

Summer 2022, Open Air Interface Workshop  
Orange Labs, Orange Gardens, Paris

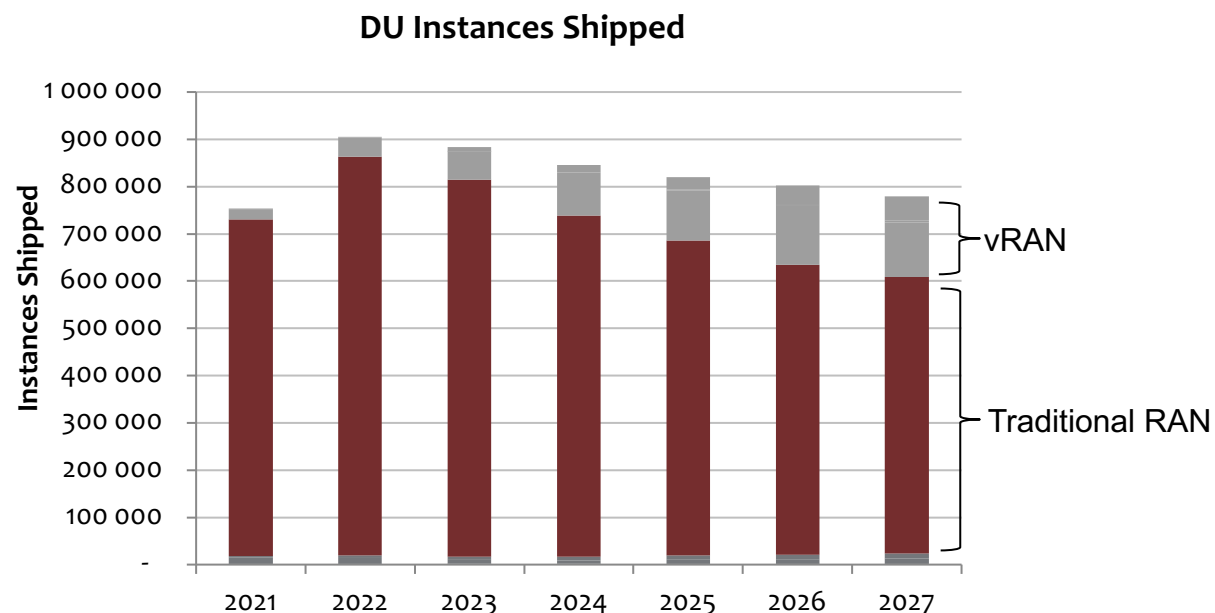
# Hardware Acceleration for Disaggregated RAN

Raghu M. Rao Ph.D.  
Director, Data Center and Communications Group

**AMD**   
together we advance\_

# vRAN Market Trends and Opportunity


## vRAN v/s Traditional RAN



Source: "Virtual RAN", April 2022, Mobile Experts

- ▲ vRAN Market is small but growing
  - ▲ Standard server + Accelerator cards
  - ▲ Larger opportunity in 6G
  
- ▲ Traditional RAN is large but shrinking
  - ▲ Vendor/platform lock-in not desirable
  
