LHC BLM

System readiness
Summary of FW & SW changes

Machine Protection Panel

Mathieu Saccani (SY-BI-BL) on behalf of the BLM team

10/06/2022

BLM LHC System readiness

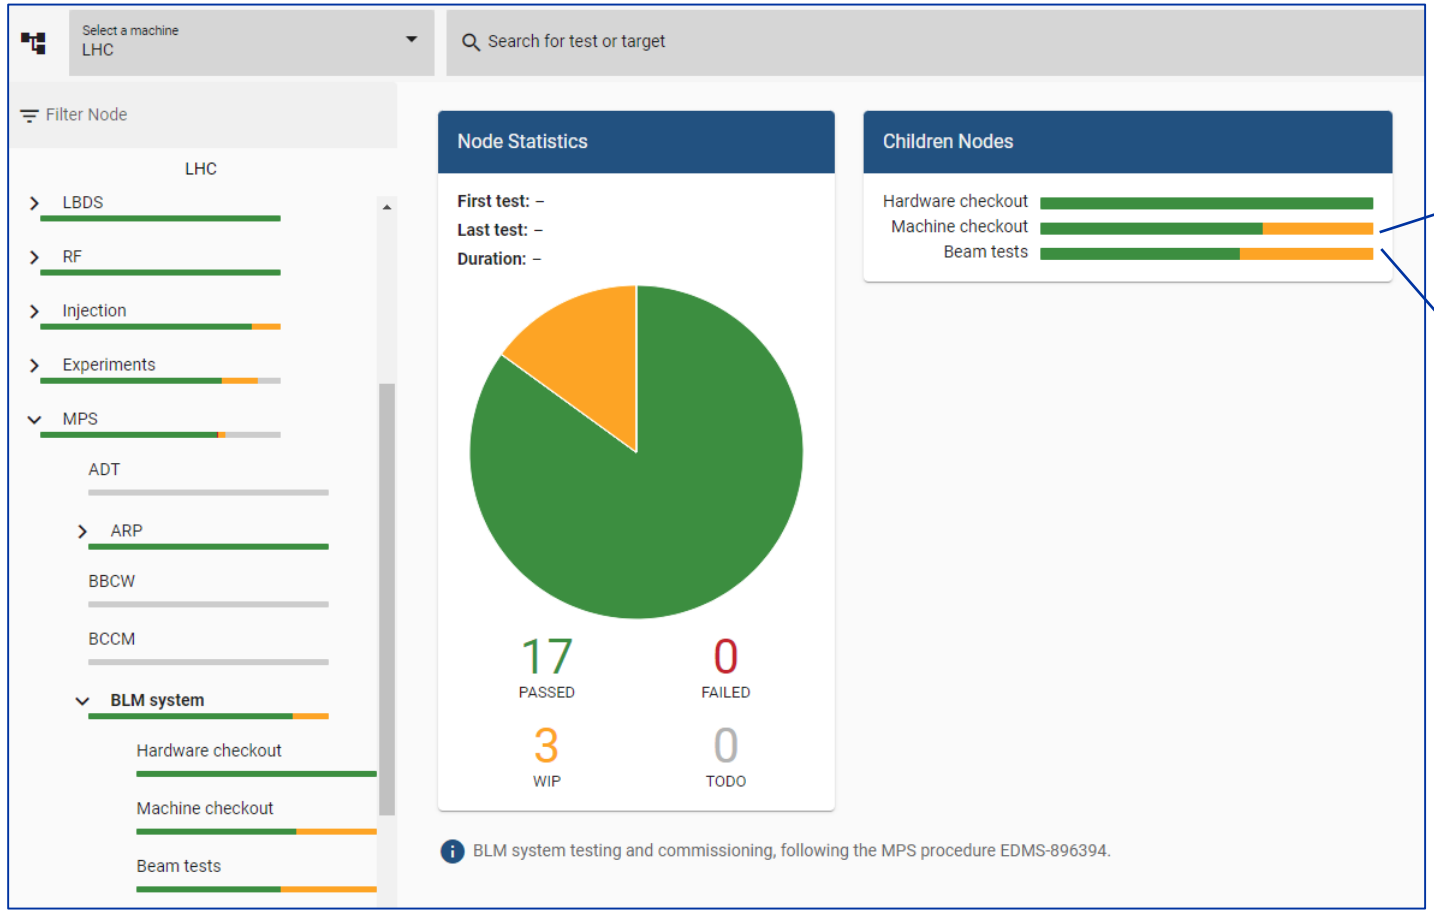

Machine Checkout

	Status
MPP/BLM - MC1 - User permit transmission	
MPP/BLM - MC2 - Threshold values change with energy	
MPP/BLM - MC3 - Missing HV detection and propagation to the SIS	

Three checks in SIS:
1- HV presence
➔ requests a dedicated test
2- Settings CRC in the electronics
3- Energy tracking

Beam Tests

	Status
MPP/BLM - BT1 - Interlock request functionality of the BLM crates	
MPP/BLM - BT2 - Interlock request functionality of the BLETC	
MPP/BLM - BT3 - Interlock request system latency	
MPP/BLM - BT4 - Test the interface of direct BLMs with the beam dumping system	
MPP/BLM - BT5 - Injection Interlock Inhibit functionality	

Already tested at the tunnel side (E. Effinger and N. Magnin).
To be done with beam.

To be done with ABT (needs 12 bunches).
Can be scheduled next week.

