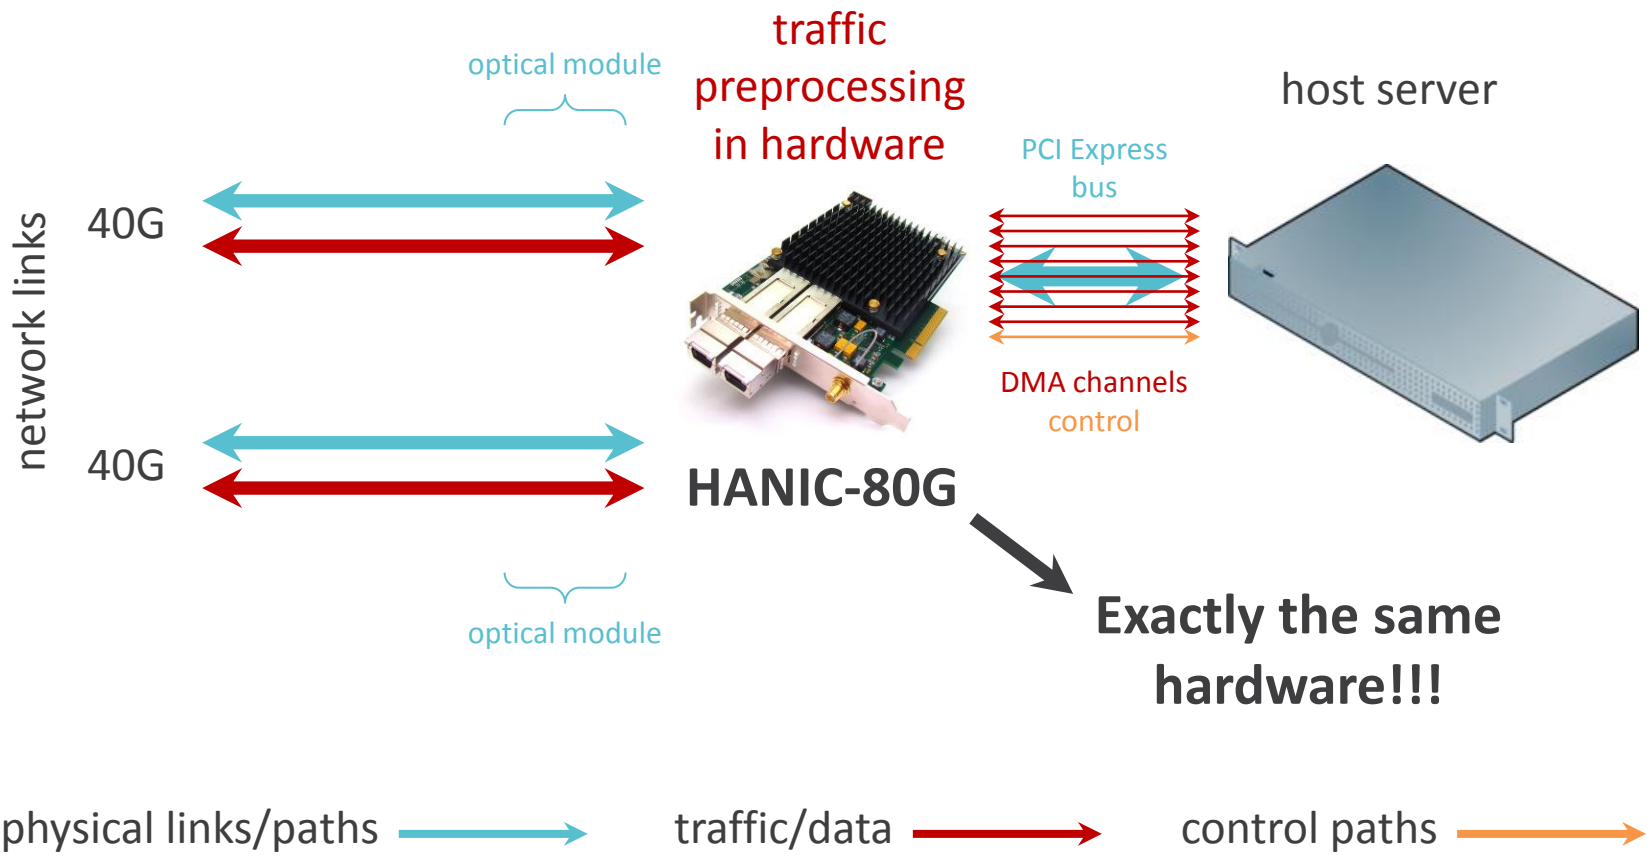

Flexibility

INVEATECH

- FPGA chip is able to be programmed even after deployment:
  - HANIC-80G example: 2× **40G** vs. 8× **10G Ethernet** only by loading different firmware!
  - HANIC-100G example: it provides architecture to support new technologies (CFP4/QSFP28 optical modules, **25G/50G Ethernet**)
- Firmwares can be simply upgraded with new features and bug fixes

Network Traffic Processing

High-speed network interfaces
