The ALICE Common Read-Out Receiver Card (C-RORC)

Heiko Engel
(IRI, Frankfurt University)

ALICE, ATLAS, CMS & LHCb Second Joint Workshop on DAQ@LHC
12.04.2016

ALICE Run 1 Read-Out Architecture

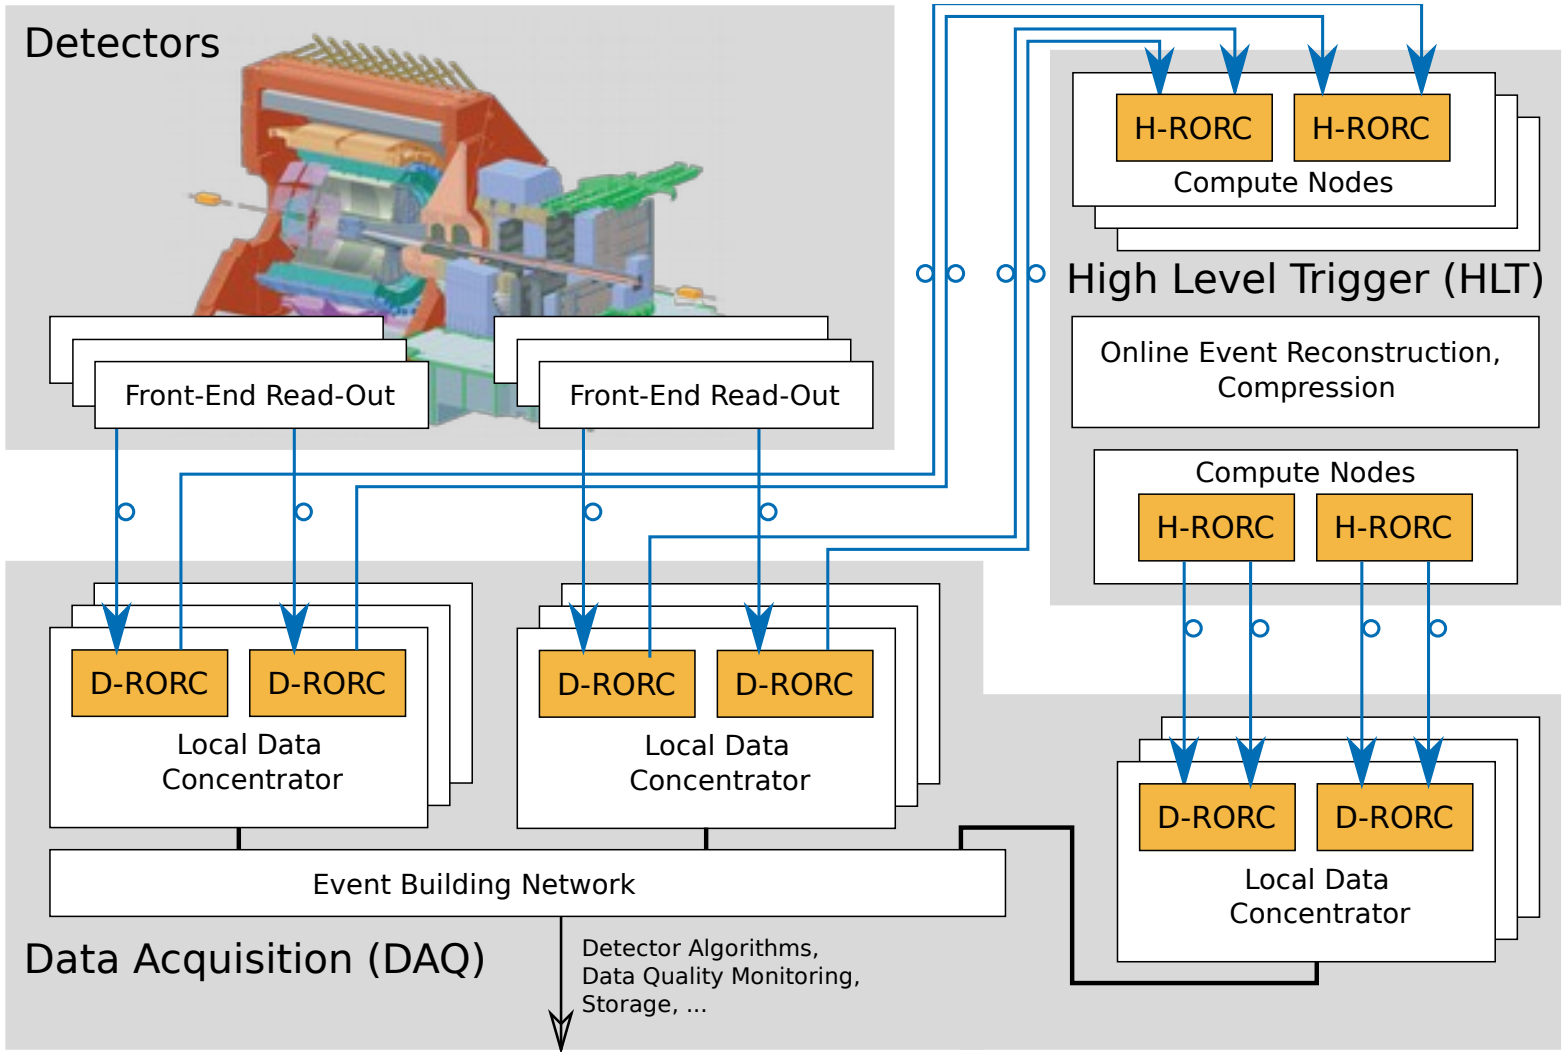

Run1
- ~500 optical links from detectors

Data Acquisition:
- ~2500 CPU Cores
- ~400 D-RORCs

High-Level Trigger
- ~2500 CPU cores
- 240 H-RORCs
- 64 GPUs

Orange boxes: custom FPGA boards:
D-RORC: DAQ Read-Out Receiver Card
H-RORC: HLT Read-Out Receiver Card