BLM Checklist

BLM LHC recent issues

1. **VME Power supply failure** on 12V (unused) in pt1
   ➔ Equipment replaced

2. **Temperature issue** (water cooling problem) in pt2 triggering optical link error interlocks.
   ➔ The alarms were not enabled, now all active
   ➔ CV has turned on the water flow to maximum
   ➔ Multiple cards have been exchanged to better resist higher temperature
   ➔ Need to change the temperature thresholds 5ºC lower (30➔~25ºC)

3. **Weak optical links** could trigger interlock if both redundant fail at the same time
   ➔ Replaced 11 BLETC at surface and a few BLECF in the tunnel (preventive maintenance)

4. **Sanity checks** issue blocking OP before injection:
   ➔ Workaround: always play the whole sequence not just a subset
   ➔ Issue in the VMW slave core: needs a BLECS FW upgrade (scheduled for YETS)

Blindable channels (inhibit at injection)

- Now present in all crates
- Disabled by default
- Acts only on maskable channels
- Programable timer per crate from injection warning (from BST + delay to be at injection)
- Inhibit the interlock output to BIC only (all running sums still active, all dump requests are logged)

Goals:

1. Test the feature with 12 bunches next week and measure the blind time needed
2. Select a first set of channels to blind (+adjust monitor factor)

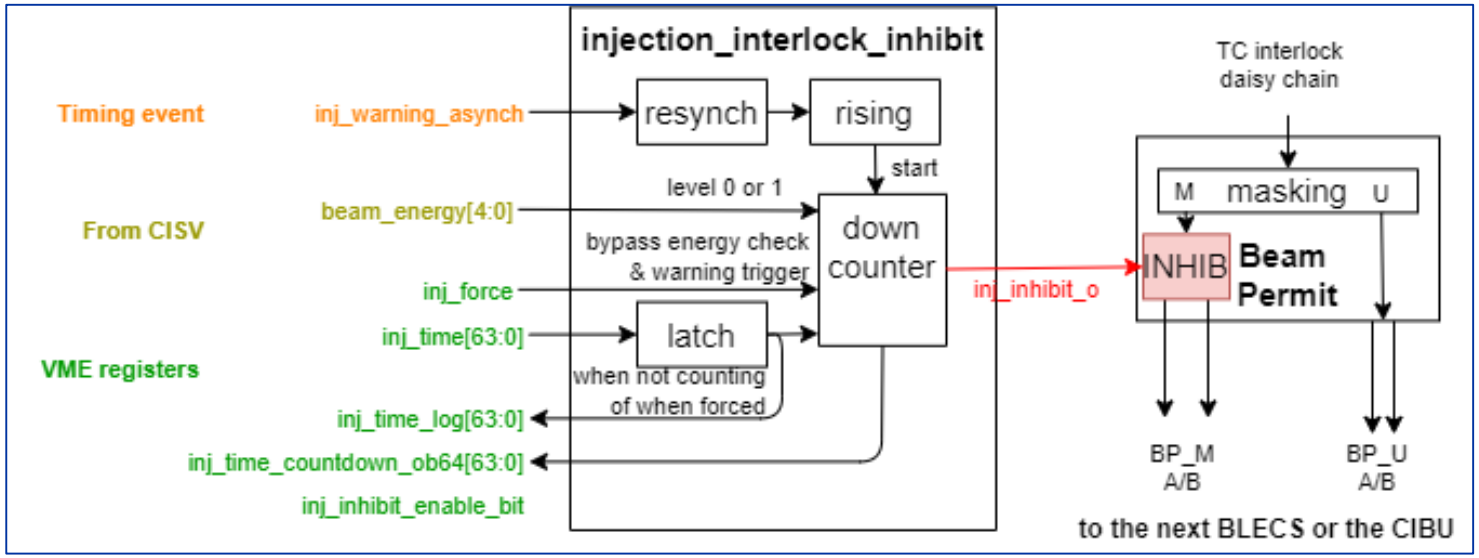

Injection Interlock Inhibit FW Implementation

BLM beam test principle

Two tests performed in parallel on the 18/05/2022 by OP:

- **Test 1: Interlock request functionality of the BLM crates**
  - Procedure written by BL and played by OP
  - Aim to trigger as many BLM crates as possible
  - 1 collimators/beam closed initially and opened using threading sequence.
  - Injection of pilot bunch – test at injection
- **Test 2: Interlock request system latency**
  - Latency must be less than 3 LHC turns (89 μs each)
  - Post-analysis performed by BI-BL from NXCALS (automatized with a Python script)

