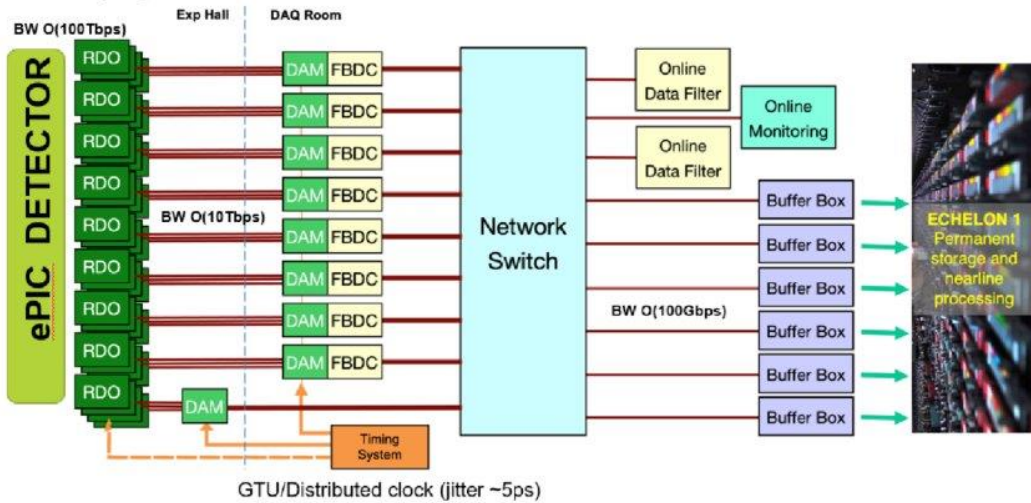



# DAM Status

**Ottorino Frezza**  
**on behalf of INFN Roma1**  
**ePIC group**

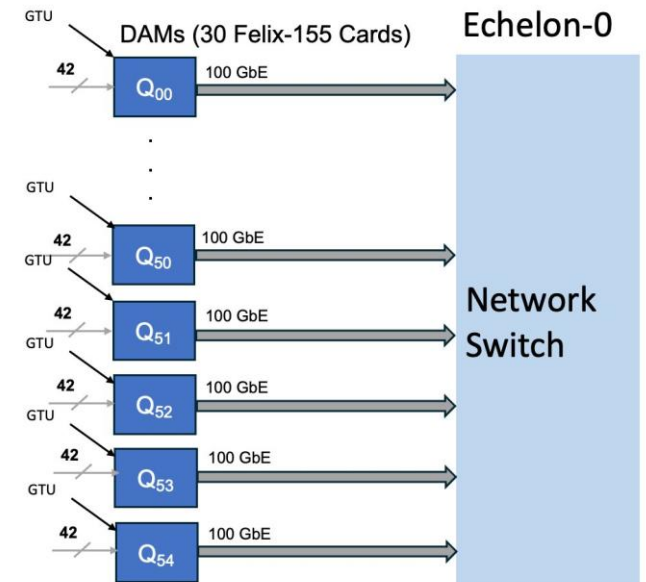
# RDO/DAM role in ePIC and dRICH PDU



## Front-End DAQ: RDO



## Back-End DAQ: DAM



## Data Aggregation Module (DAM)

- Implemented on Felix-155 board (INFN Rome)
- Aggregates data from 42 RDOs and transmit them to the ePIC buffering system (Echelon 0) through 100 GbE links.
- Provides computational capability for data handling and reduction

# How to transmit data from RDO to DAM?

**DAM Firmware:** Reuse as much as possible of the “official” FELIX firmware release.

<https://gitlab.cern.ch/atlas-tdaq-felix/firmware.git>

## FELIX Firmware

- FELIX supports different link protocols for data transfer
- The firmware flavour is defined at built time

Protocol	FELIX
GBT	TX: 4.8G, RX: 4.8G
FULL	TX: 4.8G, RX: 9.6G
IpGBT	TX: 2.58G, RX: 10.24G

### GBT :

- Protocol which provides radiation-tolerant transport of data
- Aggregate multiple lower bandwidth logic links (e-links) in one high bandwidth data link

### FULL :

- Use the GBT protocol for the links towards the FE (downlink)
- To receive data from the FE (uplink) implements a light-weight protocol (with a bandwidth of 9.6 Gb/s per optical link - 7.68Gb/s with 8b10b encoding).

### IpGBT :

- Evolution of GBT
- Two FEC codes can be used (FEC5/FEC12)

### FEC5 Uplink

	5.12 Gb/s	10.24 Gb/s
Option	7	7
Frame (bits)	128	2 x 128
Header (bits)	2	2+2
Coded header	no	no
User field (bits)	116	232
Code (bits)	10	20
16-bit multiplicity	7.250	14.5
Remainder bits	4	8
User Bandwidth (GHz)	4.64	9.28
# eLinks groups (16 bit)	7	7
eLinks bandwidth (MHz)	160/320/640	320/640/1280
#eLinks	28/14/7	28/14/7
EC bandwidth (MHz)	80	80/160
IC bandwidth (MHz)	80	80/160
Unassigned bits	0	4
Corrected (bits)	5	2 x 5
Efficiency	91%	91%

# How to transmit data from RDO to DAM?

Protocol Bandwidth (Gb/s)

Protocol	FELIX
GBT	TX: 4.8G, RX: 4.8G
FULL	TX: 4.8G, RX: 9.6G
IpGBT	TX: 2.56G, RX: 10.24G

## Design Assumption:

- RDO/DAM BW: at most 6.2Gb/s
- RDO: Use of VTRx+ Optical Link Module for Data Transmission, w/o IpGBT ASIC

## FULL:

- Less FPGA resources hungry with respect to IpGBT
- **Does not provide radiation-tolerant data transport**

## IpGBT FEC5:

- Provides **radiation-tolerant** transport of data
- FEC5 Bandwidth: 10.24Gb/s with 91% efficiency
- Scaling up to 48 IpGBT in FLX-155 could be a problem **without modifying the standard FELIX FW**
- Division of physical-link in virtual-links (e-links) is not necessary but **standard FELIX FW doesn't support single e-link configuration**

FELIX **Estimated** resource utilization

		KU115	VM1802	VP1552
GBT 24 channel	LUT	80.65%	69.60%	35.71%
	FF	77.03%	50.94%	26.13%
	BRAM	70.00%	89.45%	34.04%
	URAM		62.20%	22.14%
FULL 24 channel	LUT	52.59%	44.35%	22.75%
	FF	38.40%	33.21%	17.03%
	BRAM	40.46%	20.99%	7.99%
	URAM		62.20%	22.14%
LPGBT 24 channel	LUT	112.51%	82.94%	42.55%
	FF	52.39%	38.62%	19.81%
	BRAM	68.94%	79.52%	30.26%
	URAM		62.20%	22.14%

[https://atlas-project-felix.web.cern.ch/atlas-project-felix/user/docs/FELIX\\_Phase2\\_firmware\\_specs.pdf](https://atlas-project-felix.web.cern.ch/atlas-project-felix/user/docs/FELIX_Phase2_firmware_specs.pdf)


**FW customization required**

## AURORA:

- Xilinx low-latency data transmission protocol (up to 25 Gbps)

# Link protocol

**DISCLAIMER NOTICE:** all this is VERY



RDO "prepares" data reduction  
DAM "does" data reduction

## AWORD

50	49 48	47 46	45 39	38 30	29 27	26 24	23 22	21 9	8 0
K CODE FLAG	FEB ID	TDC ID (trailing)	Coarse (trailing)	Fine (trailing)	Col. ID	Pixel ID	TDC ID (leading)	Coarse (leading)	Fine (leading)

