

Architecture and Performance of Integrated High-speed and Versatile Embedded Networking

Ulrich Langenbach




GEFÖRDERT VOM



Bundesministerium
für Bildung
und Forschung

Outline

- Applications
 - Radar Networks
 - Embedded Video Streaming
- Requirements
- Components
 - AMD MPSoC, RFSoc
 - Corundum mqnic
 - MLE Network Protocol Acceleration Platform (NPAP)
- Solution
- Performance Tests
 - Timing
 - Bandwidth
- Conclusion



Radat Network Demonstrators

- AI based Radat Image Processing
- Radat Image Compression
- OFDM Radat, etc.
- Radat Imaging
- Front-end research

Frequency/ Trigger Distribution via

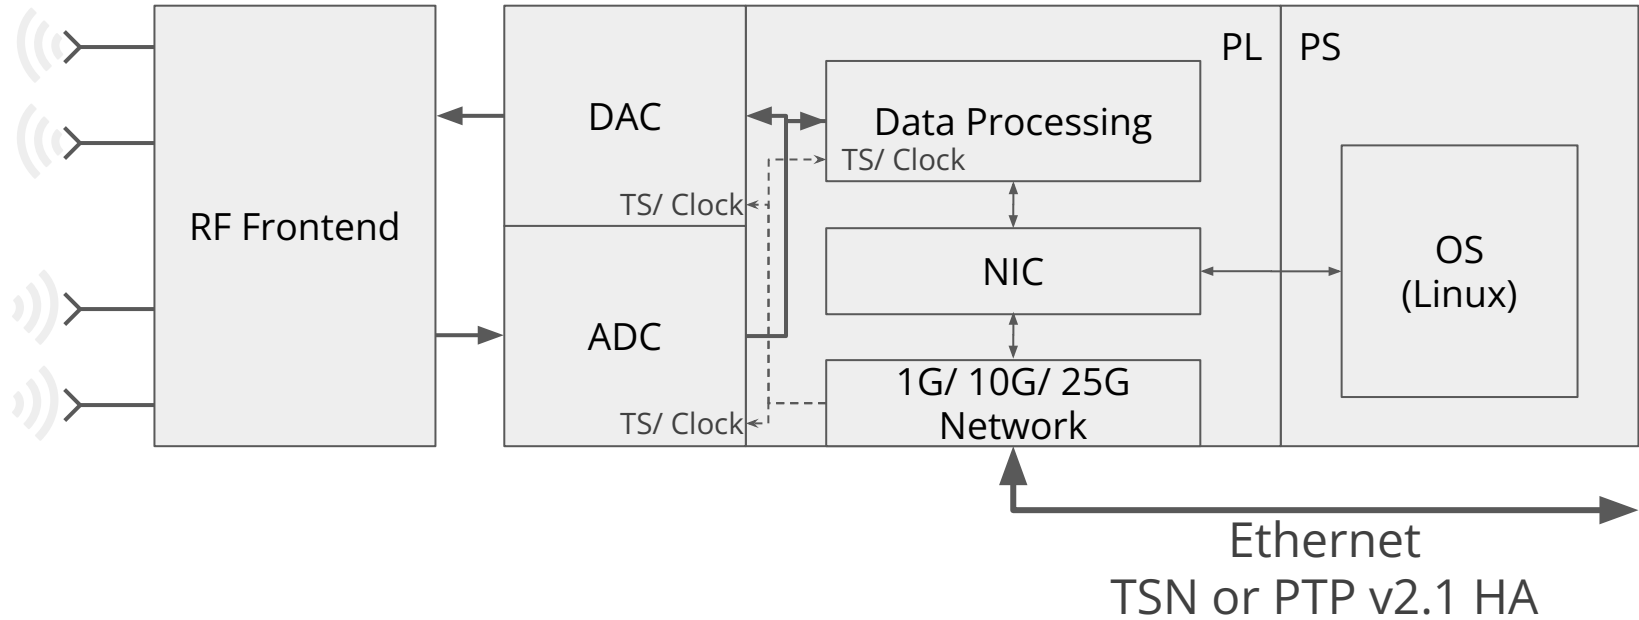
- Coax cable
- Ethernet with PTP v2.1
- Wireless LO distribution

VCXO Technologies ("Dormouse" FMC Card)

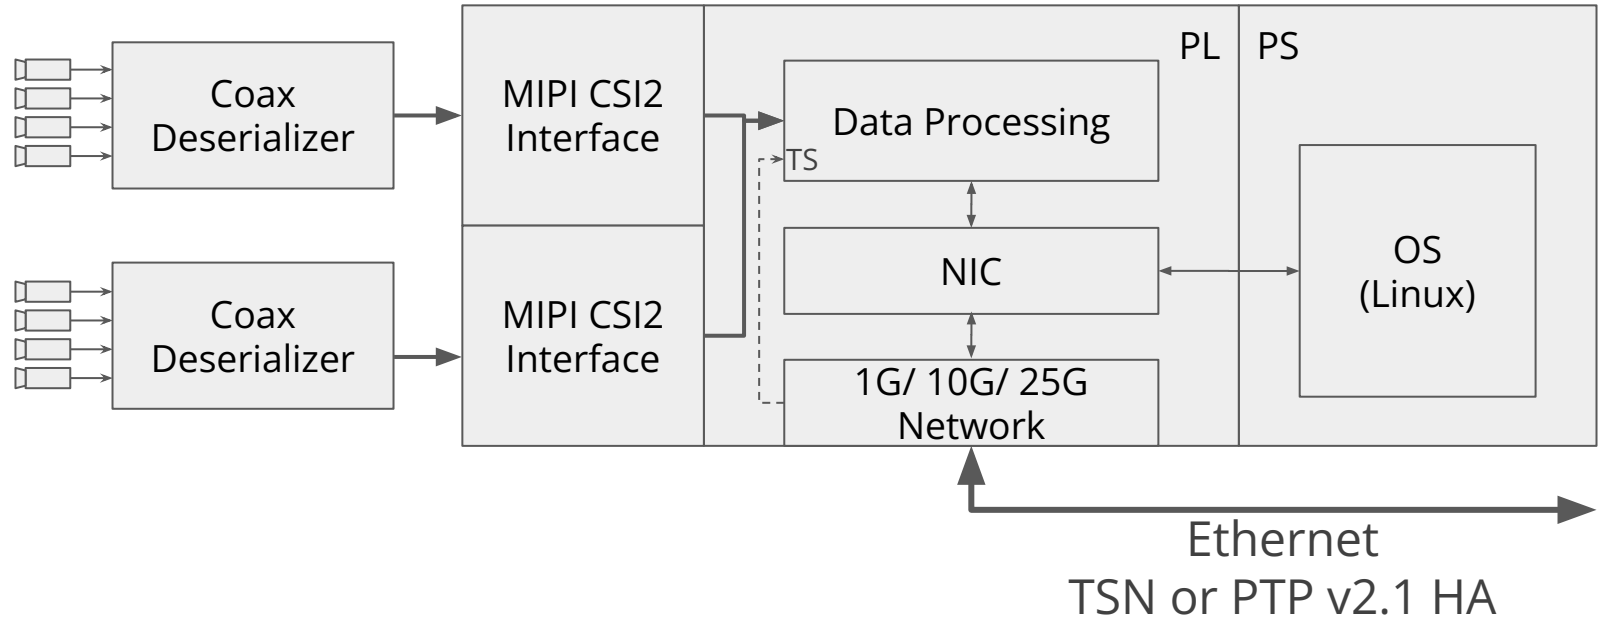
- AMD ZUP/RFSoc GTH4/GTYe4 QPLL
- SiT3521

[10]