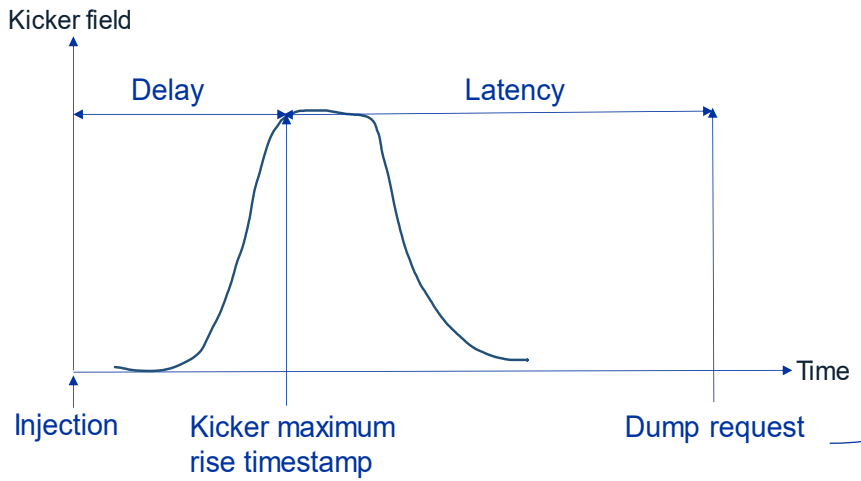

Possible to perform both tests in parallel

Latency calculated for each triggered crate

Procedure in EDMS

BLM beam test result

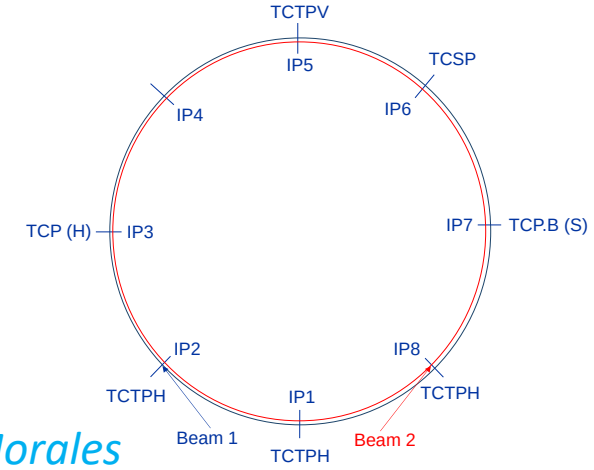

Courtesy S. Morales

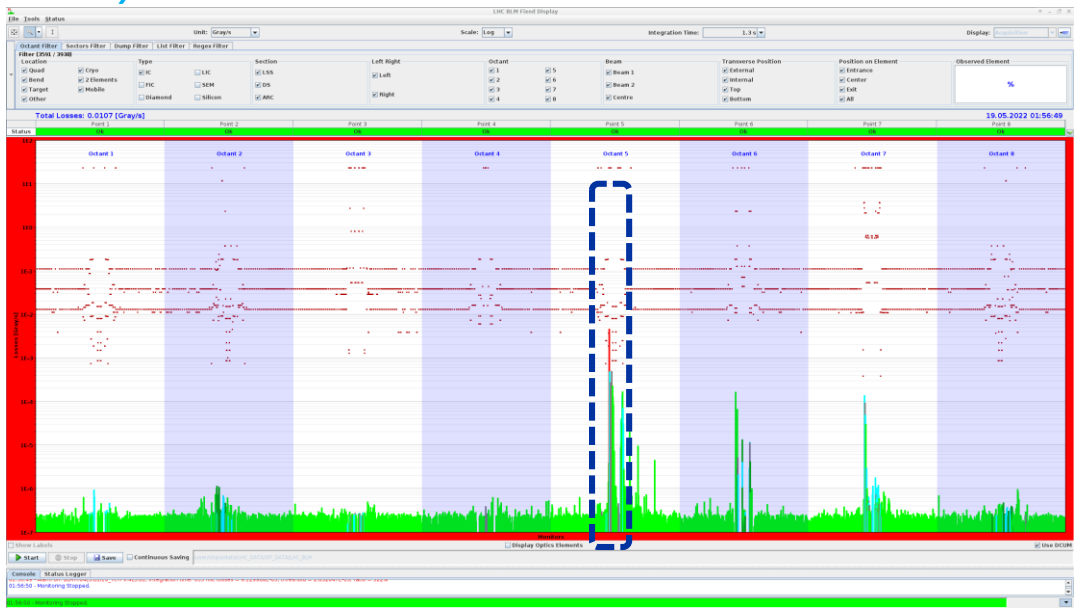

Example: Triggering of B1 Dump in IP5 at TCTPV

- Selection of collimator orientation arbitrary
- Same collimator type per point for B1 and B2
- Most dumps from BLM central crate, maskable channels (OK)
- BLM latency below 3 LHC turns – **OK**

Latency	B1	B2	Combined
(μs)	BLM_BIC	BLM_BIC	
IP1	129	127	✓
IP2	47	113	✓
IP3	78	175	✓
IP4	-	-	-
IP5	69	92	✓
IP6	85	60	✓
IP7	125	105	✓
IP8	128	183	✓

Includes transmission time of the signal through cabling from detectors to CIBUS (few km)

FW & SW changes summary

FW

On the 4 optical links reception chains on the surface processing board:

1. Add input delay constraints for all data lines
2. Improve the clock domain crossing mechanism

All the rest of the HDL code remains the same (as v1.1.7).

This new firmware v1.2.0 is deployed and tested with beam.

SW

To avoid losing CTIM events (XPOC and PM missing data) because of CPU high activity:

1. CTRP IRQ priority increase
2. FESA RT thread priority rescaled
3. In the future the CPU upgrade would give more margins (profiling & statistics under development)