	DF3	DF2	DF1	DF0	DCS (free)	AWORD	AWORD	AWORD	AWORD
FULL	255	254	253	252	251 204	203 153	152 102	101 51	50 0
lpGBT	223	222	221	220	219 204	203 153	152 102	101 51	50 0

### Protocol over optical link:

- FULL (256 bits/clock)
- lpGBT (224 bits/clock with FEC5)
- Data Flag
- Space for DCS bus

link protocol operated at **39.4 MHz** CLK (2/5 of EIC clock, close to LHC clock)

**AURORA/FULL are our current baseline choices!**

- extracts BC from bits 21-11
- extracts geographic position from RDOID (DAM is link-aware) + FEBID (49-48) + AlcorCh (29-24): geo info → data pattern
- NN decision is reached, the DAM must add to the 51-bit word the 11 bits of the RDOID → 62-bit word
- 2 bits added → PCI 64-bit words → dRWORDS (dRICH words)
- dRWORDS out in timeframes to StreamingReadOut computing

# DAM prototype on FLX182

## Firmware

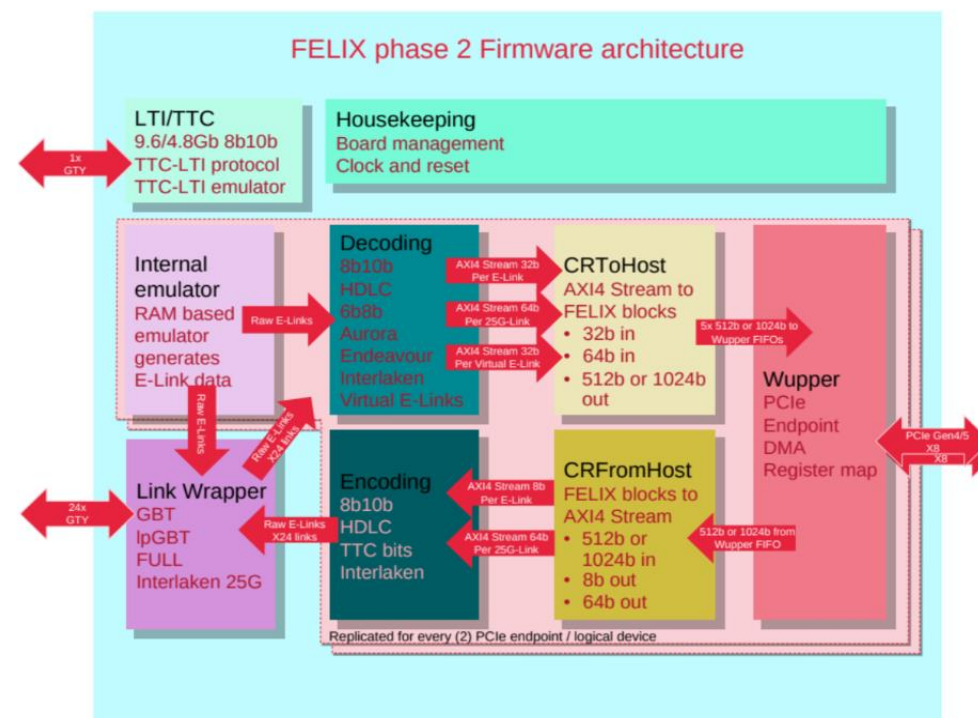
- Loading Felix official firmware onto the board failed due to a DRAM check error.
  - In current Felix FW versions the PS performs just ancillary functions (BIST,...), while all the design core functionalities are implemented in the PL.
  - We forced Versal Platform Management Controller (PMC) to ignore DRAM status check at start up.



**FW loaded and board detected!**

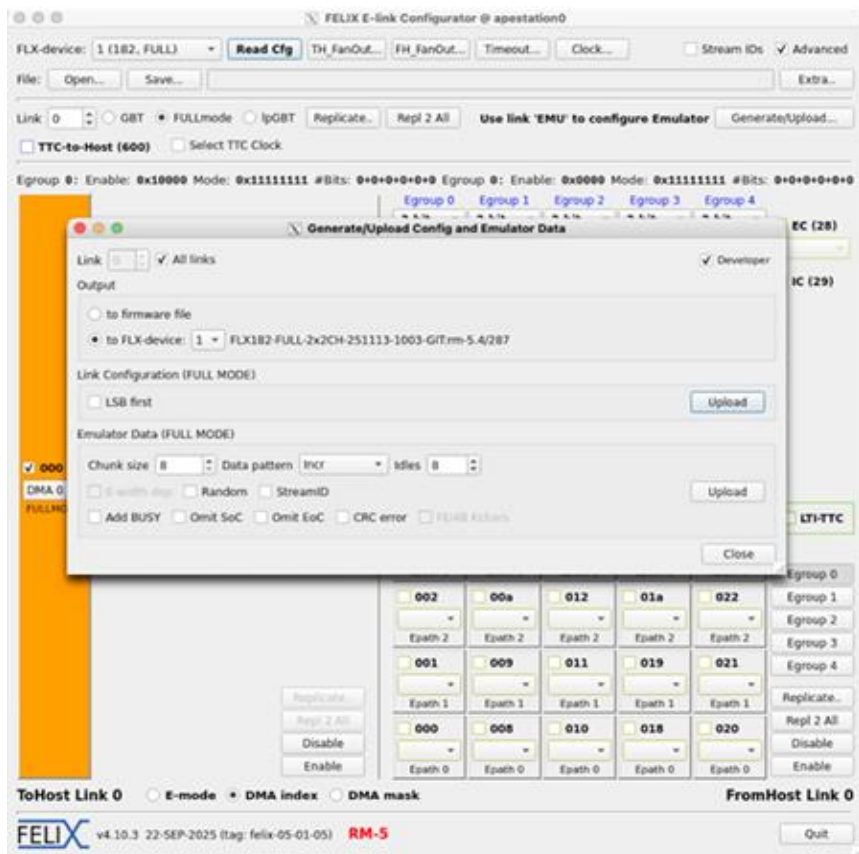
```
[locicero@apestation0 ~]$ lspci -vvvvv | grep CERN
34:00.0 Network controller: CERN/ECP/EDU Device 0428
Subsystem: CERN/ECP/EDU Device 0038
35:00.0 Network controller: CERN/ECP/EDU Device 0427
Subsystem: CERN/ECP/EDU Device 0038
```

```
lonardo@apestation0 x86_64-el9-gcc13-opt]$ bin/flx-init
2025-05-14 18:06:05 Opening card 0 (device 0)...
2025-05-14 18:06:05 Card type: FLX-182
2025-05-14 18:06:05 Firmware : FULL
2025-05-14 18:06:05 Clock : Local
```



# DAQ back-end FLX182 FULL flavour tests

- Exercised the full software/firmware stack using internal data generators to trigger DMA transactions to the host memory.



```
[lonardo@apestation0 ~]$ fdaq -d 0 -t 15 -e movesi
```

Consume FLX-device data while checking the data (blockheader and trailers), counts errors including chunk truncation, halts when the memory buffer is near overflowing. Also counts chunk CRC errors.

