ALICE

SRO@ALICE

-

Streaming readout Workshop SRO-XI

Filippo Costa
for the ALICE O^2/FLP

Filippo Costa

- Responsible for the detector readout activities in ALICE.
- ALICE Software Release Coordinator.
- Software and firmware developer for the readout card (PCIe card with FPGA).

INTRODUCTION

The presentation will describe the operations of the ALICE experiment during the first HI data taking in 2023.

(Description of the new ALICE readout concept was presented during past version of the workshop in 2021 https://indico.phy.ornl.gov/event/112/contributions/478/)

The first part describes the major ALICE O^2 components and how the experiment is controlled from a single central point.

The central part of the presentation gives details concerning dataflow, different processes running and system performance.

The last part of the presentation is dedicated to describe the major challenges in operating such a large system running at high data rate.

MOTIVATION

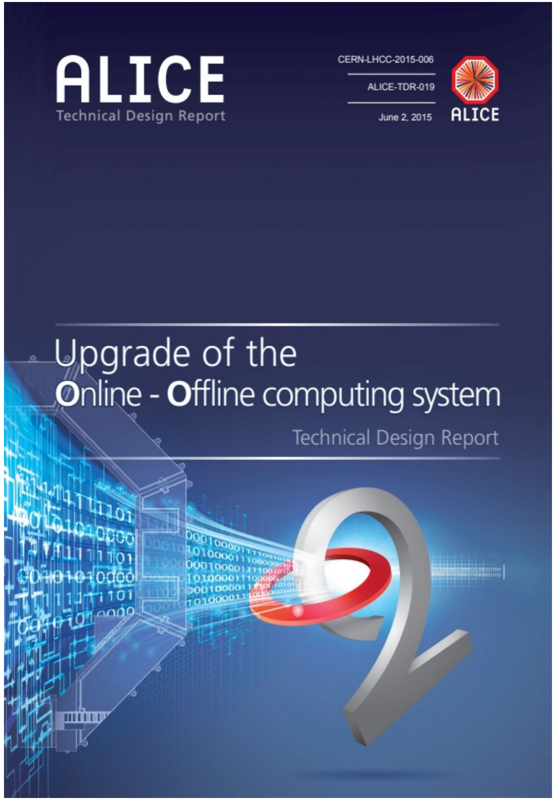

ALICE upgrade is based on the LHC running conditions after **LS2** which will deliver Pb–Pb collisions at up to L = 6 · 1027 cm−2 s −1 , corresponding to an interaction rate of **50kHz**.

In order to keep up with the 50kHz interaction rate, the TPC will also require the implementation of a continuous read-out process to deal with event pile-up and avoid trigger-generated dead time.

The resulting data throughput from the detector has been estimated to be greater than **3TB/s for Pb–Pb** events, several orders of magnitude more than in Run 1/2.

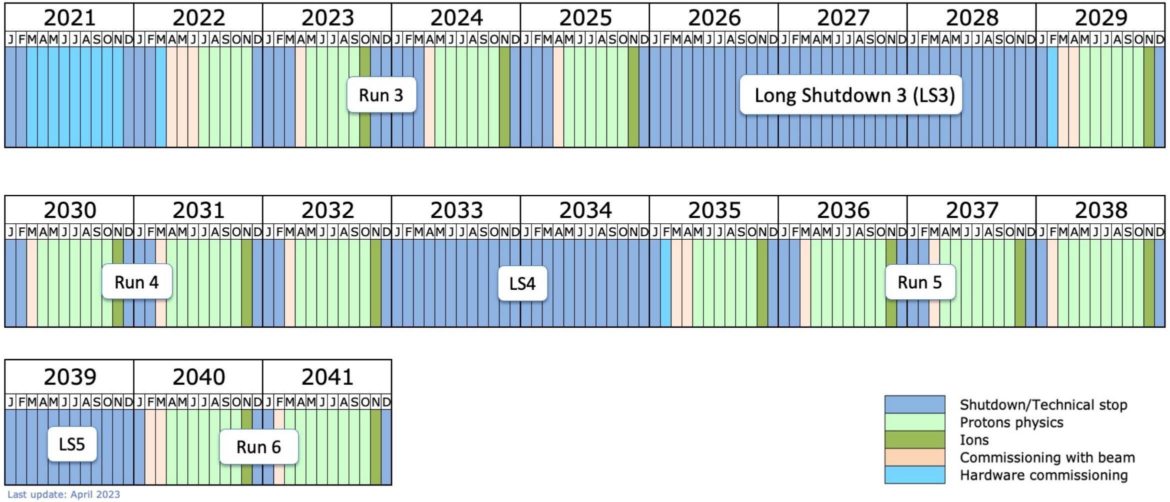

ALICE O^2 UPGRADE for the LHC RUN 3

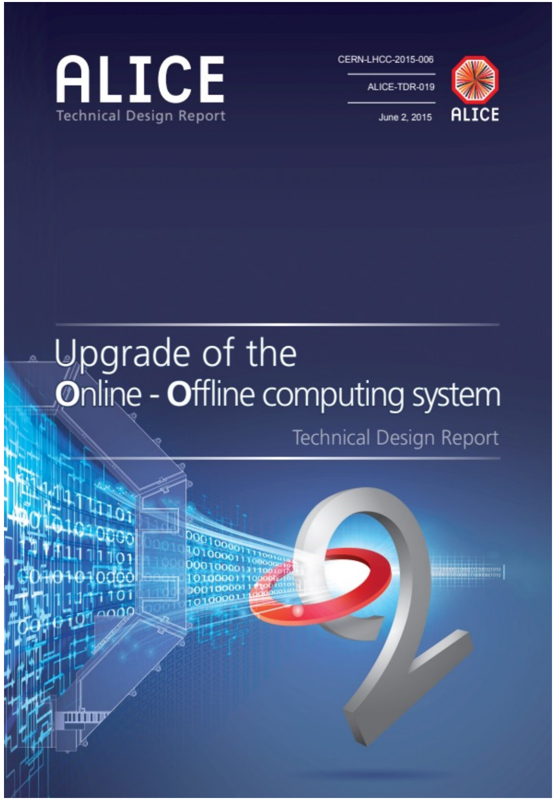

A new **ALICE Online and Offline (O^2)Computing** has been developed, the **ALICE O^2.**

The new O^2 facility provides:

- Continuous readout.
- Synchronous and asynchronous reconstruction.
- Two different categories of computing nodes, corresponding to the two data aggregation steps.
- GPU data processing.