ALICE Run1 Read-Out Receiver Cards

- ALICE Data Acquisition: D-RORCs

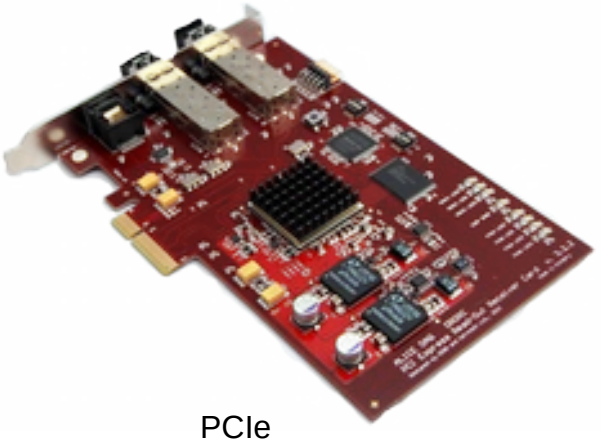

PCIe

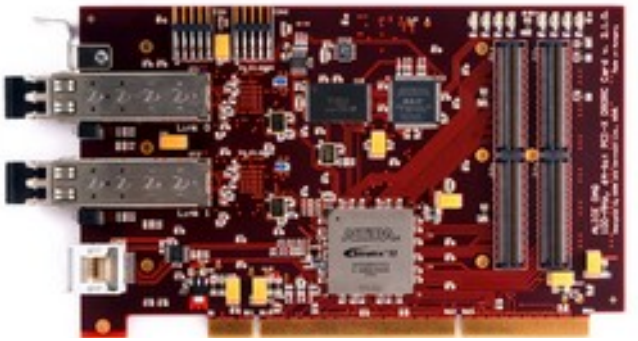

Images: cerntech.hu

PCI-X

- ALICE High-Level Trigger: H-RORCs

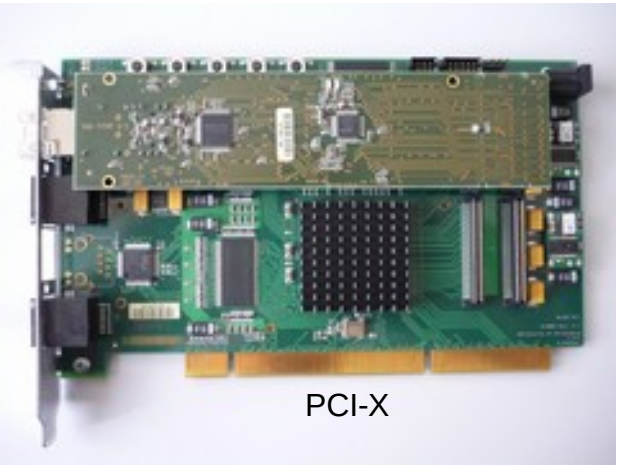

PCI-X

Read-Out Receiver Card for Run 2

- Main requirements
  - PCI-Express
  - 12 optical links
  - support for higher link speeds
  - failsafe FPGA configuration
  - DAQ: LVDS
  - HLT: onboard memory & data preprocessing capabilities
- Timeline
  - project kickoff mid 2010
  - schematics done mid 2011
  - First prototypes end 2012