A Networked Radar Node



Embedded Video Streaming



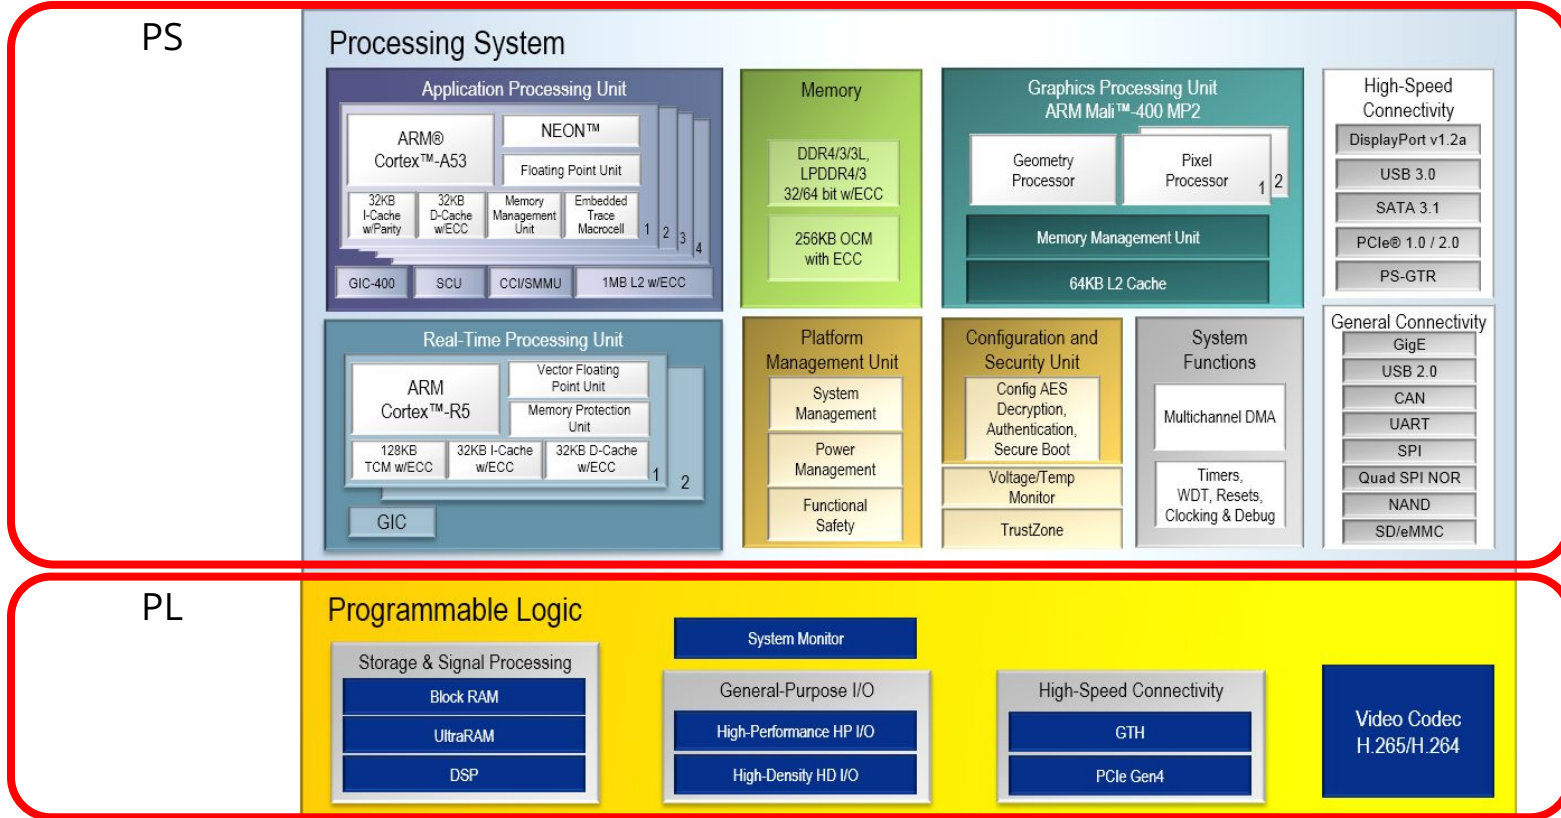
Requirements

- Standard Driven
- High-Performance: Low-Latency, High-Bandwidth
- Synchronized Time Stamping via PTP v2
- Interoperable
- Remote Operable
- Software Driven
- Interactive Usage of the System
 - Low latency connectivity for the embedded software
 - Minimize impact on bandwidth of the main data streams

Existing Components

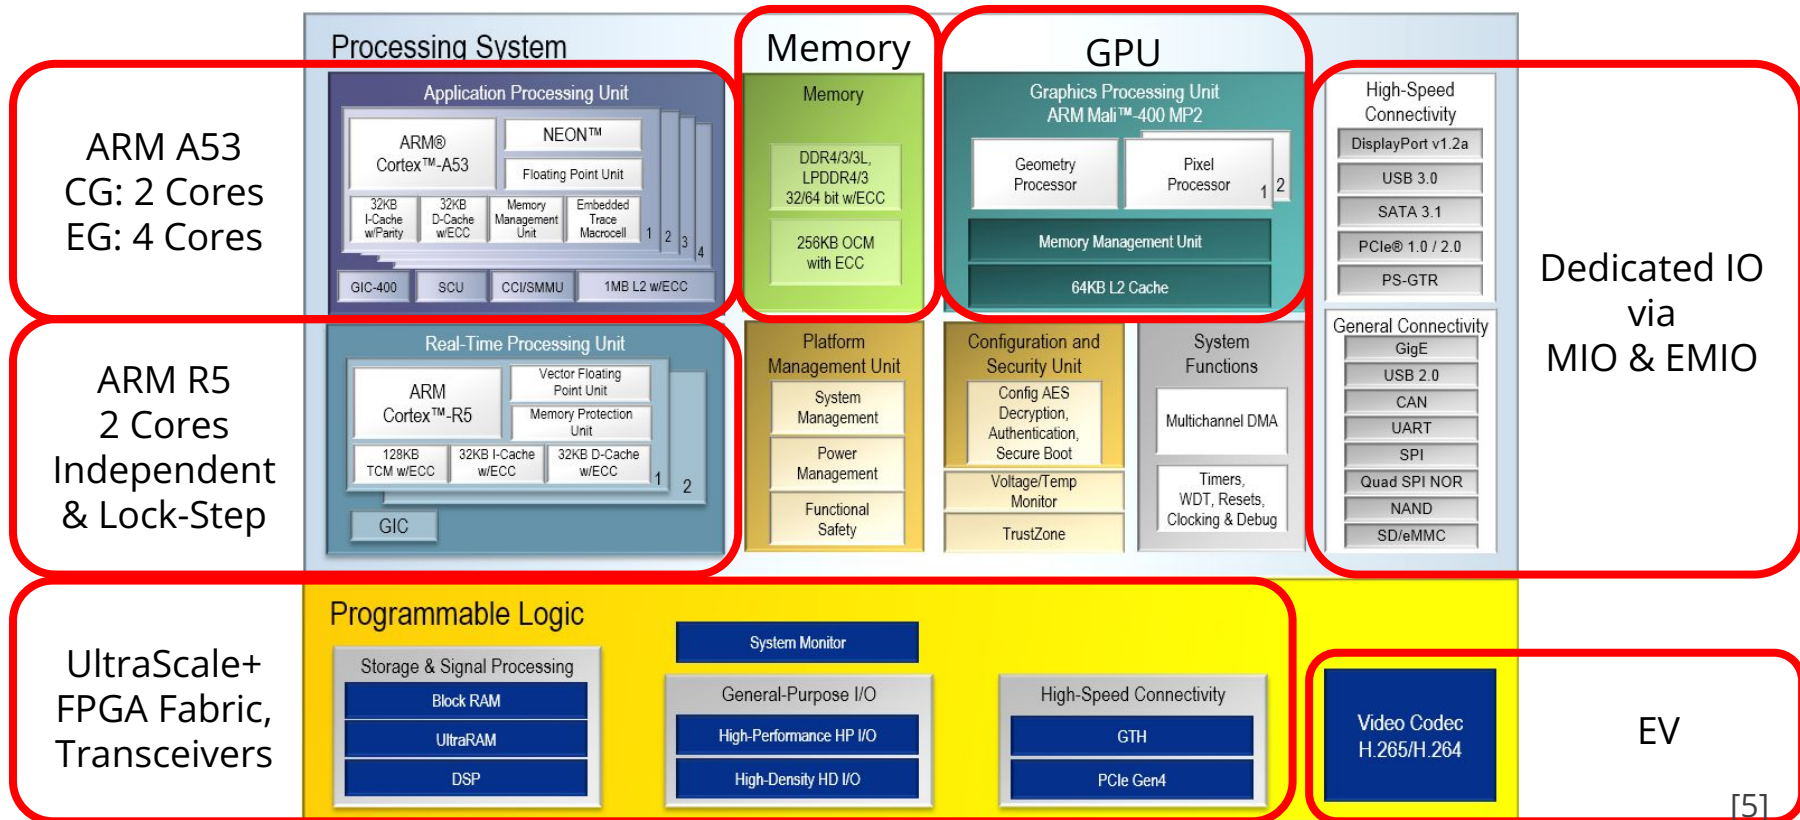
- MPSoC -> RFSoc
- Corundum mqnic, PCIe NIC -> Embedded NIC
- Network Protocol Accelerator Platform (NPAP)

AMD MPSoC (CG, EG, EV)



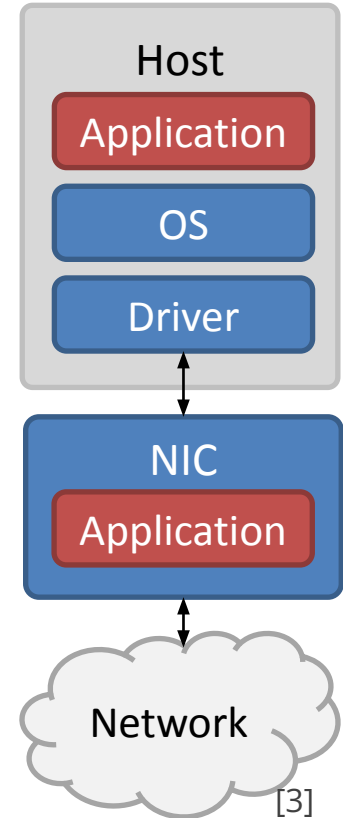
[5]

AMD MPSoC (CG, EG, EV)