Annex

More details regarding FW & SW changes

FW Changes

Optical link reception improvement

LHC BLM Architecture

- 2 redundant optical links from the tunnel to the surface electronics

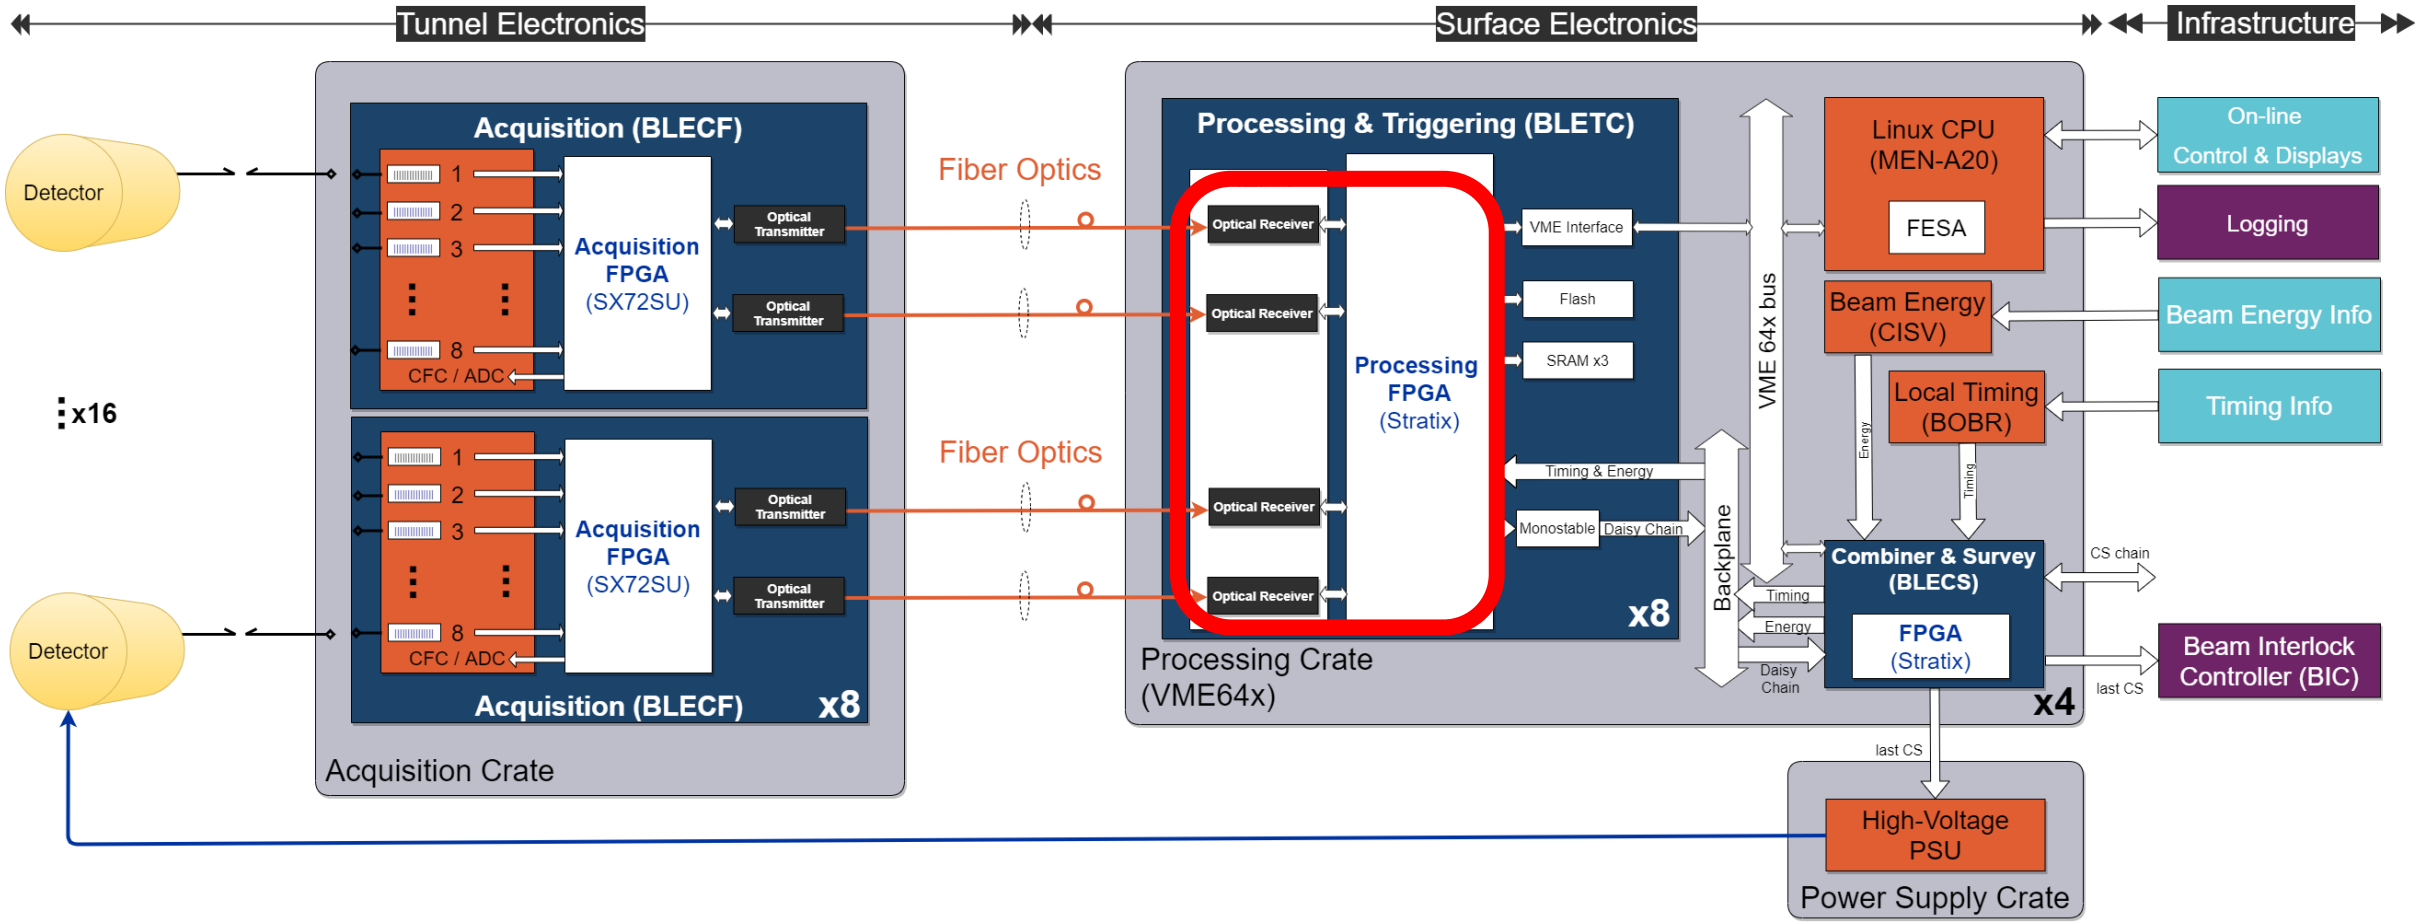

Optical links Architecture

- **FW update in the BLETC (Threshold comparator) FPGA**
- **Only in the RCC (Receive, Check and Compare) block**