C-RORC Hardware Overview

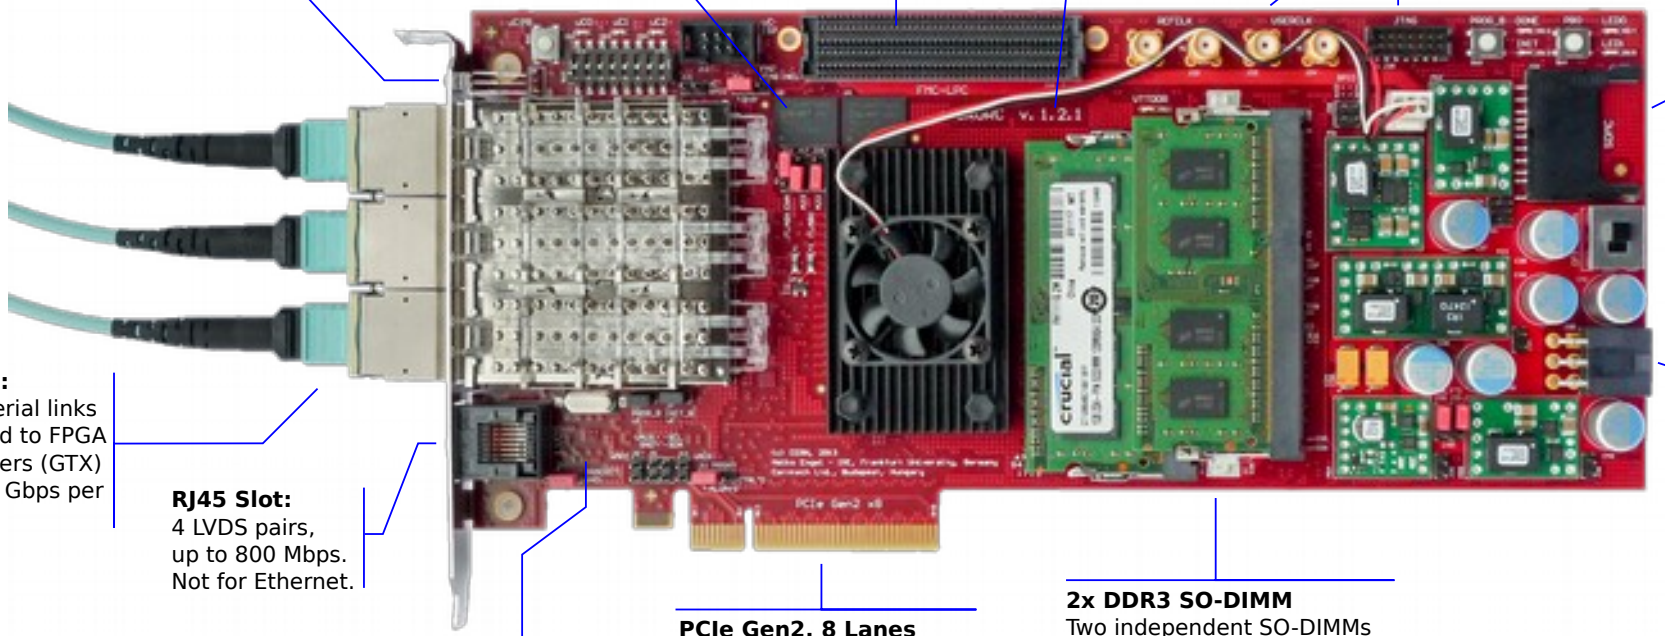

Configuration Flash
2x 128Mbit Xilinx PlatformFlash XL

FMC LPC
VITA57.1 FMC LPC Connector

GTX RefClk
Configurable GTX Reference Clock

System Clock
200 MHz

SMA Connectors
for additional GTX and System Clocks

JTAG-Connection
for FPGA Programming Cables

Bi-Color **LED** in IO-Bracket

SDCard Socket

Power Connection
via 6-Pin PCIe GPU Power Connector

3x QSFP:
12 fast serial links connected to FPGA transceivers (GTX)
Up to 6.6 Gbps per channel

RJ45 Slot:
4 LVDS pairs, up to 800 Mbps.
Not for Ethernet.

Microcontroller for Configuration monitoring and connection to Host via PCIe SMBus Lines.

PCIe Gen2, 8 Lanes
8x 5.0 Gbps, connected to Xilinx PCIe Hard Block

2x DDR3 SO-DIMM
Two independent SO-DIMMs
Up to 1066 Mbps with Xilinx MIG DDR3 Softcore

C-RORC Hardware Overview

- PCIe throughput close to physical limit

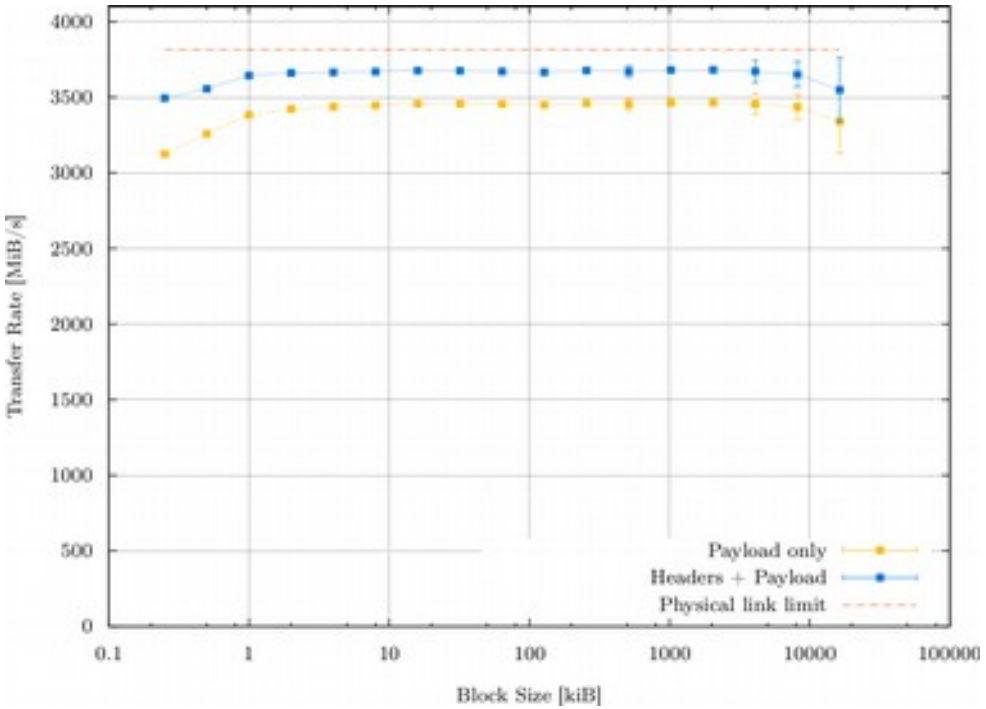

- DDR3 at 1066 / 800 / 606 Mbps

[MB/s]	DDR3-1066	DDR3-800	DDR3-606
Seq. read	8140	5940	4620
Seq. write	7320	5420	3300