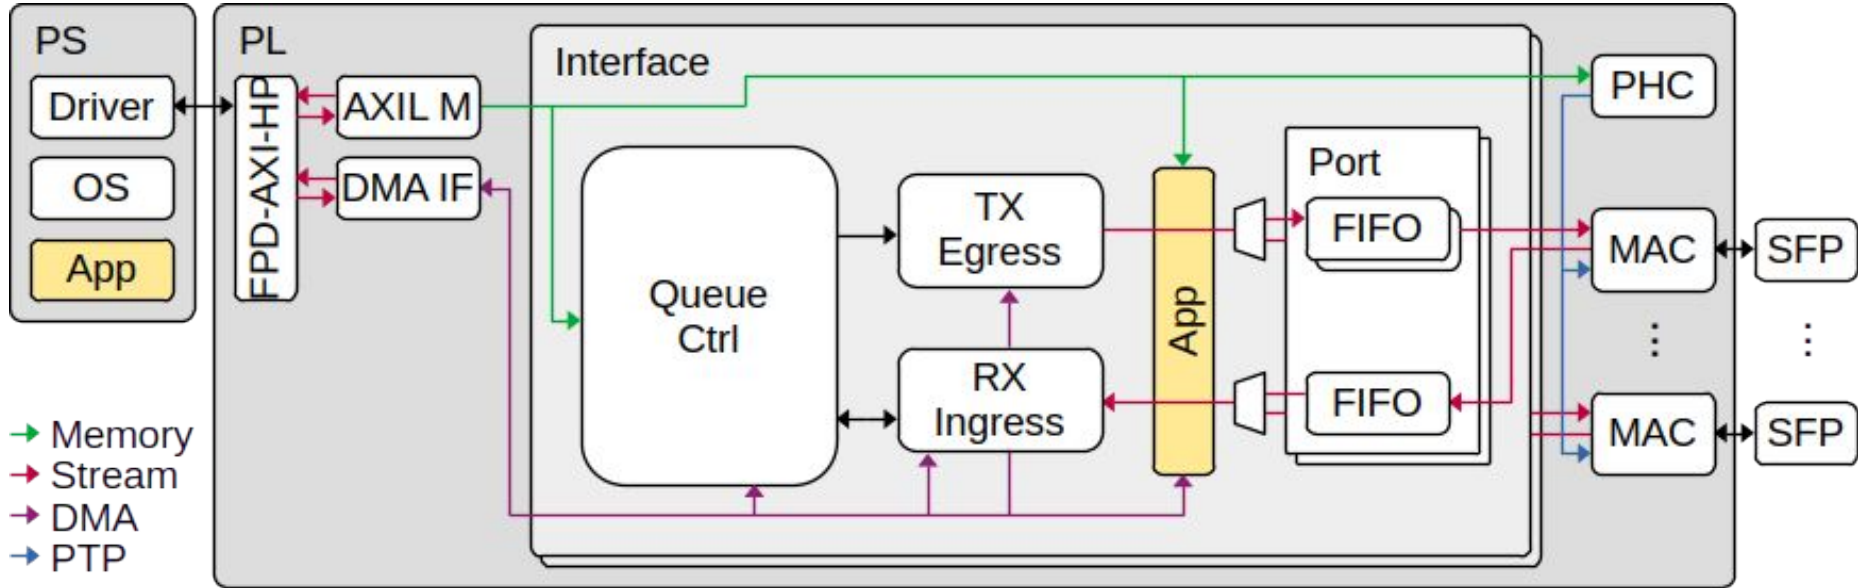


Corundum: An Open Source FPGA based NIC

- Network Interface Controller (NIC) connects software to the network
- NIC functionality is evolving
 - Line rate increases
 - Offload networking functions from CPU to NIC
- More general: in-network compute
 - Offload compute to programmable NICs, switches, etc.
 - Not limited to network stack
 - Application block with full access to all data and resources
- Open Source, includes Gateware, Driver, etc.
- Fine-grained traffic control via 10k+ hardware queues
- Provides PTP v2 timestamping and time synchronisation
- Wide AMD and Altera device support



Corundum: Simplified mqnic Block Diagram



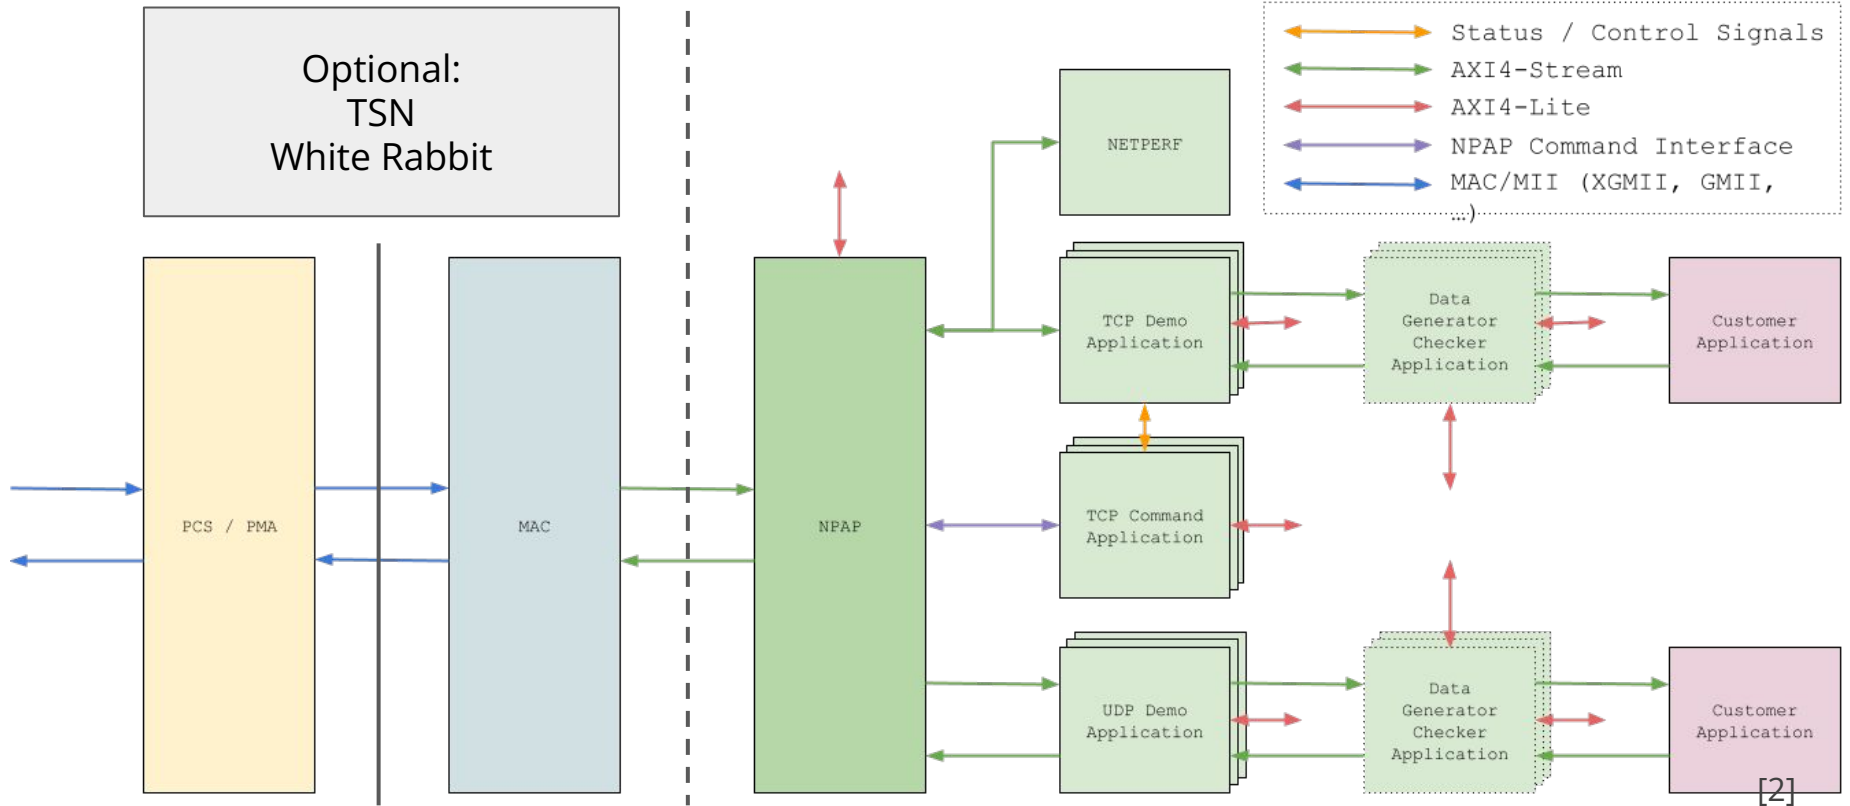
[1], [3]

Network Protocol Accelerator Platform (NPAP)

- Implements Ethernet, IP, UDP, TCP Protocol Layers in Logic
- Platform- independent HDL code, currently supported platforms include
 - AMD 7 series to Versal Platforms
 - Altera Stratix 10
 - Microchip PolarFire
- > 50 Gbps of bandwidth
- Provides RTT latencies < 10 us for 10 GbE Link speed
- AMD US+ Resource Consumption (UDP + 3 TCP Sessions) ~ 35k LUTs
- BRAM usage mostly depend on bandwidth-delay-product requirements of the system
- Includes licensed technology from Fraunhofer HHI
- Maintained and Supported by MLE

