




# **Vitis Network Example (VNx)**

## **A Lightweight UDP Network Layer for FPGA**

Lucian Petrica

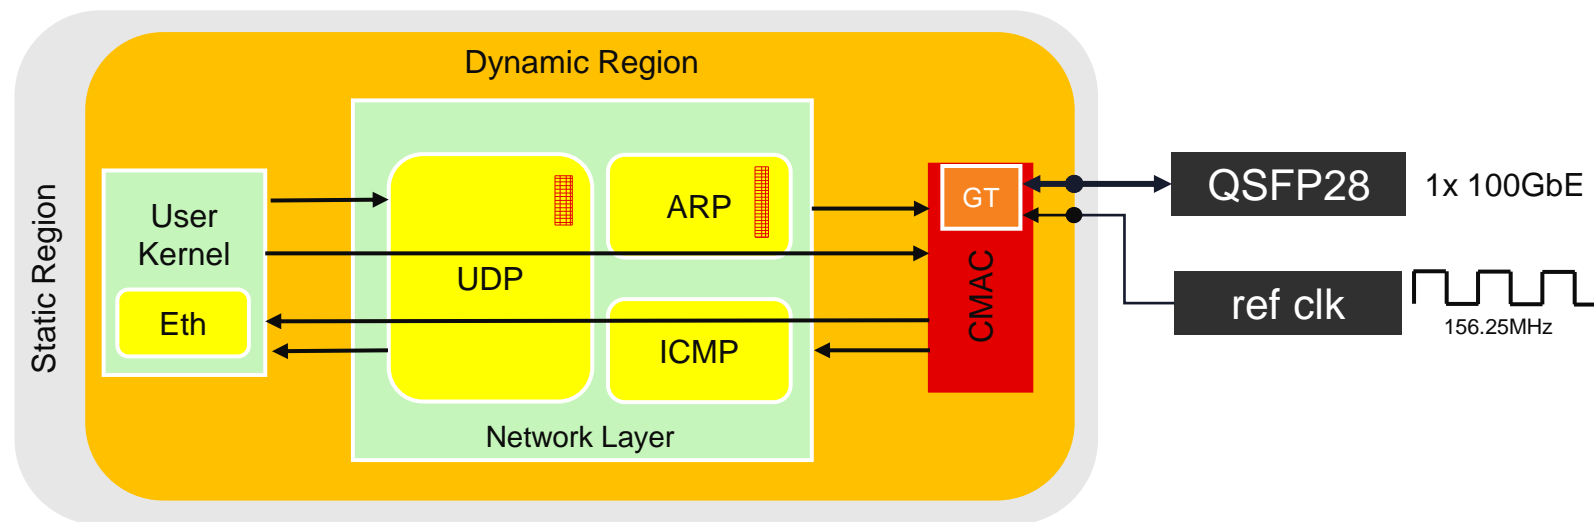
# VNx Goals

- Provide networking on Vitis platforms for Alveo cards
  - Fast stream transport over routed fabrics
  - Low resource overhead
  - Easy to link with existing designs that use streaming kernels
  - Portable across Alveo range
- Easy integration & ease of use
  - Works out-of-the-box at HACC
  - PYNQ-based Python driver
  - C++ driver (community contributed)



