

Development of FELIX for HL-LHC

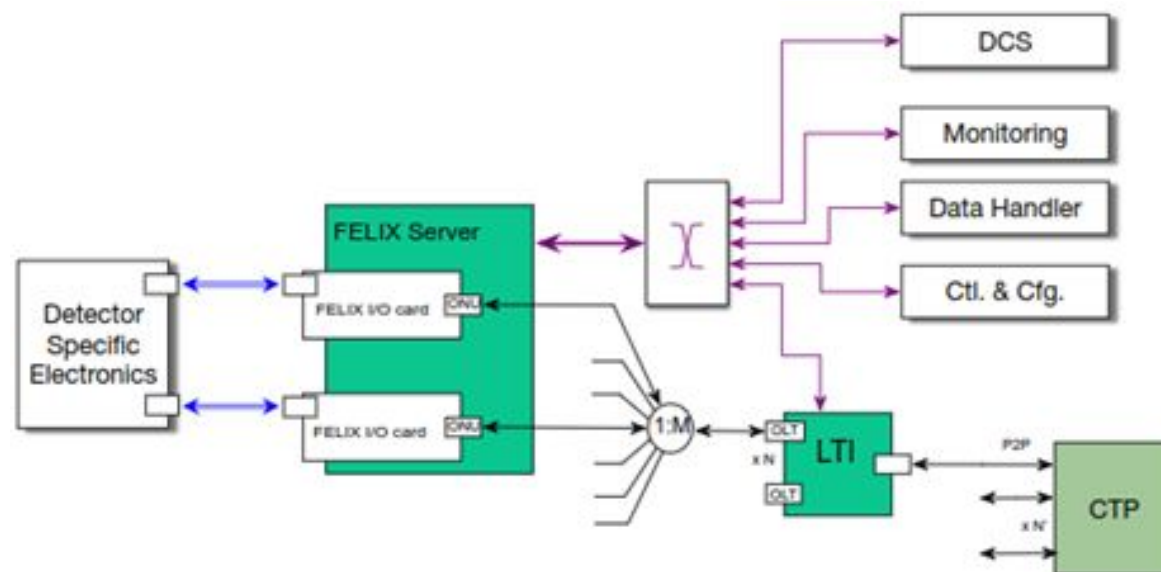
Mathieu Benoit, Filiberto Bonini, Hucheng Chen, Andy Liu, Dimitrios Matakias, Shaochun Tang,
Hao Xu, Elena Zhivun

July 21, 2022

FELIX - FrontEnd Link eXchange

FELIX: data/signal/message routing from/to FE

- Router between FE serial links and commercial network
- Data transport decoupled from data processing
- Get and distribute TTC (Timing, Trigger and Control) signals
- GBT-mode configurable e-links
- Detector independent

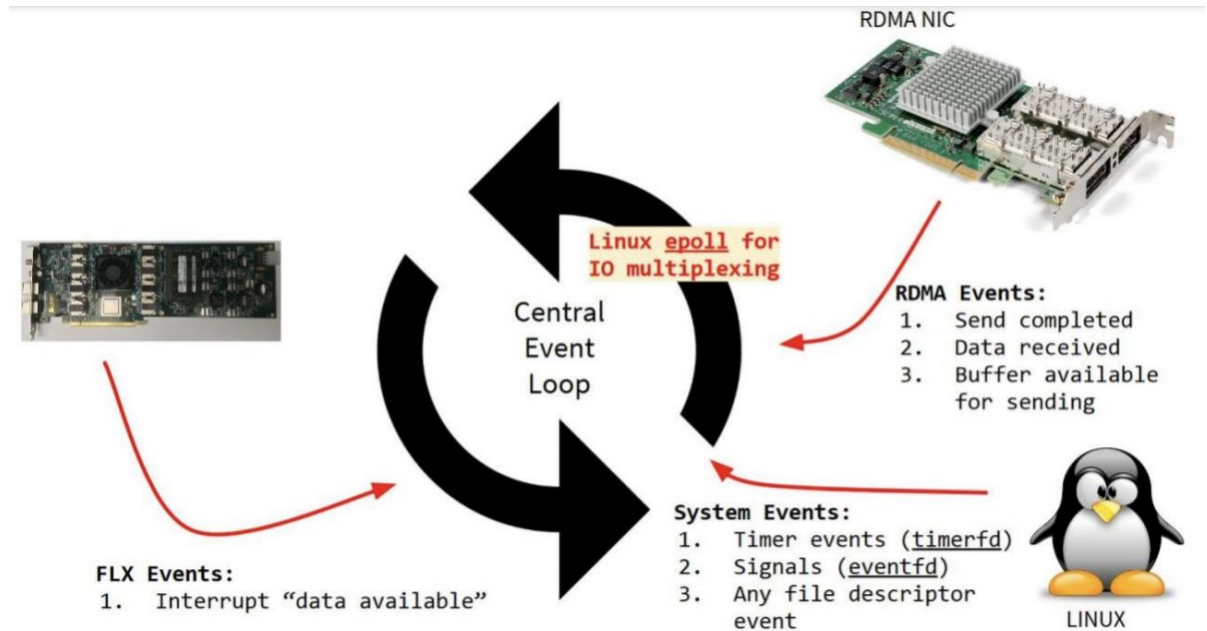


FELIX Software

- Each FELIX server hosts up to two FELIX cards and one NIC.
- Low level software has been developed for basic configuration and monitoring.
- High level software has been developed for data taking and channel monitoring.