Opened FLX-device 0, FLX182-FULL-2x2CH-251113-1003-GIT:rm-5.4/287, trailer=32bit ChunkHeaders, buffer=1024MB, DMA=0

\*START(emulator)\* using DMA #0 polling

Secs	Recv[MB/s]	File[MB/s]	Total[(M)B]	Rec[(M)B]	Buf[%]	Wraps
1	241.4	0.0	241.4	0	0	0
2	241.3	0.0	482.6	0	0	0
...						
15	241.3	0.0	3618.8	0	0	3

\*STOP\*

-> Data checked: Blocks 3533027, Errors: header=0 trailer=0

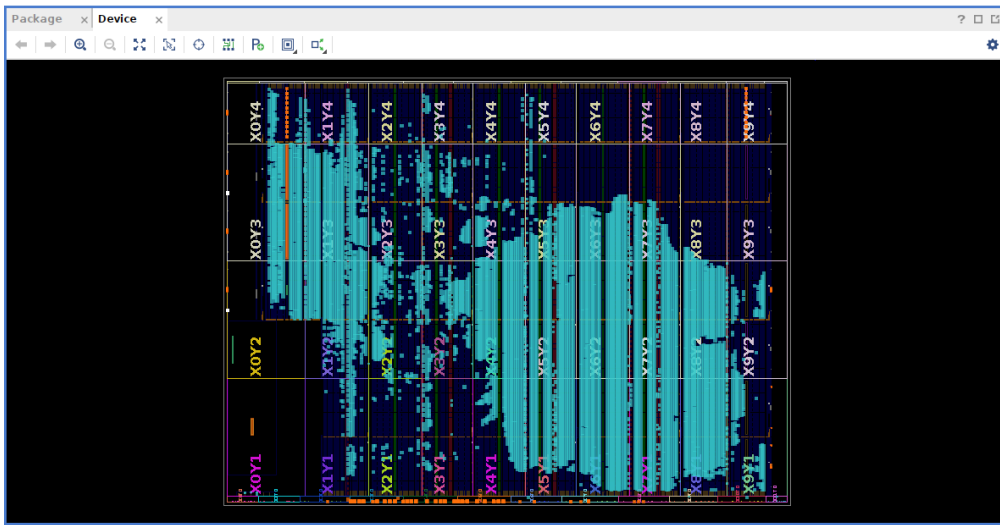
Exiting..

```
[lonardo@apestation0 ~]$ hexdump movesi-251121-122848-1.dat | head -16
```


```
00000000 0000 c0cf 0008 6000 00aa 0000 bb00 10aa 0000010 0008 6000 00aa 0100 bb00 10aa 0008 6000
00000200 00aa 0200 bb00 10aa 0008 6000 00aa 0300 0000030 bb00 10aa 0008 6000 00aa 0400 bb00 10aa
00000400 0008 6000 00aa 0500 bb00 10aa 0008 6000 0000050 00aa 0600 bb00 10aa 0008 6000 00aa 0700
00000600 bb00 10aa 0008 6000 00aa 0800 bb00 10aa 0000070 0008 6000 00aa 0900 bb00 10aa 0008 6000
00000800 00aa 0a00 bb00 10aa 0008 6000 00aa 0b00 0000090 bb00 10aa 0008 6000 00aa 0c00 bb00 10aa
00000a00 0008 6000 00aa 0d00 bb00 10aa 0008 6000 00000b0 00aa 0e00 bb00 10aa 0008 6000 00aa 0f00
00000c00 bb00 10aa 0008 6000 00aa 1000 bb00 10aa 00000d0 0008 6000 00aa 1100 bb00 10aa 0008 6000
00000e00 00aa 1200 bb00 10aa 0008 6000 00aa 1300 00000f0 bb00 10aa 0008 6000 00aa 1400 bb00 10aa
```

# FELIX Firmware – GBT flavour

## Synthesis results for 4 channels



4 ch GBT:  
4TX-4RX



→ Pci 2x8 gen3

Name	Registers (1799680)	CLB LUTs (899840)	LUT as Logic (899840)	LUT as Memory (449920)	LOOKAHEAD8 (112480)	SLICE (112480)	CLB Registers (1799680)	Block RAM Tile (967)	URAM (463)	Bonded IOB (692)	GTYE5_QUAD (11)	PCIE 40E5 (4)
▼ N felix_top	16.99%	24.43%	23.69%	1.48%	2.14%	41.38%	16.99%	24.46%	43.63%	28.32%	54.55%	50.00%

# FELIX Firmware – IpGBT flavour

## Synthesis results

- ✓ • FLX182\_LPGBT\_12CH\_LTICLK\_GIT\_phase2-master\_rm-5.4\_204\_250925\_1327
- ? • FLX182\_LPGBT\_16CH\_LTICLK\_GIT\_phase2-release-5.3\_rm-5.3\_762\_250531\_1331
  - Can't program FPGA, missing prj file (can't disable DDR calibration)


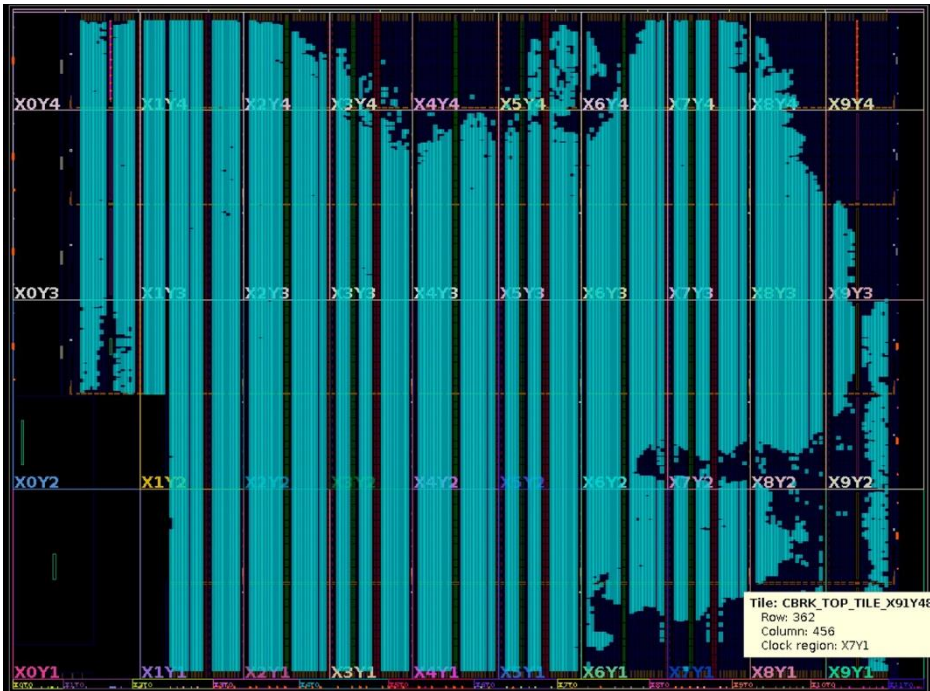
Instance	Module	Total LUTs	Logic LUTs	LUTRAMs	SRLs	FFs	RAMB36	RAMB18	URAM	DSP Blocks
felix_top	(top)	525065	512813	9156	3096	657519	400	776	294	0
xcvm1802		899840				1799680	612	1224	368	1968
% used		58,3509291				36,5353285	65,3594771	63,3986928	79,8913043	0

- ✗ • FLX182\_LPGBT\_24CH\_LTICLK\_GIT\_phase2-master\_rm-5.4\_204\_250925\_1327

[DRC UTLZ-1] Resource utilization: RAMB18 and RAMB36/FIFO over-utilized in Top Level Design (This design requires more RAMB18 and RAMB36/FIFO cells than are available in the target device. This design requires 1956 of such cell types but only 1934 compatible sites are available in the target device. Please analyze your synthesis results and constraints to ensure the design is mapped to Xilinx primitives as expected. If so, please consider targeting a larger device.)  
 [Vivado\_Tcl 4-23] Error(s) found during DRC. Placer not run.