Copyright ETH Zürich - 2022

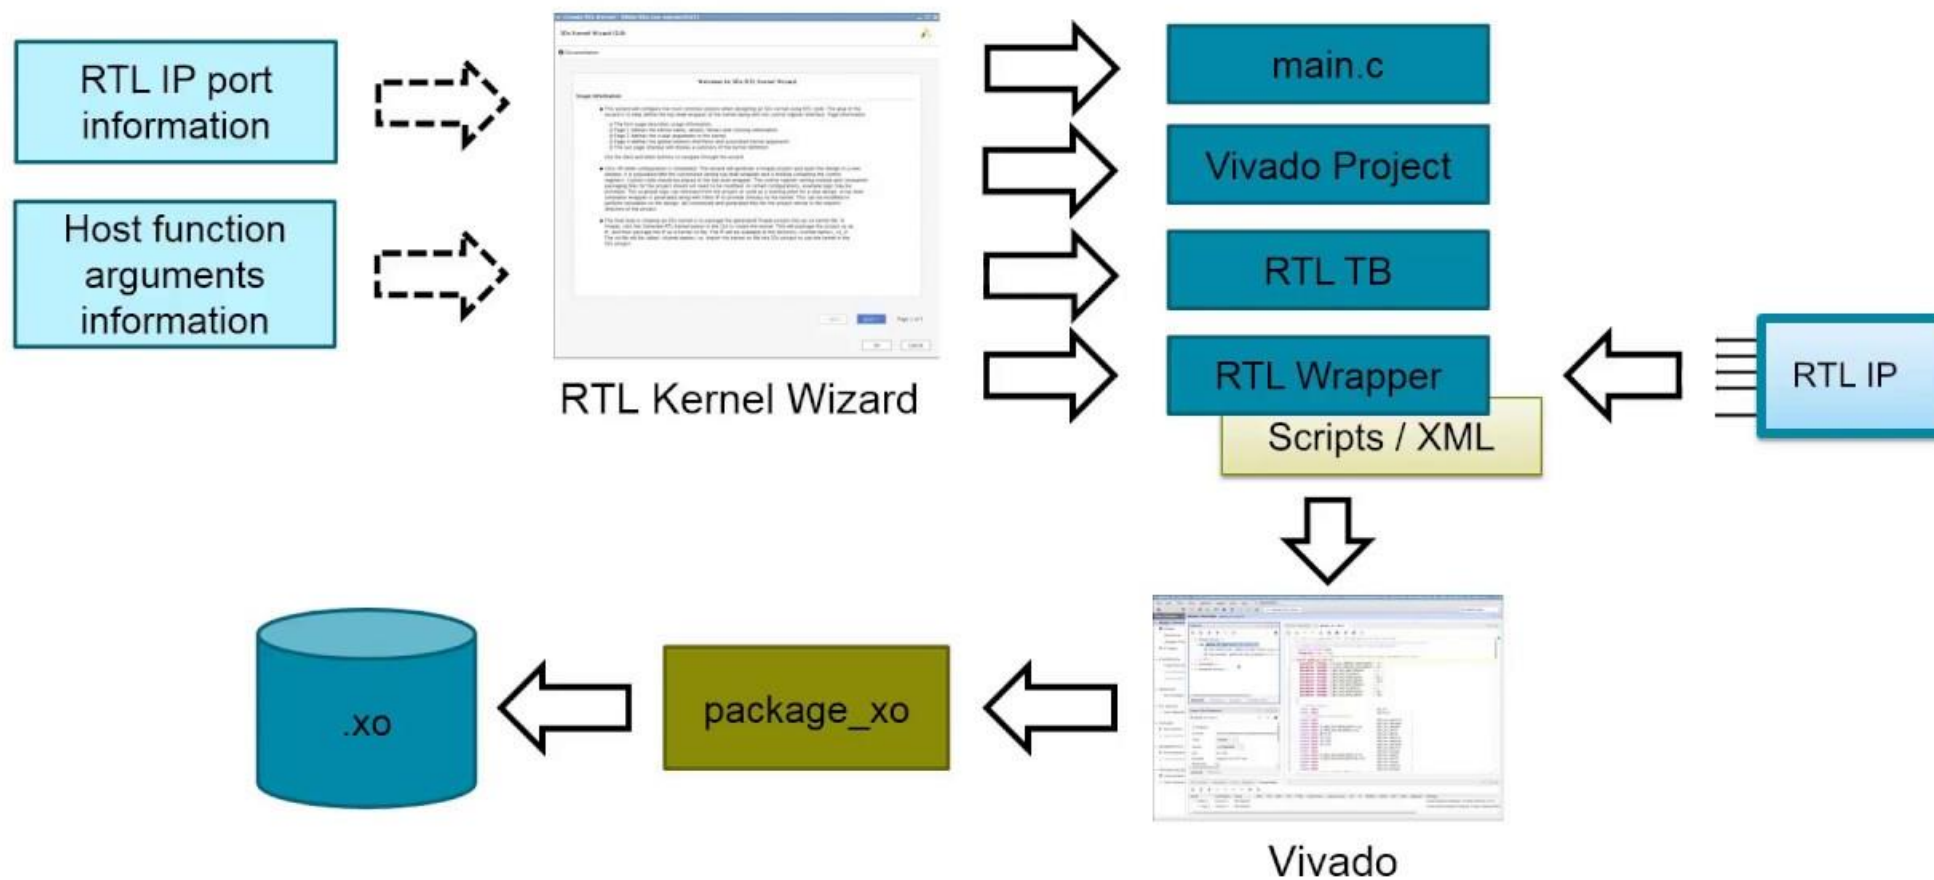
# VNx Architecture



- Support for multiple interfaces and cards
  - Any Vitis Alveo platform with GT access
  - U50, U55C, U200, U250 and U280
- Layered architecture
  - Easily adaptable to user needs

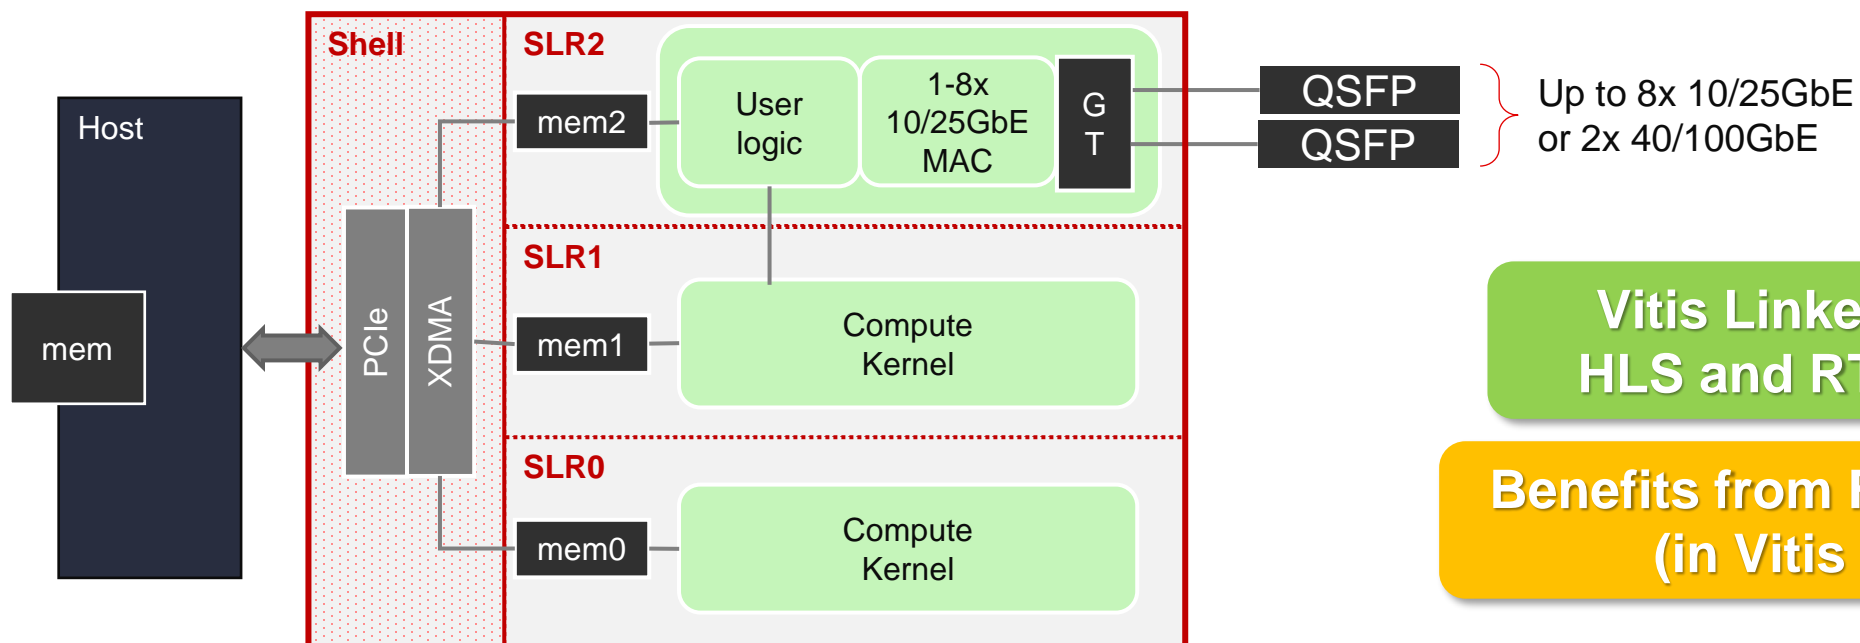
- CMAC kernel
  - Configuration is specific to Alveo card and network interface (if card is multi-port)
- Network Layer
  - UDP as transport protocol
  - Single channel application interface

# Refresher: Vitis RTL Kernels



# Vitis GT Kernels – Networking Flexibility

- RTL kernels which connect to Gigabit Transceiver (GT) pins
- Can implement your choice of MAC
- Documentation [here](#)