- ▲ O-RAN presents an opportunity to have multi-vendor solutions for RU, DU, CU
  
- ▲ Increasing interest in open-source RAN, Core and Orchestration solutions

# Drivers For Open RAN/vRAN Acceleration



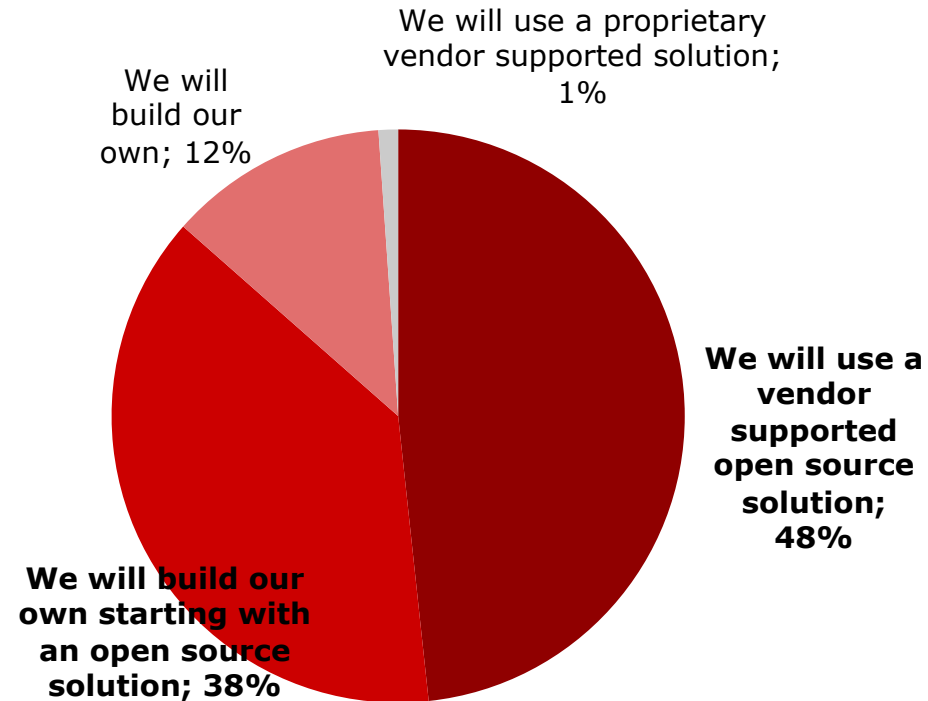
Source: Heavy Reading Webinar, May 2021



- ▲ Heavy Reading conducted a survey of operators in early 2021
- ▲ Topic- Accelerating Open RAN Platforms
- ▲ The important business drivers for Open RAN-
  - ▲ Faster roadmap and ability to bring new features
  - ▲ Reduce vendor lock-in
  - ▲ Reliability improvement and cost savings with virtualization
- ▲ The **three** key drivers for acceleration-
  - ▲ Price,
  - ▲ Power
  - ▲ Performance
    - ▲ L1 has many high compute modules that are expensive to implement on GPPs.
    - ▲ Requires many more CPU cores and is expensive and not power efficient.

<https://www.windriver.com/resource/charting-the-use-of-hardware-acceleration-and-integration-in-open-ran-deployments>

# Software Infrastructure



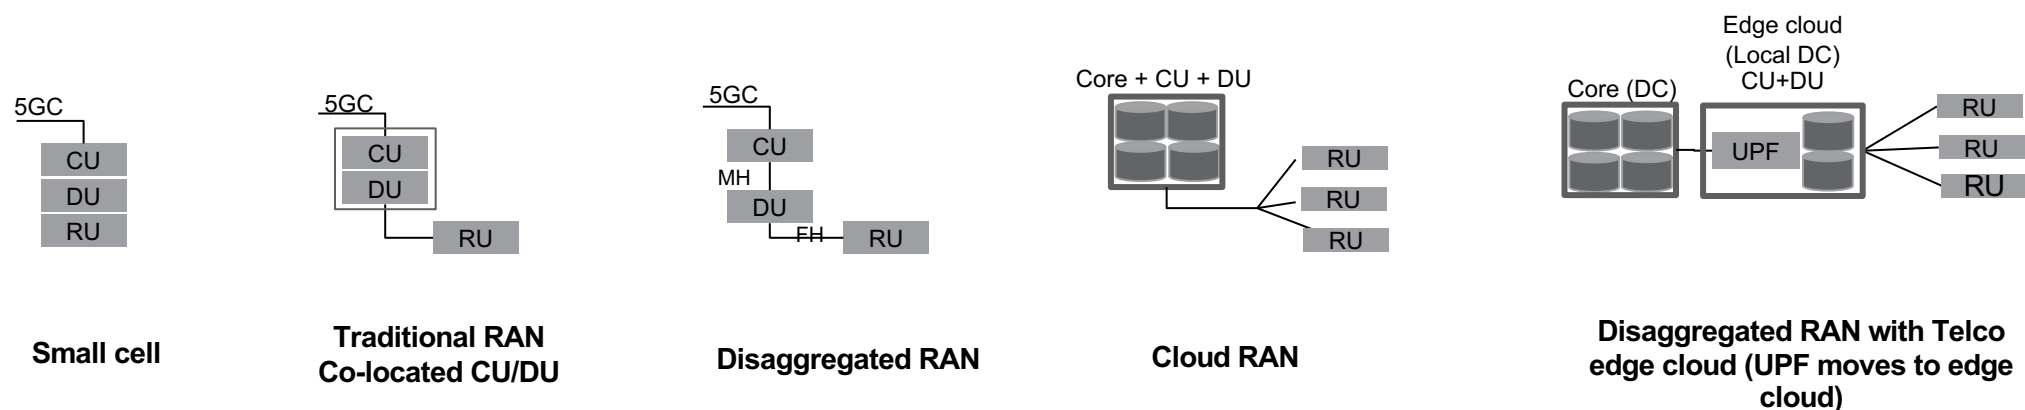
- ▲ Key drivers for moving to software and virtualizing RAN
  - ▲ Improved reliability
  - ▲ Ease of deployment
  - ▲ Maintenance
  