- Configuration control via PCIe sideband

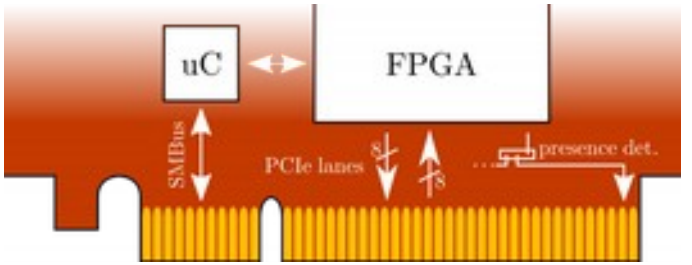

ATLAS TDAQ in Run2

ATLAS@LHC - CERN
General purpose detector
Wide physics search goals
46m long, 22m high, 7000 tons
140M channels

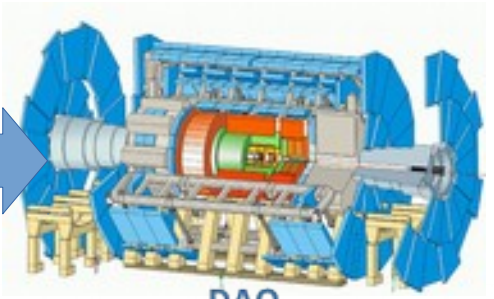

Run1 ReadOut System (ROS) Setup:
~1600 optical inputs
~150 PCs
~300 1GbE ports

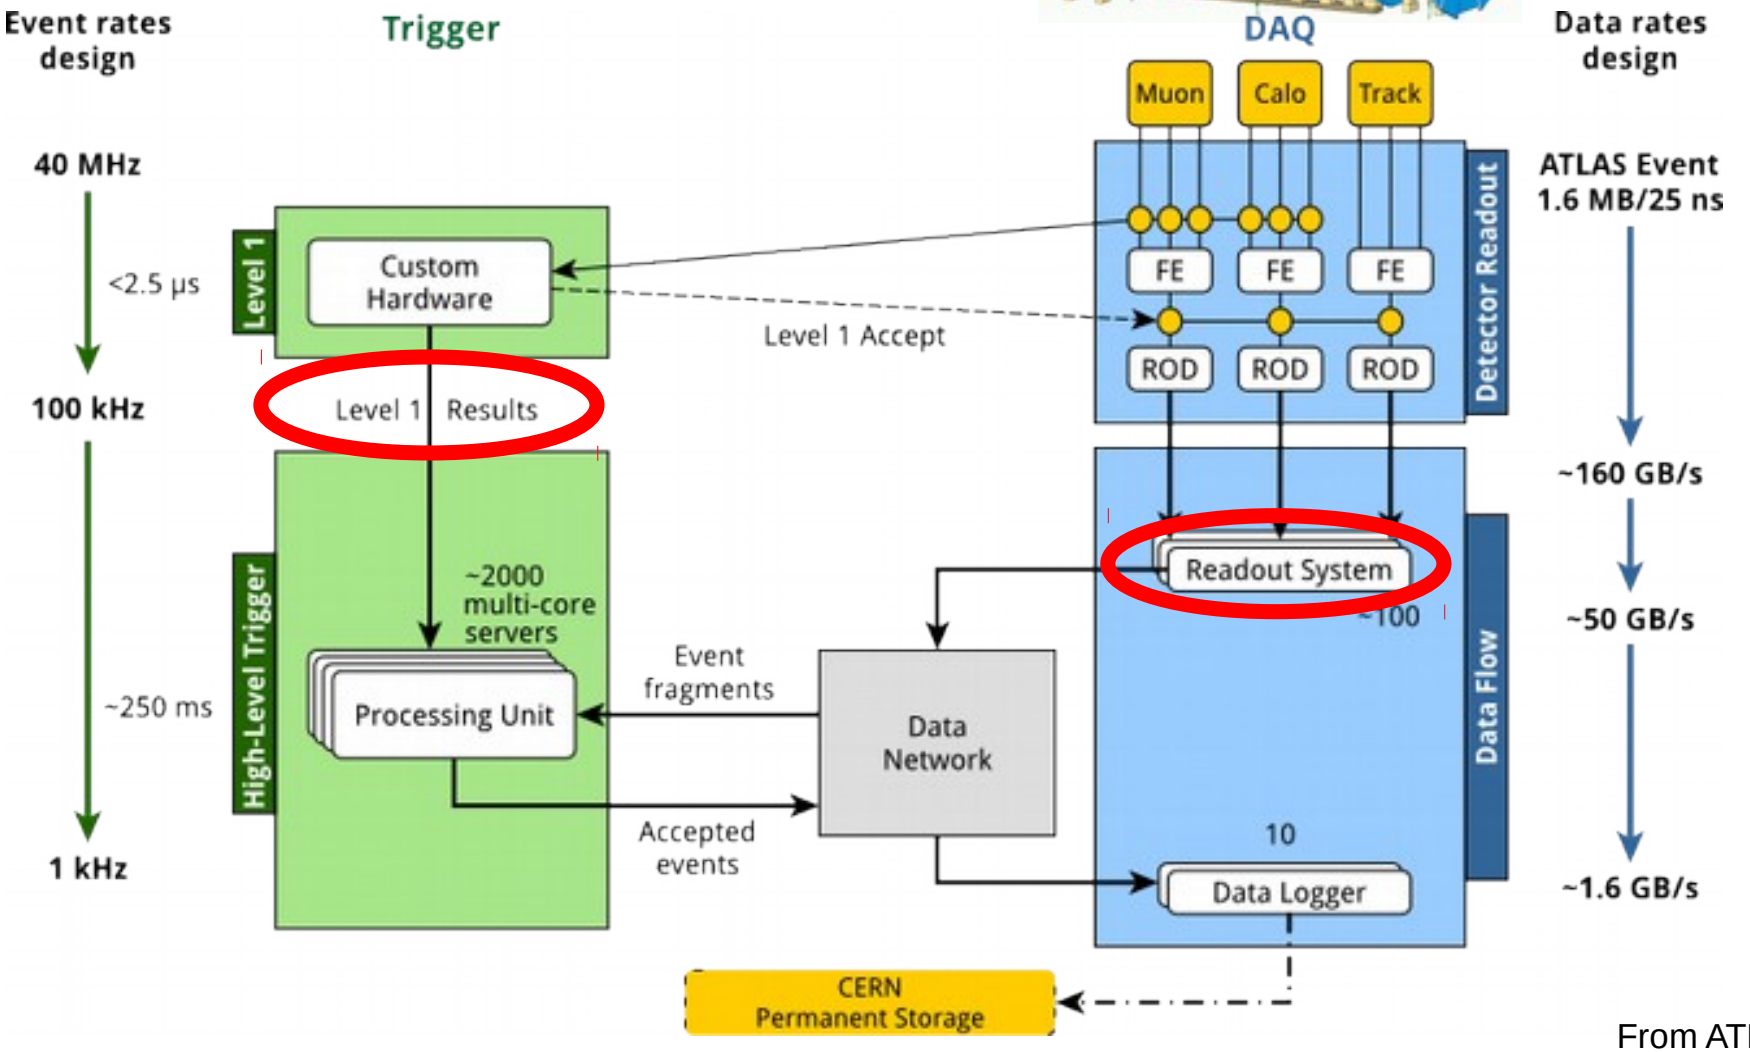

From ATLAS ROS Team

C-RORC Hardware Production

- Common production
  - Approx. 400 boards
- Common hardware test suite
  - Firmware
  - Software
  - Peripherals
- Installation into production systems Aug. 2014

- Timeline
  - Market Survey / Call for Tender Q1-2/2013
  - Contract mid 2013
  - Pre-series production early 2014
  - Series production done mid 2014

C-RORC Applications

Common hardware platform - independent firmware projects

ALICE Data Acquisition
ALICE High-Level Trigger