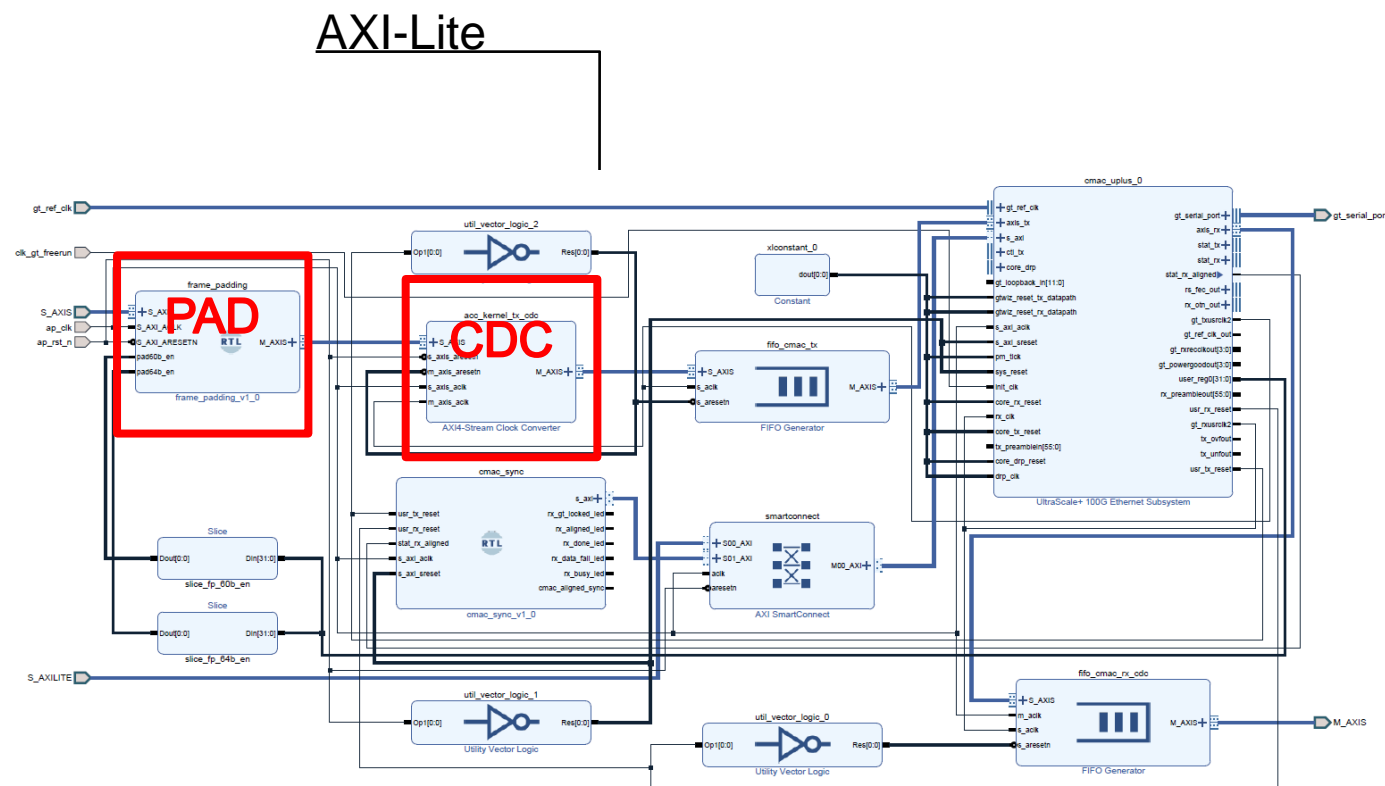
Vitis Linker can mix  
HLS and RTL kernels

Benefits from Floorplanning  
(in Vitis config)

Production support for all Alveo cards

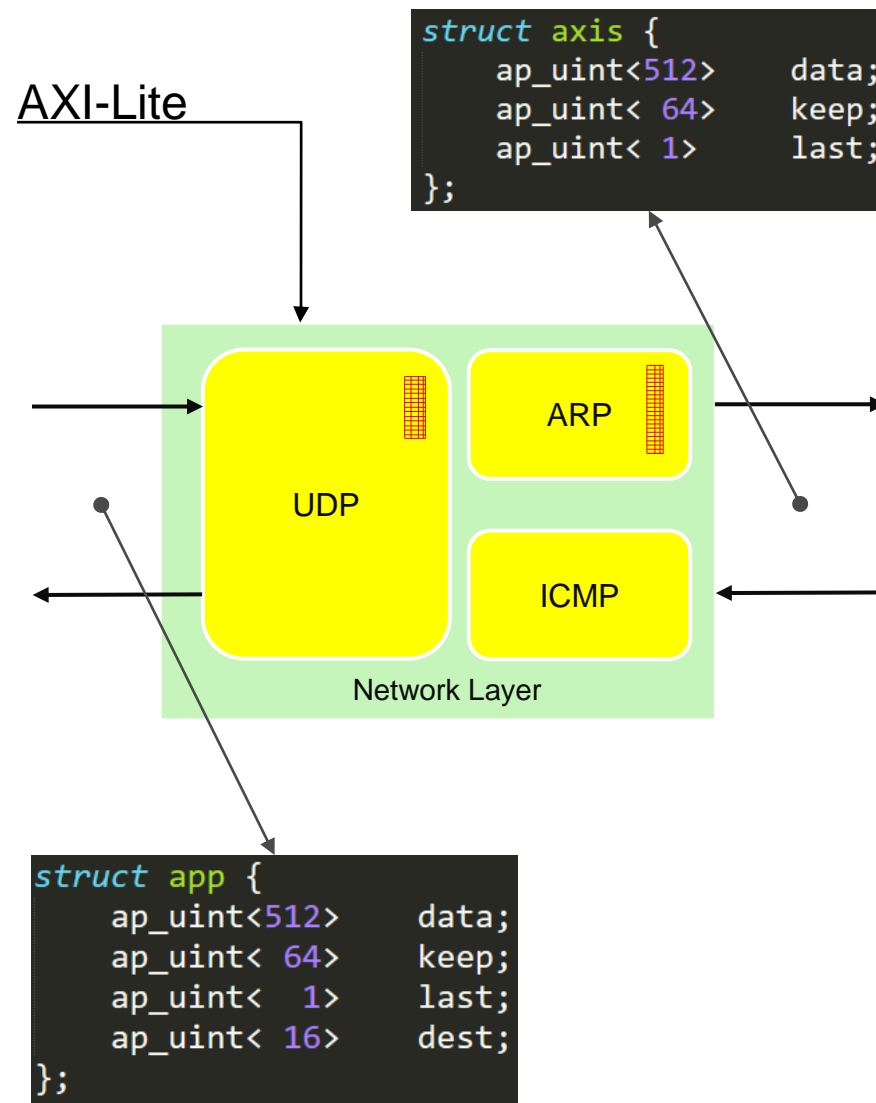
# CMAC Kernel Implementation

- **Vitis-compatible**
  - Standard AXI-Lite and AXI-Stream interfaces + GT pins
  - Single external kernel clock
  - All configuration via AXI-Lite - access to the CMAC IP register map
- **Statistics** gathering and readout via AXI-Lite
- **RS-FEC** capability included in the CMAC IP
  - optionally enabled from host code
- **Portable** across UltraScale+ Alveo cards
  - CDC included to internal clock domains
  - No fixed clocks required, due to internal CDC
- Optional **frame padding** to 60 or 64 bytes