- ▲ This is a trend continuing from the virtualization of the Core and network automation
  
- ▲ A large majority (~85%) preferred the use of open-source software, either vendor supported or internal development starting with an open-source solution

Source: Heavy Reading Webinar, May 2021




[https://www.lightreading.com/lq\\_redirect.asp?pid=1&lgid\\_docid=769527](https://www.lightreading.com/lq_redirect.asp?pid=1&lgid_docid=769527)

# The Disaggregated RAN




- ▲ The disaggregated RAN distributes the RAN functionality over the RU, DU and CU (optionally)
- ▲ Some functionality from the core (i.e., UPF) could move to the CU
- ▲ Varied customer base: Telco, Enterprises with Private 5G, etc.
- ▲ Numerous deployment scenarios lead to many product skews
- ▲ One size does not fit all
  - ▲ Optimizing for price, performance and power is difficult
  - ▲ Leads to significant challenges to equipment provider eco-system

# Hardware Acceleration Trends




## Look-aside FEC Acceleration

- ▶ Enables software centric RAN and virtualization
- ▶ High complexity modules impact performance
- ▶ Look-aside acceleration improves performance
- ▶ But limited by PCIE bandwidth



## Inline L1 Acceleration

- ▶ Offloads complete L1 functionality
- ▶ Removes PCIE BW issue
- ▶ Performance limited to thermals and power constraint of PCIE slot
- ▶ Mismatched server and card can lead to inefficiencies