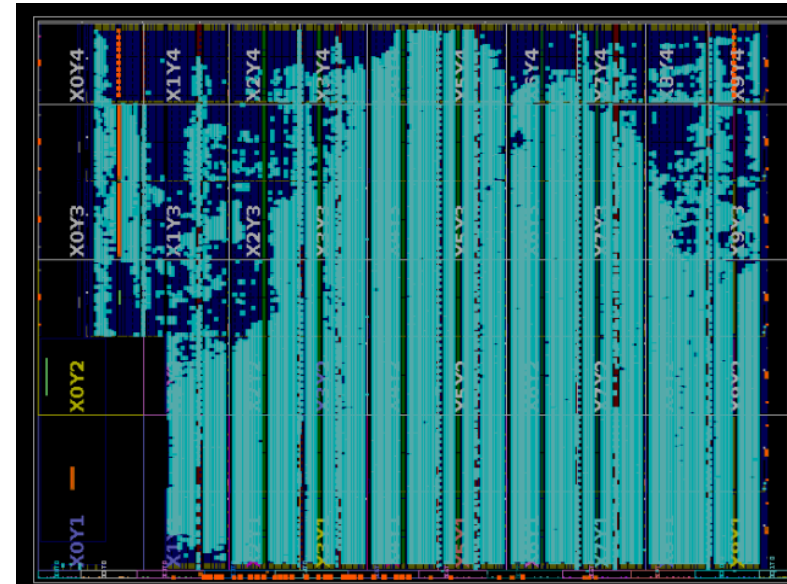

# FELIX Firmware – IpGBT flavour

## Synthesis results for 12 channels



# FELIX Firmware – FULL flavour

## Synthesis results for 24 channels



Name	Registers (1799680)	CLB LUTs (899840)	LUT as Logic (899840)	LUT as Memory (449920)	LOOKAHEAD8 (112480)	SLICE (112480)	CLB Registers (1799680)	Block RAM Tile (967)	URAM (463)	Bonded IOB (692)	GTYE5_QUAD (11)
> N felix_top	28.72%	38.39%	37.02%	2.74%	4.24%	67.93%	28.72%	66.70%	78.19%	28.32%	100.00%


# VLDB+

- **IpGBT: Highly configurable ASIC**
  - Mode of operation (TX/RX/TX-RX)
  - Uplinkdata rate
  - FEC option
  - Number of eLINK
  - ...




**11 pins and 320 registers need to be taken care of**

**VLDB+**: board designed by CERN's EP-ESE group to provide an evaluation kit for the new Versatile Link ecosystem.



- On board:
  - IpGBT
    - configuration and monitoring through the RaspberryPi 4 and piGBT web application.
  - VTRX+
    - only one Tx link is connected to the VTRx+
  - Rad-tol DC/AC regulators
  - FMC connectors for prototyping with FEB or FPGA development kit

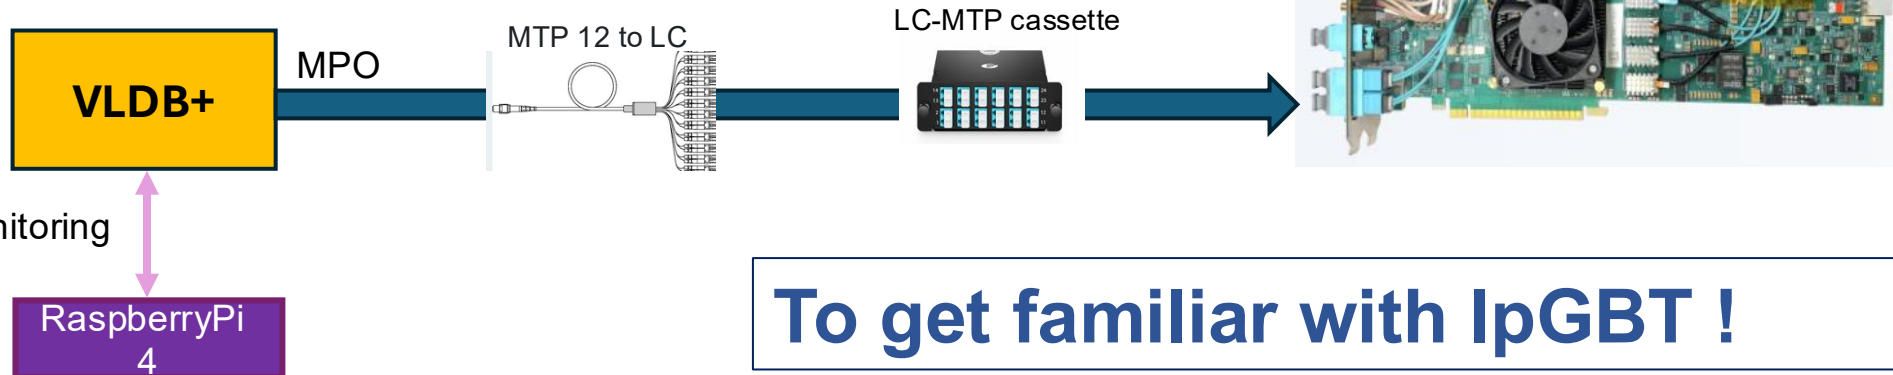
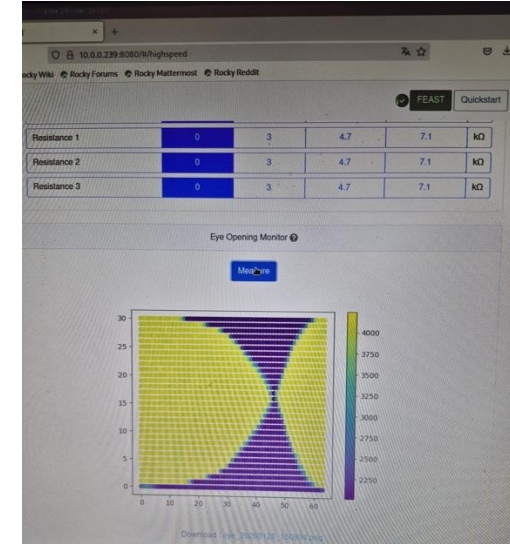


# Testbed (I) @Rome



**VLDB+** : board designed by CERN's EP-ESE group to provide an evaluation kit for the new Versatile Link ecosystem.


- On board:
  - IpGBT,
  - VTRx+ connector
  - Rad-tol DC/AC regulators
  - FMC connectors for prototyping with FEB or FPGA dev kit
- IpGBT configuration through **RaspberryPi 4** and piGBT web app.



# DAM/RDO emulator system


## FMC Adaptor card for DAM Emulation (FADE) module

- Designed for the ePIC detector test data readout and the incremental DAQ software development.



## AXAU15

- AMD Artix UltraScale+ FPGA development board in PCIe card format developed by ALINX.




FADE card plugged into the ALINX board


# DAM/RDO emulator system

## Emulator system operational mode:

- **Slave mode:** Emulator system receives System clock and sends the 'DAM' status via the GTU\_QSFP connector from/to a GTU. **It performs as a DAM emulator.**
- **Master mode:** Emulator system sends the System clock and receives the 'DAM' status via the GTU\_QSFP connector to/from a DAM.