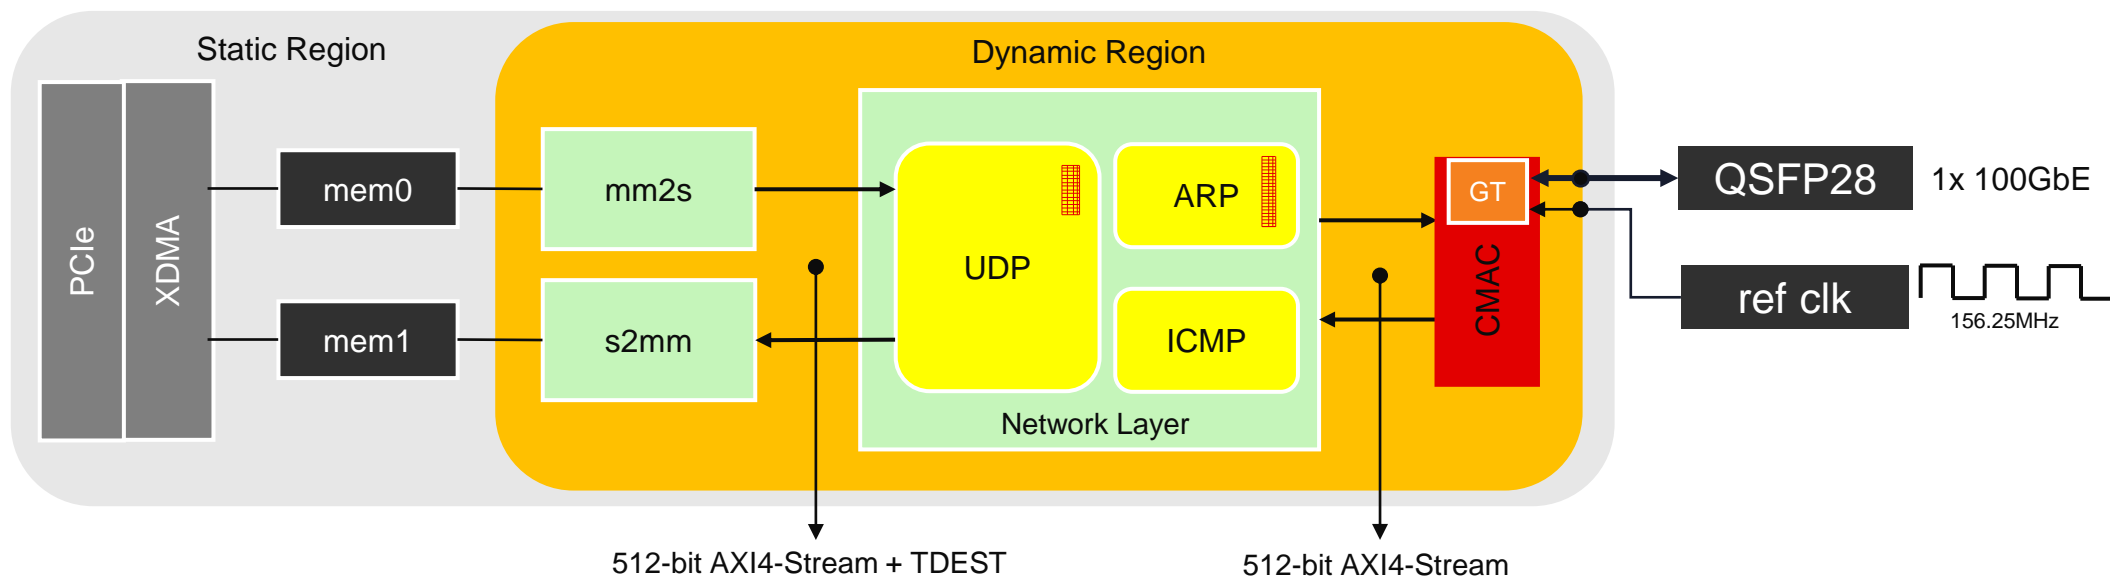


# Network Layer

- Based on Limago
- ARP
  - Translation between MAC and IP addresses
  - ARP table is accessible from host
- ICMP
  - ping functionality
- UDP
  - Transport protocol
  - Host-configured connection/socket table (16 entries)
  - TDEST in application interface specify connection ID
- AXI-Stream Interfaces
  - 64-byte data (with associated TKEEP)
  - 16-bit TDEST carries index in connection table