More information is available on FELIX website:

<https://atlas-project-felix.web.cern.ch/atlas-project-felix/>



FELIX software development


FELIX Hardware FLX-712

FLX-712 hardware

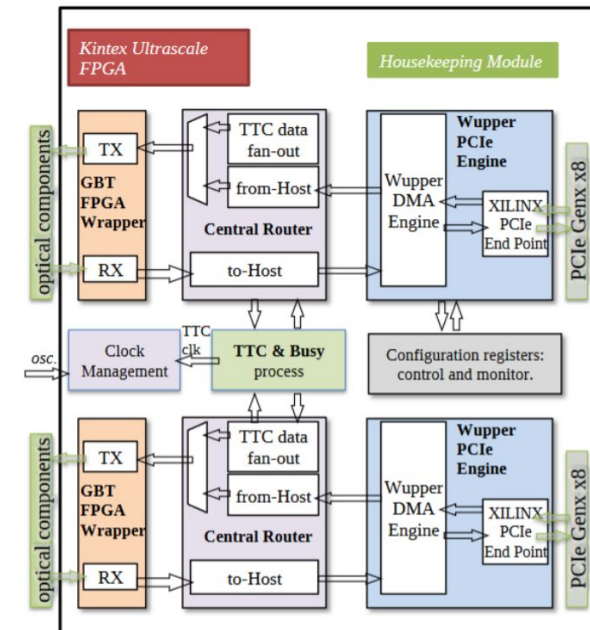
- Xilinx Ultrascale FPGA
- 24/48 duplex optical fibers for FE
- PCIe Gen3 x 16 lanes

FELIX firmware supports both of GBT mode and Full mode

- **GBT mode**
 - Firmware to interface FELIX to GBTx ASICs
 - GBTx = radiation hard chip aggregating FE electrical links, coupled to optical transceiver
 - GBT = protocol used by the GBTx
 - One link (4.8 Gb/s) divided into several e-links
 - Up to 24 GBT links per FLX-712
- **Full Mode**
 - Firmware to interface the FELIX to other FPGA-base systems
 - Up to 24 channels per FLX-712, 9.6 Gb/s each