ATLAS TDAQ Read-Out System
ATLAS TDAQ RoI Builder
ALICE Run 3 Read-Out Prototype

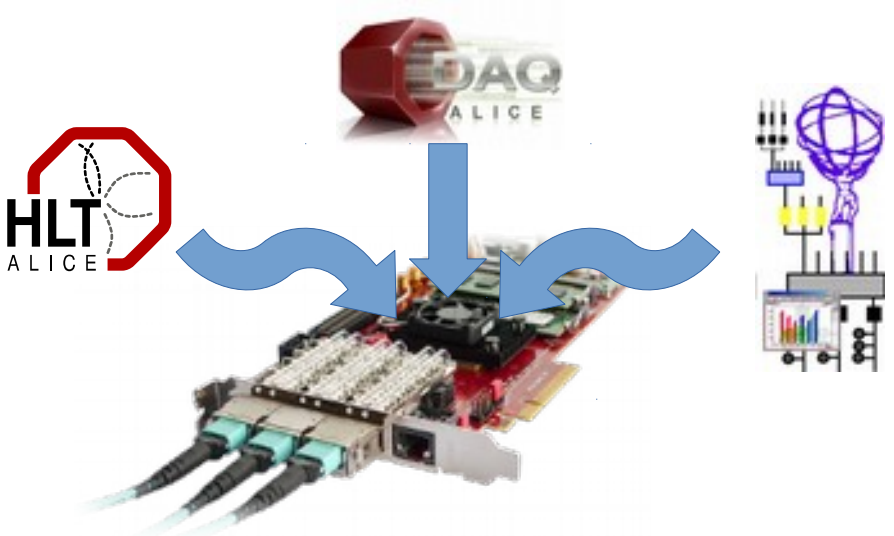

ALICE Data Acquisition Firmware

- Data Acquisition and Detector Configuration/Control protocols
  - DATA TAKING
    - RX channel used to download events from the detector electronics to the DAQ farm
    - TX channel used for flow control to avoid data overflow.
  - DETECTOR configuration:
    - TX channel used to send configuration data to the front end electronics
    - RX channel used to receive acknowledgement
- Replace TPC, TRD and HLT-to-DAQ D-RORCs
- Optical Links: Detector Data Link (DDL) Protocol,
  - Gen1 at 2.125 Gbps
  - Gen2 at: 3.125 / 4.0 / 5.3125 Gbps
  - Linkrate configurable at runtime
- DMA: Device to Host
  - Physmem buffer handling
  - PLDA DMA engine, 6 data channels
    for data taking and for detector configuration