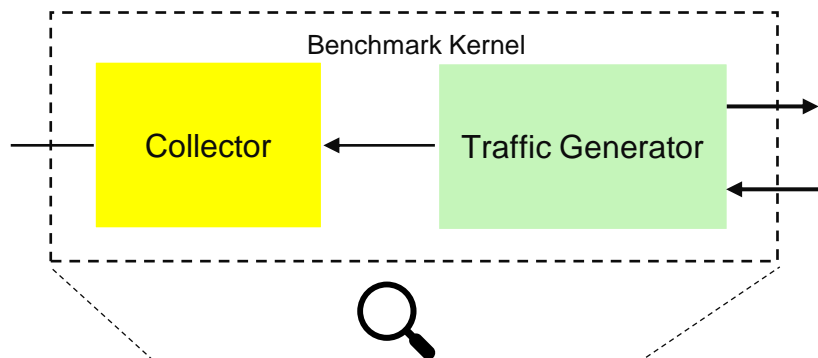
# Basic Example



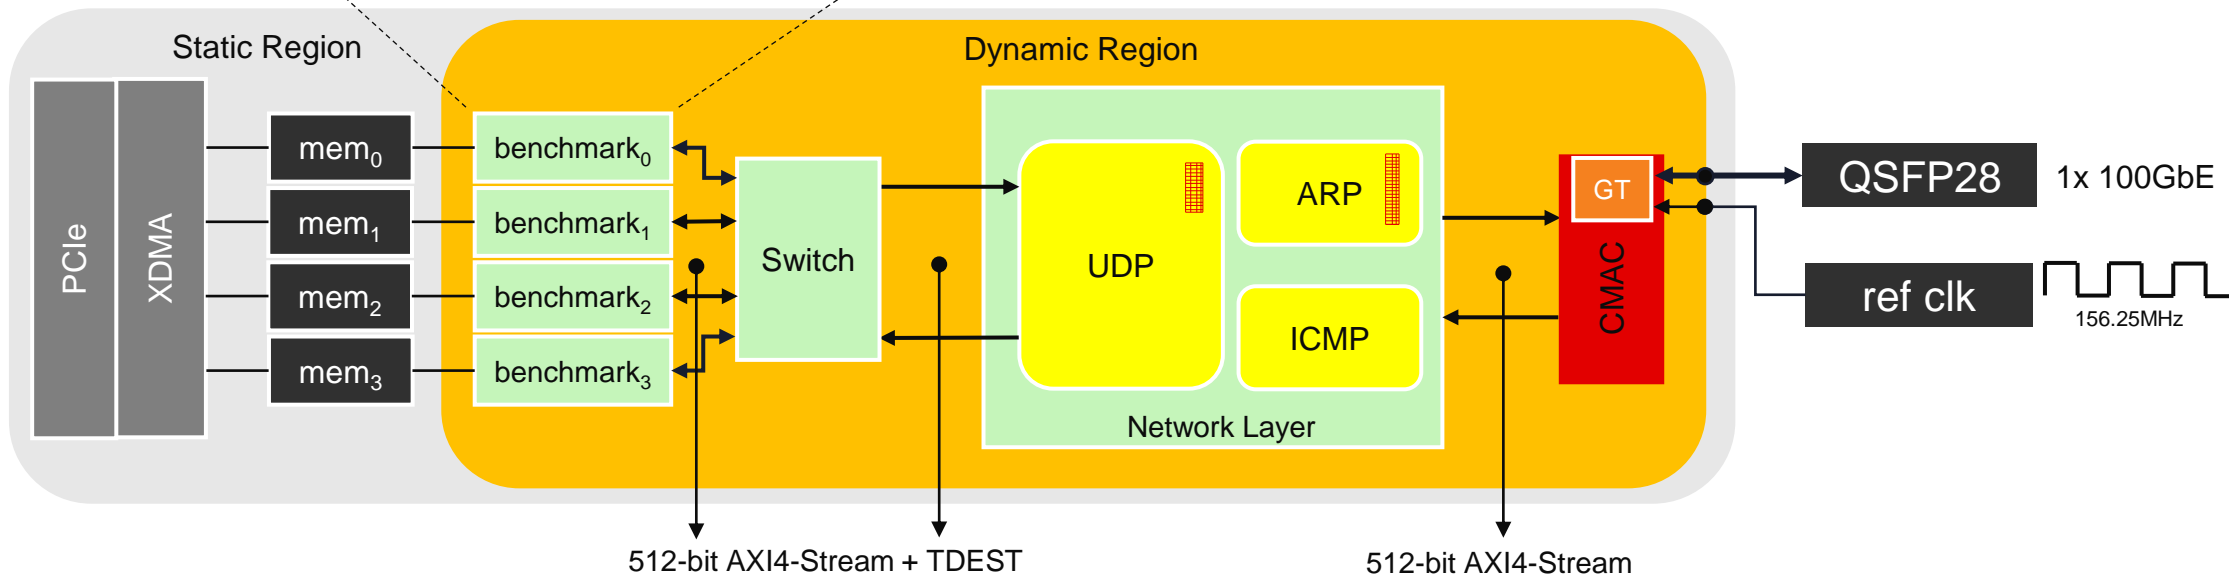
- Two user kernels
- Memory mapped to stream (mm2s)
  - Push data from host to network
- Stream to memory mapped (s2mm)
  - Pull data from network to host
- Works for 1 or 2 interfaces, depending on the board



# Benchmark Example



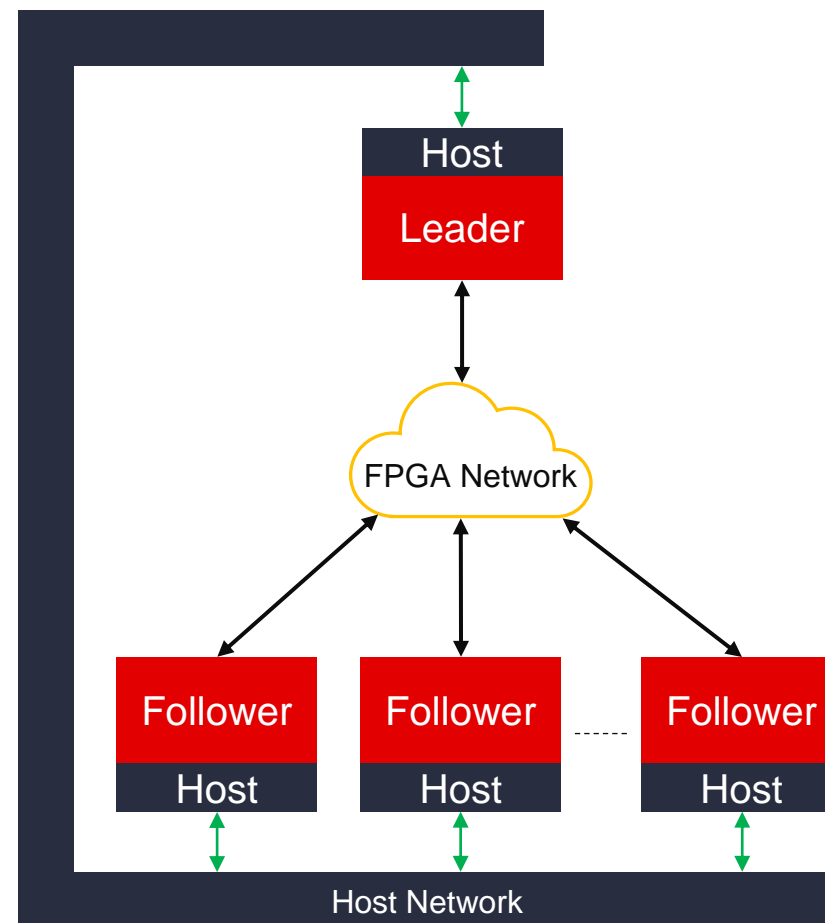
- Benchmark kernel
  - Throughput
  - Latency (RTT)



# Host Code Execution on Heterogeneous Systems

- Steps for application execution
  1. Configure the devices
    - Download FPGA configuration file
    - Configure IP addresses
    - ARP Discovery
    - Populate UDP table
  2. Allocating buffers
  3. Writing the buffers to FPGA memory
  4. Running the accelerators
  5. Reading the buffers from FPGA memory
- How to execute these steps remotely?