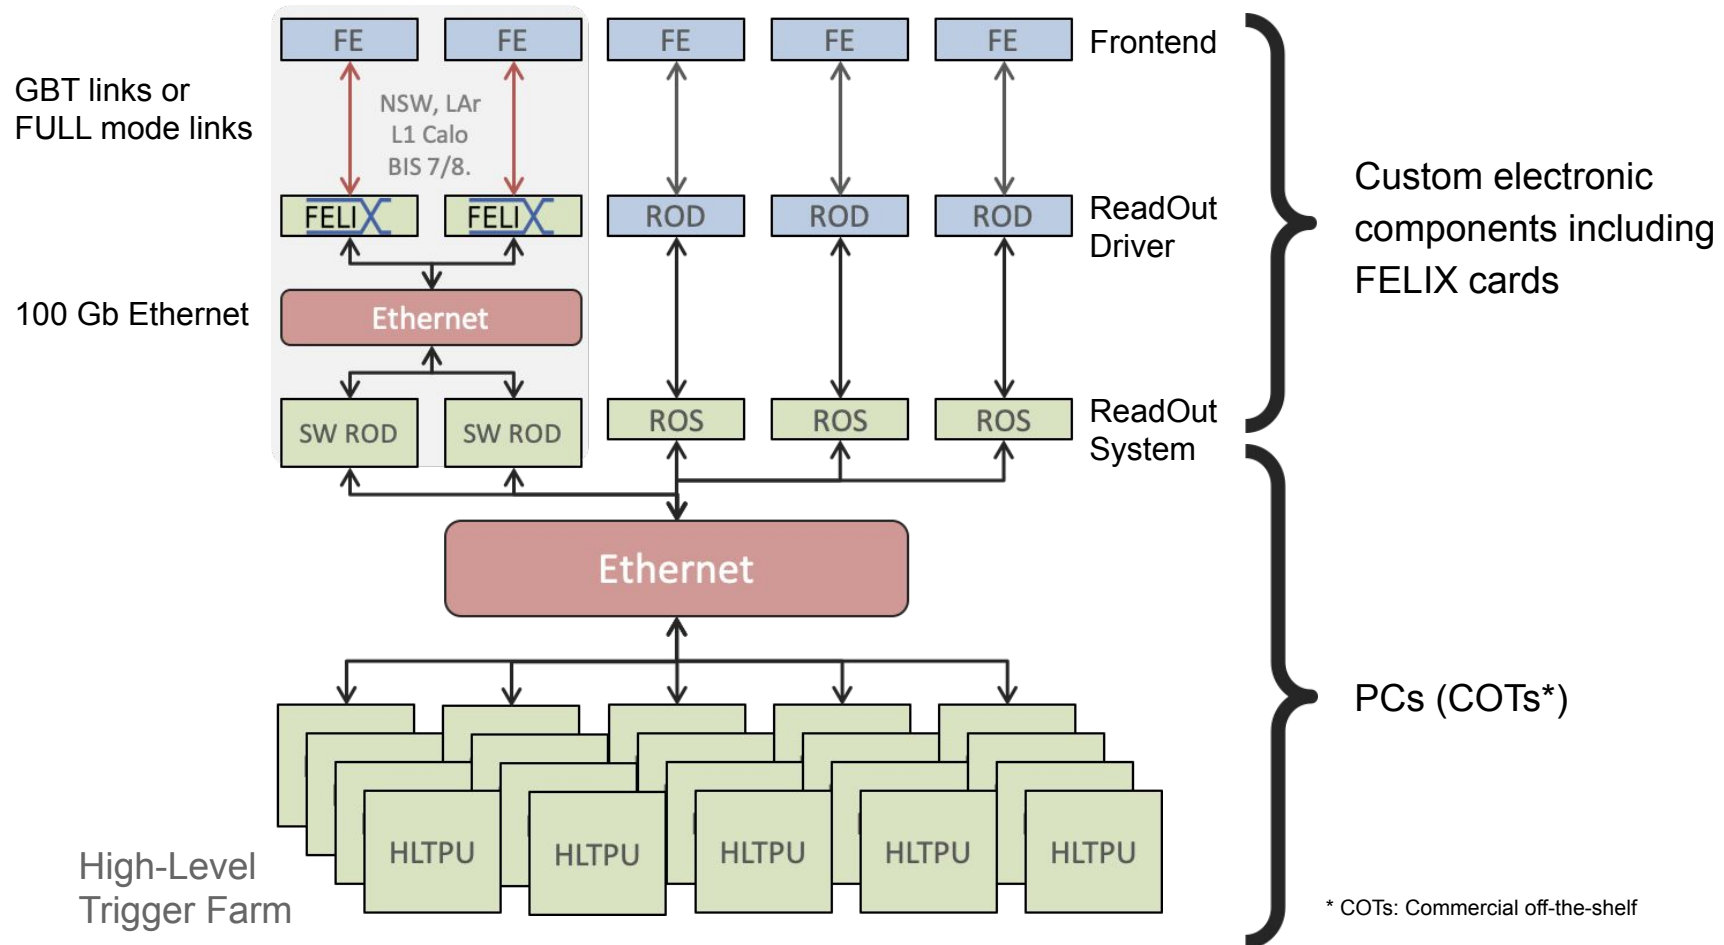


Picture of FLX-712



Block diagram of FLX-712 firmware

FLX-712 for ATLAS Phase-I Upgrade




* COTs: Commercial off-the-shelf

FLX-712 for sPHENIX DAQ

- All three sPHENIX tracking detectors use streaming readout
- Completed construction of sPHENIX FELIX DAQ interface (~50)

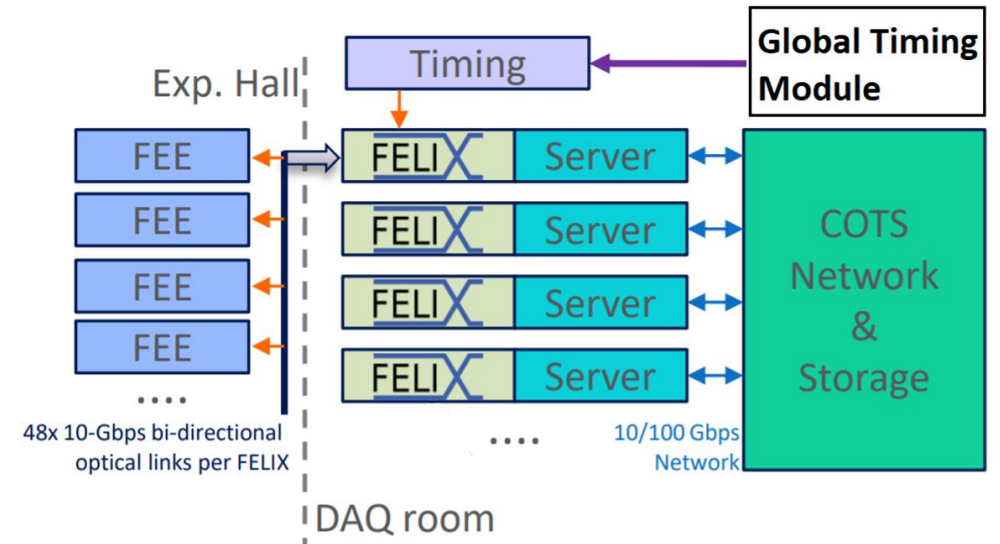
MVTX RU




INTT ROC



TPC FEE



Architecture of sPHENIX streaming DAQ



Data Flow

Requirements from ATLAS for HL-LHC

- FELIX will be used for ATLAS detector Readout system in Phase-II
 - To receive event data from detector front-end links
 - To relay Timing, Trigger and Control (TTC) information from the Phase-II TTC system to on-detector electronics
- Requirements for FELIX
 - ~12 k uplinks for Phase-II Readout system
 - Link speeds: from 2.5 Gb/s to 25 Gb/s
 - 24 channels for majority application, 48 channels are preferred for LAr LTDB and some NSW sectors
 - PCIe Gen4 (at least) for data output
 - Electrical connection for trigger signal
- A new FELIX with advanced FPGA is being developed to meet high throughput requirements

Table 2.4: Summary of Phase-II Detector Readout Link and Bandwidth Requirements. Downlink refers to data travelling toward the front-end electronics, and uplink to data travelling from the front-end toward the rest of the DAQ system. Detectors with existing FELIX installations from Phase-I will be updated with new hardware as required. Where no split between downlink and uplink is presented, it is assumed that the number of downlinks will at most match the number of uplinks, with a downlink bandwidth of 2.5 Gb/s for IpGBT and 5 Gb/s for GBT mode (and implicitly FULL mode). For 25 Gb/s uplinks the corresponding downlink speed has yet to be specified, but will likely be 5 or 10 Gb/s.

Detector	Number of FELIX boards	Number of Links	Bandwidth (Gb/s)	Top-level Link Protocol
ITk Pixel downlink	220	1564	2.5	IpGBT
ITk Pixel uplink		4684	10	
ITk Strips downlink	76	1552	2.5	IpGBT
ITk Strips uplink		1824	10	
LAr LASP downlink	5	200	5	GBT
LAr LASP uplink	48	1136	10	FULL
LAr LDPB downlink	8	31	5	FULL
LAr LDPB uplink		155	10	
LAr LATS downlinks	6	26	5	IpGBT
LAr LTDB downlinks	16	620	5	GBT
L0Calo downlink	8	16	5	FULL
L0Calo uplink		120	10	
NSW downlink	96	1728	5	GBT
NSW uplink		2880	5	
NSW Trigger Processor	4	96	5	GBT
RPC downlink (incl BI)	4	32	5	FULL
RPC uplink (incl BI)		96	10	
CTP downlink	1	0	5	FULL
CTP uplink		12	10	
MUCTPI downlink	1	2	5	FULL
MUCTPI uplink		8	10	
MDT Trigger Processor	32	64	2.5	IpGBT
MDT Trigger Processor		768	10	
Global Trigger GEP	7	50	25	Interlaken
Global Trigger MUX & gCTPI	4	74	25	IpGBT
Tile	8	161	10	FULL
TGC Endcap Sector Logic	8	192	10	FULL
HGTD DAQ Path	48	1152	10	IpGBT
PLR	4	32	10	IpGBT
BCM*	2	12	10	IpGBT
LUCID	1	4	10	IpGBT
ZDC	1	9	10	IpGBT
AFP	2	32	10	IpGBT