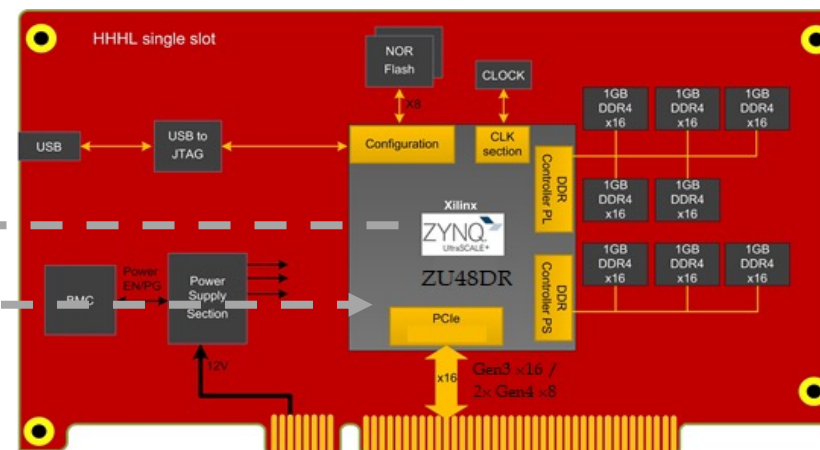
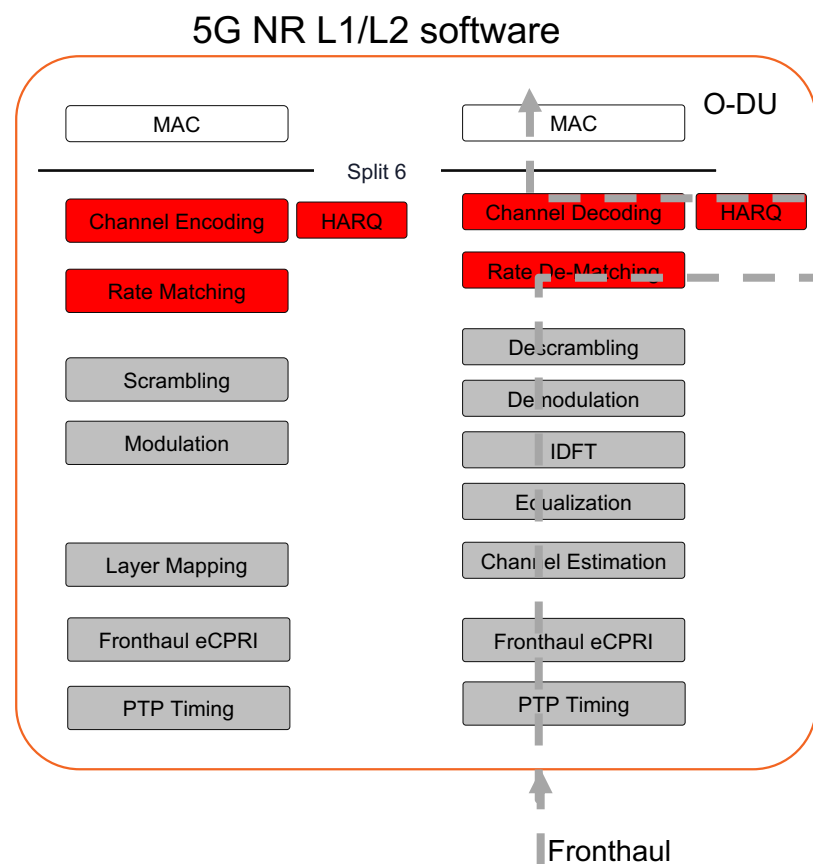


- ▶ One size does not fit all!
- ▶ Solution needs to-
  - Scalable
  - Support various deployment scenarios
  - Cost effective across varying deployment and customer scenarios
  - Matched CPU and acceleration capability
  - Not constrained by PCIE BW
  - Not constrained by thermal requirements
  - Energy efficient