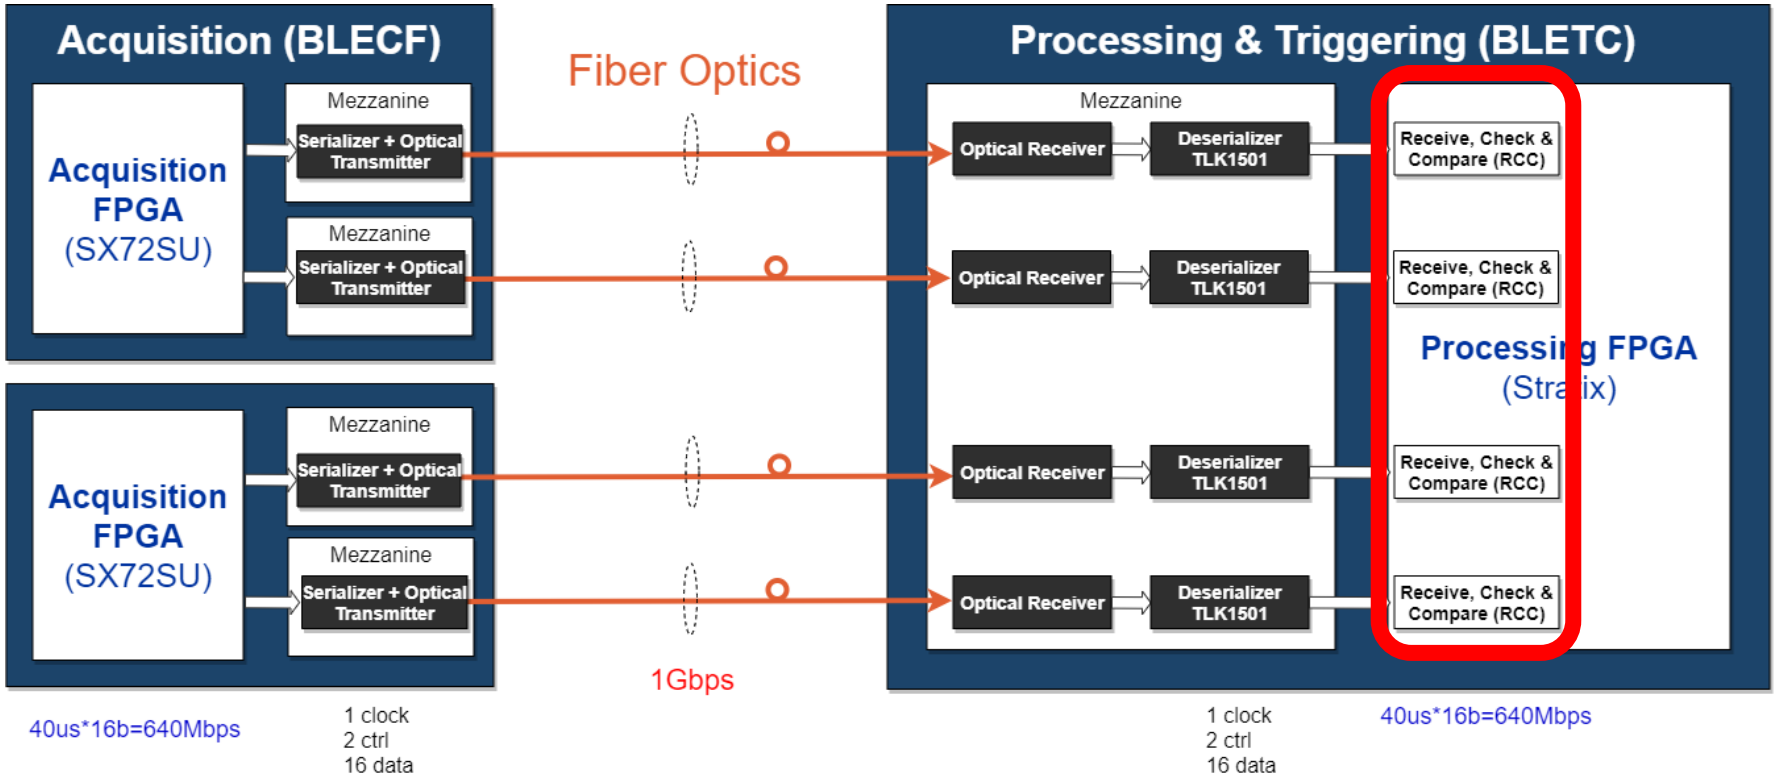

Optical Link Data Reception

- **Redundant link**
- **Protected by CRC32**
- **Frame ID counter**

CID (card identity number)	
STATUS 1	
STATUS 2	
Count 1	ADC 1
ADC 1	Count 2 / ADC 2
ADC 2	Count 3
ADC 3	Count 4
Count 4	ADC 4
Count 5	ADC 5
ADC 5	Count 6 / ADC 6
ADC 6	Count 7
ADC 7	Count 8
Count 8	ADC8
FID (frame identity number)	
DAC1	DAC2
DAC3	DAC4
DAC5	DAC6
DAC7	DAC8
CRC	
CRC	

40us frame
*(20*16bits word)*

- **TLK1501 deserializer**
  - 40MHz clock
  - 2 control bits: StartOfFrame (SOF) + RestOfFrame (ROF)
  - 16 data bits