Target FPGA

- FPGA: Xilinx Versal Prime XCVM1802-1MSEVSVA2197
 - ~900k LUTs, ~1.8M FF
 - GTY Transceivers
 - 44 transceivers in 11 Quads
 - Up to 26.5625 Gb/s for -1M devices
 - 4 x PCIe Gen4 x8 end-points
- PCIe Architecture
 - PL Blocks for PCIe
 - PCIe Gen4, up to 8 GT lanes concurrently active
 - Soft IP subsystems available
 - Easiest migration for legacy designs
 - CPM Block for PCIe and CCIX
 - Up to Gen4 x16 link or two Gen4 x8 links
 - Two hardened ports for PCIe

Resource	FLX-712 / KU115	FLX-182 / VM1802
LUTs	663,360	899,840
FlipFlops	1,326,720	1,799,680
BlockRAM 36kb	2160	967
UltraRAM 288kb	-	463
Transceivers GTH < 16.3 Gb/s	64	-
Transceivers GTY < 32.75 Gb/s	-	44
PCIe interface	Gen3	Gen4

Pre-Prototype: FLX-181

FPGA: Xilinx Versal Prime XCVM1802-1MSEVSVA2197 ES

- High speed optical links
 - 12 x FE-Links: 1 pair of Samtec FireFly 14 Gb/s 12-ch modules
 - 16 x GTY links @ 25 Gb/s on FMC+
- Dual PCIe Gen4 x8, up to 256 GT/s
 - 16 x GTY links
- 3 x Mini-UDIMM DDR4 modules
 - Accessible by both PL and PS through NoC
- FMC+ mezzanine card
 - 34 x differential pairs from XPIO banks
 - 16 x GTY links
- Peripherals:
 - Micro SD 3.0 and QSPI Flash for system boot
 - USB - I2C/UART




Photo of assembled FLX-181

Status of FLX-181

Hardware

- Three boards have been produced for firmware development
 - One is being used for ATLAS firmware development at Nikhef
 - One is being used for DUNE firmware development at CERN
 - One is kept at BNL for hardware and firmware development

Firmware

- Different firmware flavours have been developed

Flavour	Channels	LUTs	FF	BRAM	URAM	Power
FULLMODE	24	58%	38%	38%	69%	
GBT	24	58%	45%	94%	73%	
LPGBT	24	84%	39%	87%	69%	31.5 W
STRIP	24	65%	45%	76%	90%	30.4 W


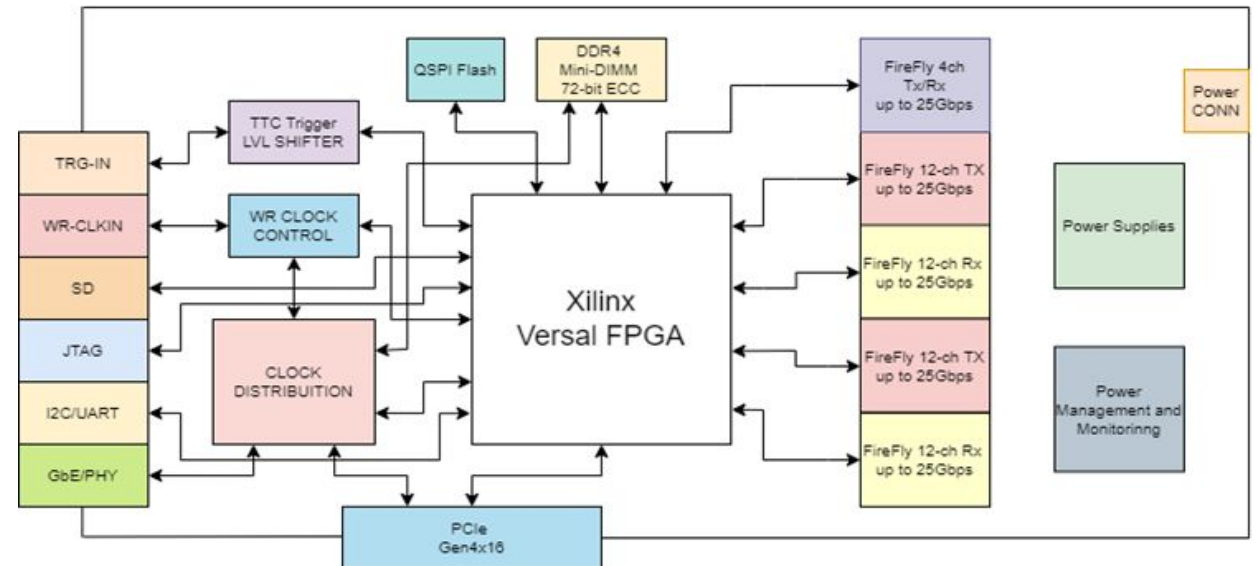


Photo of FLX-181 with 25Gbps FireFly FMC+

Prototype: FLX-182

FPGA: Xilinx Versal Prime XCVM1802-1MSEVSVA2197
production device

- PCIe Gen4 x16: PL and CPM compatible
- 24 FireFly links with 3 possible configurations
 - 24 links @25 Gb/s
 - 24 links @10 Gb/s (CERN-B-Y12)
 - 12 links @25 Gb/s + 12 links @10 Gb/s
- 4 FireFly links with 2 possible configurations
 - LTI interface
 - 100GbE
- Electrical signals on front panel
 - 3 inputs and 3 outputs
- 1 DDR4 Mini-UDIMM
- USB-JTAG/USB-UART