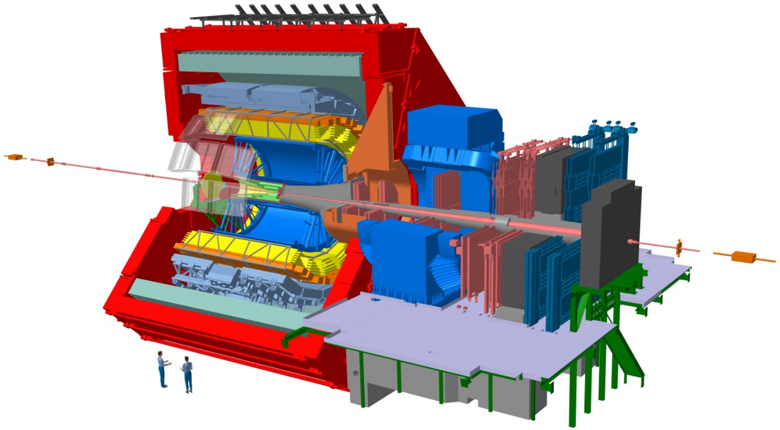

ALICE O^2 FARM

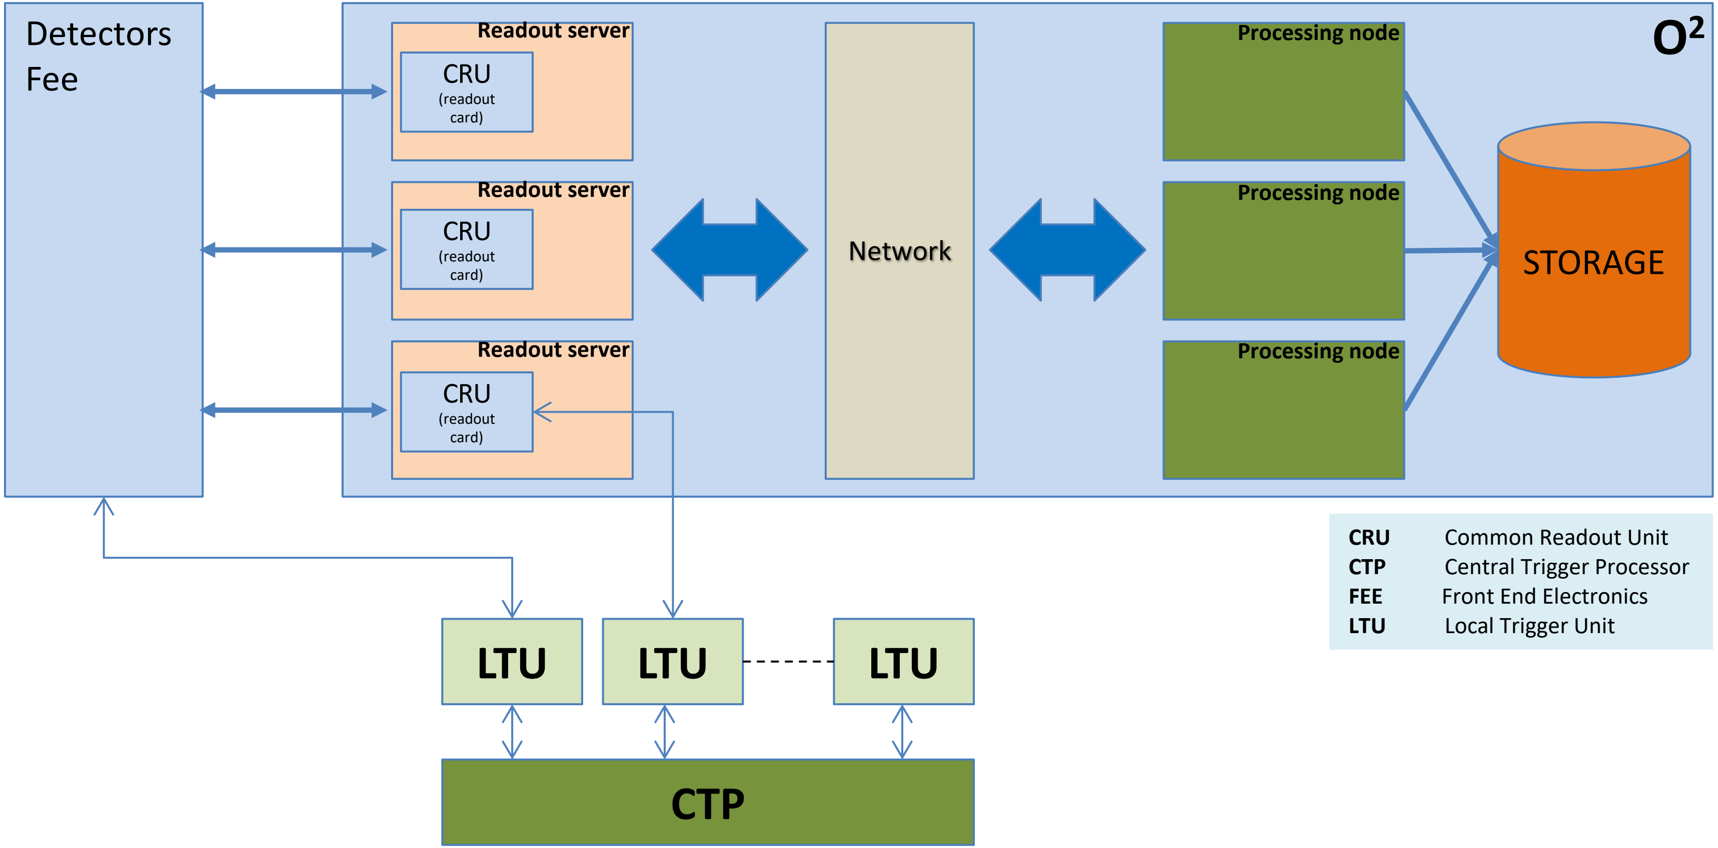

CRU	Common Readout Unit
CTP	Central Trigger Processor
FEE	Front End Electronics
LTU	Local Trigger Unit

ALICE O^2 FARM

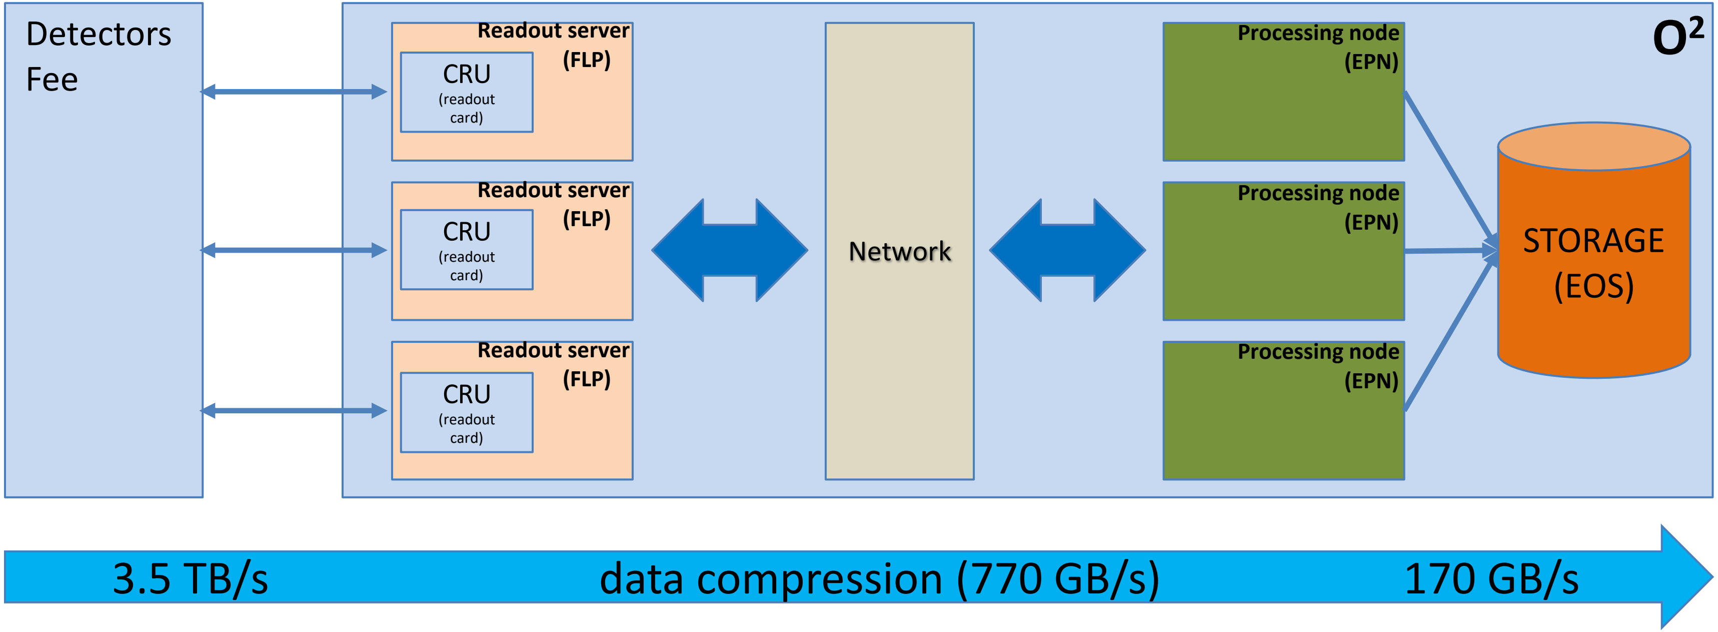

- 8000 links connect the detectors to O^2 farm. Total data rate > 3 TB/s
- 200 FLPs receive data from the detectors FEE. Total data rate ~770 GB/s
    - 500 readout cards.
- 350 EPN collect and store data on EOS. Total data rate 170 GB/s