[2]

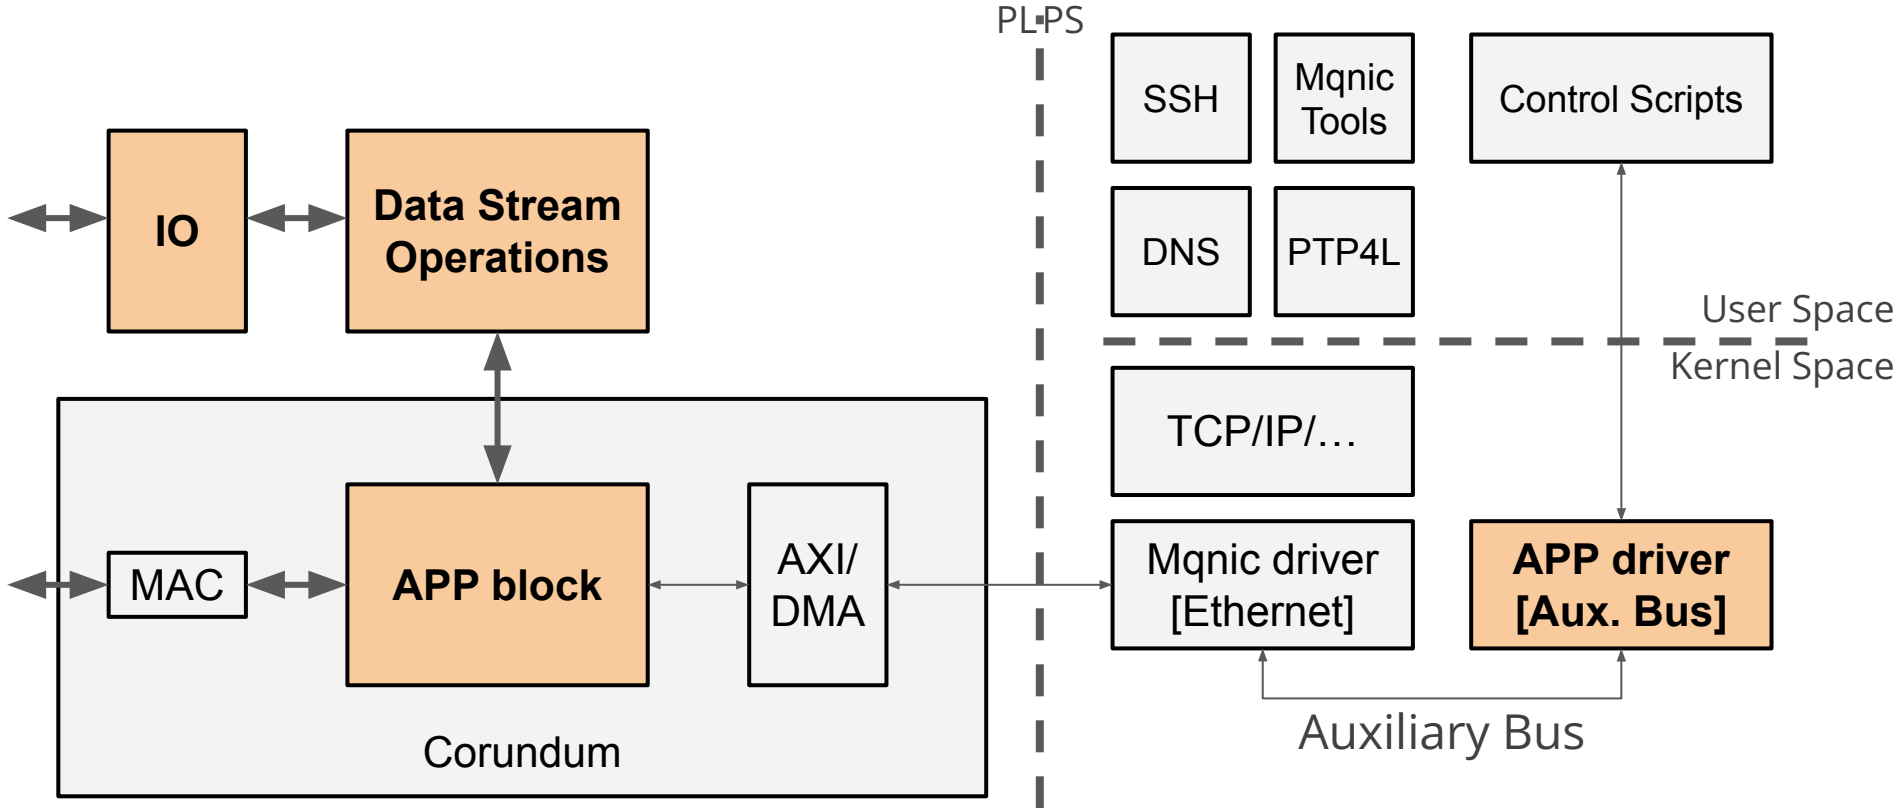
NPAP Subsystem



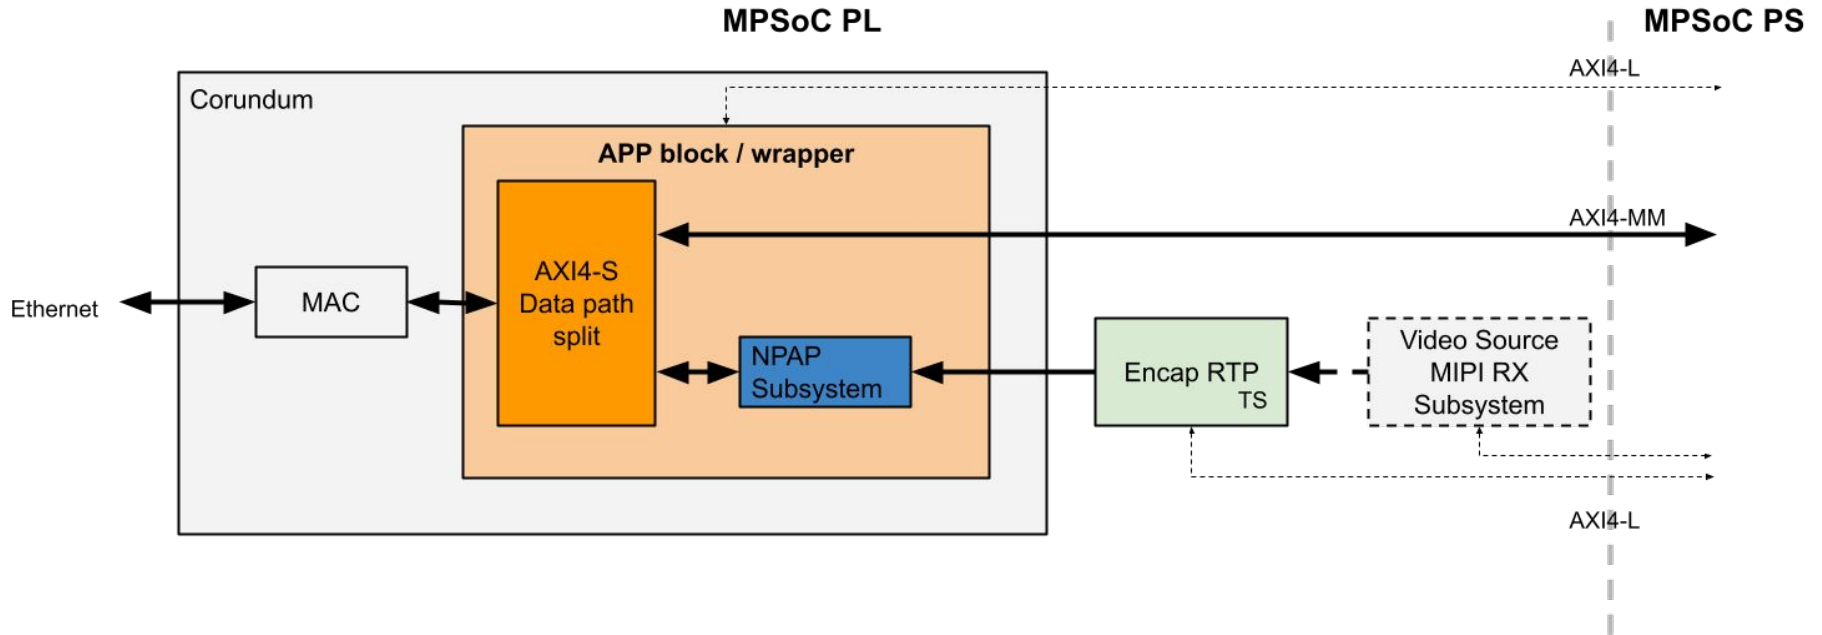
Solution

- Architecture Overview
- Video Solution Architecture
- Solution Platforms

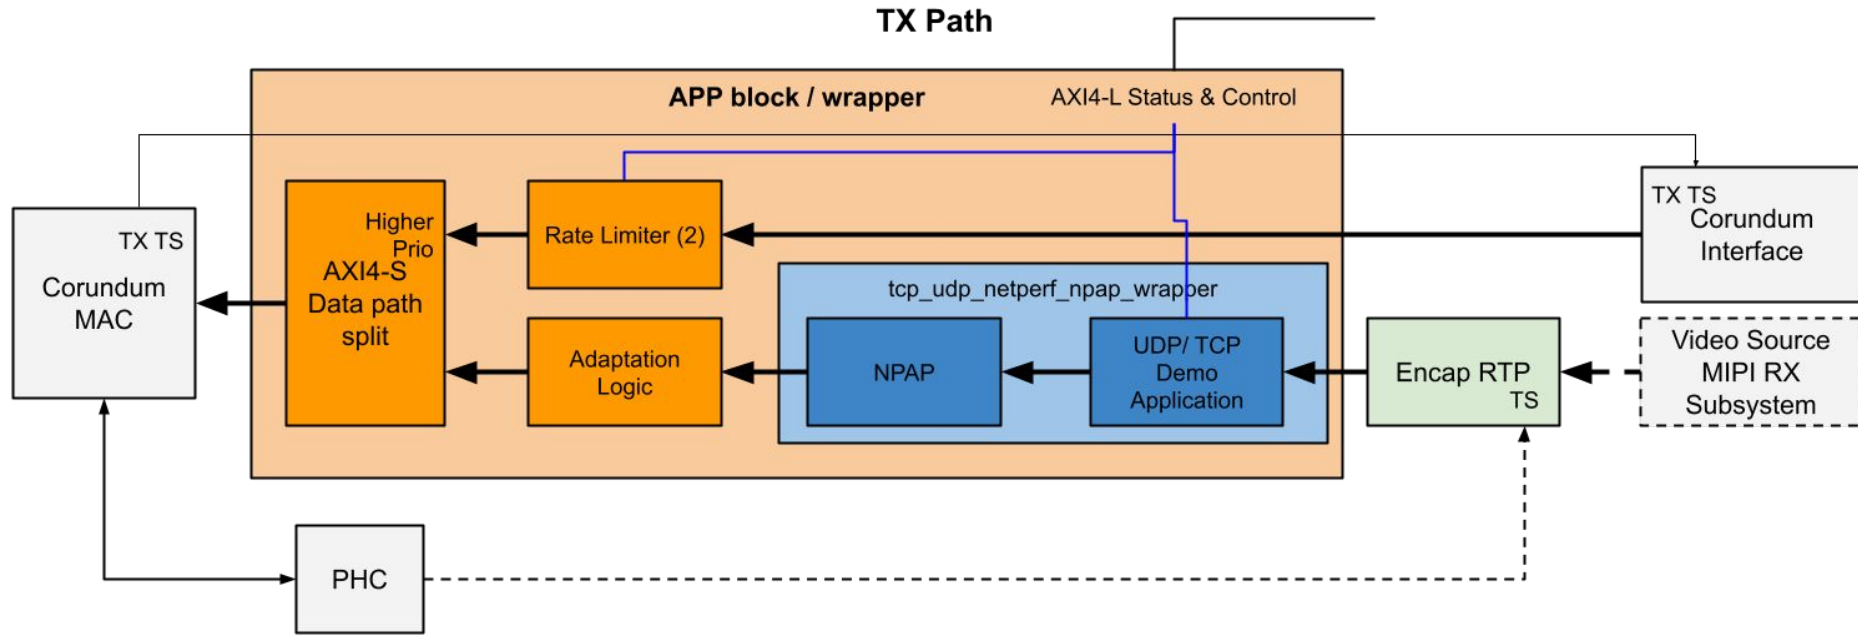
Solution: Architecture Overview



Video Solution: Architecture Overview



Video Solution: Architecture Details (Transmit)




Solution Platforms

AD-GMSL2ETH-SL
(AMD KR26 + AD Carrier)



[7]

AMD ZCU102





[8]

Performance Tests

- PTP Synchronization -> PPS Jitter
 - Each node provides a PPS output
 - Oscilloscope Measurement of the outputs of the nodes
 - Digital Phosphor View
- Bandwidth