Block diagram of FLX-182

Architecture and Interfaces

- PCIe Gen 4 x 16 lanes
- Transceiver
 - Transceiver Type: Samtec FireFly transceiver
 - Transceiver Speed: up to 10 Gb/s (“CERN-B”) or 25 Gb/s
- Number of Optical Connectors per Card
 - At least 24 bi-directional connections to front-end electronics
 - A separate bi-directional connection to the TTC/BUSY system
- Configuration
 - Boot from JTAG/QSPI/SD card
 - Remote FPGA configuration from Multiple Flash Partitions
- DDR4/Flash Memory/SD card
- I2C
- External Electrical Interface
- Voltage Protection
- Temperature Protection

Thermal Simulation



- Power estimate
 - FPGA: ~60W
 - Whole board: ~133W
- Heatsink
 - FPGA
 - Passive or fan-sink
 - Temperature is below 70°C @70W
 - FireFly
 - 14Gbps module: standard pin-fin heatsink
 - 25Gbps module: high performance pin-fin heatsink
- Geometric model for simulation with COMSOL
 - Main power-dissipating / air-flow conditioning components:
 - Power modules
 - VM1802 with heatsink
 - FireFly modules
 - DDR4 UDIMM w/ socket
 - Front Panel



FLX-182 thermal simulation

FLX-182 Status

- Design passed FELIX review, will be sent out for fabrication in this week
- First assembled board is expected to be delivered in early September 2022
- 7 boards will be produced if there's no big design issues, by December 2022
- Small production for more boards is possible once FPGA is available



Plan for 48-ch FELIX

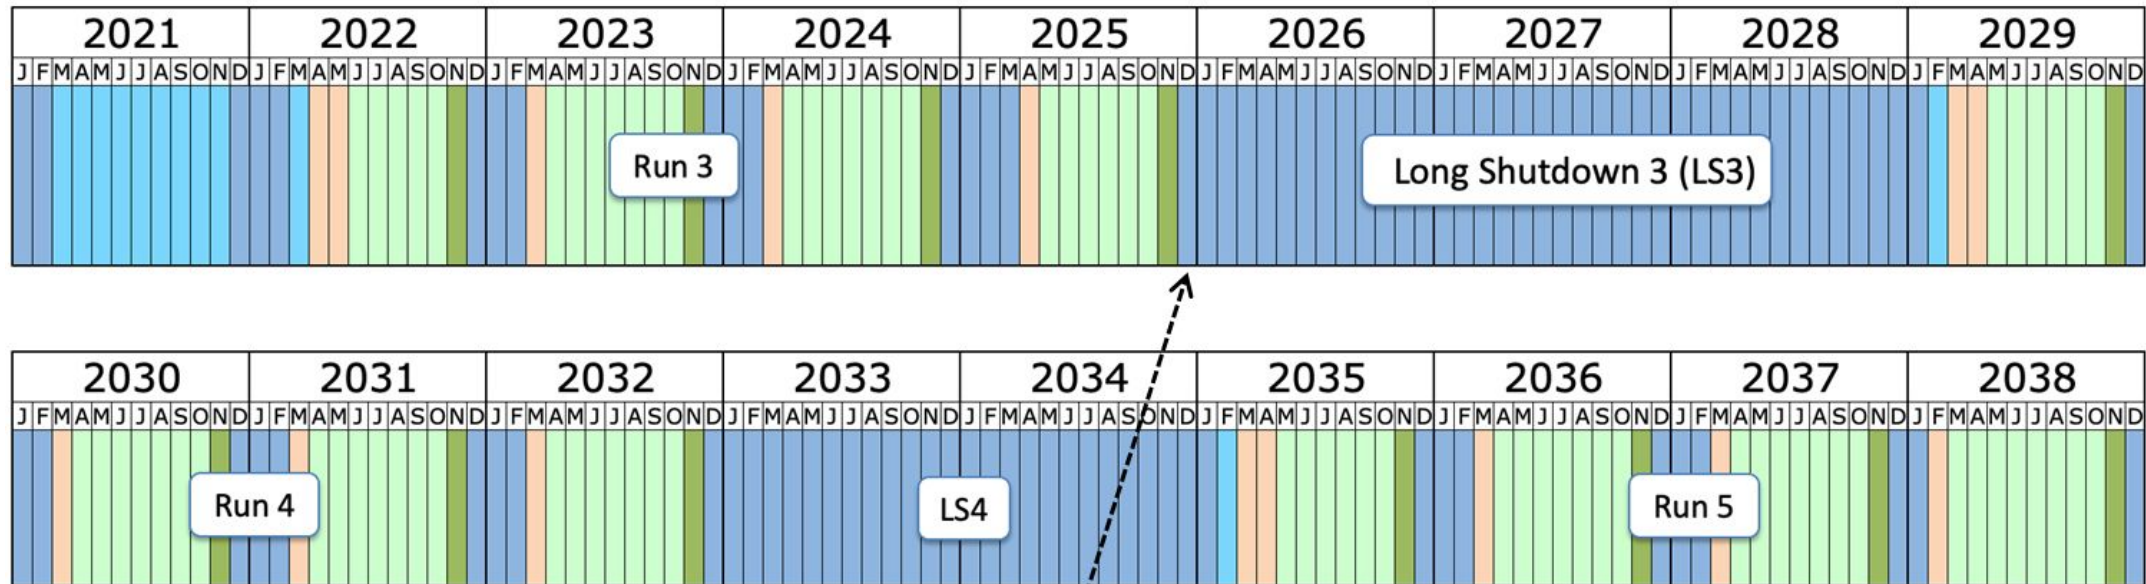
- FPGA: Versal Premium, e.g. VP1552
- Transceivers: Up to 100+ GTYP/GTM
- PCIe Gen 5 up to 16 lanes
- If FPGA is available as planned, design will start in Q1 of 2023, first board is expected to be available in Q3 2023.