CRU	Common Readout Unit
FLP	First Level Processors
EPN	Event Processing Nodes
EOS	EOS Open Storage

CHALLENGES:
- Control from a central system
- DATA processing (no backpressure)
- DATA quality (must be good)
- High efficiency

AliECS global overview

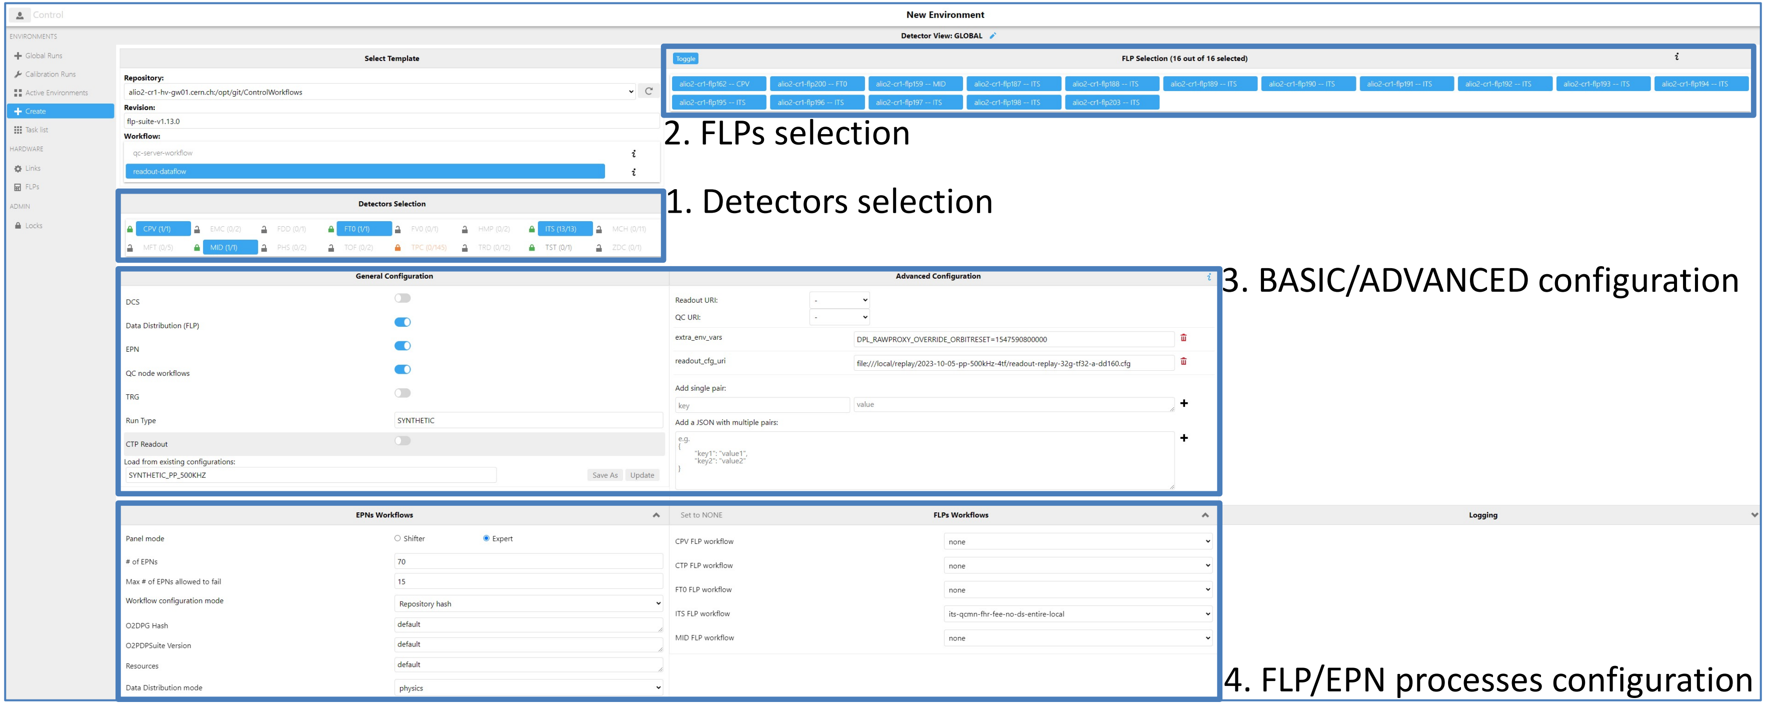

ALICE CONTROL SYSTEM (AliECS)

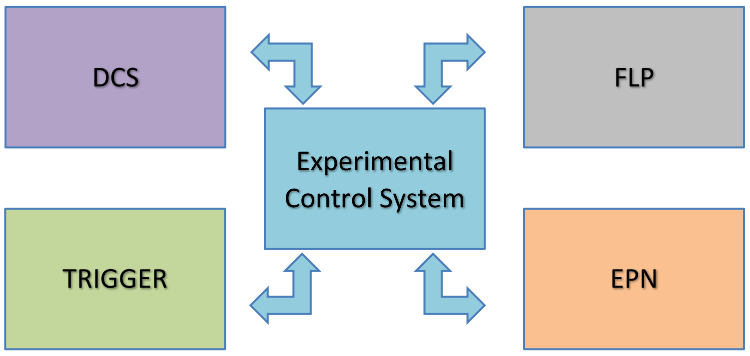

AliECS core

Control software to operate all the central components.

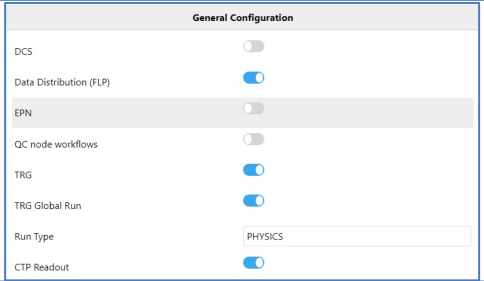

AliECS GUI

Graphic Interface to control the **CORE.**

It allows the shifter to operate the system and to start data taking with a few mouse clicks.

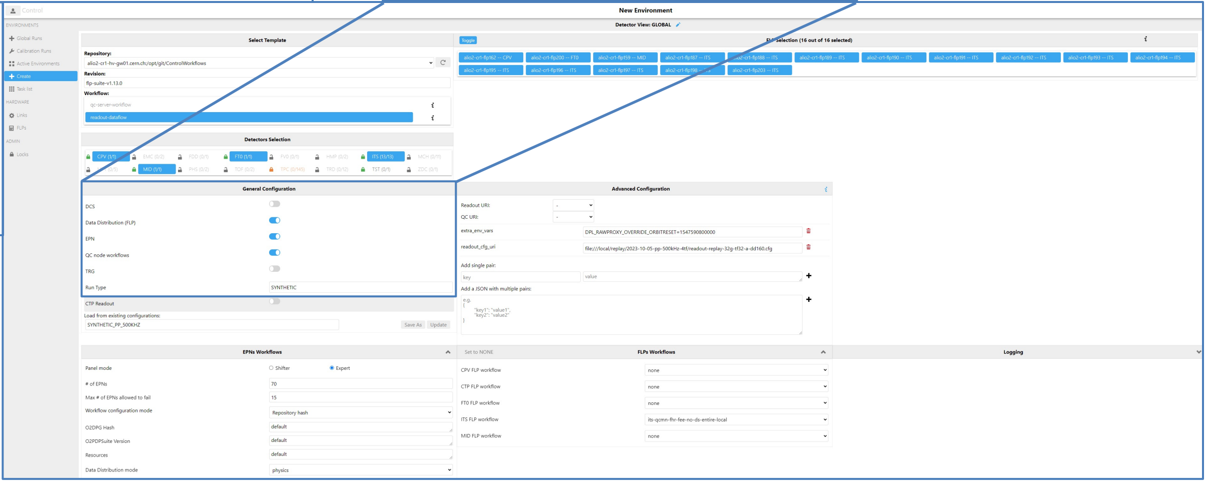

ENVIRONMENT CRATE/CONFIGURE

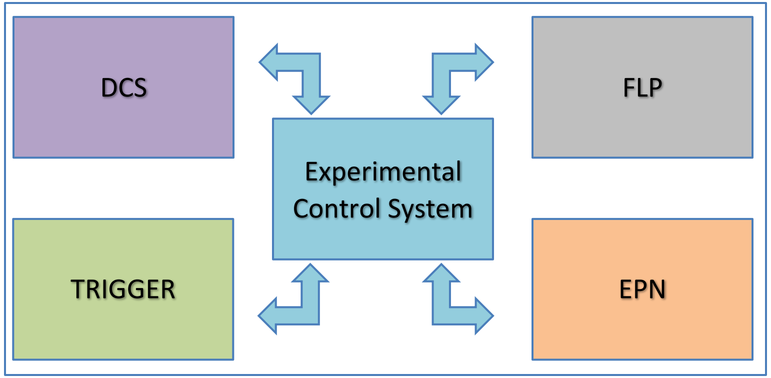

The ENVIRONMENT creation sets:

- list of detectors included in data taking,
- list of FLPs,
- number of EPNs,
- configuration of processes running on FLP/EPN.

All the operations are executed in parallel.

DCS : Prepare For Run. Detectors are configured via DCS.
TRIGGER : the list of detectors is sent to CTP to prepare the trigger configuration.
FLP : memory is allocated, readout cards are configured for data taking.
EPN : memory is allocated, detectors processes are started.

START (DCS SOR)

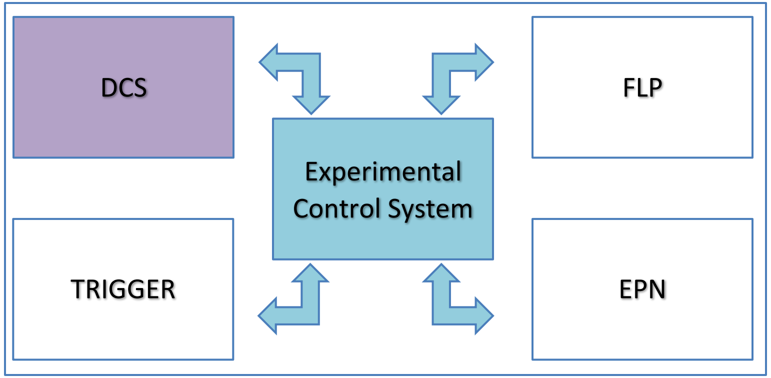

1. **DCS** : DCS SOR. Detectors that don’t have PFR are configured now

START (FLP EPN READY)

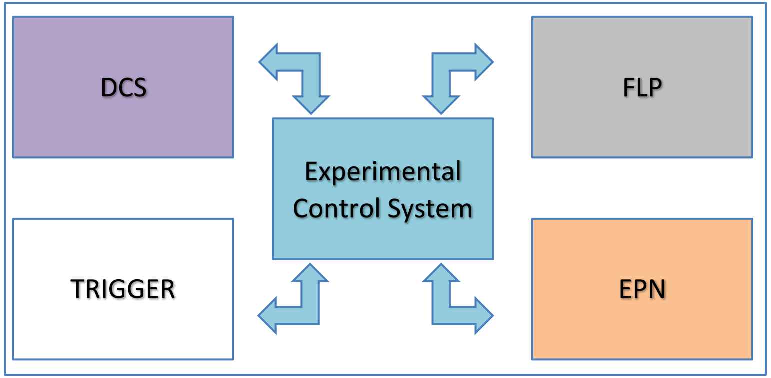

1. **DCS** : DCS SOR. Detectors that don't have PFR are configured now
2. **FLP** : DMA is enabled. **EPN** : ready to receive data from FLP

START (TRIGGER SOR)

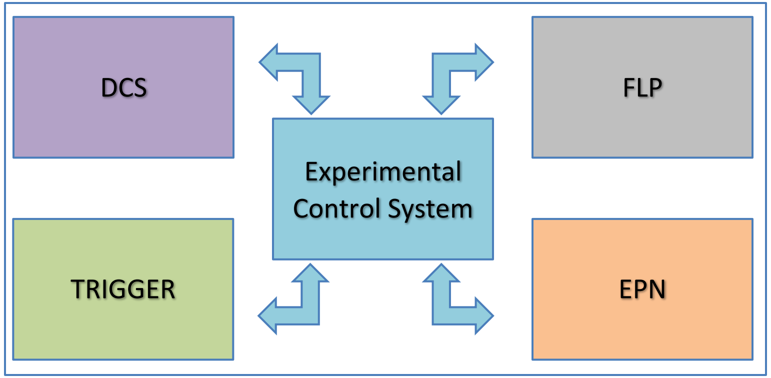

1. **DCS** : DCS SOR. Detectors that don't have PFR are configured now
2. **FLP** : DMA is enabled. **EPN** : ready to receive data from FLP
3. **TRIGGER** : Start of RUN trigger is sent to all the detectors

DATA QUALITY CONTROL?

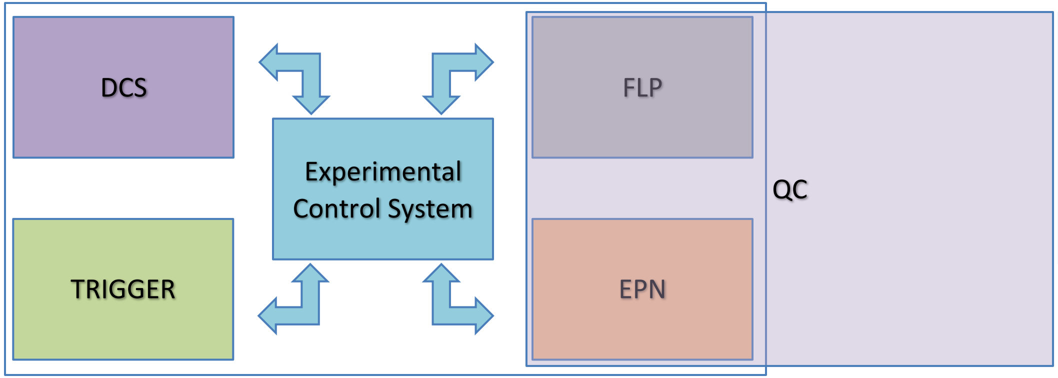

QC is controlled by ECS and it is connected to both FLP and EPN.

It receives data whenever a new RUN is started and it provides detector specific information on the data quality.