 - Software logs to show the bandwidths


Performance Test Setup: PTP v2 Timing

- 2x ZCU102
- Directly connected, point-to-point
- Linux running PTP4L
- Unloaded Network



Performance Test Setup: PTP v2 PPS Jitter

- < 30 ns relative PPS jitter
- No global synchronisation to UTC
- Unloaded network




Performance Test Setup: Bandwidth

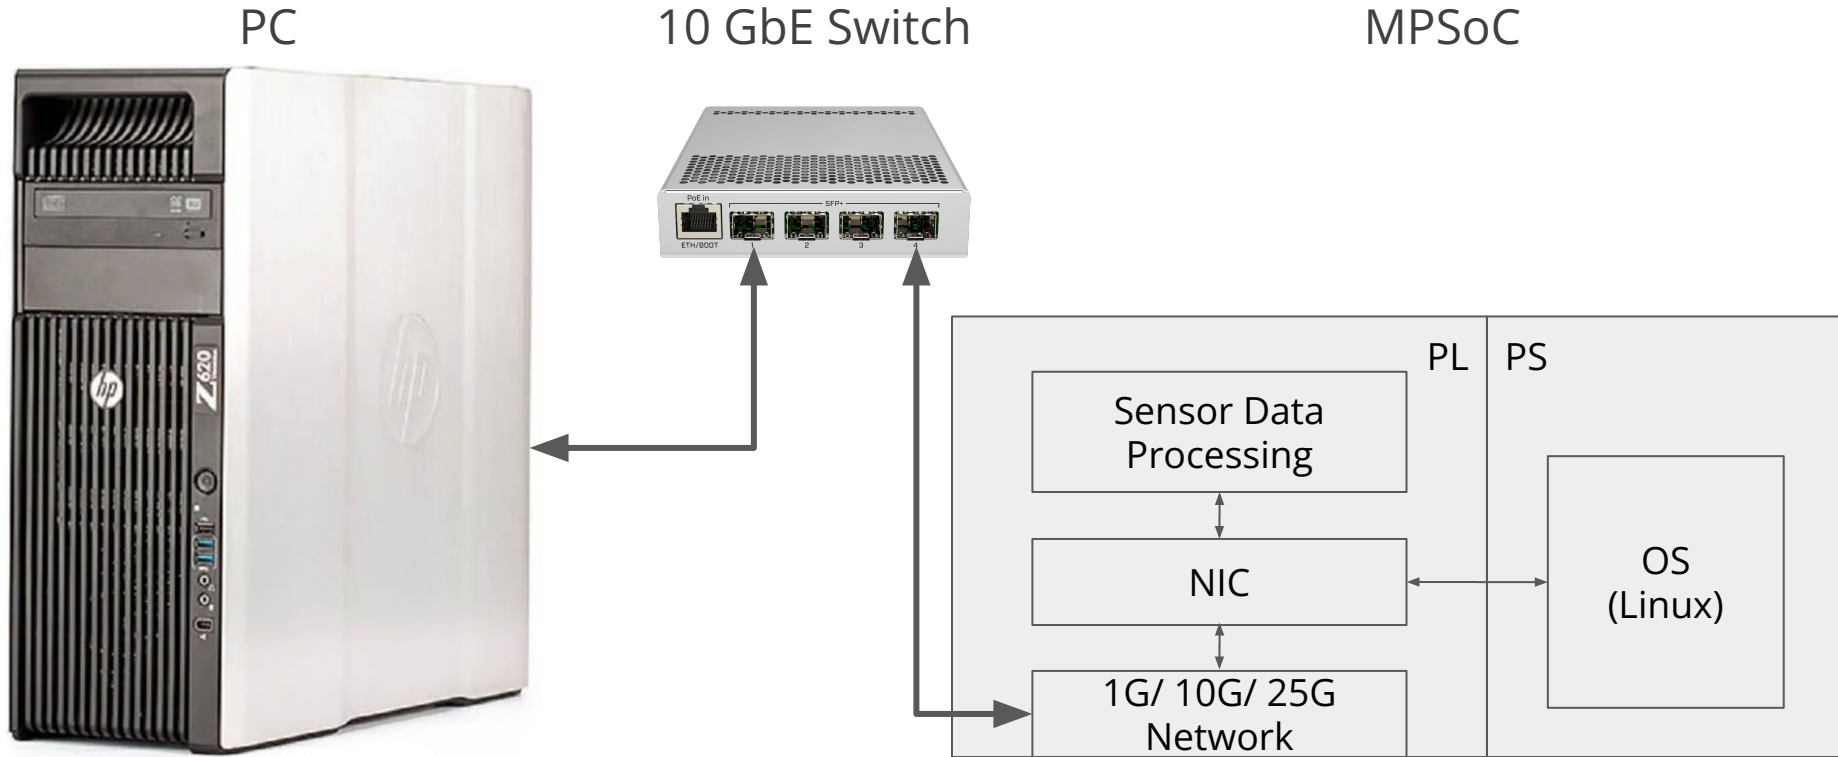
PC

10 GbE Switch

ZCU102 DevKit



Bandwidth Performance Test Setup: Data Streams



Bandwidth Performance Test Setup: Data Streams

	PC	PS	PL
PC	-	1x TCP	-
PS	1x TCP	-	-
PL	4x TCP	-	-

All of these data streams are concurrently active.