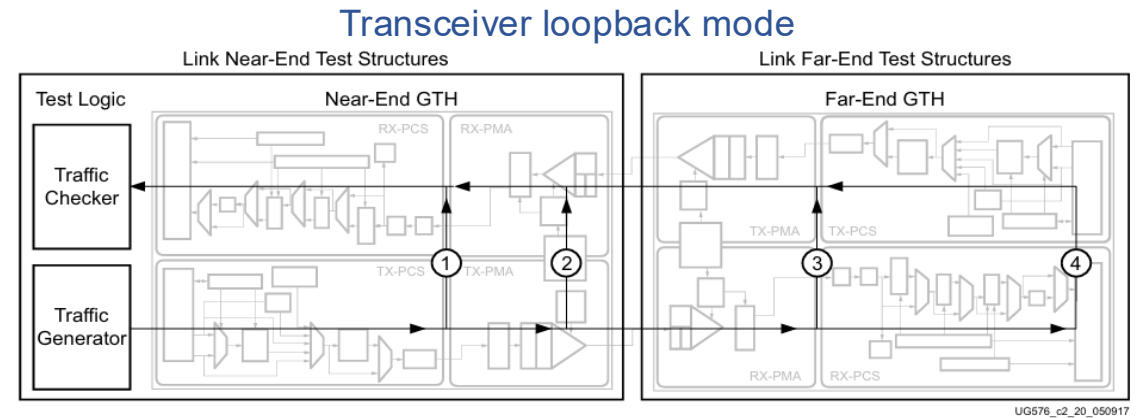
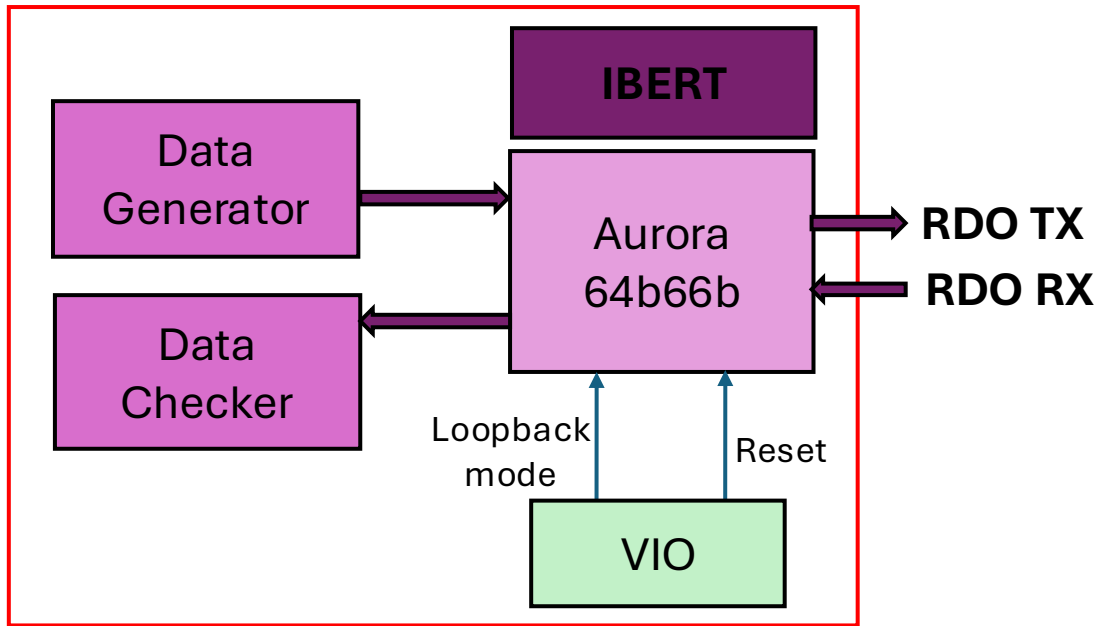


- **RDO emulator**



# Testbed (II) @Rome

## Alinx preliminary FW – Aurora reference design



Test with Transceiver loopback mode=2



Data received with NO ERROR

- Starting from Aurora Reference Design
- Useful to validate the data transmission

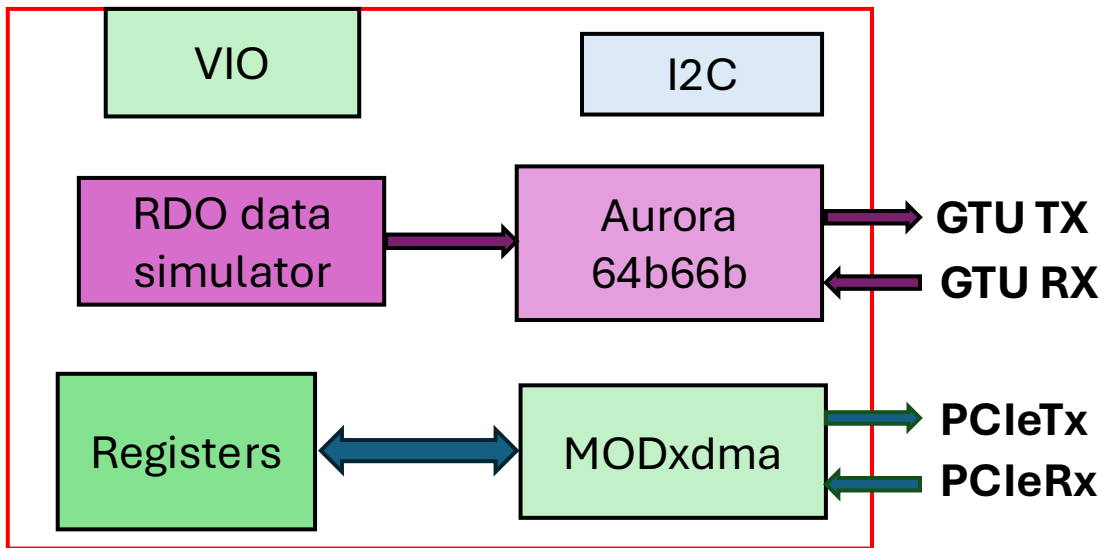
**NEXT STEP:**

- Connect Alinx to Felix board evaluating the overall link quality (IBERT) - Loopback mode 3,4

# Testbed (II) @Rome

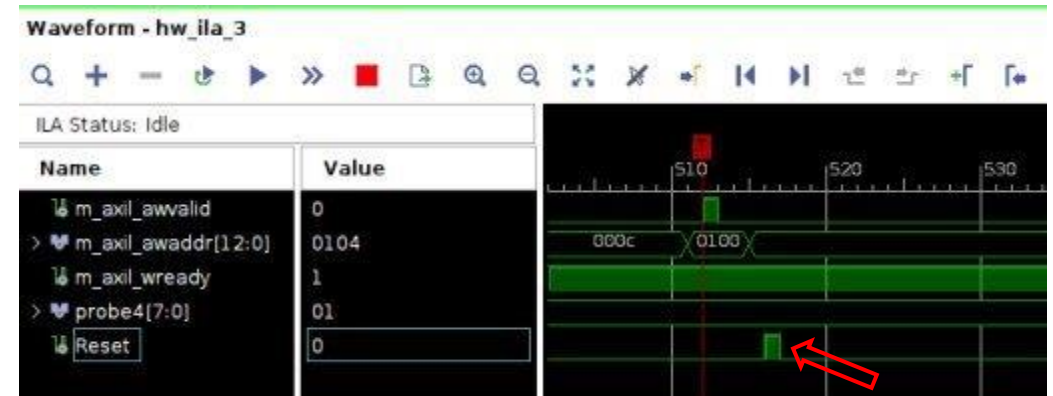
## Alinx preliminary FW – Register access via PCI

Thanks to William Gu (JLab)



```

[locicero@apecyc02 ~]$ cd Epic_Felix/FADE_APE/dma_ip_drivers/XDMA/linux-kernel/tools
[locicero@apecyc02 tools]$ sudo ./reg_rw /dev/xdma0_user 0x100 w 0x00000fff
device: /dev/xdma0_user, address: 0x100 (0x0+0x100), access write,
access width: word (32-bits)
character device /dev/xdma0_user opened.
Memory 0x0 mapped at address 0x7f6120730000.
Write 32-bits value 0x00000fff to 0x100 (0x0x7f6120730100)
    
```