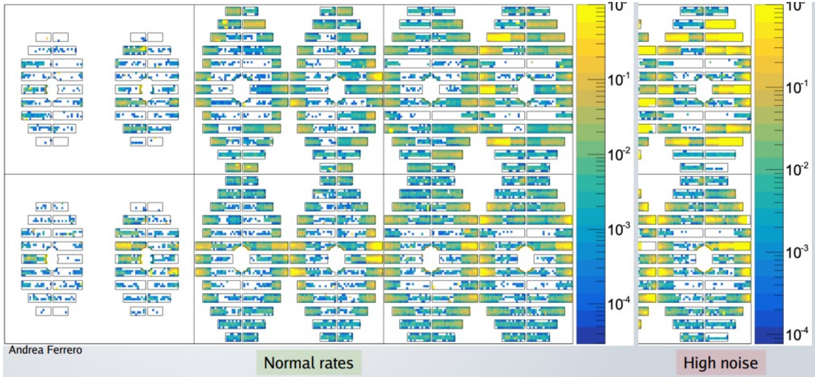

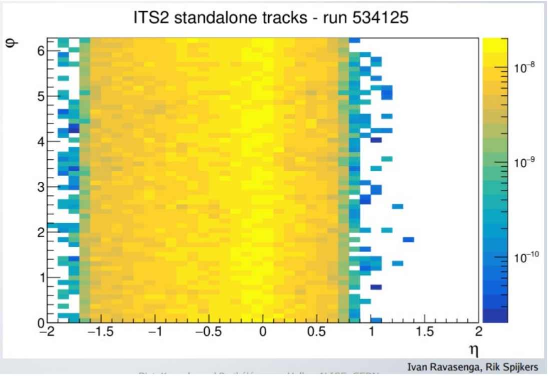

DATA FLOW- from detector to disk

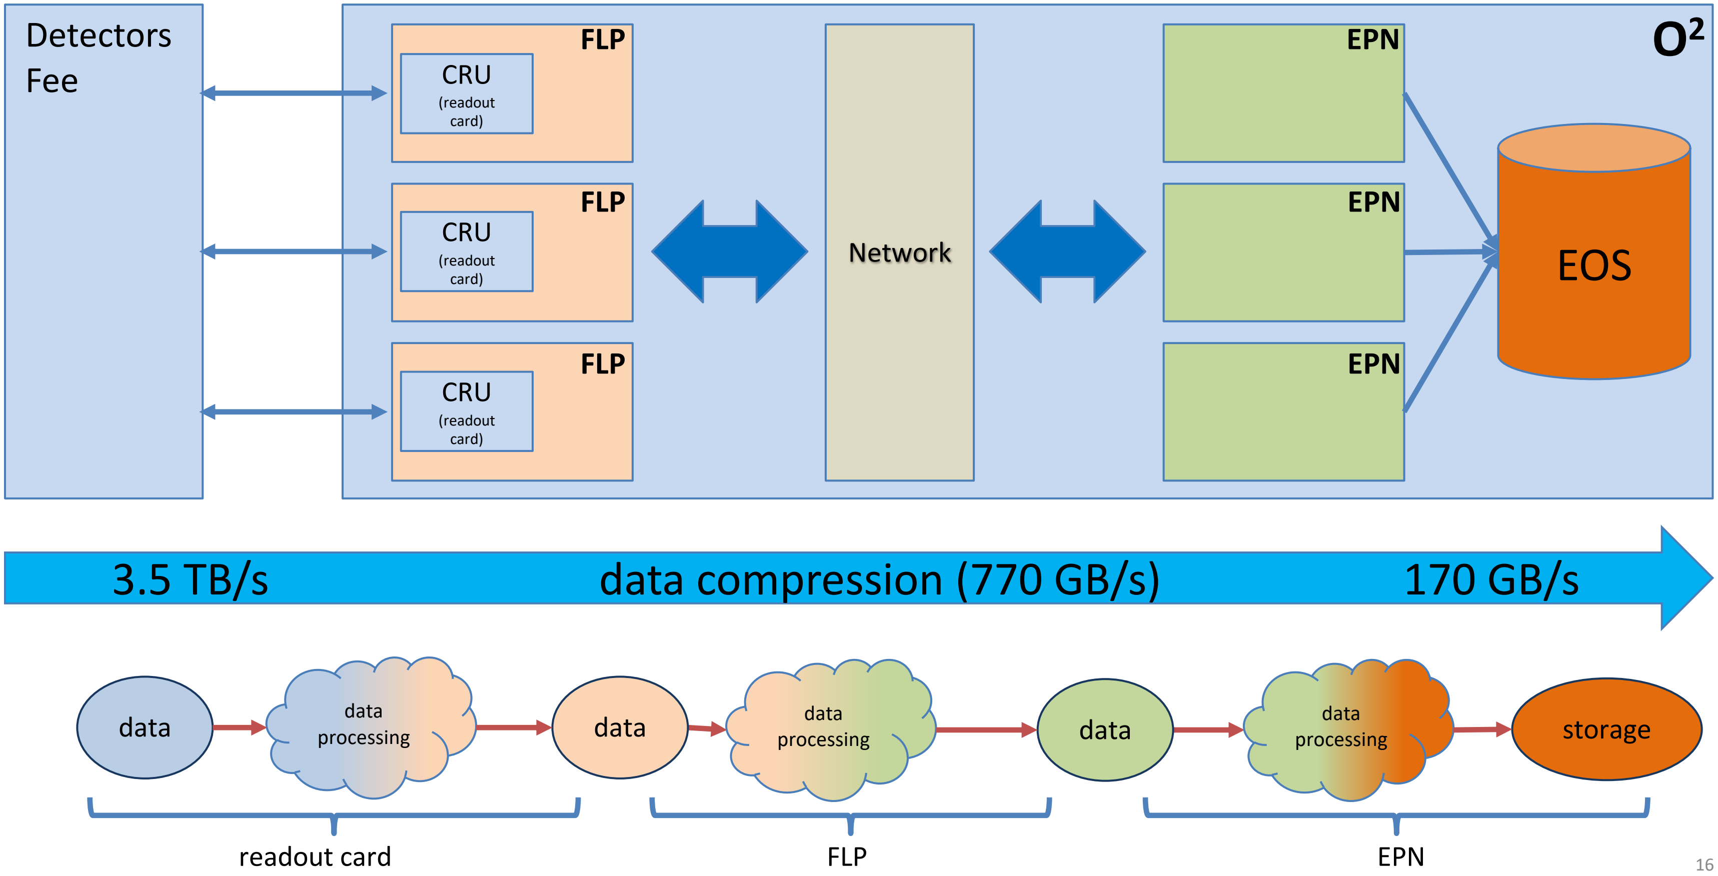

DATA PROCESSING (CRU)