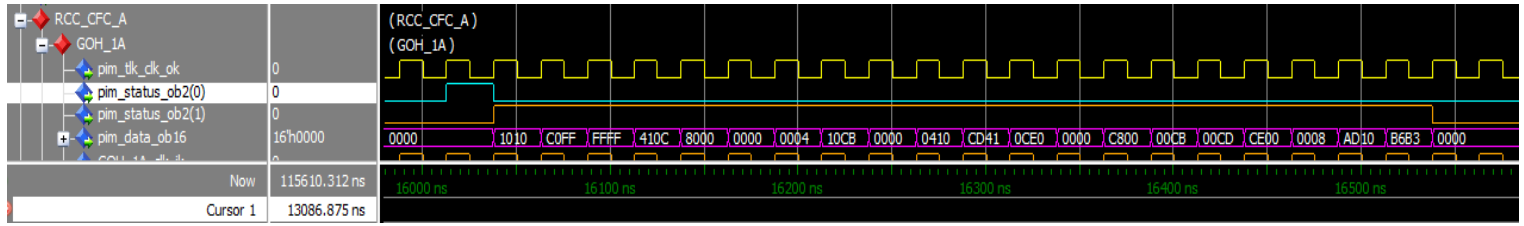

Simulation example

- **Each of the four TLK1501 generates its own 40MHz**
- **The TLK1501-FPGA lines not skew compensated (neither on the carrier nor on the mezzanine v4.0)**

Optical Link Data Reception

- Most of the installed mezzanines are v4.0
- Skew compensation introduced from mezzanine v5.0

FPGA (Dab64x side) EDA-00998-V3			TLK (Mezzanine side) EDA-00780-V3		
PIM C (1B) = LINK_1 CFC_B	length mm	Jn4	J1	length	TLK1501_IC2
PIM_StatusC[0]	174	PIN_F1 -to PIM_IO[1]	1	15	RX_2ER
PIM_StatusC[1]	173	PIN_F2 -to PIM_IO[2]	2	14	RX_2DV
PIM_DataC[15]	160	PIN_F4 -to PIM_IO[3]	3	20	RX_2D-15
PIM_DataC[14]	155	PIN_G1 -to PIM_IO[4]	4	15	RX_2D-14
PIM_DataC[13]	163	PIN_G2 -to PIM_IO[5]	5	19	RX_2D-13
PIM_DataC[12]	160	PIN_G3 -to PIM_IO[6]	6	16	RX_2D-12
PIM_DataC[11]	160	PIN_G4 -to PIM_IO[7]	7	24	RX_2D-11
PIM_DataC[10]	170	PIN_G6 -to PIM_IO[8]	8	20	RX_2D-10
PIM_DataC[9]	175	PIN_H1 -to PIM_IO[9]	9	24	RX_2D-9
PIM_DataC[8]	164	PIN_C2 -to PIM_IO[10]	10	20	RX_2D-8
PIM_TLK_clkC	189	PIN_H3 -to PIM_IO[11]	11	28	RX_2CLK
PIM_DataC[7]	153	PIN_H4 -to PIM_IO[12]	12	27	RX_2D-7
PIM_DataC[6]	165	PIN_H5 -to PIM_IO[13]	13	29	RX_2D-6
PIM_DataC[5]	161	PIN_H6 -to PIM_IO[14]	14	25	RX_2D-5
PIM_DataC[4]	189	PIN_H7 -to PIM_IO[15]	15	33	RX_2D-4
PIM_DataC[3]	165	PIN_J2 -to PIM_IO[16]	16	29	RX_2D-3
PIM_DataC[2]	67	PIN_J3 -to PIM_IO[17]	17	35	RX_2D-2
PIM_DataC[1]	70	PIN_J6 -to PIM_IO[18]	18	33	RX_2D-1
PIM_DataC[0]	73	PIN_J7 -to PIM_IO[19]	19	41	RX_2D-0

1B = LK1 CFCB	total length mm
	188.71
	187.12
	179.63
	170.11
	182.20
	176.03
	183.98
	190.24
	198.90
	184.24
	217.19
	179.82
	193.78
	186.32
	221.51
	193.70
	102.33
	102.54
	113.82

Lines length for link 1 CFC-B

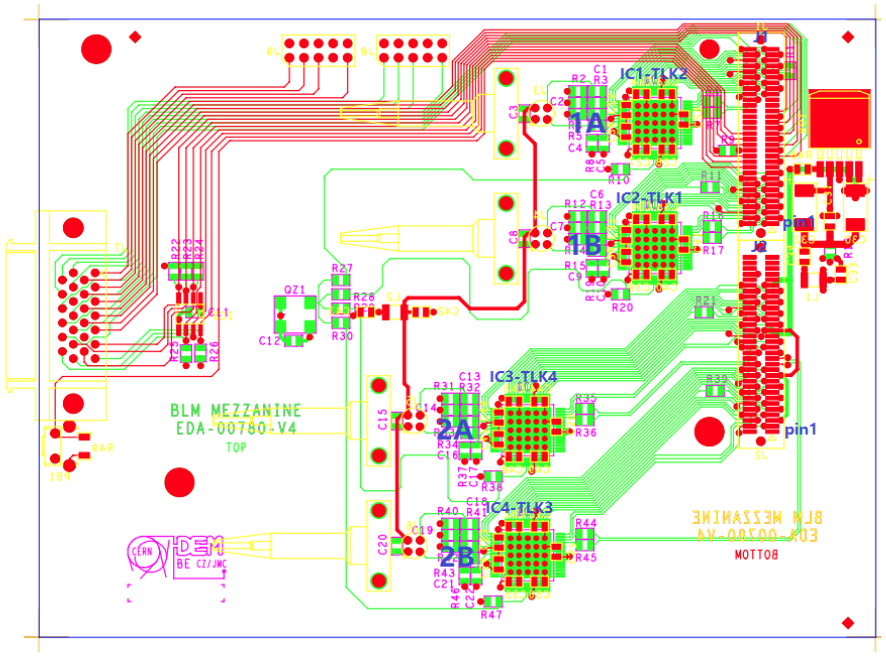

Mezzanine v4.0 layout

- Skew now partially compensated in the FPGA (input delay)

```
# Link 1B:
set_instance_assignment -name stratix_decrease_input_delay_to_internal_cells -to PIM_IO[11] on
set_instance_assignment -name decrease_input_delay_to_input_register -to PIM_IO[17]  off
set_instance_assignment -name decrease_input_delay_to_input_register -to PIM_IO[18]  off
```