ALICE DAQ Run2 Setup
130 Nodes / ~2900 cores
60 C-RORCs
235 D-RORCs
10G Ethernet Interconnect

From Filippo Costa, ALICE DAQ

ALICE High-Level Trigger: Firmware

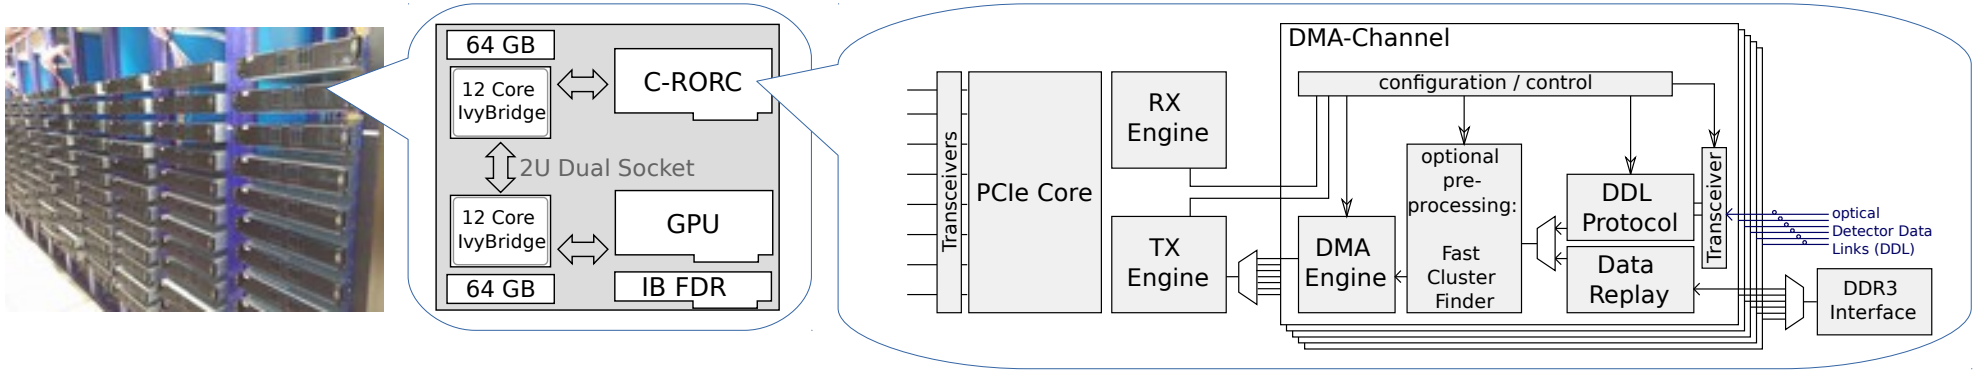

- Replacing all previous HLT boards
- Up to 12 links / board:
  - Data input from DAQ: DMA-to-Host, DDL at 2.125 / 3.125 / 4 Gbps
  - Data output to DAQ: DMA-to-Device, DDL at 5.3125 Gbps
- FPGA-based data pre-processing: cluster finding on TPC raw data
- Custom DMA engine for scatter-gather DMA
- DDR3 for DataReplay

ALICE HLT Run2 Setup
180 Nodes / ~4300 cores
180 GPUs
74 C-RORCs
Infiniband FDRx4 (56 Gbps)

ATLAS ROS/RoIB Architecture

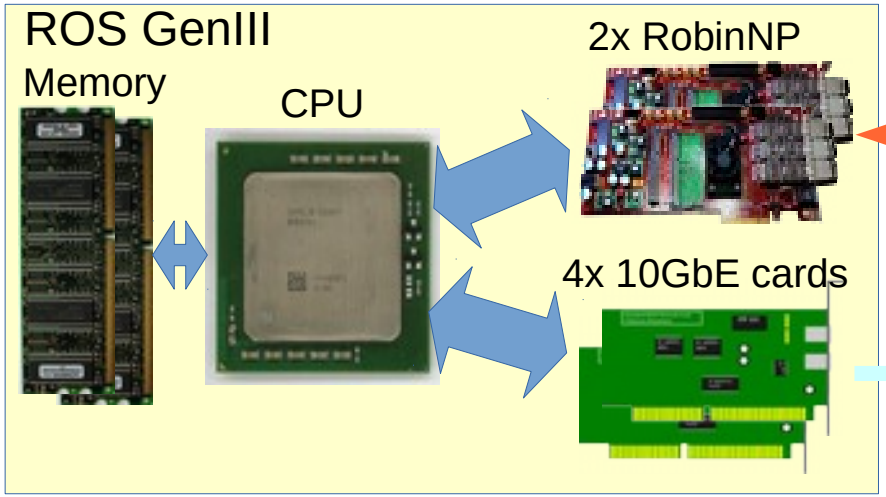

24x optical link

4x 10GbE

At 50% readout
→up to **~16 Gbps**
- RobinNPs to memory
- memory to network cards