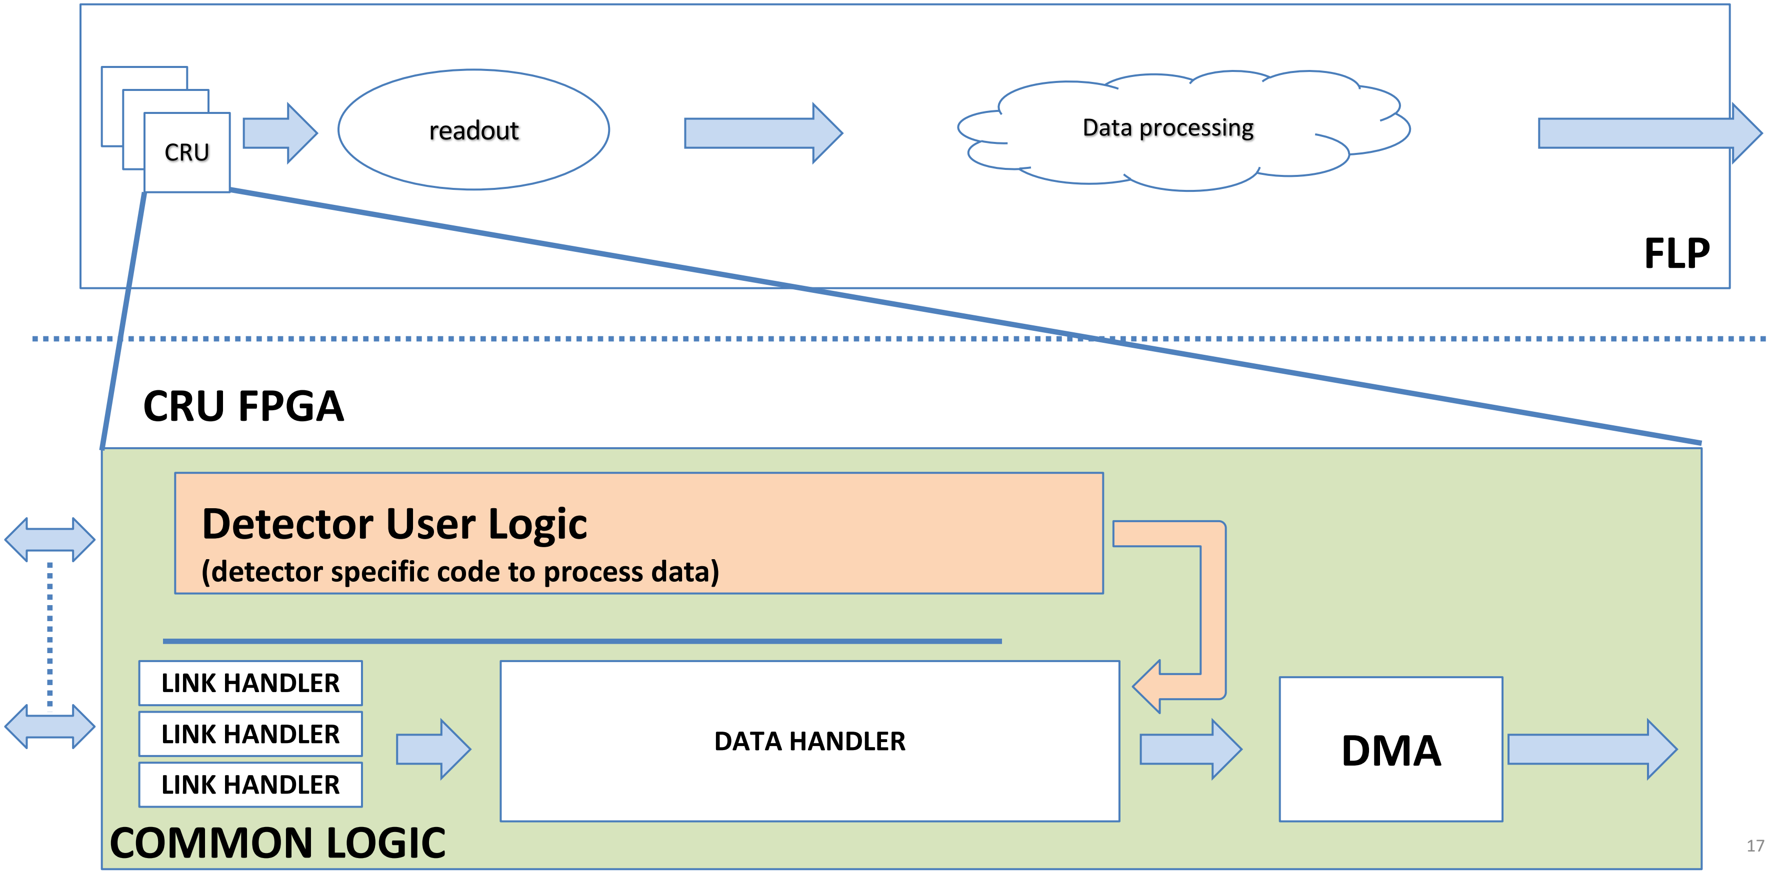

DATA PROCESSING (FLP)

ALICE

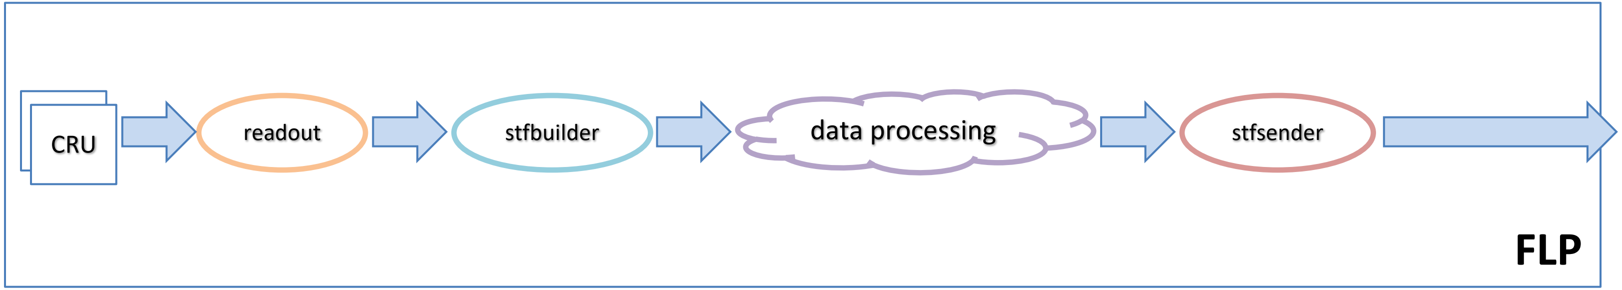

alio2-cr1-flp178

Name	PID	Locked	Status	State
readout	4449		ACTIVE	CONFIGURED
stfbuilder	4450		ACTIVE	CONFIGURED
stfsender	4451		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-internal-dpl-clock	4457		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-readout-proxy	4463		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t0	4470		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t1	4477		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t2	4487		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t3	4495		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t4	4509		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t5	4508		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t6	4517		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t7	4522		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t8	4527		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t9	4557		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t10	4565		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t11	4572		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t12	4580		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t13	4586		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-tof-compressor-0_t14	4603		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-Dispatcher	4614		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-qc-task-TOF-TaskRaw	4655		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-TOF-TaskRaw-proxy	4680		ACTIVE	CONFIGURED
jit-a6368d3a765462d9c2db1085fb5e5a128dfd9e1b-dpl-output-proxy	4711		ACTIVE	CONFIGURED

List of processes varies from detector to detector

DATA PROCESSING (EPN)

ALICE

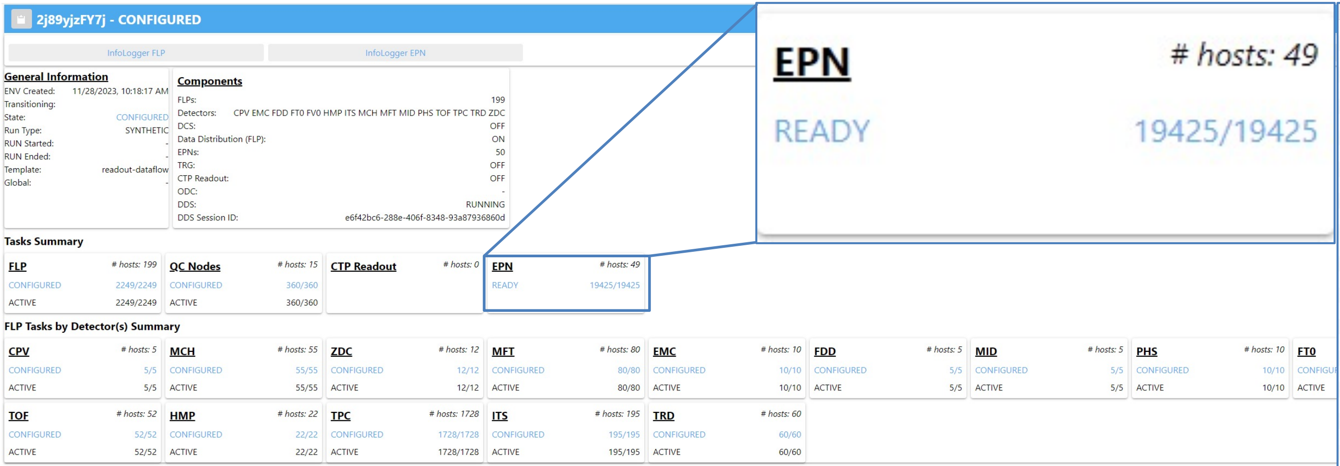

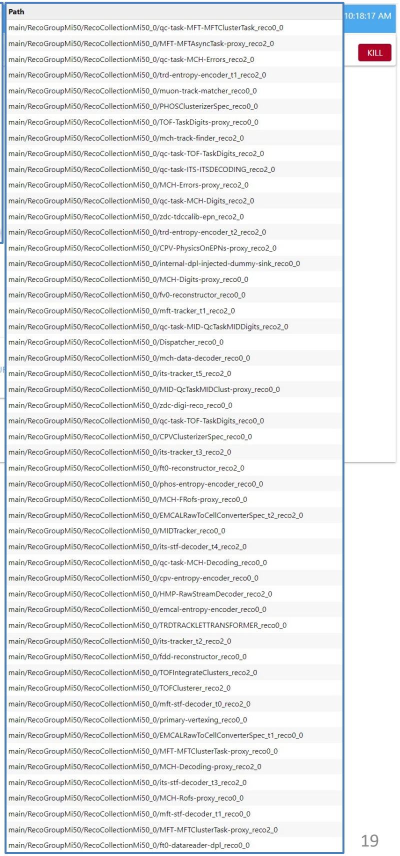

To be able to process data fast enough we had a processing farm consisting of a total of 350 EPN nodes and 2800 GPUs (Without GPUs, more than **2000 64-core servers** would be needed for online processing!).
All the EPNs have the same list of processes as there is no DETECTOR-EPN specific, but there are some dedicated to CALIBRATION.