Bandwidth and stability (~1 d runtime)

CPU Load

PL -> PC

PC -> PS

PS -> PC

ZCU102

The screenshot displays a multi-panel terminal window with the following sections:

- Top Row (CPU Load):** Four terminal windows showing system logs for 'root@zcu102-5 - Konsole'. The logs include system boot information, network interface status (eth0, eth1), and various system messages.
- Middle Row (Network Statistics):** Four terminal windows showing network traffic statistics for 'root@zcu102-5 - Konsole'. Each window displays a table of received and transmitted data for various IP addresses and ports.
- Bottom Row (Switch Configuration):** A terminal window titled 'root@switch - Konsole' showing network configuration for a switch. It lists interfaces (eth0, eth1, eth2, eth3) and their corresponding IP addresses and speeds.

Interface	IP Address	Speed
eth0	1500	1592 99.0 Gbps
eth1	1500	1592 99.0 Gbps
eth2	1500	1592 99.0 Gbps
eth3	1500	1592 99.0 Gbps

PC

Switch

Bandwidth and stability (~1 d runtime)

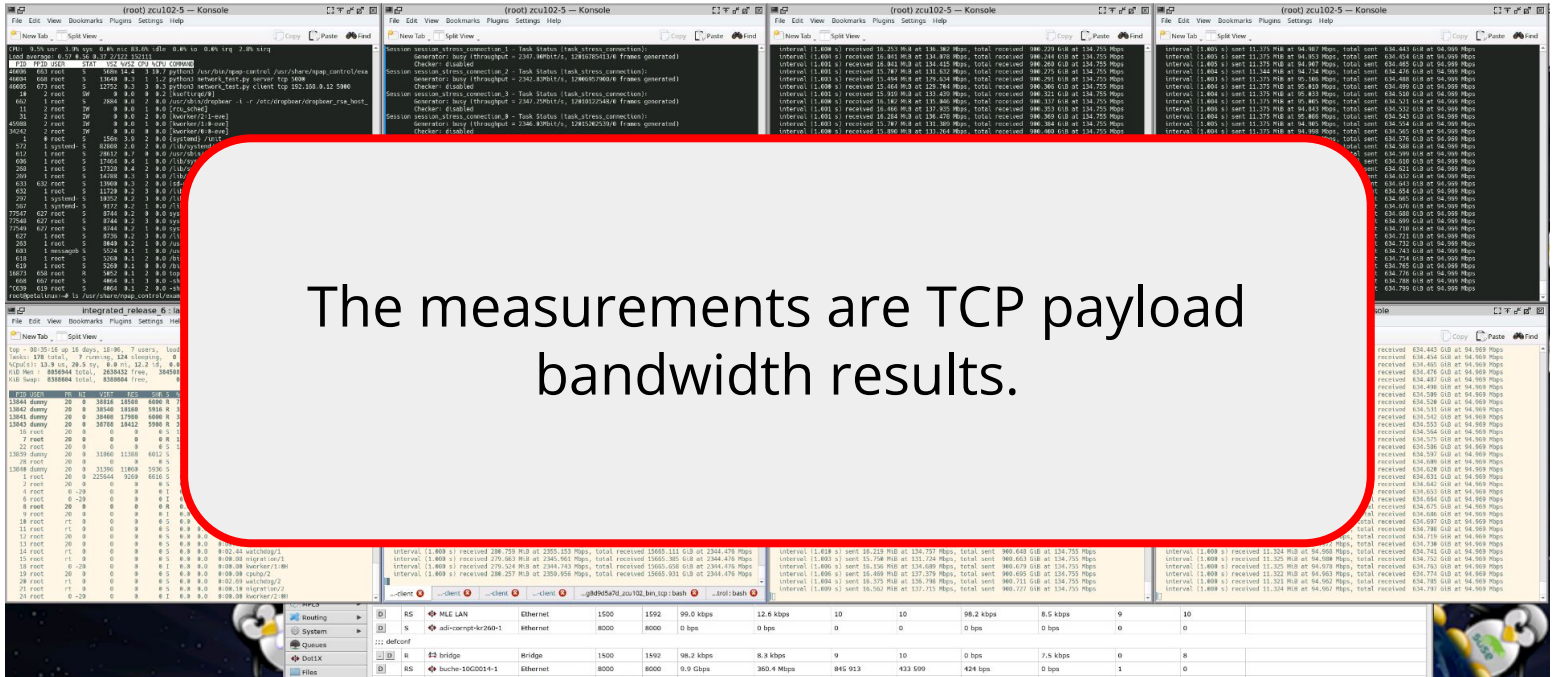
CPU Load

PL -> PC

PC -> PS

PS -> PC

ZCU102



The measurements are TCP payload bandwidth results.

PC

Switch

Bandwidth CPU Utilization

ZCU102

PC

```
(root) zcu102-5 — Konsole
CPU: 9.5% usr 3.9% sys 0.0% nic 81.6% idle 0.0% io 0.0% irq 2.8% sirq
Load average: 0.57 0.56 0.37 2/122 152111
PID PPID USER STAT VSZ %VSZ CPUS COMMAND
46006 663 root S 568m 14.4 19.7 python3 /usr/bin/npap-control /usr/share/npap-control/eva
46004 668 root S 1364m 3.3 1.2 python3 network_test.py server tcp 500#
46005 673 root S 12752 0.3 3 0.3 python3 network_test.py client tcp 192.168.0.12 5900
1# 2 root SW 0 0.0 0 0.2 [ksftirq/0#]
662 1 root S 2884 0.0 2 0.0 /usr/sbin/dropbear -i -r /etc/dropbear/dropbear_rsa_host_
11 2 root IW 0 0.0 1 0.0 [rcu_sched]
31 2 root IW 0 0.0 2 0.0 [kworker/2:1#eve]
45988 2 root IW 0 0.0 1 0.0 [kworker/1:0#eve]
34242 2 root IW 0 0.0 0 0.0 [kworker/0:0#eve]
1 0 root S 156m 3.9 2 0.0 {systemd} /unit
572 1 systemd- S 82808 2.0 2 0.0 /lib/systemd/systemd-timesyncd
612 1 root S 28612 0.7 0 0.0 /usr/sbin/strf-agent --L-
606 1 root S 17464 0.4 1 0.0 /lib/systemd/systemd-logind
268 1 root S 17328 0.4 2 0.0 /lib/systemd/systemd-journald
289 1 root S 14788 0.3 3 0.0 /lib/systemd/systemd-udev
633 632 root S 13900 0.3 2 0.0 [sd-pam]
632 1 root S 11720 0.2 3 0.0 /lib/systemd/systemd --user
297 1 systemd- S 10352 0.2 3 0.0 /lib/systemd/systemd-networkd
567 1 systemd- S 9172 0.2 1 0.0 /lib/systemd/systemd-resolved
77547 627 root S 8744 0.2 0 0.0 systemd-userwork
77548 627 root S 8744 0.2 3 0.0 systemd-userwork
77549 627 root S 8744 0.2 1 0.0 systemd-userwork
627 1 root S 8736 0.2 3 0.0 /lib/systemd/systemd-userdbd
263 1 root S 8040 0.2 1 0.0 /usr/sbin/haveged -w 1024 -v 1 --Foreground
603 1 messageb S 5524 0.1 1 0.0 /usr/bin/dbus-daemon --system --address=systemd: --fork
618 1 root S 5260 0.1 2 0.0 /bin/login -p --
619 1 root S 5260 0.1 0 0.0 /bin/login -f
18873 658 root R 5052 0.1 2 0.0 top
668 667 root S 4864 0.1 1 0.0 -sh
^C639 619 root S 4864 0.1 2 0.0 -sh
root@petalinux:~# ls /usr/share/npap_control/example_stress_4connection.cfg
```

```
integrated_release_6 : labgrid-client — Konsole
top - 08:35:16 up 16 days, 10:06, 7 users, load average: 2.00, 2.73, 2.66
Tasks: 170 total, 7 running, 124 sleeping, 0 stopped, 0 zombie
%Cpu(s): 13.9 us, 20.5 sy, 0.0 ni, 12.2 id, 0.0 wa, 0.0 hi, 53.4 si, 0.0 st
KiB Mem : 8056944 total, 2638432 free, 384508 used, 5034004 buff/cache
KiB Swap: 8388604 total, 8388604 free, 0 used, 7332292 avail Mem

PID USER PR NI VIRT RES SHR S %CPU MEM TIME+ COMMAND
13844 dummy 20 0 38816 18508 6000 0 75.4 0.2 715:39.39 python3
13842 dummy 20 0 38540 18160 5916 0 33.2 0.2 353:35.41 python3
13841 dummy 20 0 38408 17900 6000 0 32.2 0.2 303:20.16 python3
13843 dummy 20 0 38788 18412 5900 0 31.2 0.2 312:59.50 python3
16 root 20 0 0 0 0 0 0.0 0.0 167:46.81 ksftirq/1
7 root 20 0 0 0 0 0 0.0 0.0 378:29.87 ksftirq/0
22 root 20 0 0 0 0 0 0.0 0.0 145:15.55 ksftirq/2
13800 dummy 20 0 13800 6012 5 9.3 0.1 88:34.17 python3
0 0 0 0 0 0 0.0 0.0 31:34.56 ksftirq/3
0 0 0 0 0 0 0.0 0.0 18:59.18 python3
269 6616 S 0.0 0.1 0:29.40 systemd
0 0 0 0 0 0 0.0 0.0 0:08.18 kthreadd
0 0 0 0 0 0 0.0 0.0 0:08.08 kworker/0:0#H
0 0 0 0 0 0 0.0 0.0 0:08.08 mm_percpu_wq
0 0 0 0 0 0 0.0 0.0 1:30.11 rcu_sched
0 0 0 0 0 0 0.0 0.0 0:08.08 rcu_bh
10 root rt 0 0 0 0 0.0 0.0 0:08.07 migration/0
11 root rt 0 0 0 0 0.0 0.0 0:02.69 watchdog/0
12 root 20 0 0 0 0 0.0 0.0 0:08.08 cpulp/0
13 root 20 0 0 0 0 0.0 0.0 0:08.08 cpulp/1
14 root rt 0 0 0 0 0.0 0.0 0:02.44 watchdog/1
15 root rt 0 0 0 0 0.0 0.0 0:08.08 migration/1
18 root 0 -20 0 0 0 0.0 0.0 0:08.08 kworker/1:0#H
19 root 20 0 0 0 0 0.0 0.0 0:08.08 cpulp/2
20 root rt 0 0 0 0 0.0 0.0 0:02.69 watchdog/2
21 root rt 0 0 0 0 0.0 0.0 0:08.18 migration/2
24 root 0 -20 0 0 0 0.0 0.0 0:08.08 kworker/2:0#H
```