→ traffic preprocessing in FPGA chip
→ high-speed transfer to host computer
→ packet processing on multicore CPUs

All advantages of processing on commodity servers are preserved!

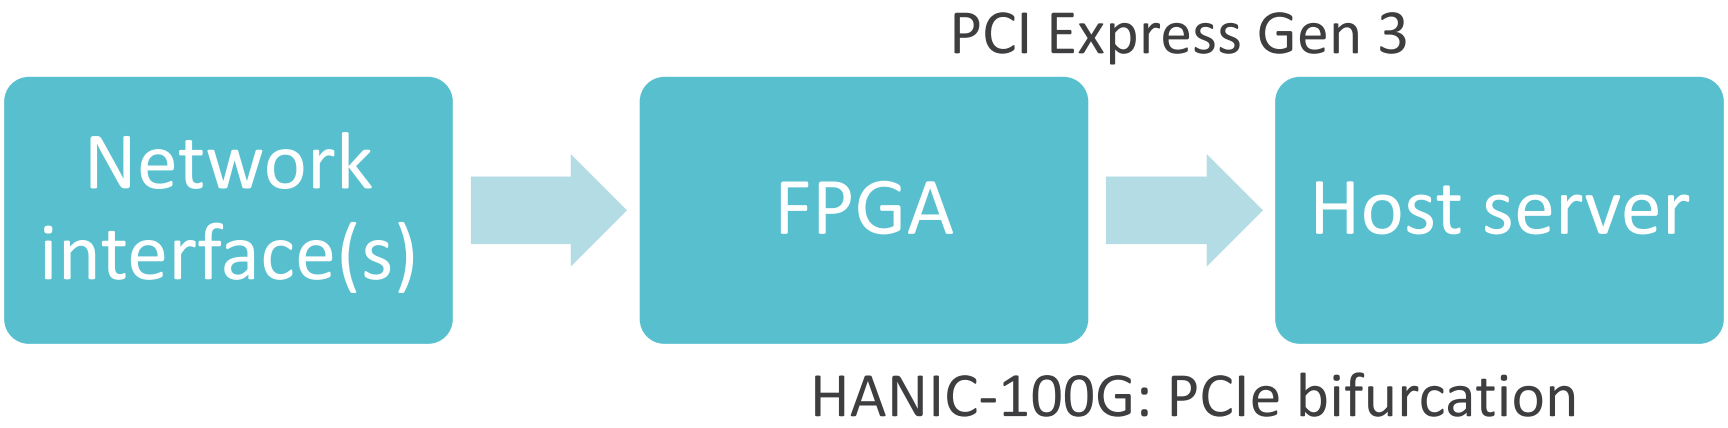

Usability

- High-performance requires special processing on both HW and SW side
- Device drivers provide **transparent user interface**
  - Standard application interface (PCAP)
  - Optimized for throughput
  - Special techniques for 100GE (PCIe bifurcation)