- PHYSICS
- CALIBRATION
- SYNTHETIC/REPLAY

ALICE EVENT IDENTIFICATION

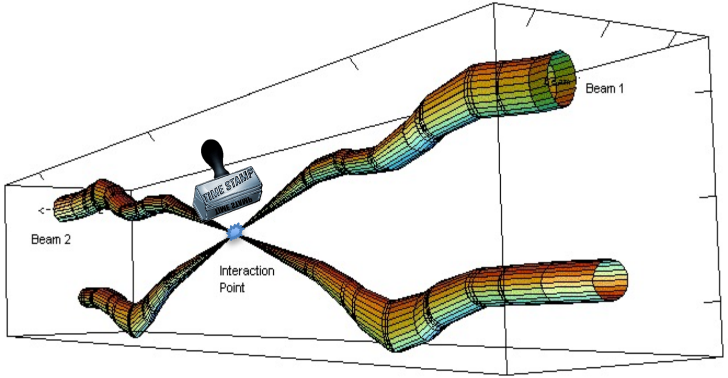

LHC has a global **MACHINE clock** that runs @ 40 MHz.
2 numbers are used to identify each interaction:

- **Bunch Crossing (BC)** : a number from 0 to 3563.
- **ORBIT** : a 32 bit number.

Every time BC overlaps, ORBIT+1.
The combination of the 2 gives a unique time stamp every ~4days

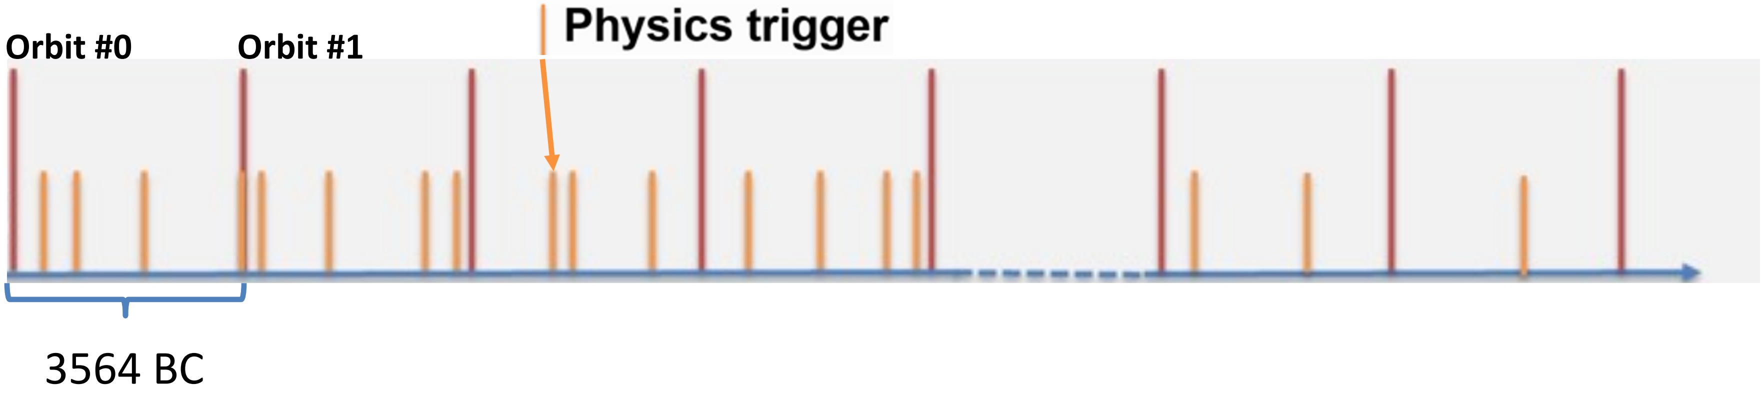

HB TRIGGER and TIME FRAME

LHC clock 40 MHz
3564 Bunch Crossing in 1 ORBIT
ORBIT rate ~10 KHz
Time Frame = 128 Orbits

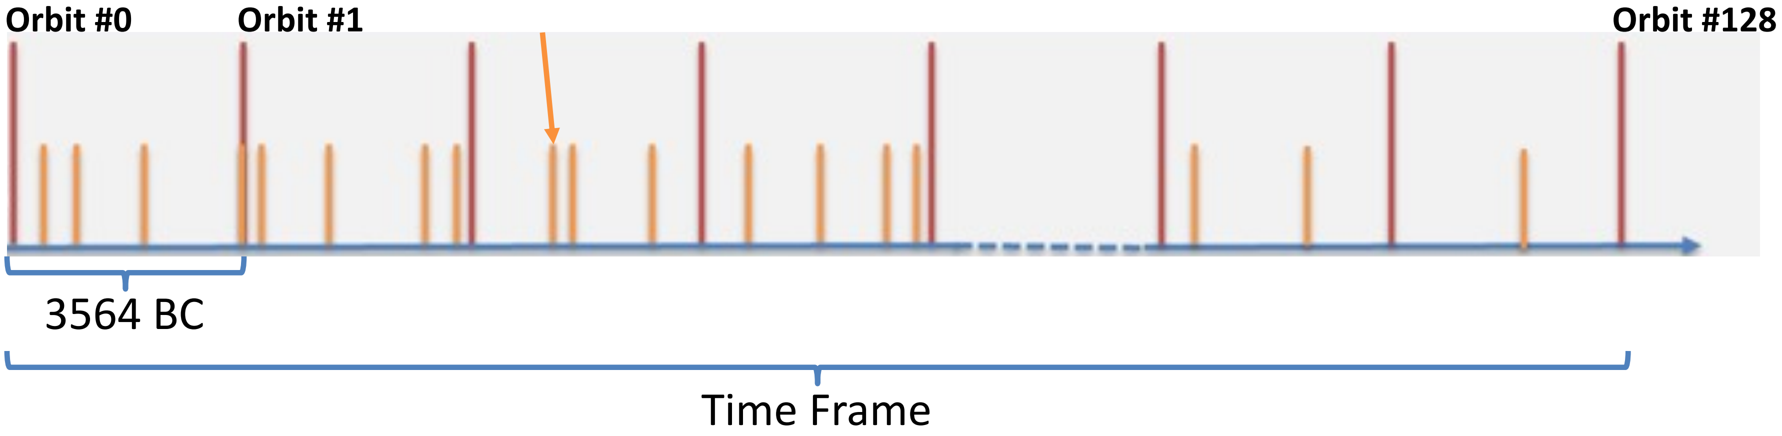

TF length, why it is important

ALICE

LONGER TF, less statistic we lose,
SHORTER TF, less memory we need.

- **Time Frame = 256 orbits**
  - This was done to optimize the data transfer and reduce the percentage of data at the boundary of each TF. Every TF reaches a different EPN.
- **Time Frame = 128 orbits**
  - EPN has 512 GB of RAM. We reduced the size as 256 required more memory than available. Even in this configuration we were on the edge of the memory (with only 30 GB of 512 GB left).
- **Time Frame = 32 orbits**
  - GPUs are processing individual time frames, using a TF or 32 ORBIT was the best choice for performance and memory usage.
  - Synchronous processing: EPN farm build for synchronous processing!
  - Asynchronous reprocessing : More detectors with significant computing contribution (more memory requirements)
  - Even with this configuration on the GRID sites we run out of memory (under investigation).

PERFORMANCE – how well did ALICE operate

A challenging year

We were getting ready for our HI commissioning:

- IP8 Inner Triplet incident on 17/07
- 7 weeks without beam during a crucial period to prepare the HI run
- No testing of detectors stability
- SW upgrades every week, but only partial validation with synthetic and cosmics runs

A challeng

We were gettin

- IP8 Inner Tri
- 7 weeks with
- No testing of
- SW upgrades ... nd cosmics runs

MAJOR CHALLENGE

Running an experiment with 10’000 links, ~500 servers, thousands of processes (more then 1000 in each EPN) comes with some technical challenges.