Solution compatible with v5.0

- Timing Analysis constraints added per individual lines (set_input delay)

```
# min delay: Th (from TCLK docs) + data trace delay - clock delay
set_input_delay -clock_fall -clock {virt_cfcB_1k1_clk_40} -min [expr $TLK_TCO_min + $cfcB_1k1_ctrl_0_delay  - $cfcB_1k1_clk_delay] [get_ports {PIM_IO[1]}]
set_input_delay -clock_fall -clock {virt_cfcB_1k1_clk_40} -min [expr $TLK_TCO_min + $cfcB_1k1_ctrl_1_delay  - $cfcB_1k1_clk_delay] [get_ports {PIM_IO[2]}]
set_input_delay -clock_fall -clock {virt_cfcB_1k1_clk_40} -min [expr $TLK_TCO_min + $cfcB_1k1_data_0_delay  - $cfcB_1k1_clk_delay] [get_ports {PIM_IO[3]}]
set_input_delay -clock_fall -clock {virt_cfcB_1k1_clk_40} -min [expr $TLK_TCO_min + $cfcB_1k1_data_1_delay  - $cfcB_1k1_clk_delay] [get_ports {PIM_IO[4]}]
```

Data reception in HDL

- One clock domain per link for the frame detection and CRC check
- One BRAM per link for clock domain crossing
- A common error check and selection per redunded link
- Main modification => use synchronous SOF
  - Avoid asynchronous reset of counters in the main 40MHz domain
  - Reduces the number of resynchronization registers (anti-metastability)
  - Tested in simulation, in the lab for several days, then deployed in LHC on 17/05/2022

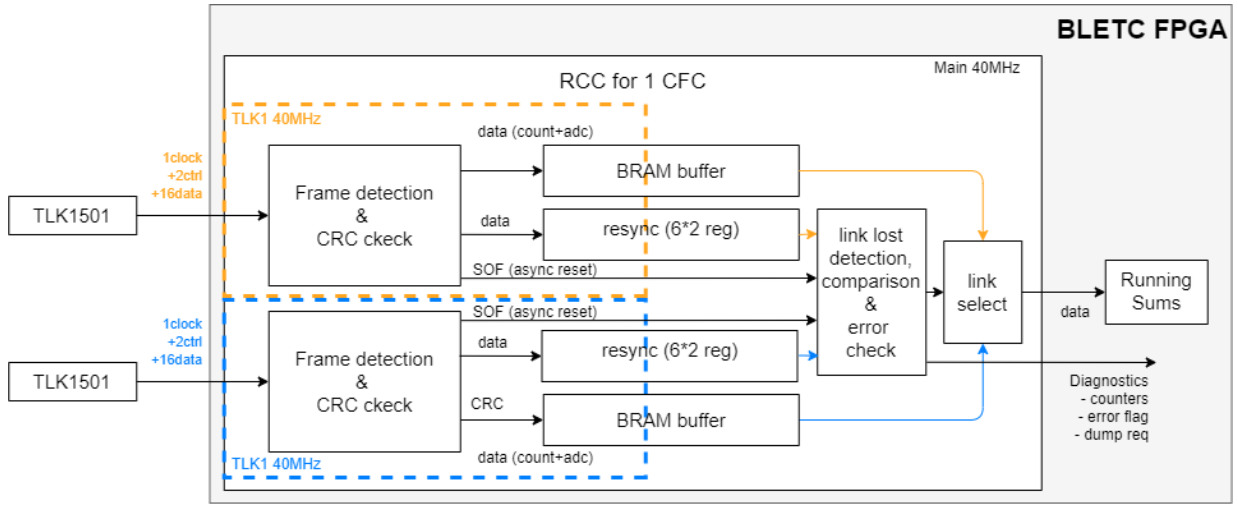

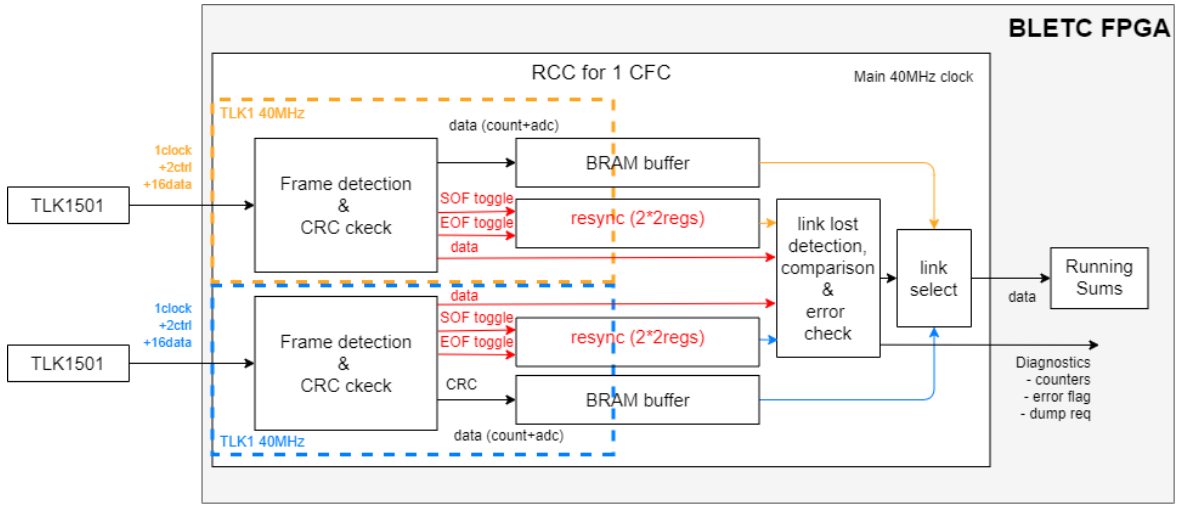

SW Changes

FESA + OS IRQ management

Issue Description

Issue:

