

# Looking Back at 2024

---

And Ahead to 2025

**Alessandro**

on behalf of the CT team



Science and  
Technology  
Facilities Council

# Just one year ago....

## Looking ahead to 2024 ...

- Some significant slip in our software plan since just May 2023
  - Already a ~3 month delay over the last 6 months
- Overall focus for 2024 remains: finish software foundation, and build up functionality from there
- Prepare for necessary Production and Procurement readiness reviews (most systems ~~Nov 2024, earlier for timing~~) for FD
- Preparation of ND TDR
- Operations of the FD ProtoDUNEs! And Near detector prototype (2x2) and other teststands and ...
  - Testing functionality with data from an operating detector → iteration on our design and solidifying our specifications

	2023				2024			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
dunedaq-v4.0.0								
Separation of FD/ND software releases complete								
Transition to ethernet readout for all WIBs complete								
DAQ in Kubernetes systematically in use at EHN1								
Network requirements complete (incl any flow control, protocols, buffer req)								
Configuration framework and application interface (schema, data distribution, API) complete								
Op Monitoring complete incl code review								
Review and revision of appfwk complete (cmdlib,appfwk, iomanager, conf I/F, opmon I/F)								
Timing SW/FW complete including code reviews								
DAQ Deployment model demonstrated at EHN1								
Database editors complete								
Resources management complete incl code review								

12 11 Dec 2023 W. Ketchum | DUNE DAQ 2023



## ProtoDUNE LAr procurement just been announced

## ... and beyond

- FD detectors are becoming a reality, and happening soon
  - FD2 TDR done and posted!
  - Detailed installation planning as cavern excavation makes good progress
  - **We cannot expect further delays in the detector**
- We should anticipate
  - PRRs in 2024
  - Procurements to start shortly after in 2025
  - Installation activities to begin in summer of 2026 @ SURF
  - Support for detector installation @ SURF to start in fall of 2026
- ...all while continuing to make progress on the near detectors of course



13 11 Dec 2023 W. Ketchum | DUNE DAQ 2023





# An Unprecedented Challenge

- Pursuing critical developments, supporting operations at CERN, preparing technical documentation for procurements
- Intense planning period at the start of 2024
  - ▶ Tracking the evolution of the NP04 (and NP02) operation plan
  - ▶ Setting up new development + production workflows
  - ▶ 3-days planning workshop in May

**Planning Workshop Day 3**  
Thursday May 16, 2024, 7:00 AM → 10:25 AM US/Central

**Planning Workshop Day 2**  
Wednesday May 15, 2024, 7:00 AM → 10:40 AM US/Central

**Planning Workshop Day 1**  
Tuesday May 14, 2024, 7:00 AM → 10:30 AM US/Central

Speakers: Alessandro Thea (STFC Rutherford Appleton Laboratory), Asher Kaboth (RHUL), Roland Sipoš (CERN), Wesley Ketchum (Fermi National Accelerator Laboratory)

7:00 AM → 7:15 AM **Welcome and Introduction**  
Speaker: Alessandro Thea (STFC Rutherford Appleton Laboratory)