- ▶ Numerous interesting proposals on the table

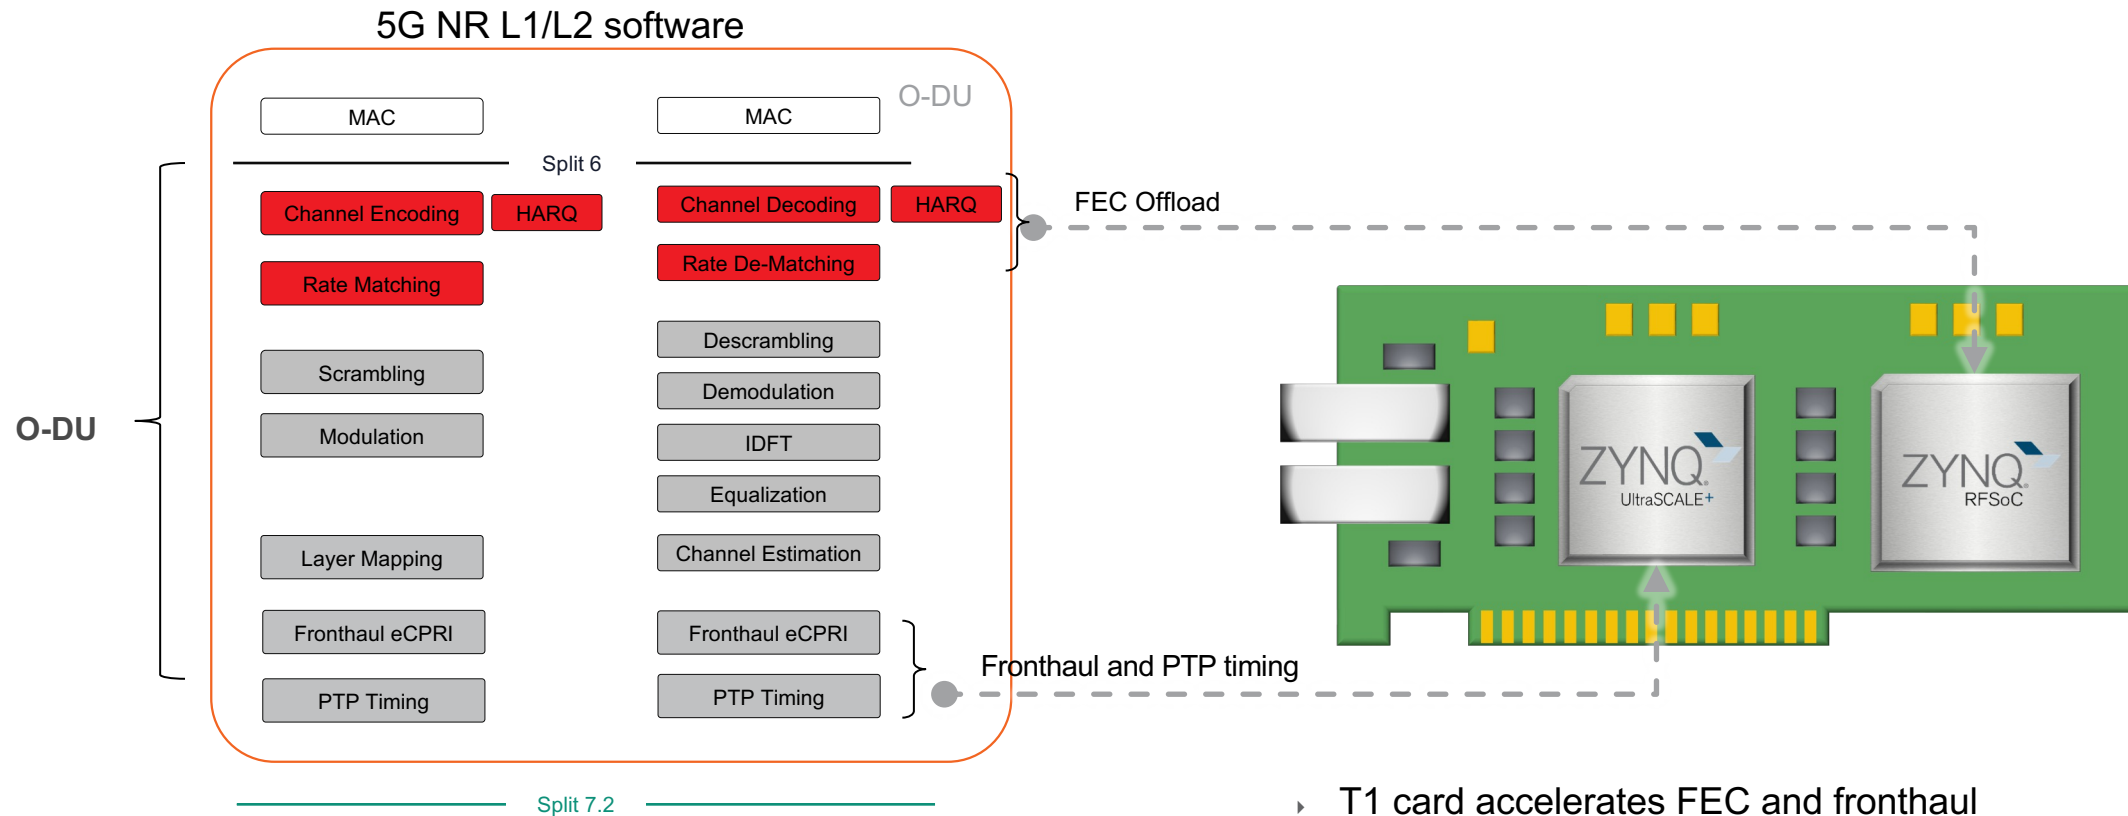
# T2 Telco Accelerator Card - Look-aside Acceleration