"Reset" register writing

- Useful for learning how to access (read/write) registers via PCI in order to configure the firmware (e.g. operational mode)

- **NEXT STEPS:**

- Define data simulator format
- Connect Alinx to Felix board

# Testbed @Rome




FLX-182

2 x cassette MTP to 24 LC ports


VLDB+

Alinx FPGA + FADE board

FLX-182: 2 MTP TX/RX  
24 channels ports

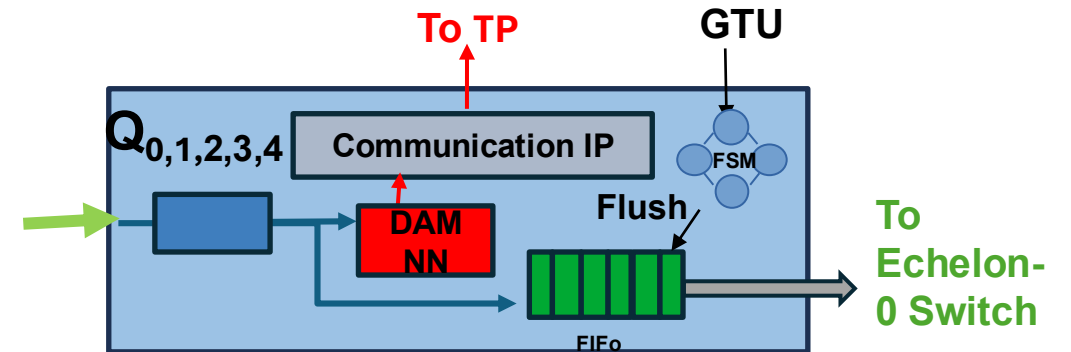
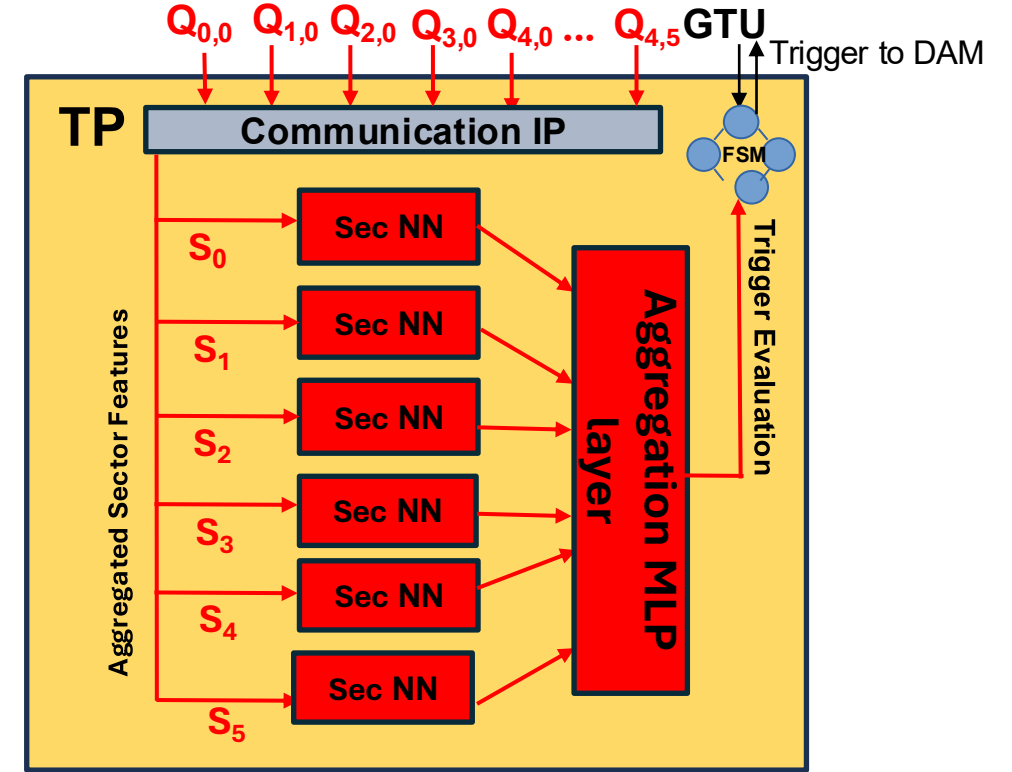
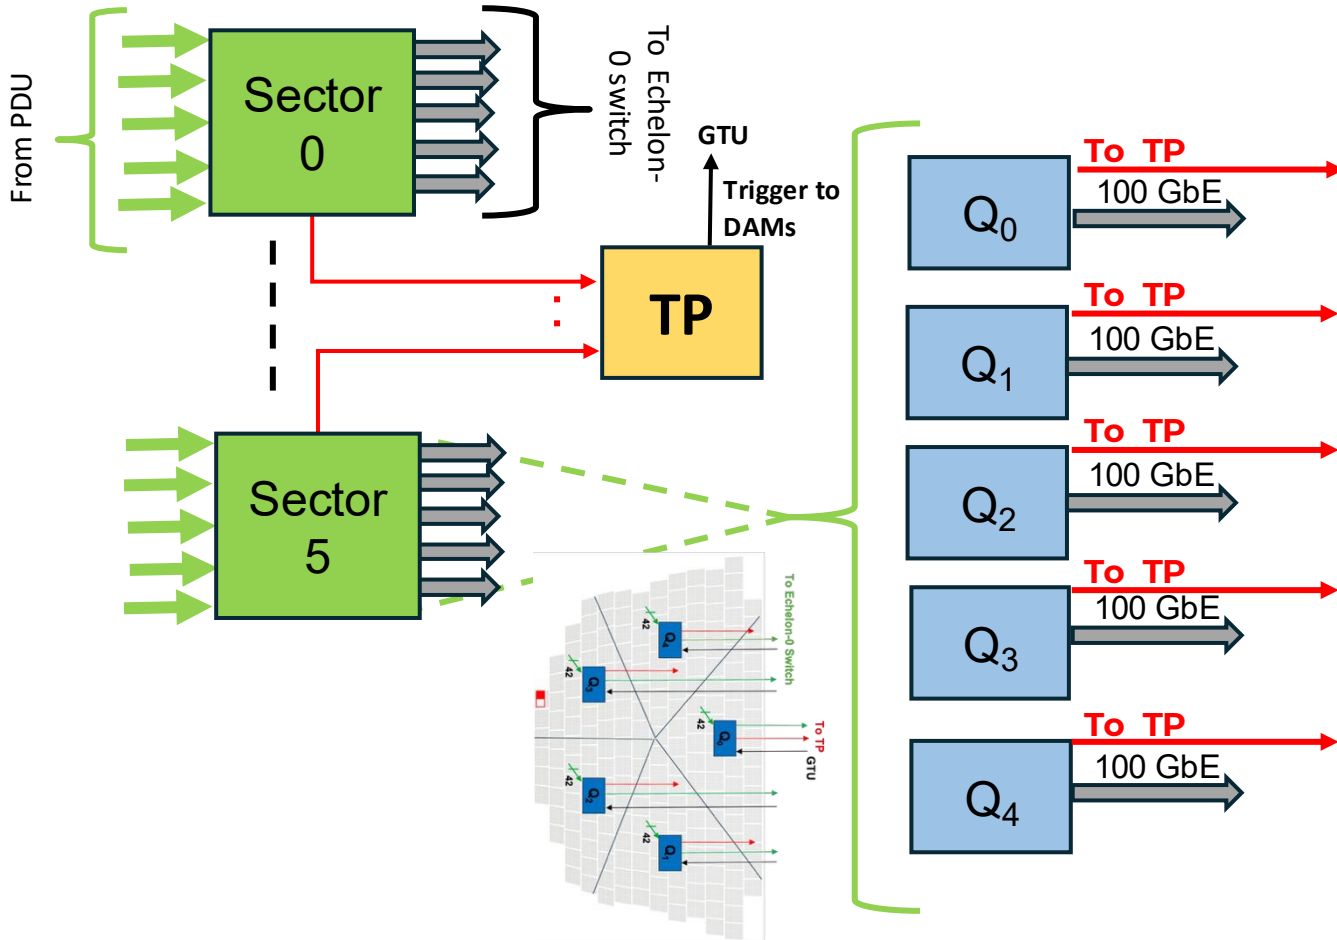