Millisecond event (PM and XPOC) missing from time to time in some random BLM crates.
JIRA-TIMING-4011
Other systems (SY-EPC-CCS) seemed to face kind of the same problem
JIRA-TIMING-4027

Root cause:

- The CTR loses some CTIM events (separated by 125us) when the CPU activity is high even if the CTR event queue is not full.
- Was reproduced by BE-CEM&CSS
- No notification of interrupt loss in CTR driver (gateware problem?)
- The issue appeared after the migration to FESA3, where priorities are defined on a range [0-100] and no more by category/offset (LOW|NORMAL|HIGH / -2|-1|0|+1|+2)

Special thanks to Marine Gourber-Pace, Michel Arruat, Frederic William Hoguin & Stephane Deghaye

Workaround 1: Priorities

Increase the priority of the CTRP IRQ kernel thread

- JIRA-BIBML-2373
- This solution was implemented on the 01/06/2022
- It seems to solve the problem for now
- Need to be confirmed by N. Magnin (SY-ABT-BTC) with his XPOC logs

Implemented solution:

1. CTRP IRQ increase to 88 (instead of 87)
2. FESA RT thread priority rescaled in range [0:25] (instead of [0:70]), but the FESA precedence remains (same behavior)

```
ps -eLo comm,rtprio| grep irq
irq/9-acpi                87
irq/23-ehci_hcd           87
irq/23-uhci_hcd           87
irq/8-rtc0                87
irq/28-eth0               87
irq/19-i801_smb           87
irq/14-ata_piix           87
irq/15-ata_piix           87
irq/16-serial             87
irq/17-vme_brid           87
irq/16-ctrp.02:           ~~87~~ 88
```

CTR IRQ priority increased

```
ps -eLo comm,tid,rtprio | grep BLMLHC_DU_M
BLMLHC_DU_M      2990     ~~70~~ 25
BLMLHC_DU_M      2993     ~~70~~ 25
BLMLHC_DU_M      2994      1
BLMLHC_DU_M      3005     ~~70~~ 25
BLMLHC_DU_M      3006     ~~70~~ 25
BLMLHC_DU_M      3015      5
BLMLHC_DU_M      3016      1
BLMLHC_DU_M      3017      7
BLMLHC_DU_M      3018     10
BLMLHC_DU_M      3019      9
BLMLHC_DU_M      3020      6
BLMLHC_DU_M      3021      8
BLMLHC_DU_M      3022      5
BLMLHC_DU_M      3023      8
BLMLHC_DU_M      3024     11
BLMLHC_DU_M      3032     10
BLMLHC_DU_M      3033      1
BLMLHC_DU_M      3034      6
…
```

FESA RT threads rescaled

Special thanks to Stephen Jackson

Workaround 2: CPU upgrade

Move the LHC BLM systems to MenA25

The CPU upgrade would drastically reduce the CPU activity

- 4 cores instead of 2
- More RAM
- Faster MBLT VME data throughput

➔ A continuous profiling on operational A20s is developed and will be compared to A25 in the lab.

The upgrade of BLM CPUs could be done during the next YETS

- The new CPU behaviour will be first fully characterised in the lab
- The exchange can be done quite quickly and easily
- Will give more margins in processing time
- Will ease the future system maintenance and upgrades
- 30 CPUs (27 BLM FECs in LHC + 2 lab crates + 1 spare for piquet)

MENA20

- Launch date Q3-2006
- Intel Core2 Duo L7400 (1.5GHz)
- 2 cores / 2 threads
- 64-bit VMEbus (TSI148)
- 1 GB RAM
- 1 Gb Ethernet in front
- 2 PMC/XMC slots
- Serial console in front

MENA25

- Launch date Q2-2016
- Intel Pentium D1519 (1.5GHz, turbo 2.1GHz)
- 4 cores / 8 threads
- 64-bit VMEbus (open-source bridge)
- 8 GB DRAM
- 1 Gb Ethernet in front
- 2 USB 3.0
- 1 PMC/XMC slot
- Serial console in front

Thank you for your attention!

Questions?

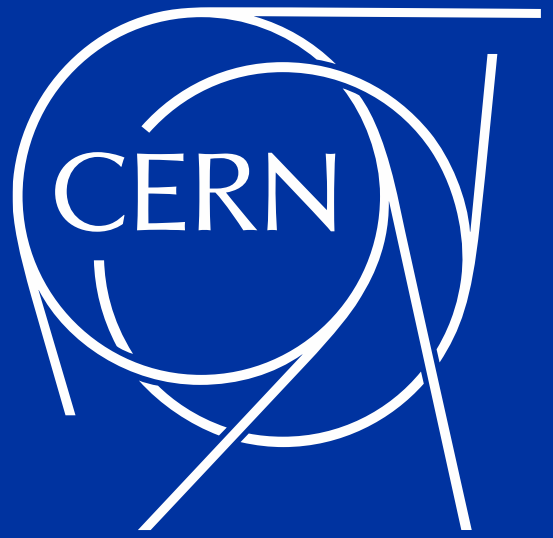
CERN