ALICE adopted a readout system without BUSY, so you can’t slow down 3.5 TB/s and you don’t want to lose data.

ALICE efficiency 76%

- HI run max IR 47 kHz.
- Data rate into EPN = 770 GB/s
- Data rate into STORAGE = 170 GB/s

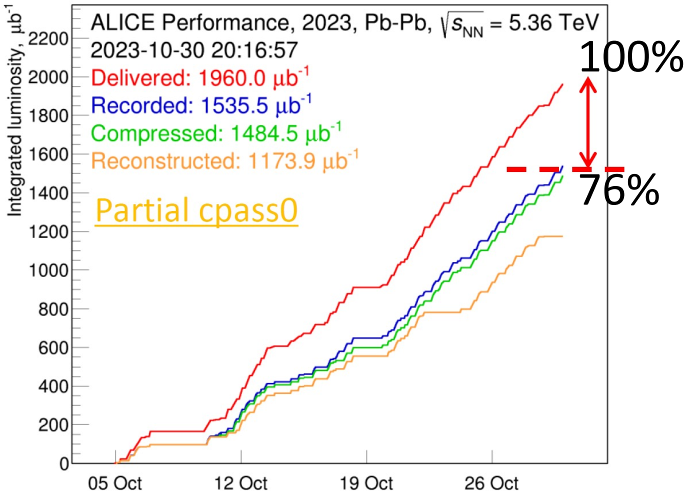

24% of inefficiency is caused by runs stopped due to :
- user error,
- Detector problem (misconfiguration, SEU, …),
- Process crashed,
- Clock phase changes,
- … and more.

detector	% data stored on disk
CPV	0.00515878866767245
CTP	0.00882093634283997
EMC	0.0965574200869967
FDD	0.0127348827326441
FT0	0.0401928297161023
FV0	0.0149288995582678
HMP	0.019687999041066
ITS	3.94027096898005
MCH	1.5866397777917
MFT	1.24892675680083
MID	0.0404086086670641
PHS	0.0180259765773141
TOF	0.328620040146893
TPC	90.9539527948842
TRD	1.33096730644395
ZDC	0.354106013562415

Simple task, complex infrastructure

To achieve a simple task like "data taking", ALICE O^2 runs many components (hardware and software), developed by different teams, that have to communicate with each other at the same time, in a well coordinated way.

This makes the system rather complex.

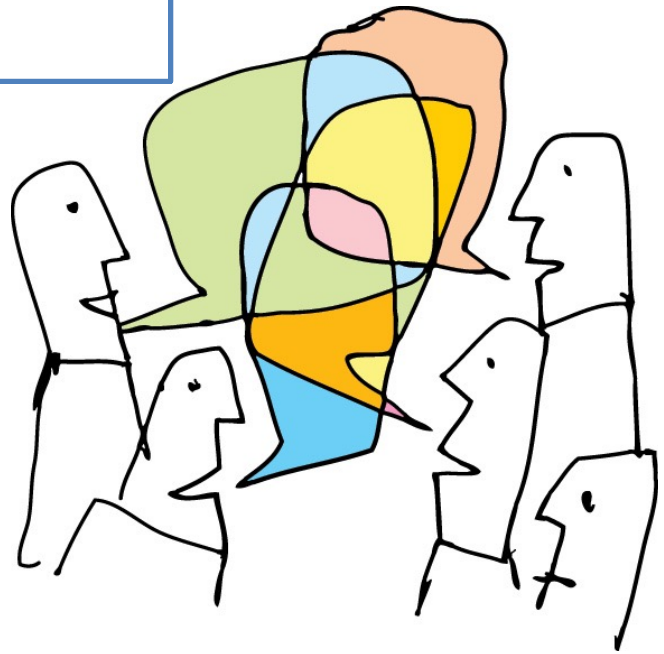

User eXperience

From the moment you click START, the user is flooded with information. Every process prints messages in the logging software, generating several hundreds of information per seconds.

Currently it is not easy to have clear overview of the status of the system.

Error detection

There are several components that run together in a defined sequence.
Interference, connection timeout, wrong software configuration result in an ENVIRONMENT in ERROR.
While this is clearly notified to the shifter, to identify the source of the problem is one of the major challenge even for the experts.

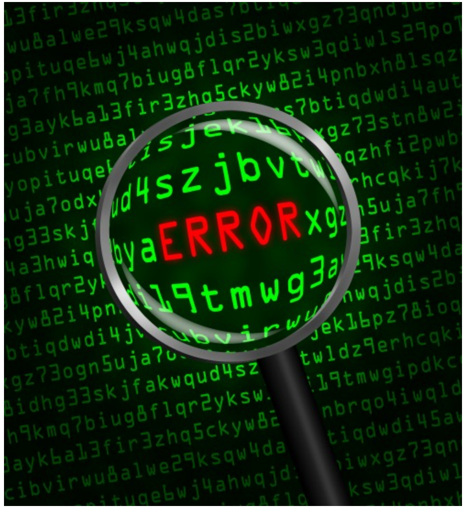

Error detection, fast reaction

The creation time of the ENVIRONMENT last several minutes (~7 minutes).
In 2023 was very important to identify possible source of errors as early as possible to reduce the number of environment creation to the minimum.
All the time spent in creating a new ENVIRONMENT is beam time lost.

Resource sharing

EPNs are equipped with powerful hardware for data processing (GPU). When they are not used in online for data taking, they are used in offline to process the data collected.

Currently the move of EPN from online to offline is a lengthy operation reducing the optimization of resource usage.

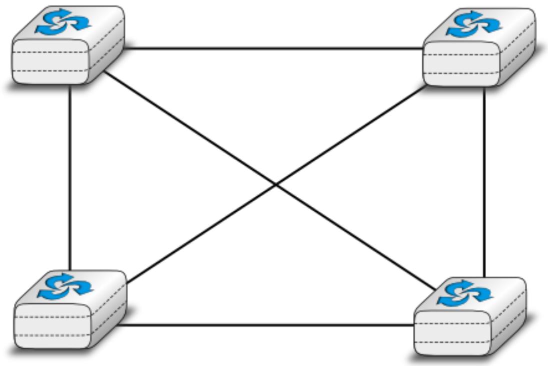

SW verification

Every software deployment brings in many new commits. Although we have different systems to test the software, the final verification is done in **PRODUCTION**.
During Stable Beam period it makes difficult to deploy new software and it requires coordination between different experts.

During HI in 2023 we spent a lot of effort to deploy only selected commits in time window where the BEAMs were not circulating.

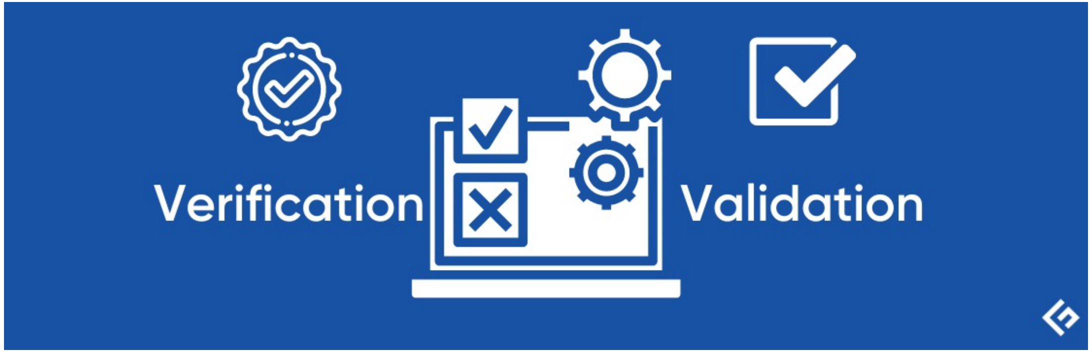

CONCLUSIONS

ALICE

- The O^2 framework worked well. ALICE could process data coming from the detectors during HI without introducing backpressure.
- ALICE was writing data into the storage at 170 GB/s.

- We are currently working to improve the software to have a system
  - Faster.
  - Better user experience.
  - More stable.