- ROS / ROI-Builder
  - Same firmware and host PC
  - ROS: 12 optical links / board
  - RoIB: 10 optical links / board
  - separate control & fragment building software for ROS and RoIB

ATLAS ROS Run2 Setup:
- 98 ROS PCs
- 1860 ROLs
- 178 C-RORCs
- 392 10GbE links

From ATLAS ROS Team

ATLAS RobinNP Firmware and Software

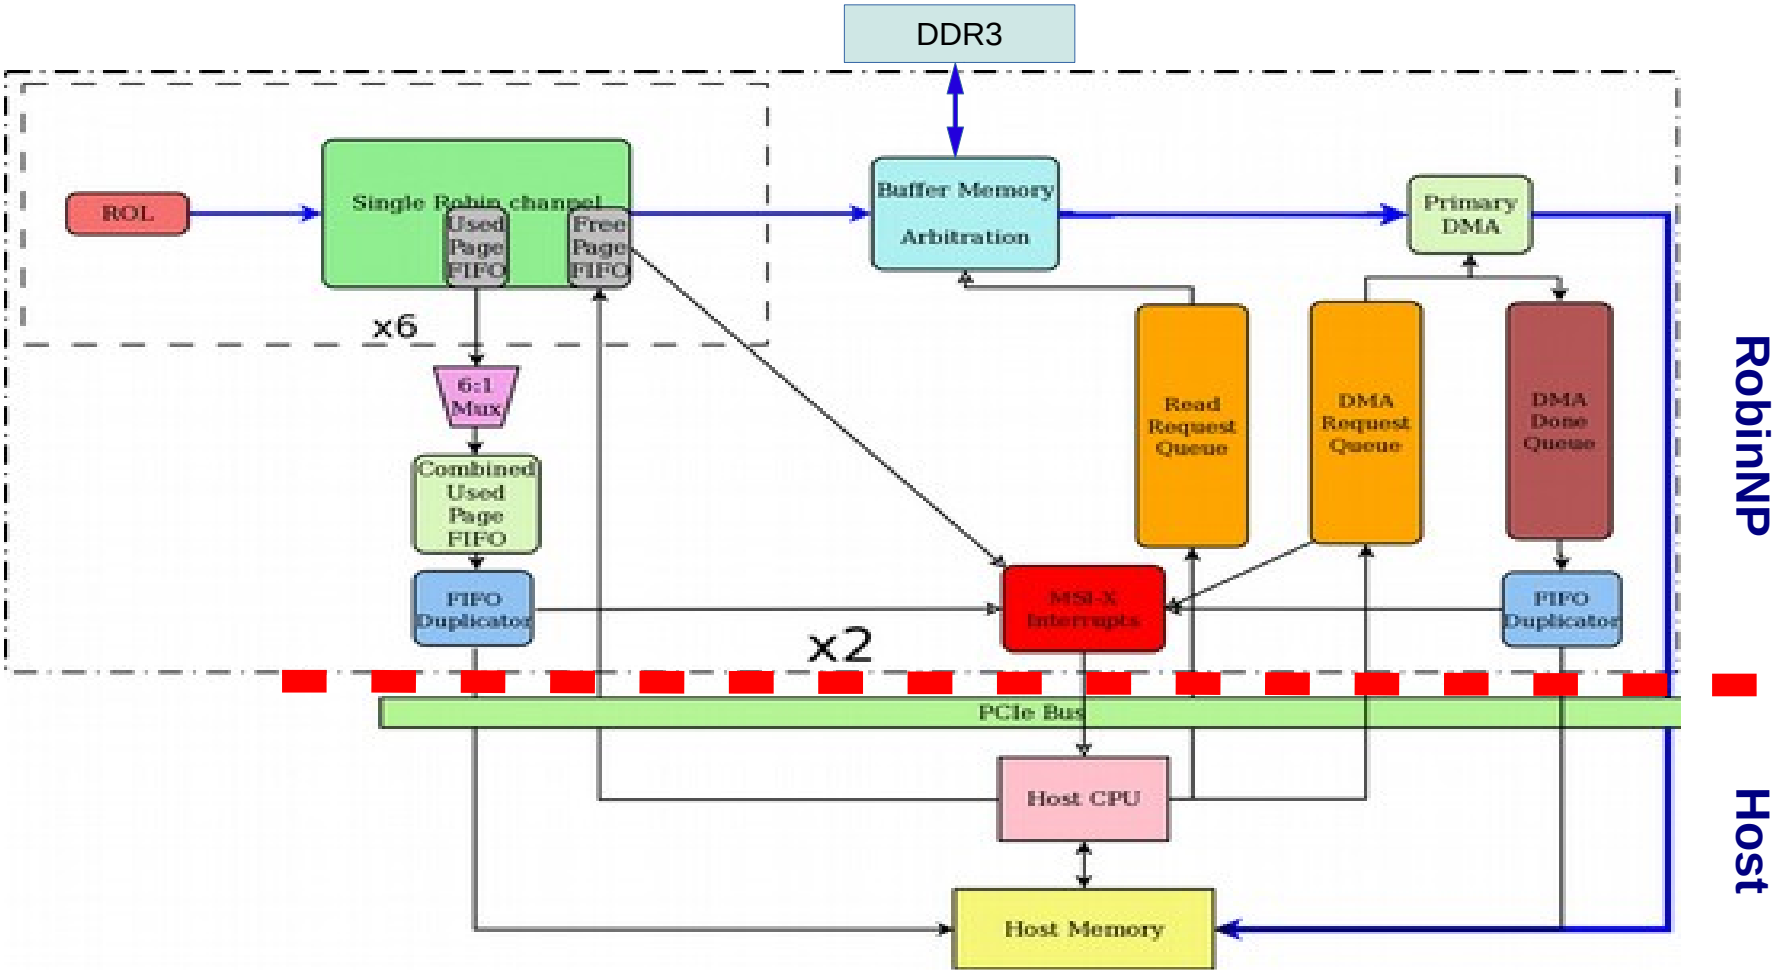

- C-RORC as RobinNP (No Processor)
  - based upon Robin firmware, but offload tasks to the host CPU
  - PLDA DMA engine
- Key innovations at the boundary between the RobinNP and the host

From ATLAS ROS Team

C-RORC as Prototype for Run 3