U.S. ATLAS TDAQ Phase II Upgrade Project Schedule


- Prototype firmware complete, [July 2023](#)
- Technical evaluation of FLX-182 and 48-ch prototype will be used in coordination with detector requirements to determine the best design candidate for production, ATLAS PDR Follow-up, [August 2023](#)
- The selected design will be review in Felix FDR, [October 2023](#)
- Pre-production ready, [December 2024](#)
- Production complete, [May 2025](#)
- DOE CD-4 complete, [December 2028](#)
- U.S. ATLAS project CD-4 complete, [Q1 FY 2031](#)

FELIX Technical Support

Technical support will be provided through the HL-LHC operation



Last updated: January 2022



Long Shutdown 3 (LS3) – a 3-year shutdown and installation period beginning in Jan 2026 which drives the upgrade construction completion schedule.

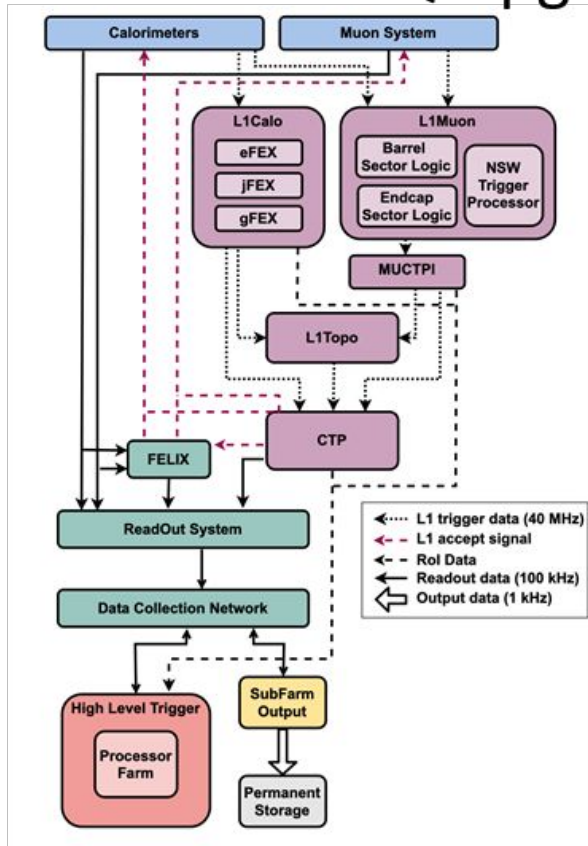
Outside of ATLAS

- FELIX is a generic platform for high throughput readout
 - Integration of hardware, firmware and software
 - Both firmware and software are open source development
 - Official technical support will be available in 2030s
- Collaboration outside of ATLAS is welcome
 - FLX-712 is being used for DAQ systems in sPHENIX@BNL, CBM@GSI, NA62@CERN, ProtoDUNE@CERN
 - And several smaller scale testbeam experiments

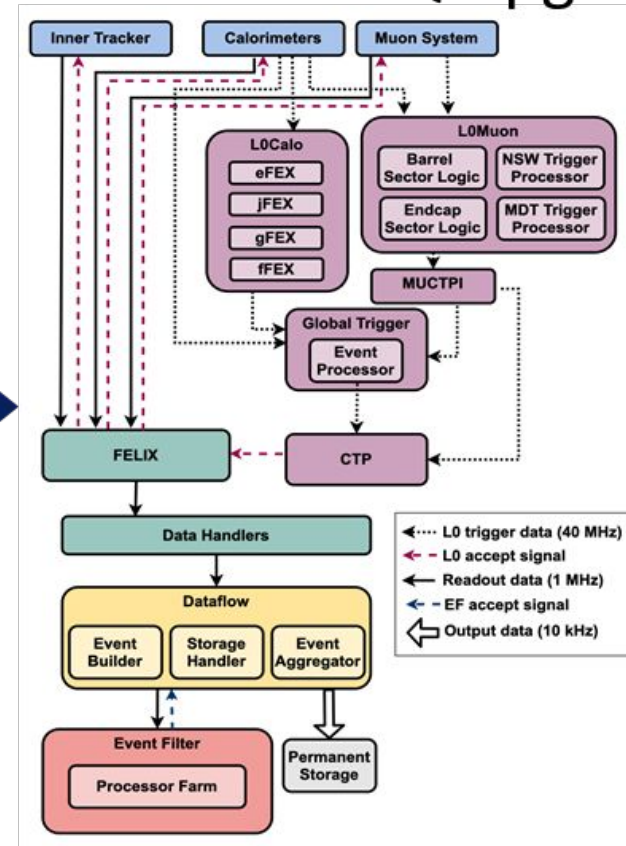
Backup

ATLAS TDAQ Upgrade for HL-LHC

Phase-I TDAQ Upgrade



HL-LHC TDAQ Upgrade



Architecture of TDAQ

- Trigger rate: 1MHz
- Data rate: 5.2TB/s
- 17093 optical links from detectors

FELIX:

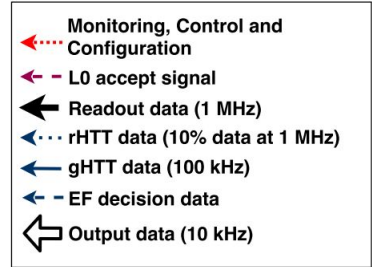
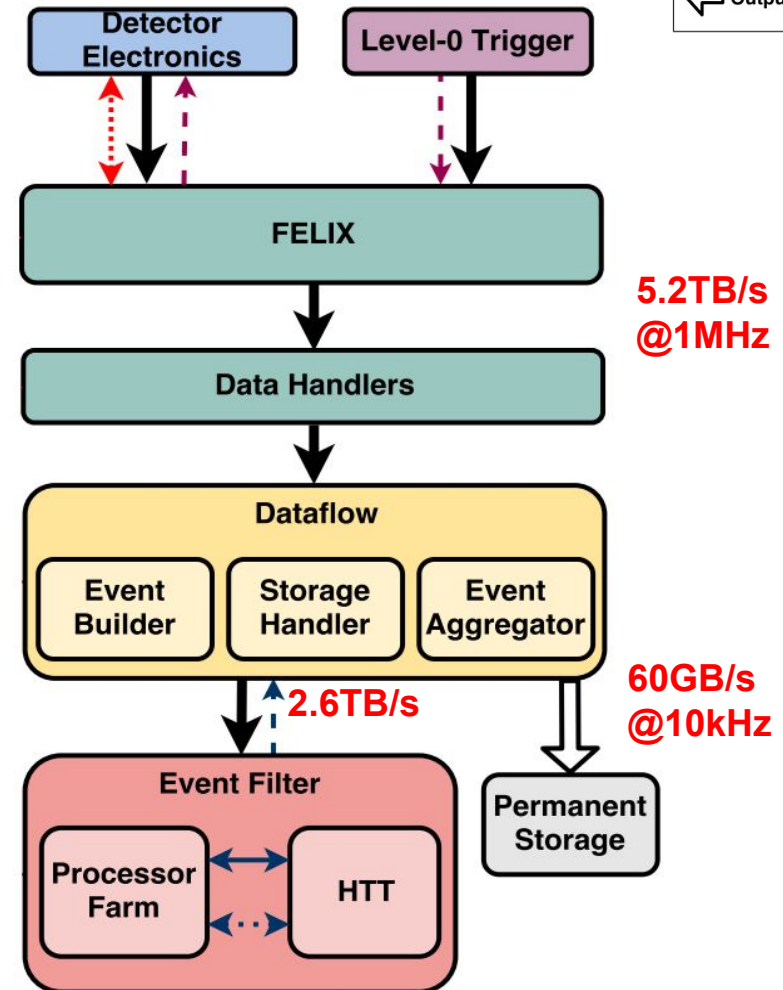
- 639 FLX cards for HL-LHC
- Distribute trigger & command signals to the Front-Ends
- Transmits the full detector data up to the Data Handlers

Data Handlers:

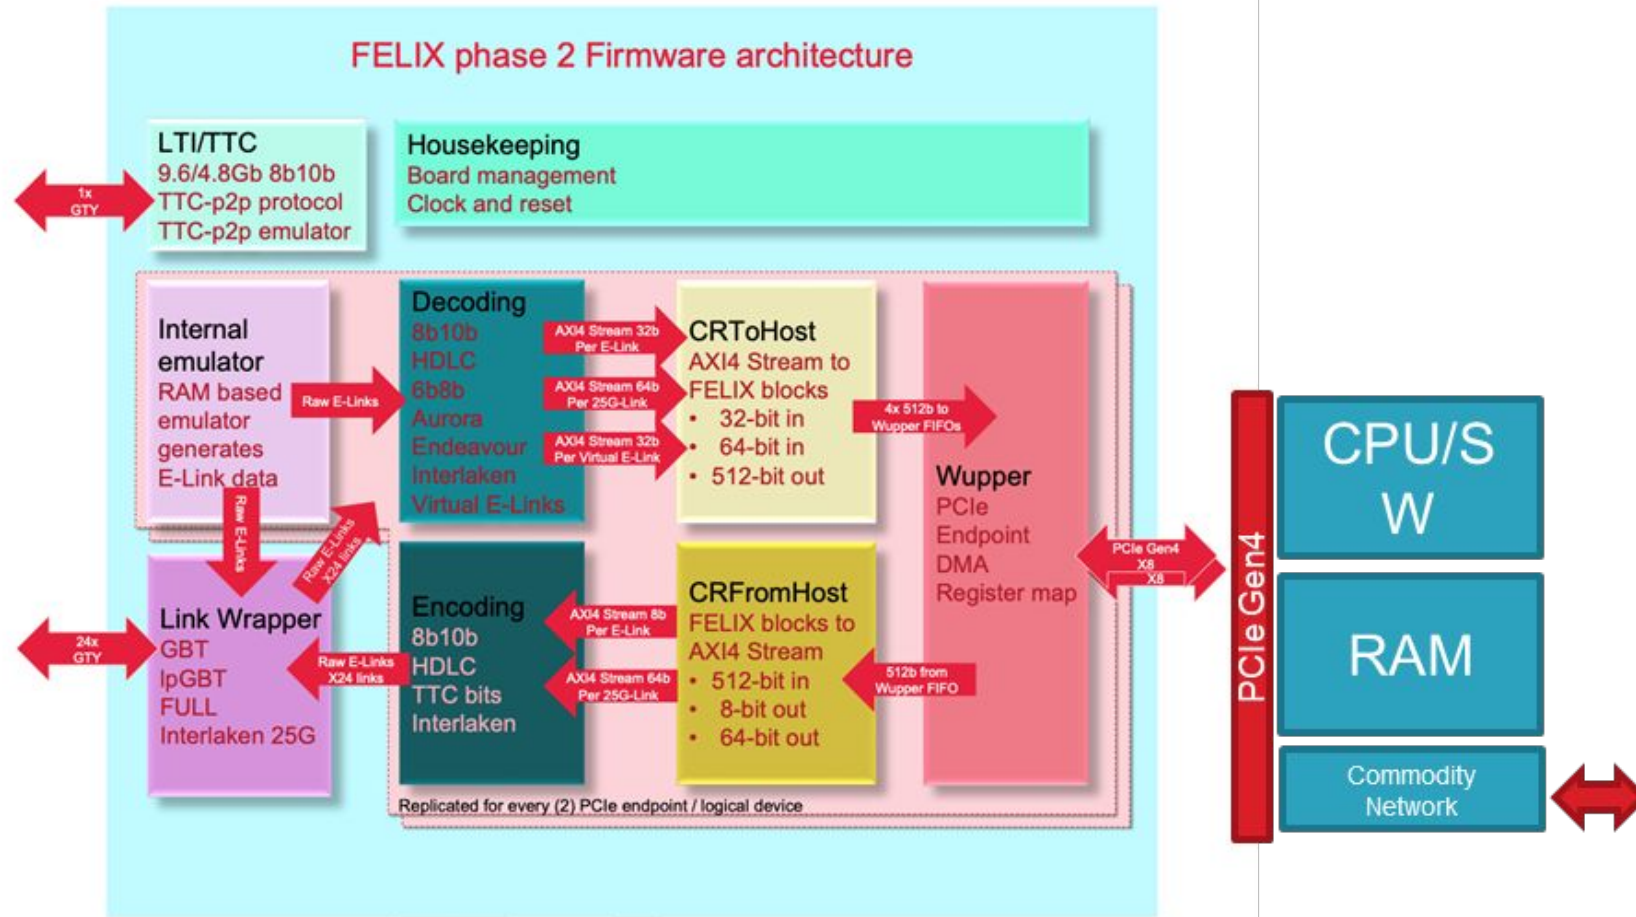
- Receives the data from FELIX servers over the network at 1 MHz
- Performs data formatting and send data fragments to the Dataflow system