} setup



# Dask-on-PYNQ

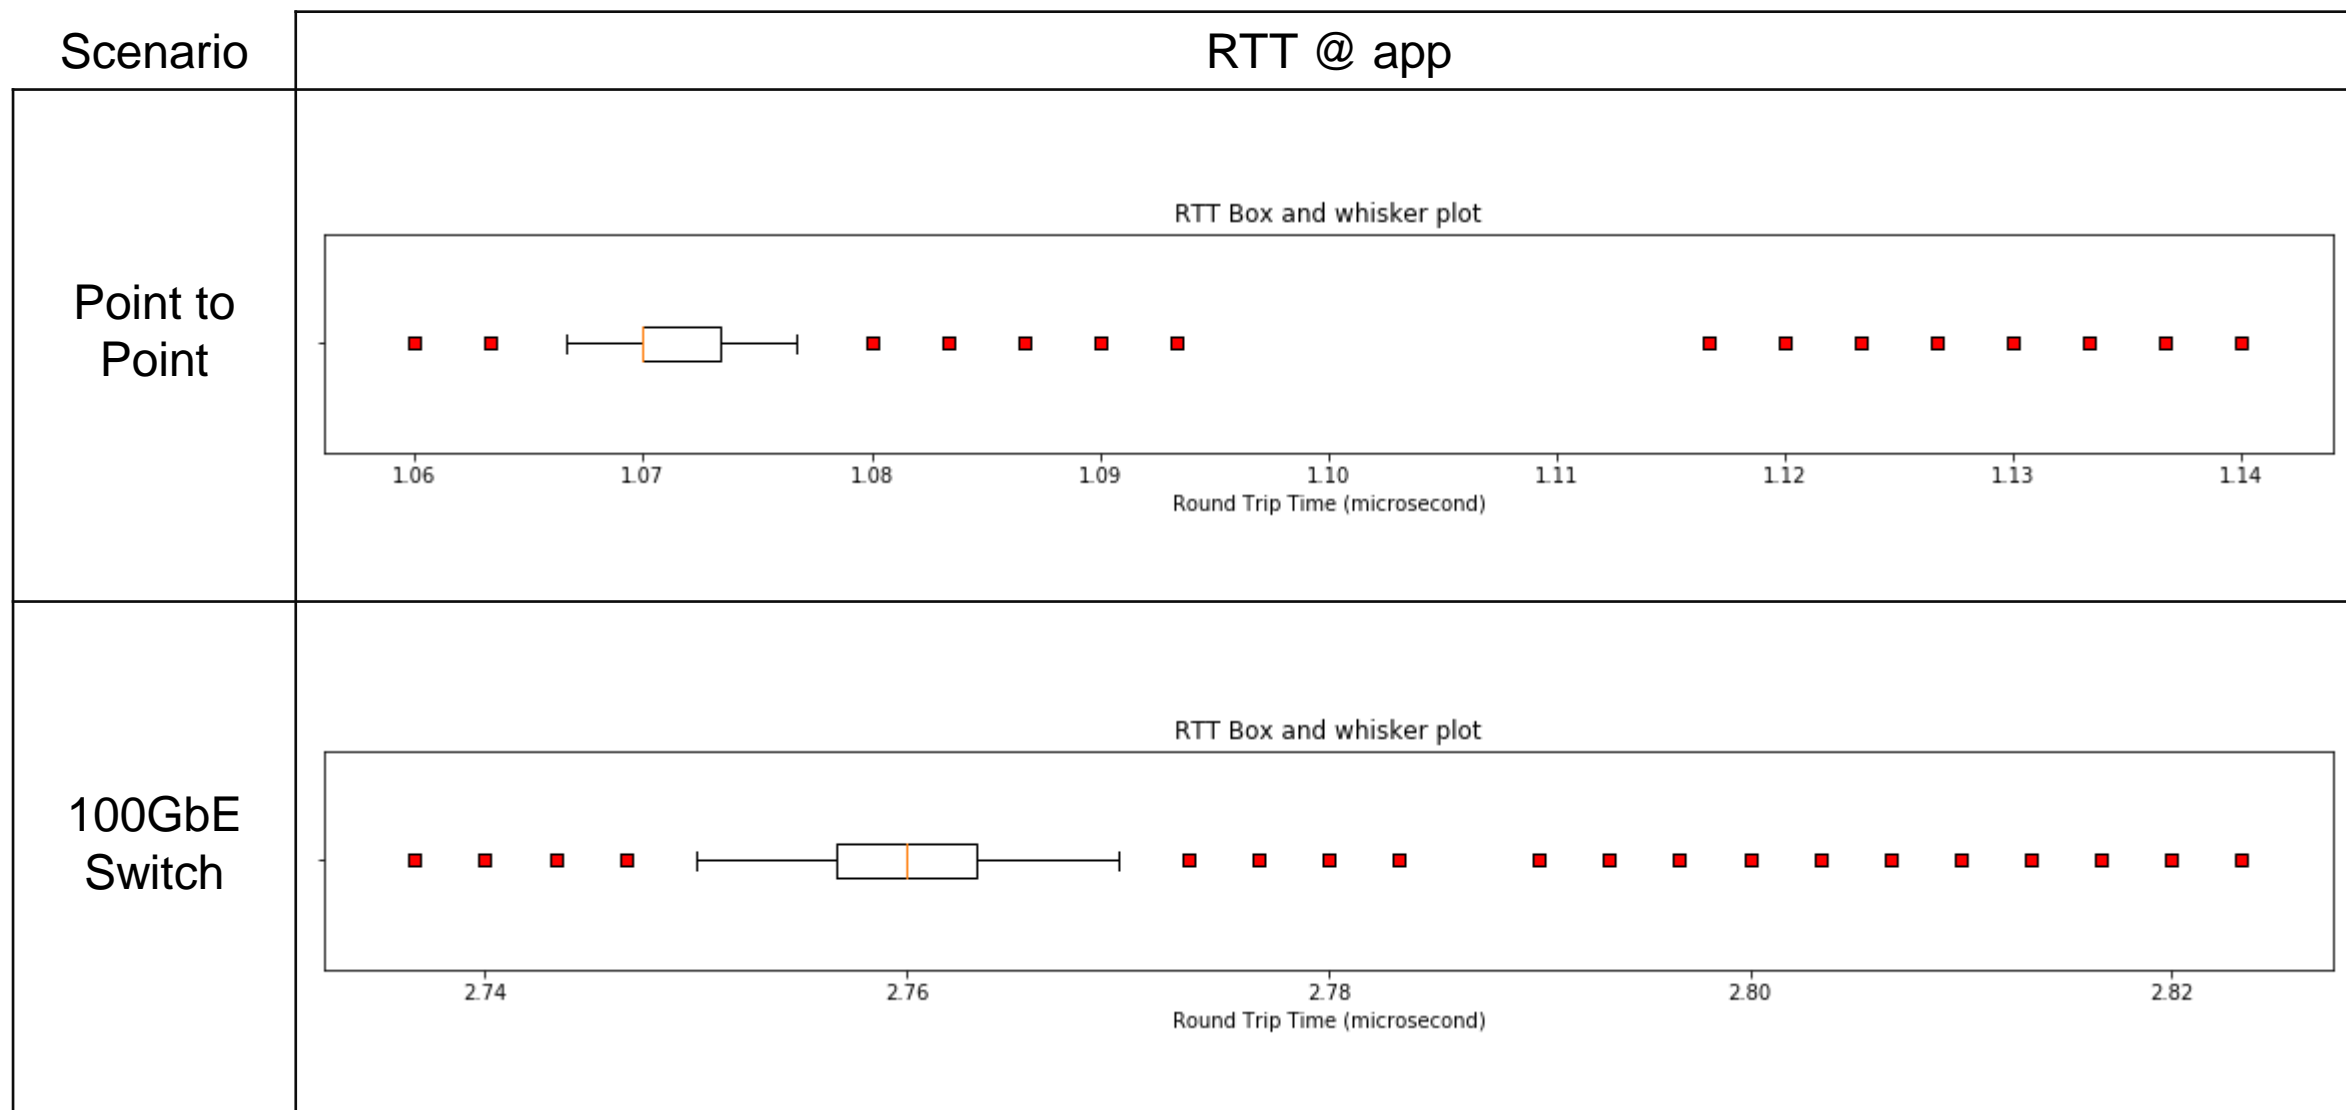
- Dask is a flexible library for parallel computing in Python
- PYNQ-on-Alveo as host code
  - Notebooks on Jupyter Lab
- Dask class wrapper
  - Distributed configuration
  - Remote buffer allocation
  - Remote task execution
  - All the same functionality as PYNQ but distributed!