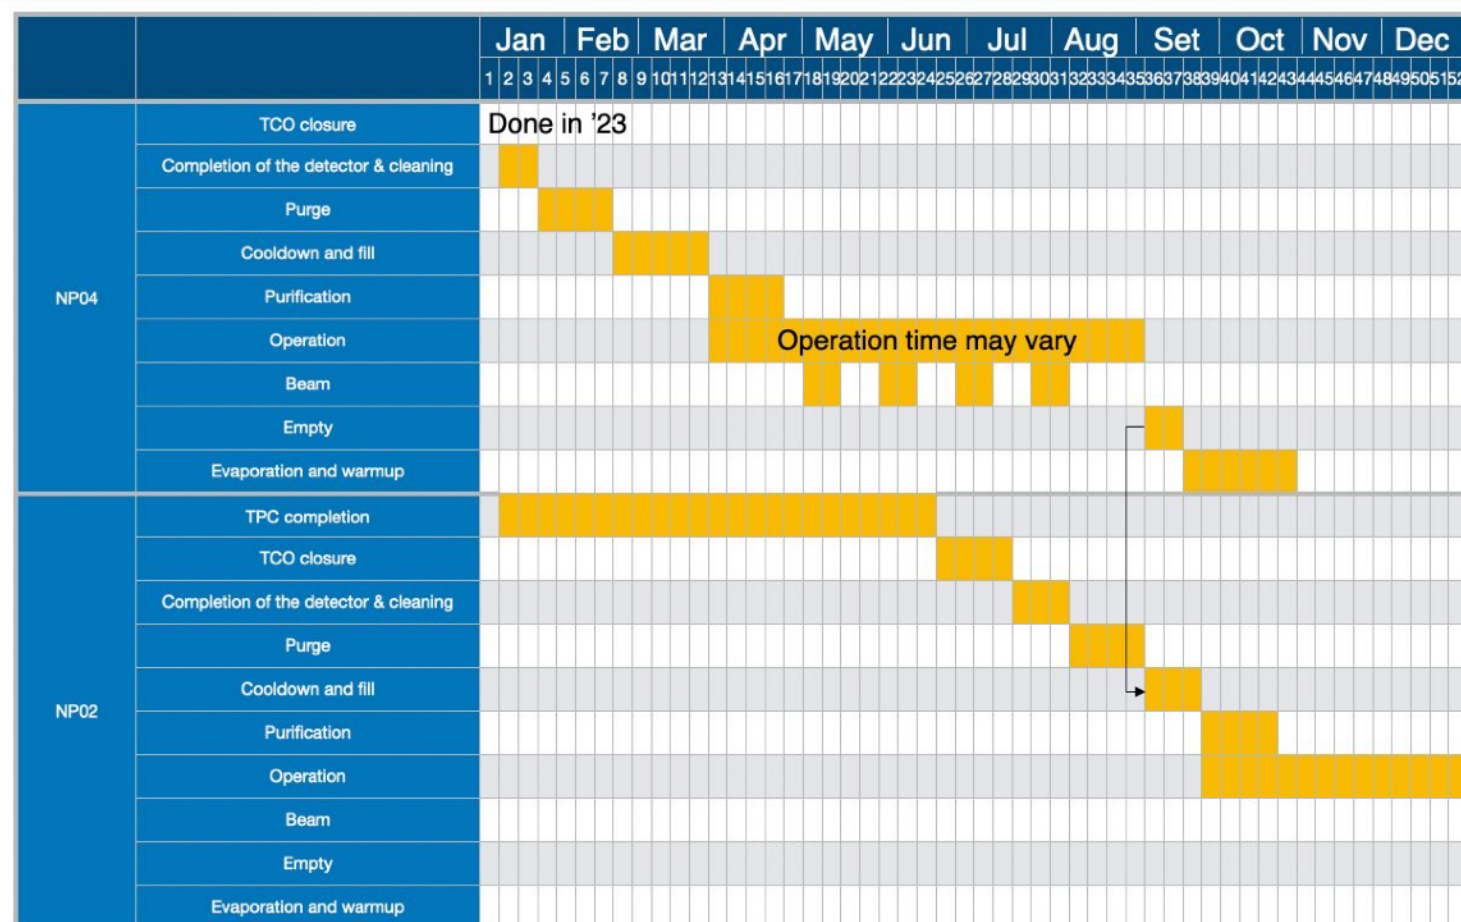
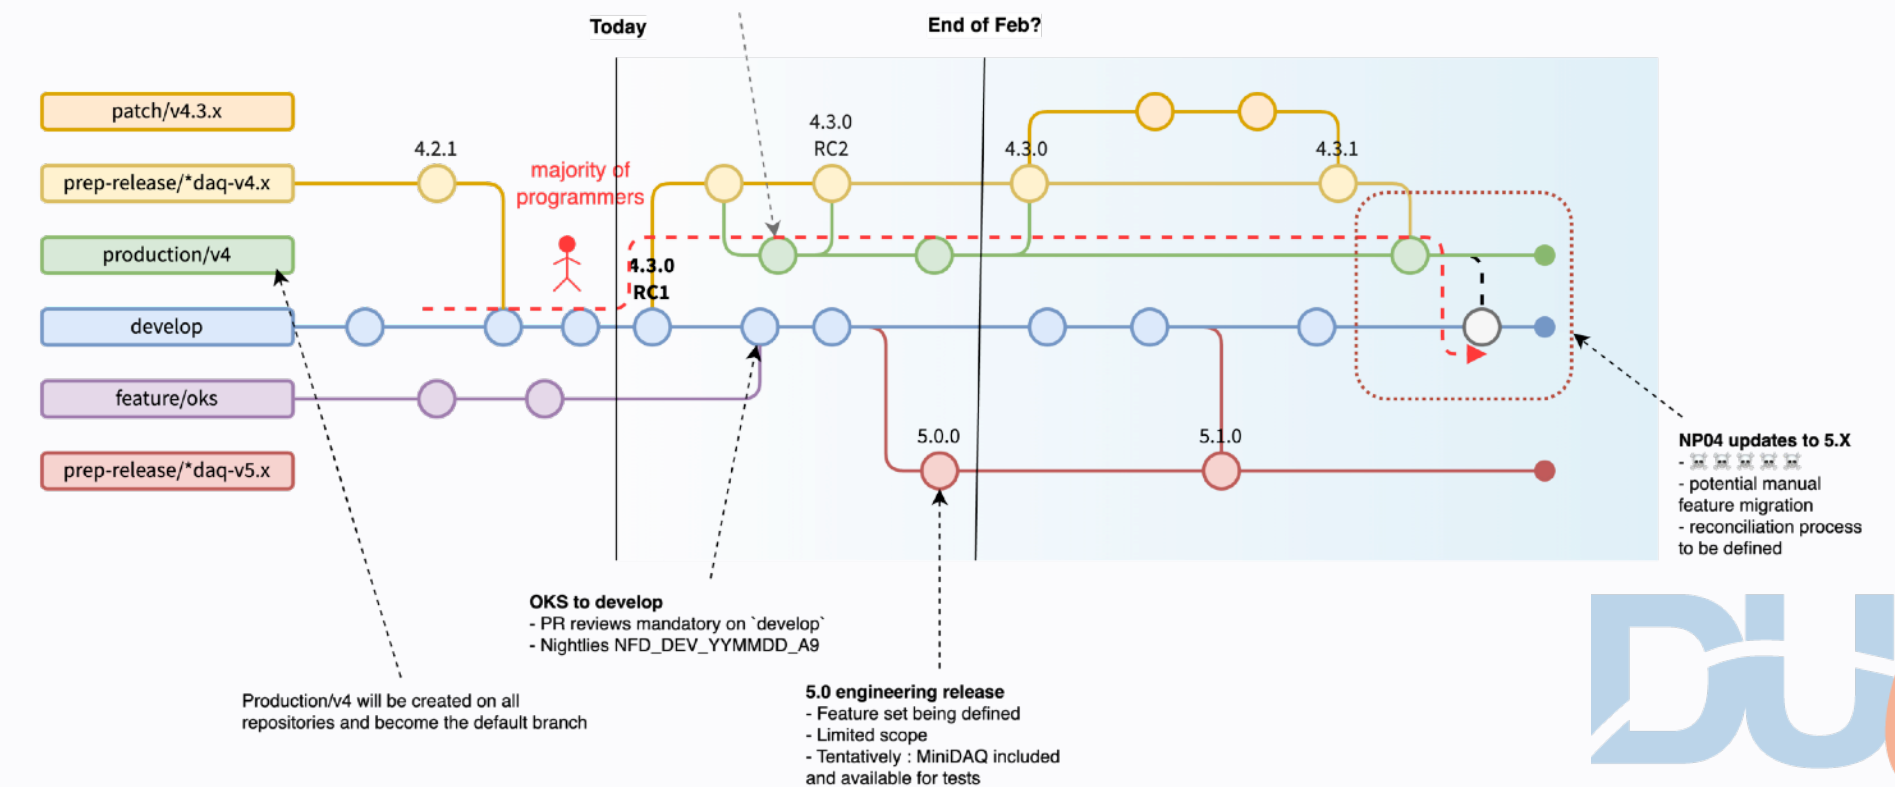
7:15 AM → 7:45 AM **NP04**  
Speaker: Wesley Ketchum (Fermi National Accelerator Laboratory)