10 %
1x 70 %
3x 32 %

Bandwidth: PL - PC (Sensor Data)

ZCU102

PC (4x)

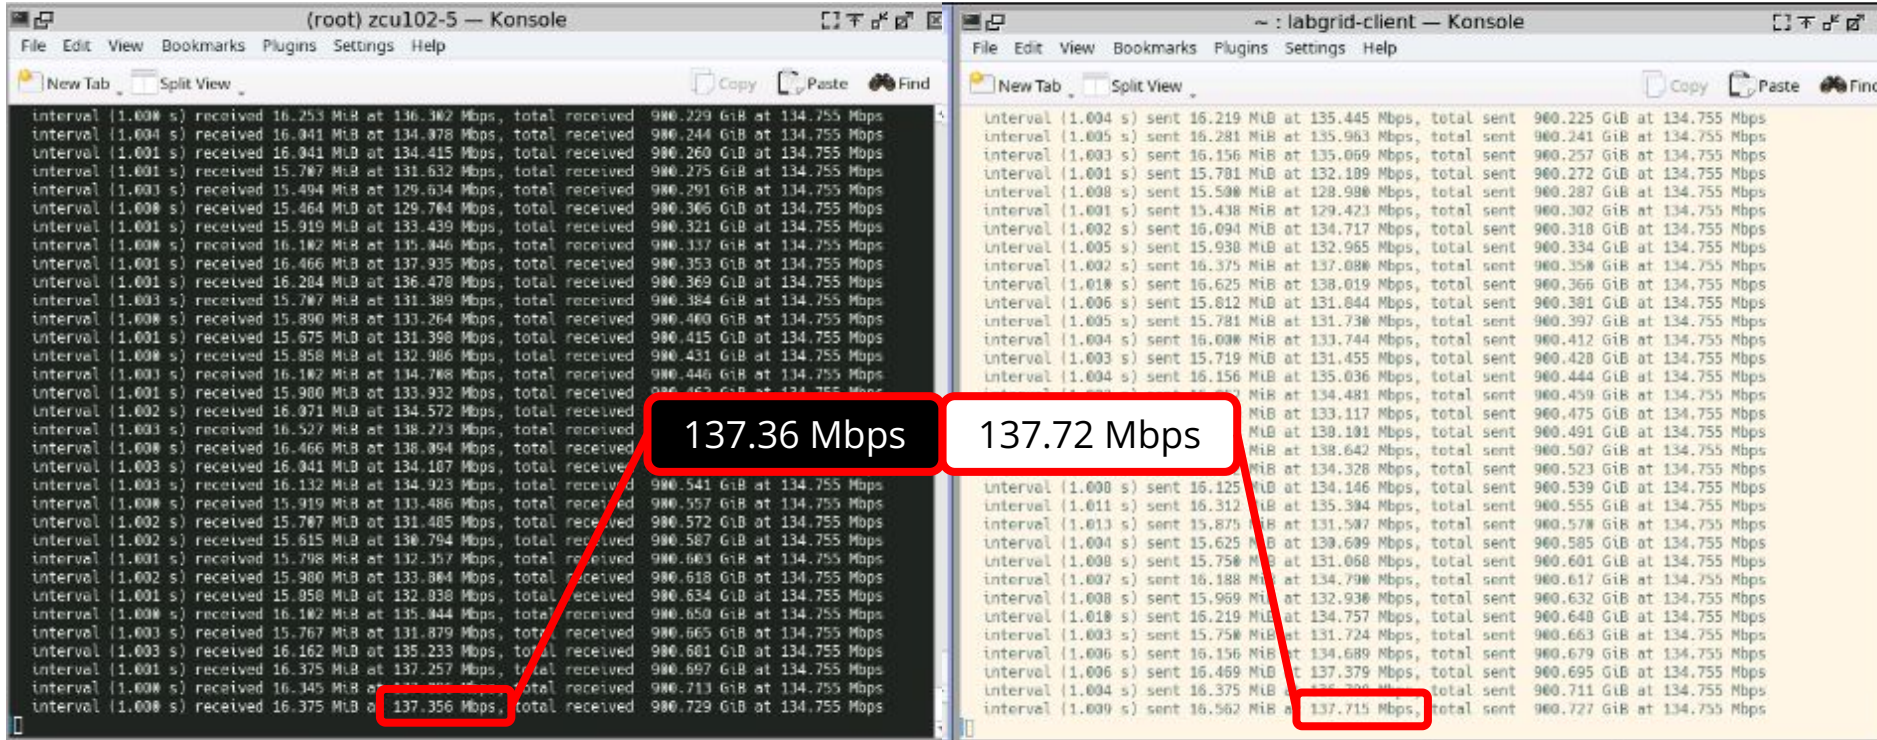
```
(root) zcu102-5 — Konsole
File Edit View Bookmarks Plugins Settings Help
New Tab Split View Copy Paste Find
Session session_stress_connection_1 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2347.90Mbit/s, 12016785413/0 frames generated)
Checker: disabled
Session session_stress_connection_2 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2342.83Mbit/s, 12006957900/0 frames generated)
Checker: disabled
Session session_stress_connection_3 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2347.25Mbit/s, 12010122548/0 frames generated)
Checker: disabled
Session session_stress_connection_0 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2346.03Mbit/s, 12015202539/0 frames generated)
Checker: disabled
Session session_stress_connection_1 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2342.57Mbit/s, 12016844029/0 frames generated)
Checker: disabled
Session session_stress_connection_2 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2327.92Mbit/s, 12007016446/0 frames generated)
Checker: disabled
Session session_stress_connection_3 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2333.09Mbit/s, 12010100919/0 frames generated)
Checker: disabled
Session session_stress_connection_0 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2340.01Mbit/s, 12015261149/0 frames generated)
Checker: disabled
Session session_stress_connection_1 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2346.39Mbit/s, 12016902907/0 frames generated)
Checker: disabled
Session session_stress_connection_2 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2361.89Mbit/s, 12007075077/0 frames generated)
Checker: disabled
Session session_stress_connection_3 - Task Status [task_stress_connection]:
Generator: busy (throughput = 2347.17Mbit/s, 12010239453/0 frames generated)
Checker: disabled

~ : labgrid-client — Konsole
File Edit View Bookmarks Plugins Settings Help
New Tab Split View Copy Paste Find
Interval (1.000 s) received 279.428 MiB at 2344.613 Mbps, total received 15657.735 GiB at 2344.475 Mbps
Interval (1.000 s) received 279.471 MiB at 2344.276 Mbps, total received 15658.008 GiB at 2344.475 Mbps
Interval (1.000 s) received 280.414 MiB at 2352.205 Mbps, total received 15658.281 GiB at 2344.475 Mbps
Interval (1.000 s) received 279.705 MiB at 2346.246 Mbps, total received 15650.555 GiB at 2344.475 Mbps
Interval (1.000 s) received 279.677 MiB at 2346.096 Mbps, total received 15658.828 GiB at 2344.475 Mbps
Interval (1.000 s) received 280.263 MiB at 2350.951 Mbps, total received 15659.101 GiB at 2344.475 Mbps
Interval (1.000 s) received 279.351 MiB at 2343.222 Mbps, total received 15659.374 GiB at 2344.475 Mbps
Interval (1.000 s) received 279.519 MiB at 2344.770 Mbps, total received 15659.647 GiB at 2344.475 Mbps
Interval (1.000 s) received 280.235 MiB at 2350.562 Mbps, total received 15659.921 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.593 MiB at 2345.375 Mbps, total received 15660.194 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.062 MiB at 2347.633 Mbps, total received 15660.467 GiB at 2344.476 Mbps
Interval (1.000 s) received 280.374 MiB at 2351.851 Mbps, total received 15660.741 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.476 MiB at 2344.305 Mbps, total received 15661.014 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.730 MiB at 2346.383 Mbps, total received 15661.287 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.131 MiB at 2341.350 Mbps, total received 15661.560 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.538 MiB at 2344.840 Mbps, total received 15661.833 GiB at 2344.476 Mbps
Interval (1.000 s) received 280.159 MiB at 2342.993 Mbps, total received 15662.105 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.112 MiB at 2347.201 Mbps, total received 15662.379 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.459 MiB at 2344.213 Mbps, total received 15662.652 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.665 MiB at 2345.361 Mbps, total received 15662.925 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.837 MiB at 2347.339 Mbps, total received 15663.198 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.579 MiB at 2345.362 Mbps, total received 15663.471 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.535 MiB at 2344.898 Mbps, total received 15663.745 GiB at 2344.476 Mbps
Interval (1.000 s) received 200.759 MiB at 2355.153 Mbps, total received 15664.018 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.663 MiB at 2345.961 Mbps, total received 15664.291 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.524 MiB at 2345.362 Mbps, total received 15664.564 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.535 MiB at 2344.898 Mbps, total received 15664.837 GiB at 2344.476 Mbps
Interval (1.000 s) received 200.759 MiB at 2355.153 Mbps, total received 15665.111 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.663 MiB at 2345.961 Mbps, total received 15665.385 GiB at 2344.476 Mbps
Interval (1.000 s) received 279.524 MiB at 2345.362 Mbps, total received 15665.658 GiB at 2344.476 Mbps
Interval (1.000 s) received 200.257 MiB at 2350.956 Mbps, total received 15665.931 GiB at 2344.476 Mbps
```