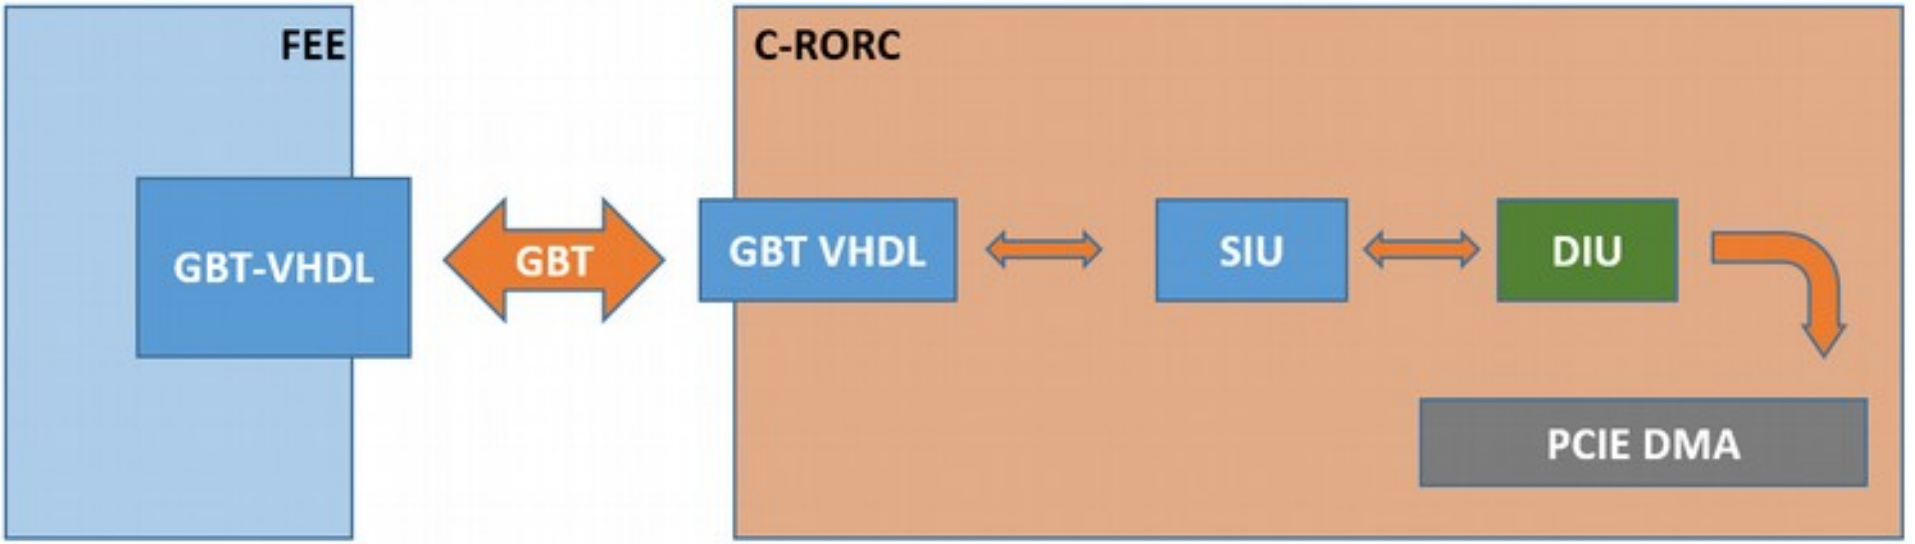

The new part pushes data to the SIU as a FEE interface

The current readout chain has been preserved.

- Run 3 hardware not available yet
- C-RORC can readout GBT link speed
- Prototype system with GBT connected to existing read-out chain

From Filippo Costa, ALICE DAQ

Summary

- Common read-out hardware for ALICE & ATLAS
- Increased link density, up-to-date interfaces
- Installed in production systems since mid. 2014
- Four applications in Run 2
  - Separate firmware projects, but with shared code
- ALICE prototyping platform for Run 3