- Not specific of VNx



```
daskdev_w0 = DaskDevice(client, workers[0])
daskdev_w1 = DaskDevice(client, workers[1])

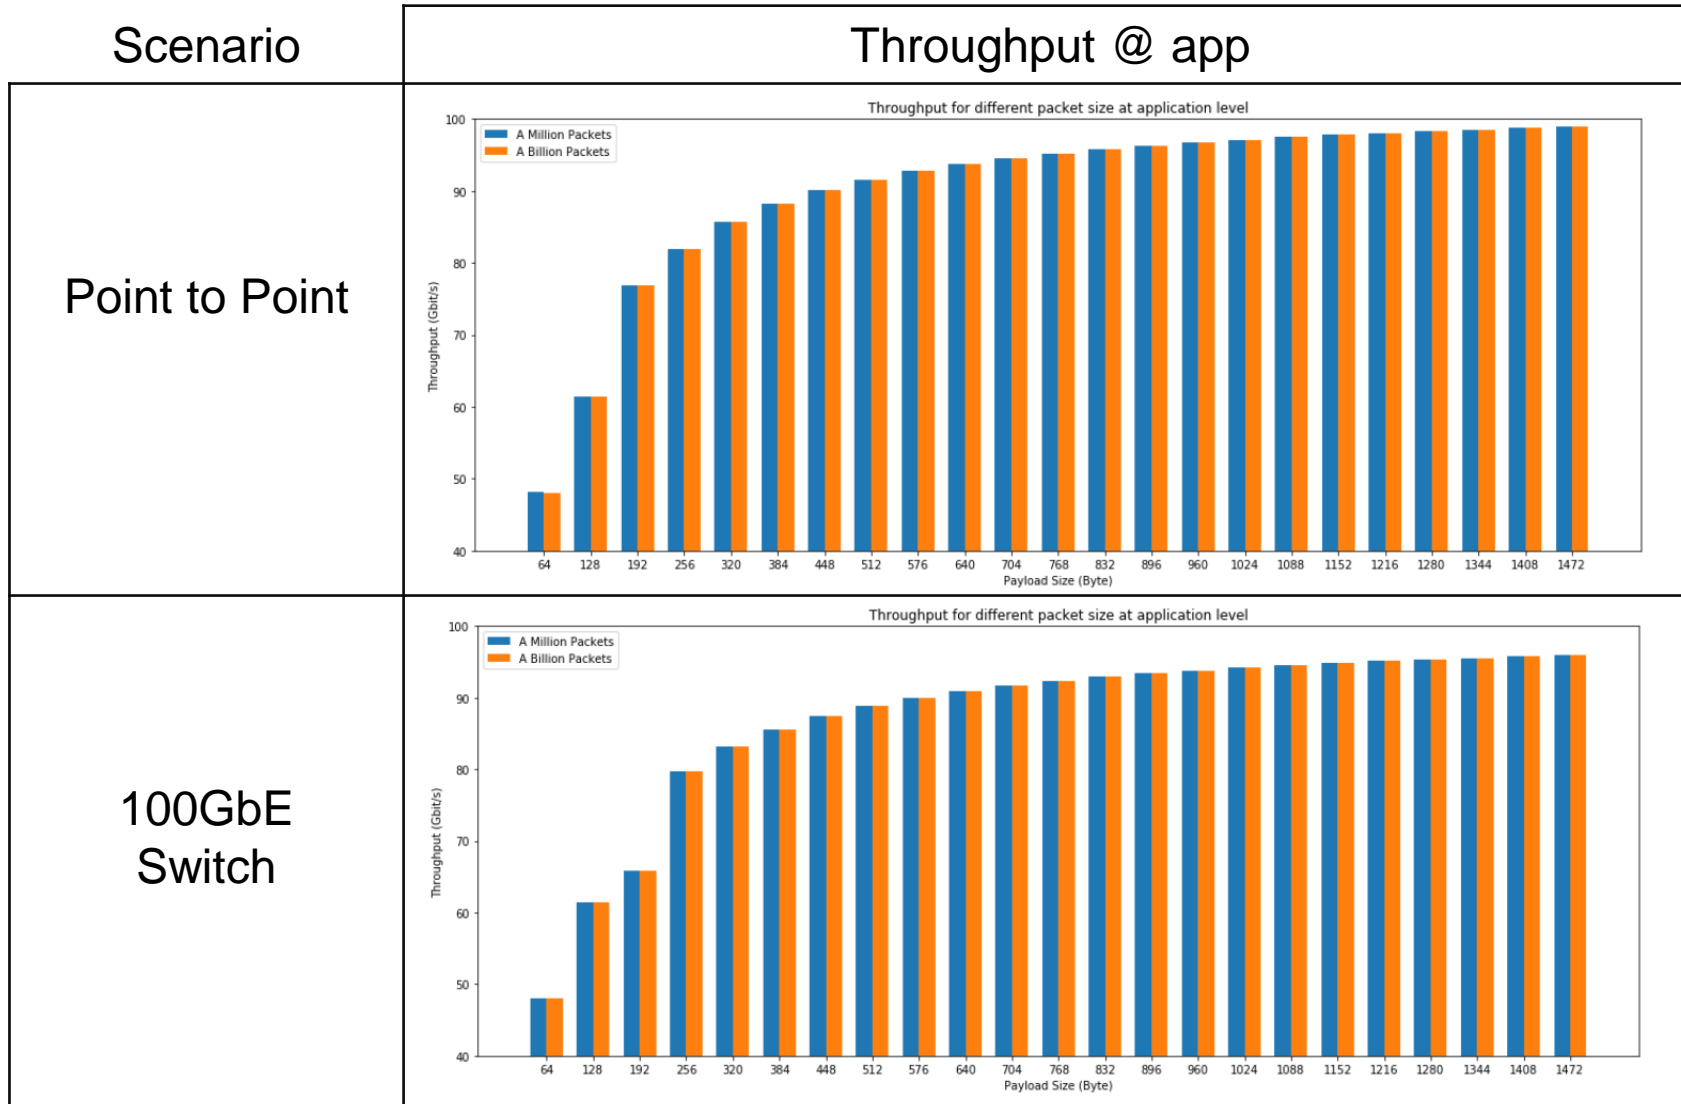
xclbin = 'vnx_benchmark_if3.xclbin'
ol_w0 = pynq.Overlay(xclbin, device=daskdev_w0)
ol_w1 = pynq.Overlay(xclbin, device=daskdev_w1)
```