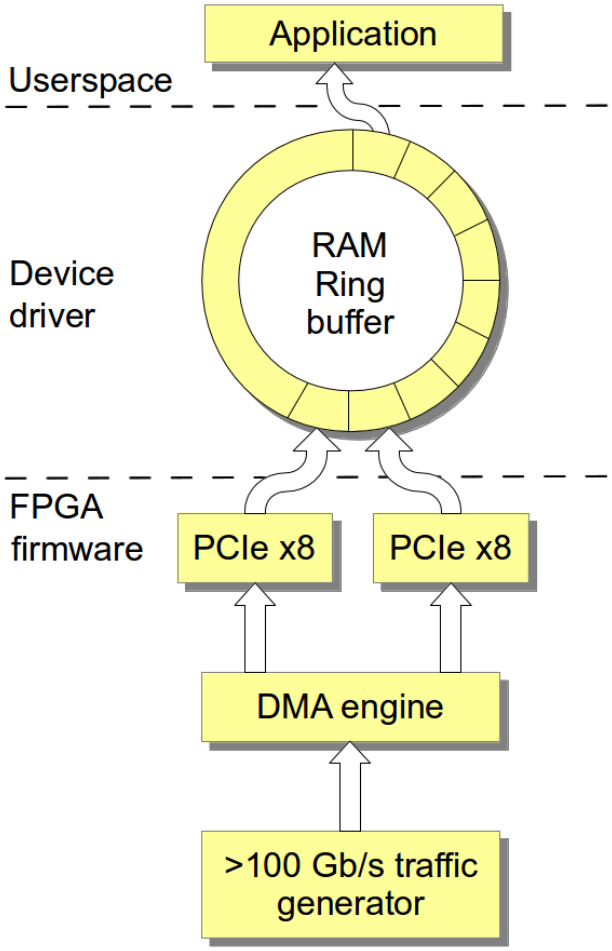

Standard tools

- Easy integration with standard tools:
  - Tcpdump  … traffic recording
  - Wireshark  … traffic analysis
  - Snort  … IDS/IPS
  - FlowMon  … traffic analysis

Summary

- Benefits of INVEA-TECH FPGA solutions:
  - **Flexibility** and **scalability** of **FPGA technology**
  - **High performance**: wire-speed traffic processing
  - **HW/SW codesign**: performance vs. time-to-market
  - Broad support of network technologies:
    - 1G, 10G, **40G & 100G Ethernet**
    - Future: 25G, 50G Ethernet
  - **High density**
    - multiple 10G links on one adapter
    - multiple adapters in one rack-mountable server

We enable you to use your standard tools with the only difference – they run faster!

Reference


upc
AVG
T-Mobile
Belnet
dedicated connectivity
SLOANE PARK
Government of the Czech Republic
Ostrava Airport
T-Com
UNIQA
Allianz
KONICA MINOLTA
CZECH STATISTICAL OFFICE
THE ACADEMY OF SCIENCES OF THE CZECH REPUBLIC
NK
IBP
Institute of Biophysics
VEOLIA
AEGON
e·on IT
SIEMENS
TELECOM Bretagne
eustream
pitcom
dm
MVV Energie
Matador
PRAGUE STOCK EXCHANGE
BURZA CENNÝCH PAPÍRŮ PRAHA
ihp
zenit elektro
OUTSOURCING SOLUTION
ATComputers
kt
AVE
CSIRT-MU
Marius Pedersen
THOMAYEROVA NEMOCNICE PRAHA
FAKULTNÍ NEMOCNICE OLOMOUC
UNIVERSITAS MASARYKIANA BRUNENSIS
ČD-Telematika
CASABLANCA INT
GEANT2
CARNet
Tomas Bata University
grnet
SURFNET
SWITCH
TÜBİTAK ULAKBİM
UNIVERSITY OF TWENTE.
Olomoucký kraj
Kraj Vysočina
OLOMOUC

INVEATECH

High-Speed Networks Technology Partner

Denis Matoušek
matousek@invea.com
+420 511 205 264

INVEA-TECH a.s.
U Vodárny 2965/2
616 00 Brno, Czech Republic
www.invea.com