Alinx+FADE: 2 QSFP+  
ports + 1 SFP port




# dRICH data reduction integrated in the DAM

## Neural network distributed over 31 Felix-155

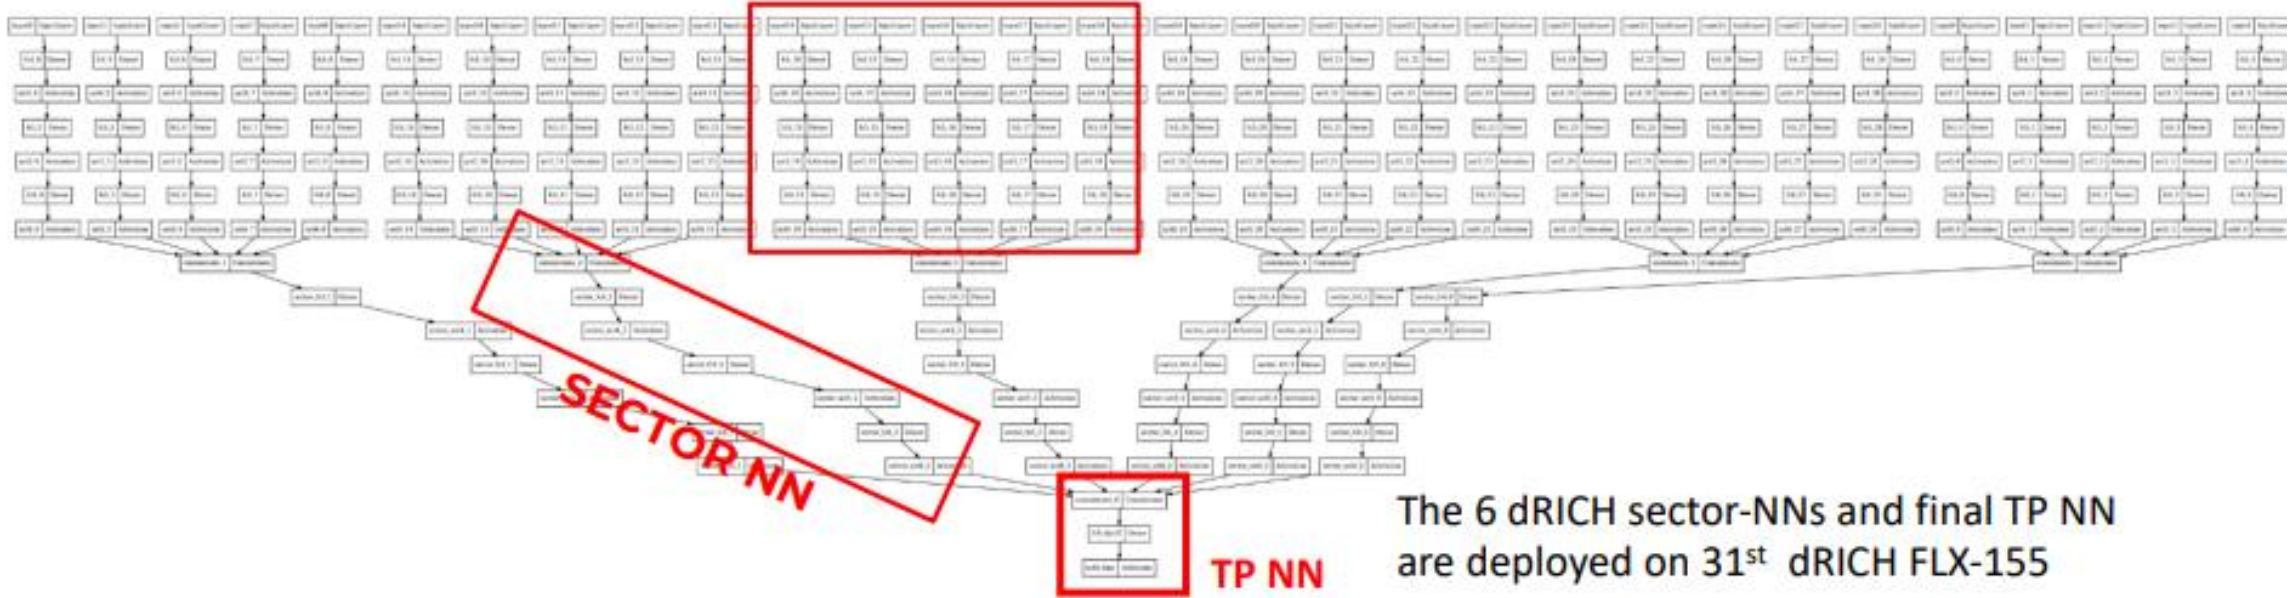


# NN deployment model and DAMs

6 dRICH sectors – 5 DAMs/sectors



5 DAM NN



The 6 dRICH sector-NNs and final TP NN are deployed on 31<sup>st</sup> dRICH FLX-155

# Conclusion

## DAM

- FLX182 board procured and installed in the APE Lab Rome test-bed.
- Study and test of the FELIX firmware and software stack, covering FULL, GBT and IpGBT flavours to drive the DAQ design and development.

## Ongoing Activities

- **DAM Tests**
  - Study of the IpGBT protocol using the VLDB+ board
  - Use Alinx + FADE to
    - emulate the GTU module for clock and run-control distribution to DAM
    - Prototype the RDO-DAM communication
- **Procurement of 2 pre-production FLX-155 (expected delivery by the end of 2026)**

The procurement process proceeding in four stages and is handled by the CERN Procurement Office:

  - **Purchase of all the FPGAs**
    - 2 AMD XCVP1552-2MSEVSVA3340 FPGAs
  - Purchase of the optical transceivers,
  - Purchase of the PCB,
  - Payment of the assembly.

# Backup

## 2 pre-production FLX-155 card procurement (end of 2026)



The procurement process will proceed in four stages separated in time and handled by the CERN Procurement Office:

- **Purchase of all the FPGAs**
  - 2 AMD XCVP1552-2MSEVSVA3340 FPGAs
  - The lead time is between 17 and 20 weeks.
- Purchase of the optical transceivers,
- Purchase of the PCB,
- Payment of the assembly.

Table 1: Breakdown of FLX-155 configurations.