# Results: Benchmark Application Latency



\*Results are included in the Notebooks

# Results: Benchmark Application Throughput



\*Results are included in the Notebooks

# VNx Summary

- 100 Gb/s UDP networking on Alveo cards
- Lightweight but unreliable
  - Ideal for point to point or controlled environment
- Simple interface
  - Easy to integrate (compute) kernels
- Two examples designs with companion PYNQ notebooks
- Dask class for distributed configuration
  - Not specific of VNx
- Python driver on top of PYNQ
- C++ driver, community contributed
- Open source
  - [https://github.com/Xilinx/xup\\_vitis\\_network\\_example](https://github.com/Xilinx/xup_vitis_network_example)

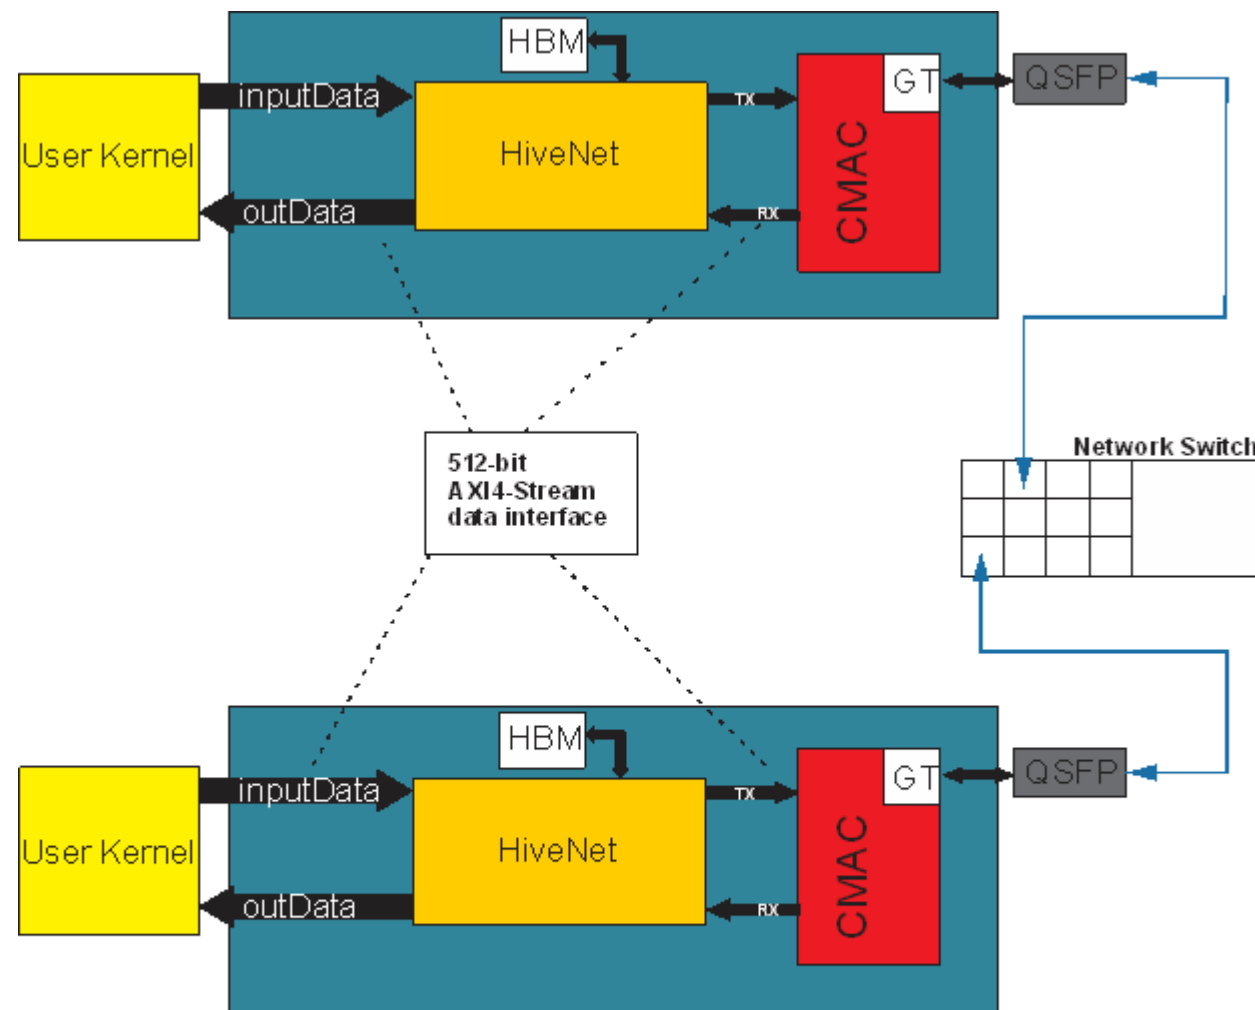
	VNx
<b>Platform(s)</b>	U50/U250/U280
<b>Host code</b>	pynq (+ dask or MPI)
<b>Transport protocol</b>	UDP (unreliable)
<b>Packet reorder</b>	No
<b>#100G Interfaces</b>	Up to 2, depends on card
<b>#Memory Banks</b>	none
<b>#Connections</b>	16
<b>Application Interface</b>	Single AXI-Stream bus
<b>Throughput</b>	Max (@ 256B+/packet)
<b>RTT</b>	2.75 $\mu$ s
<b>LUT*</b>	22,773 (1.75%)
<b>FF*</b>	48,914 (1.88%)
<b>BRAM*</b>	58.5 (2.9%)

\* Resource consumption for Network Layer only, U280 (xilinx-u280-xdma-201920.3)

# Preview: HiveNet ROCE Kernel

- ROCE protocol enables reliable transport over UDP with retransmission
- HiveNet is a Vitis-compatible ROCE kernel currently under validation
- User interface similar to VNx (AXI-Streams with 13-bit TDEST for connection ID)
- Can use same CMAC kernel as VNx
- Not as portable currently, utilizes fixed clocks only available on select platforms e.g. U55C
- Higher resource utilization compared to VNx, requirement for HBM access

<https://github.com/Xilinx/AlveoLink>



# Disclaimer & Attribution

Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change.

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

©2023 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Athlon, CDNA, EPYC, Infinity Fabric Radeon, RDNA, ROCm, Ryzen, Ryzen Threadripper, Xilinx, the Xilinx logo, Alveo, Artix, Kintex, Spartan, Versal, Vitis, Virtex, and Zynq and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft is registered trademark of Microsoft Corporation in the US and other jurisdictions. SPEC®, SPECrate®, SPECint and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See [www.spec.org](http://www.spec.org) for more information. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.



**AMD** 