Dataflow:

- **Event Builder** builds event records and manages the storage volume of the Storage Handler system
- **Storage Handler** buffers event data before and during processing by the Event Filter
- **Event Aggregator** collects, formats and transfers the output to CERN permanent storage



FLX-182 Firmware



Block diagram of firmware architecture

Firmware Flavors for Phase II

Firmware Flavours

- **GBT (+LTDB)**
 - 24 bidirectional GBT 4.8Gb/s links
 - 8b10b + HDLC decoding
 - 8b10b, HDLC and TTC encoding
 - Functionally equivalent to Phase I GBT, but aiming for 24 fully configurable
- **FULL**
 - 9.6Gb/s 8b10b ToHost / Uplinks
 - FromHost / Downlinks like GBT
 - Functionally equivalent to Phase I FULL
- **PIXEL**
 - 24 IpGBT links (10Gb/s Up, 2.5Gb/s Down)
 - Aurora decoding for RD53A and ITkPixv1
 - Command encoding
 - HDLC Encoding / decoding for EC/IC
- **STRIP**
 - 24 IpGBT links (10Gb/s Up, 2.5Gb/s Down)
 - 8b10b decoding 320Mb / 640Mb
 - LCB encoding (6b8b) / trickle memory
 - R3L1 encoding
 - Endeavour for EC, HDLC for IC
- **LPGBT**
 - 24 IpGBT links (10Gb/s Up, 2.5Gb/s Down)
 - 320, 640, 1280 Mb/s 8b10b + 80 Mb IC/EC
 - Includes HGTD encoding / decoding
- **INTERLAKEN**
 - 25.78125 Gb/s Interlaken Up
 - 25.78125 Gb/s, 9.6Gb/s 8b10b or 4.8Gb/s GBT down

FPGA Resources Estimate for Phase II

Resource estimation for all flavours

		KU115	VU37P	VM1802	VP1552
GBT 24 channel	LUT	80.65%	48.04%	69.60%	35.71%
	FF	77.03%	35.16%	50.94%	26.13%
	BRAM	70.00%	42.91%	89.45%	34.04%
	URAM		30.00%	62.20%	22.14%
FULL 24 channel	LUT	52.59%	30.61%	44.35%	22.75%
	FF	38.40%	22.92%	33.21%	17.03%
	BRAM	40.46%	10.07%	20.99%	7.99%
	URAM		30.00%	62.20%	22.14%
LPGBT 24 channel	LUT	112.51%	57.25%	82.94%	42.55%
	FF	52.39%	26.66%	38.62%	19.81%
	BRAM	68.94%	38.14%	79.52%	30.26%
	URAM		30.00%	62.20%	22.14%
PIXEL 24 channel	LUT	82.40%	41.93%	60.75%	31.17%
	FF	62.04%	31.57%	45.74%	23.46%
	BRAM	61.20%	29.86%	62.25%	23.69%
	URAM		30.00%	62.20%	22.14%
STRIP 24 channel	LUT	67.04%	34.11%	49.42%	25.35%
	FF	49.94%	25.41%	36.81%	18.88%
	BRAM	121.43%	50.10%	104.45%	39.75%
	URAM		70.00%	145.14%	51.65%
INTERLAKEN 8 channel	LUT		6.31%	9.15%	4.69%
	FF		5.44%	7.89%	4.05%
	BRAM		19.39%	40.43%	15.39%
	URAM		0.00%	0.00%	0.00%

- Choice of FPGA based on overall optimisation of channel density, space, power and cost.
- Design with more (or fewer) than 24 channels will be considered against these criteria
- Limited configuration options are considered for GBT and LPGBT Flavours
- Versal Prime VM1802
 - Routing fails when LUT count exceeds 50%. Newer tools may improve this.
- Versal Premium VP1552
 - More than 24 channels possible