- T2 is a high bandwidth Look-aside FEC Accelerator
  - Gen4 x8 PCIE – 200Gbps
  - Capable of offloading 4-sector 64TRx (16DL/8UL) massive MIMO
  - Power dissipation - ~40W-45W



- › Look-aside acceleration reduces the compute requirement of L1 by ~75% for smaller configurations (4T4R/8T8R)
- › Results in significant cost and power savings
- › Massive MIMO will require further offloading of complex functionality related to beamforming
- › For many deployment scenarios especially Private 5G and Enterprise offloading FEC is good enough

# T1 Telco Accelerator Card – Hybrid Inline and FEC Acceleration



- ▶ T1 card accelerates FEC and fronthaul
- ▶ Has IEEE 1588 timing circuitry onboard
- ▶ Two devices with 100Gbps chip-to-chip interconnect enables inline acceleration (4T4R)
- ▶ Good for the Enterprise 5G market

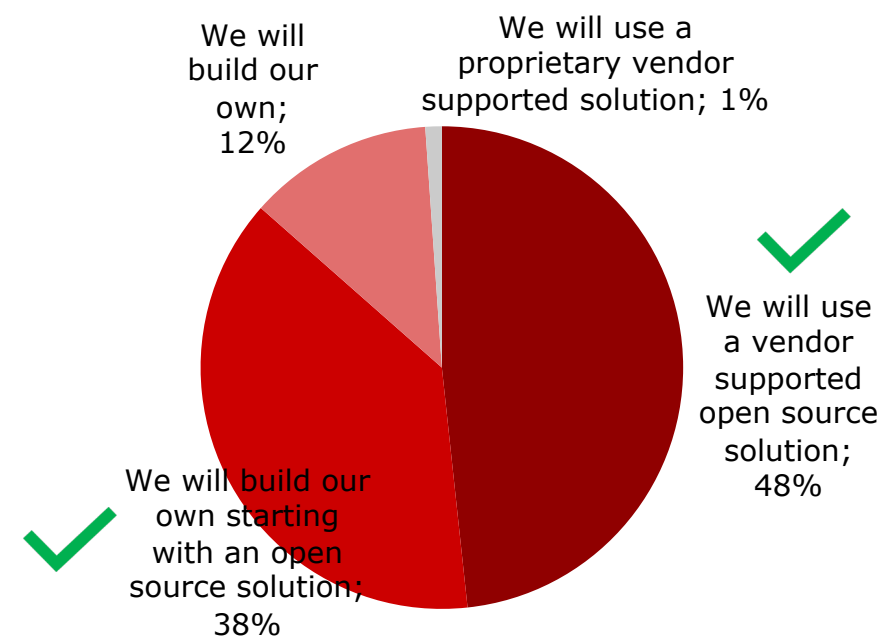
# The 5G NR Software Stack

- ▲ Much attention given to 4G/5G Layer 1
- ▲ With inline acceleration this problem may be solved
  
- ▲ But Layer 2, Layer 3 (RRC) and 5G Core are equally important
- ▲ Portions of Layer 2 (PDCP) especially in the disaggregated CU require acceleration of air-crypto and security functions
- ▲ Support for Massive MIMO
- ▲ Many protocols between disaggregated RAN and Core elements required
- ▲ Radio Intelligent Controller plays an important role
- ▲ In general O-RAN architecture and protocols needs to be supported

# Where Can OAI Contribute?

- ▲ Open-source is gaining a lot of attention
  - ▲ Success of open-source projects such as Kubernetes is bringing attention to the 5G NR stack
- ▲ Deployable 5G NR-
  - ▲ (Layer 1 may be solved by inline acceleration)
  - ▲ **High quality Layer 2/3 required**
    - ▲ **Focus on testing realistic Telco traffic models**
    - ▲ **Complete all the features**
- ▲ Fully embrace the O-RAN architecture

## Recap



---

# Disclaimer and Attribution

---

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 information. 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. GD-18

© Copyright 2021 Advanced Micro Devices, Inc. All rights reserved. Xilinx, the Xilinx logo, AMD, the AMD Arrow logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Thank you

**AMD** 



Summer 2022  
OpenAirInterface Workshop Paris