FLX-155 Configuration	FireFly Optical Transceivers					
	B04		Y12		R24	
	Quantity	Gb/s	Quantity	Gb/s	Quantity	Gb/s
A	1	14	2	16	-	-
B	1	14	2	25	-	-
C	1	14	1	16	1	16
D	1	14	4	16	-	-

# dRICH output bandwidth



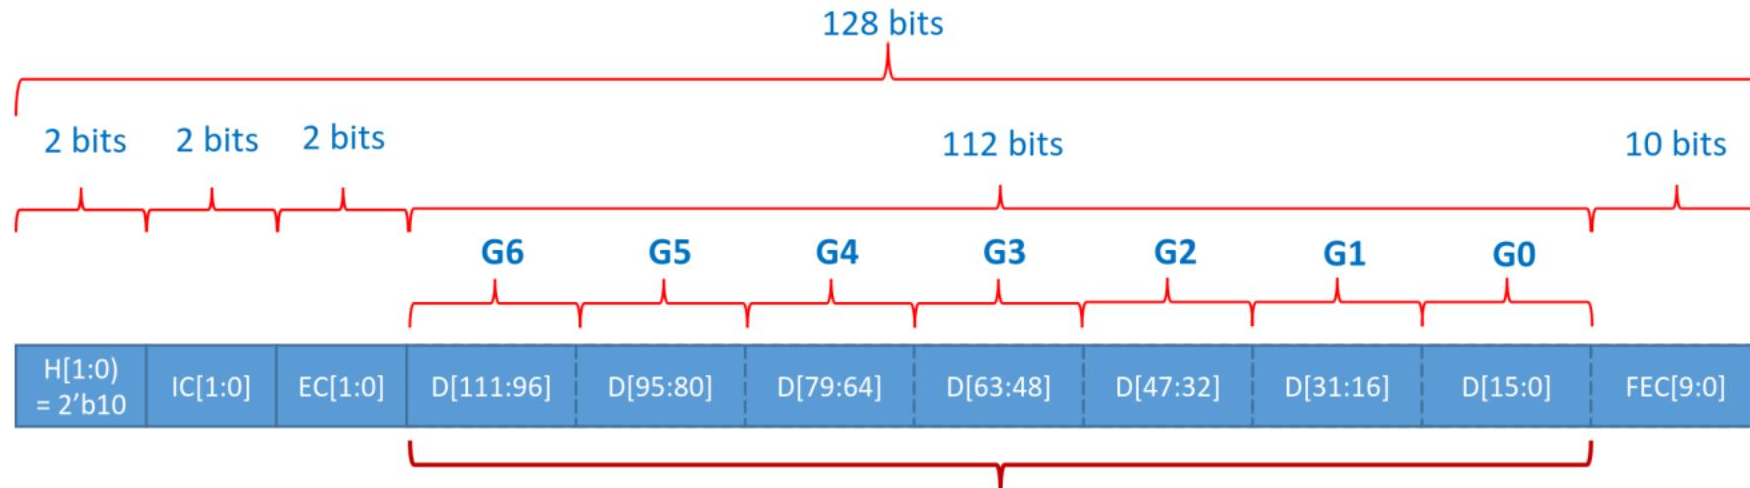
- Without data reduction with a 100 kHz DCR we are close to DAM bandwidth limit
- A data reduction factor 5 allows us to stay safe up to the 300 kHz limit
- Data reduction factor five can be achieved via shutter or provided by NN or ext. trigger or a combination of them.

# Uplink – FEC 5

- Dual data rate:
  - 5.12 Gb/s
  - 10.24 Gb/s
- Frame:
  - Header:
    - 2 / 4 – bits
    - 2'b10 (high) & 2'b01 (low)
  - IC field:
    - 2 – bit
    - 80 Mb/s
  - EC field:
    - 2 – bit
    - 80 Mb/s
  - Data field:
    - 112 / 224 – bit
    - 4.48 / 8.96 Gb/s
  - FEC field:
    - 10 / 20 – bit
- Error correction:
  - FEC:
    - Interleaving: 1 / 2
    - Symbol width: 5 – bit
    - Number of wrong symbols: 1 ( × 1 / × 2)
  - Up to 5 / 10 consecutive bits
  - Efficiency: 91%
- eLinks:
  - Data
    - 28 eLinks @ 160 / 320 Mb/s
    - 14 eLinks @ 320 / 640 Mb/s
    - 7 eLinks @ 640 / 1280 Mb/s
  - EC
    - 1 eLink @ 80 Mb/s

FEC5		
	5.12 Gb/s	10.24 Gb/s
Option	7	7
Frame (bits)	<b>128</b>	<b>2 x 128</b>
Header (bits)	<b>2</b>	<b>2+2</b>
Coded header	<b>no</b>	<b>no</b>
User field (bits)	116	232
Code (bits)	10	20
16-bit multiplicity	7.250	14.5
Remainder bits	4	8
User Bandwidth (GHz)	4.64	9.28
# eLinks groups (16 bit)	7	7
eLinks bandwidth (MHz)	160/320/640	320/640/1280
#eLinks	28/14/7	28/14/7
EC bandwidth (MHz)	80	80/160
IC bandwidth (MHz)	80	80/160
Unassigned bits	0	4
Corrected (bits)	<b>5</b>	<b>2 x 5</b>
Efficiency	<b>91%</b>	<b>91%</b>

# Example: 5.12 Gbps FEC5 Uplink Frame



Frame	Function	I/O Group
FRMUP[9:0]	FEC[9:0]	
FRMUP[25:10]	Data[15:0]	0
FRMUP[41:26]	Data[31:16]	1
FRMUP[57:42]	Data[47:32]	2
FRMUP[73:58]	Data[63:48]	3
FRMUP[89:74]	Data[79:64]	4
FRMUP[105:90]	Data[95:80]	5
FRMUP[121:106]	Data[111:96]	6
FRMUP[123:122]	EC[1:0]	EC
FRMUP[125:124]	IC[1:0]	
FRMUP[127:126]	H[1:0] = 2'b10	HFH[1:0] = 2'b10

- 7 groups of 4 input e-Ports  
Number of data ports:
- 28 eLinks @ 160 Mbps
  - 14 eLinks @ 320 Mbps
  - 7 eLinks @ 640 Mbps

Note: This is how you will see the uplink frame after it has been processed by the lpGBT-FPGA receiver but not how it is actually transmitted by the lpGBT ...

# Uplink – FEC 12

- Dual data rate:
  - 5.12 Gb/s
  - 10.24 Gb/s
- Frame:
  - Header:
    - 2 / 4 – bits
    - 10 (high) & 01 (low)
  - IC field:
    - 2 – bit
    - 80 Mb/s
  - EC field:
    - 2 – bit
    - 80 Mb/s
  - Data field:
    - 96 / 192 – bit
    - 3.84 / 7.68 Gb/s
  - FEC field:
    - 24 / 48 – bit
- Error correction:
  - FEC:
    - Interleaving: 3 / 6
    - Symbol width: 4 – bit
    - Number of wrong symbols: 1 ( × 3 / × 6)
  - Up to 12 / 24 consecutive bits
  - Efficiency: 78% (if accounting for the unassigned bits)
- eLinks:
  - Data
    - 24 eLinks @ 160 / 320 Mb/s
    - 12 eLinks @ 320 / 640 Mb/s
    - 6 eLinks @ 640 / 1280 Mb/s
  - EC
    - 1 eLink @ 80 Mb/s

FEC12		
	5.12 Gb/s	10.24 Gb/s
Option	28	28
Frame (bits)	<b>128</b>	<b>2 x 128</b>
Header (bits)	<b>2</b>	<b>2+2</b>
Coded header	<b>no</b>	<b>no</b>
User field (bits)	102	204
Code (bits)	24	48
16-bit multiplicity	6.375	12.75
Remainder bits	6	12
User Bandwidth (GHz)	4.08	8.16
# eLinks groups (16 bit)	6	6
eLinks bandwidth (MHz)	160/320/640	320/640/1280
#eLinks	24/12/6	24/12/6
EC bandwidth (MHz)	80	80/160
IC bandwidth (MHz)	80	80/160
Unassigned bits	2	8
Corrected (bits)	<b>12</b>	<b>2 x 12</b>
Efficiency	<b>80%</b>	<b>80%</b>