Bandwidth PC -> PS

ZCU102

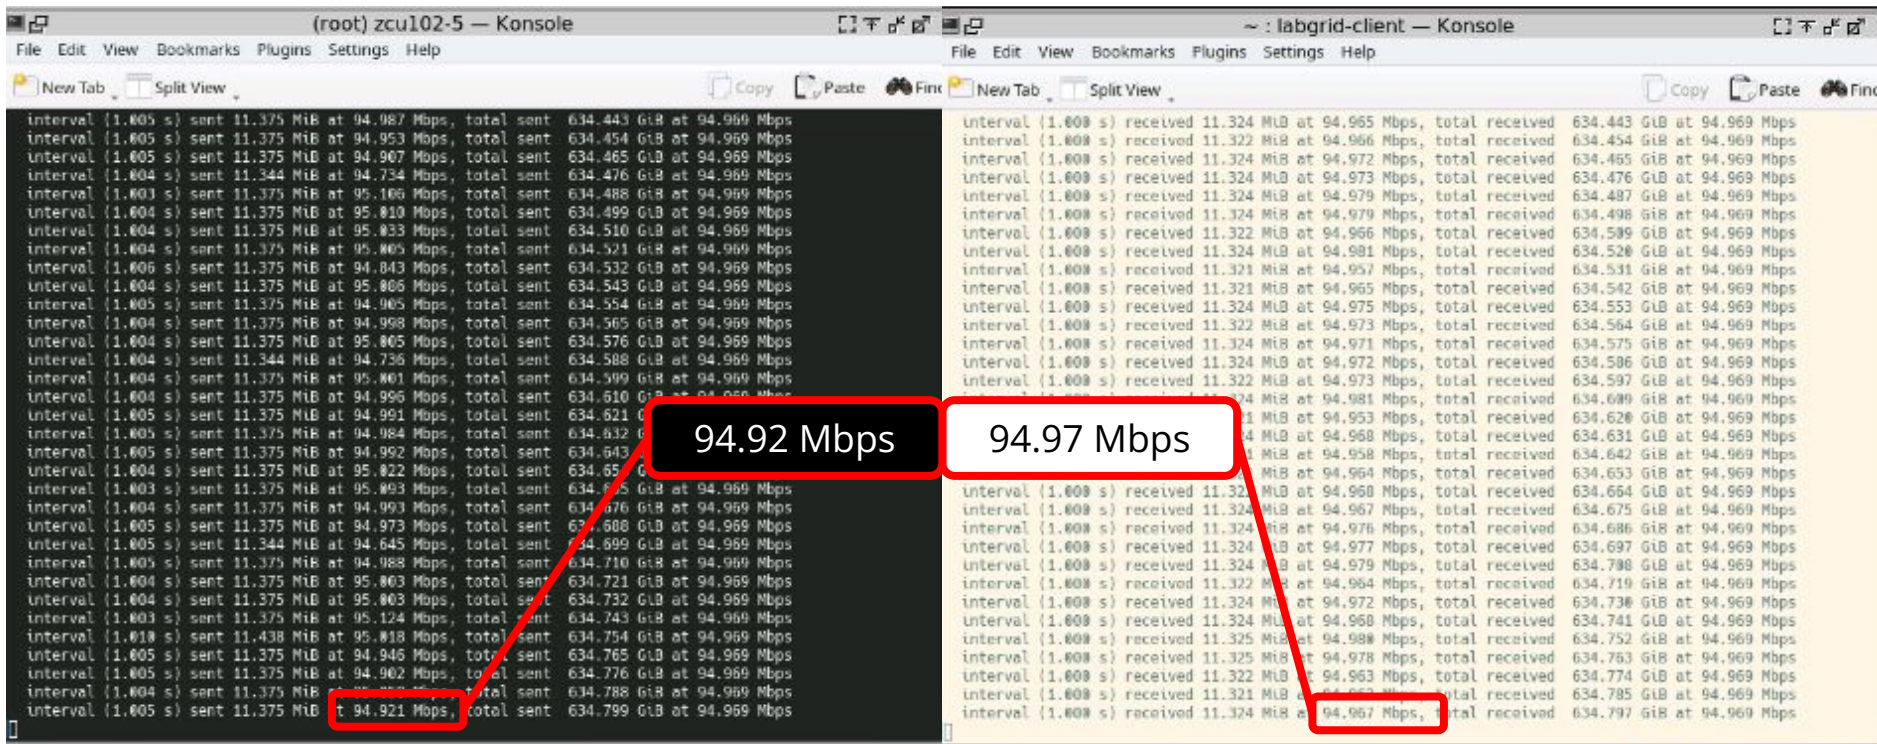
PC



Bandwidth PS -> PC (rate limited to ~ 100 Mbps)

ZCU102

PC



Total Bandwidth (TCP Payload Bandwidth)

	PC RX: ~ 9495 Mbps	PS RX: 135 Mbps	PL RX: - Mbps
PC TX: 135 Mbps	-	135 Mbps (1x TCP)	-
PS TX: 95 Mbps	95 Mbps 1x TCP	-	-
PL TX: 9400 Mbps	4x 2350 Mbps 4x TCP	-	-

Switch

	↕ bridge	Bridge	1500	1592	98.2 kbps	8.3 kbps
PC	↕ buche-10G0014-1	Ethernet	8000	8000	9.9 Gbps	360.4 Mbps

Conclusion

The described solution shows that it is possible to

- drive a 10 GbE network using an embedded system
- with integrate precision timing and to
- enable interactive use of the embedded operating system

by using a single, high-performance network interface.

Contact:

Ulrich Langenbach

ulrich.langenbach@missinglinkelectronics.com



Missing Link Electronics

Industriestraße 10
89231 Neu-Ulm

www.missinglinkelectronics.com

References

- [1] Corundum mqnrc, A. Forencich, UC San Diego, <https://github.com/corundum/corundum>
- [2] Network Protocol Acceleration Platform - NPAP datasheet, Missing Link Electronics, https://www.missinglinkelectronics.com/wp-content/uploads/2024/03/MLE_NPAP_Datasheet_for_TCP_IP_Stack-20240311.pdf
- [3] Corundum, From a NIC to a Platform for In-Network Compute, FOSDEM 2022, U. Langenbach, A. Forncich, <https://archive.fosdem.org/2022/schedule/event/corundum/>
- [4] Zynq RFSoc DFE Backgrounder, AMD, <https://www.xilinx.com/content/dam/xilinx/publications/backgrounders/zynq-rfsoc-dfe-backgrounder.pdf>
- [5] Zynq UltraScale+ MPSoC, AMD, <https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html>
- [6] Zynq UltraScale+ RFSoc, AMD, <https://www.xilinx.com/products/silicon-devices/soc/rfsoc.html>
- [7] AD-GMSL2ETH-SL, Analog Devices, <https://wiki.analog.com/resources/eval/user-guides/ad-gmsl2eth-sl>
- [8] AMD ZCU102, AMD, <https://www.xilinx.com/products/boards-and-kits/ek-u1-zcu102-g.html>

References

- [9] Light Rabbit, 13th White Rabbit Workshop 2024, Missing Link Electronics,
https://ohwr.org/project/white-rabbit/wikis/uploads/11edb649569e256ae62815615285481a/12_Light_Rabbit_1.pdf
- [10] “VERANO Projekt-Übersicht” Infineon
<https://www.infineon.com/cms/en/product/promopages/VERANO-Projekt/#Projekt-Uebersicht>