7:45 AM → 8:00 AM **NP02 Overview**  
Speaker: Wesley Ketchum (Fermi National Accelerator Laboratory)

8:00 AM → 8:15 AM **OKS review**  
Speakers: Eric Flumerfelt (Fermilab), Giovanna Lehmann (CERN), Gordon Crone (University College London), Pierre Lasorak (Imperial College London)

8:15 AM → 10:00 AM **OKS planning discussion**

Production v4 is created across DUNE-DAQ repos  
 - default branch changed to 'production/v4' on all repositories  
 - nightlies are enabled (NFD\_PROD4\_YYMMDD\_A9)  
 - PR reviews mandatory on 'production/v4'



# *How Did It All Turn Out?*





An aerial, top-down view of a winding asphalt road that curves through a dense forest of evergreen trees. The road is dark and contrasts with the lighter, textured canopy of the trees. A small, dark-colored car is visible on the road in the upper left quadrant. The overall scene is captured in a monochromatic, slightly desaturated style, giving it a professional and serene appearance.

***DAQ Operations at EHN1***



# EHN1 Activities

---

- Of course, main activity was the ProtoDUNE-II run
  - ▶ DAQ performance highly stable over course of ProtoDUNE-II HD beam run
    - ◆ Sustaining up to **40 Hz (~5.7 GB/s)** instantaneous rates during beam spill (**~15 Hz averaged trigger rate**)
    - ◆ Fully integrated TPC, PDS streaming, PDS self-triggered, CRT, laser calibration, and beam instrumentation trigger
  - ▶ Exercised many important aspects of DAQ important for FD DAQ
    - ◆ Ethernet-based readout, long-readout-windows, and of course software-based triggering
- Many other activities at EHN1 over the year, including multiple VD and HD coldbox runs and gearing up for NP02 operations running
  - ▶ Included integration of PDS in VD coldbox, and PNS triggered data in spring run





# Timing System

---

- **ProtoDUNE II HD supported by PD-I DTS prototype hardware**
  - ▶ DTS clock locked to beam instrumentation clock
  - ▶ Smooth operation with no incidents
- **DTS hardware used for HSI**
  - ▶ DAPHNE calibrations
  - ▶ CRT input
- **DTS hardware also used to provide synchronisation interface for CRT**
- **TLU continues to provides synchronisation for ICEBERG**
- **Delivered timing hardware to support CRP factories**
- **ND interface specification under development**

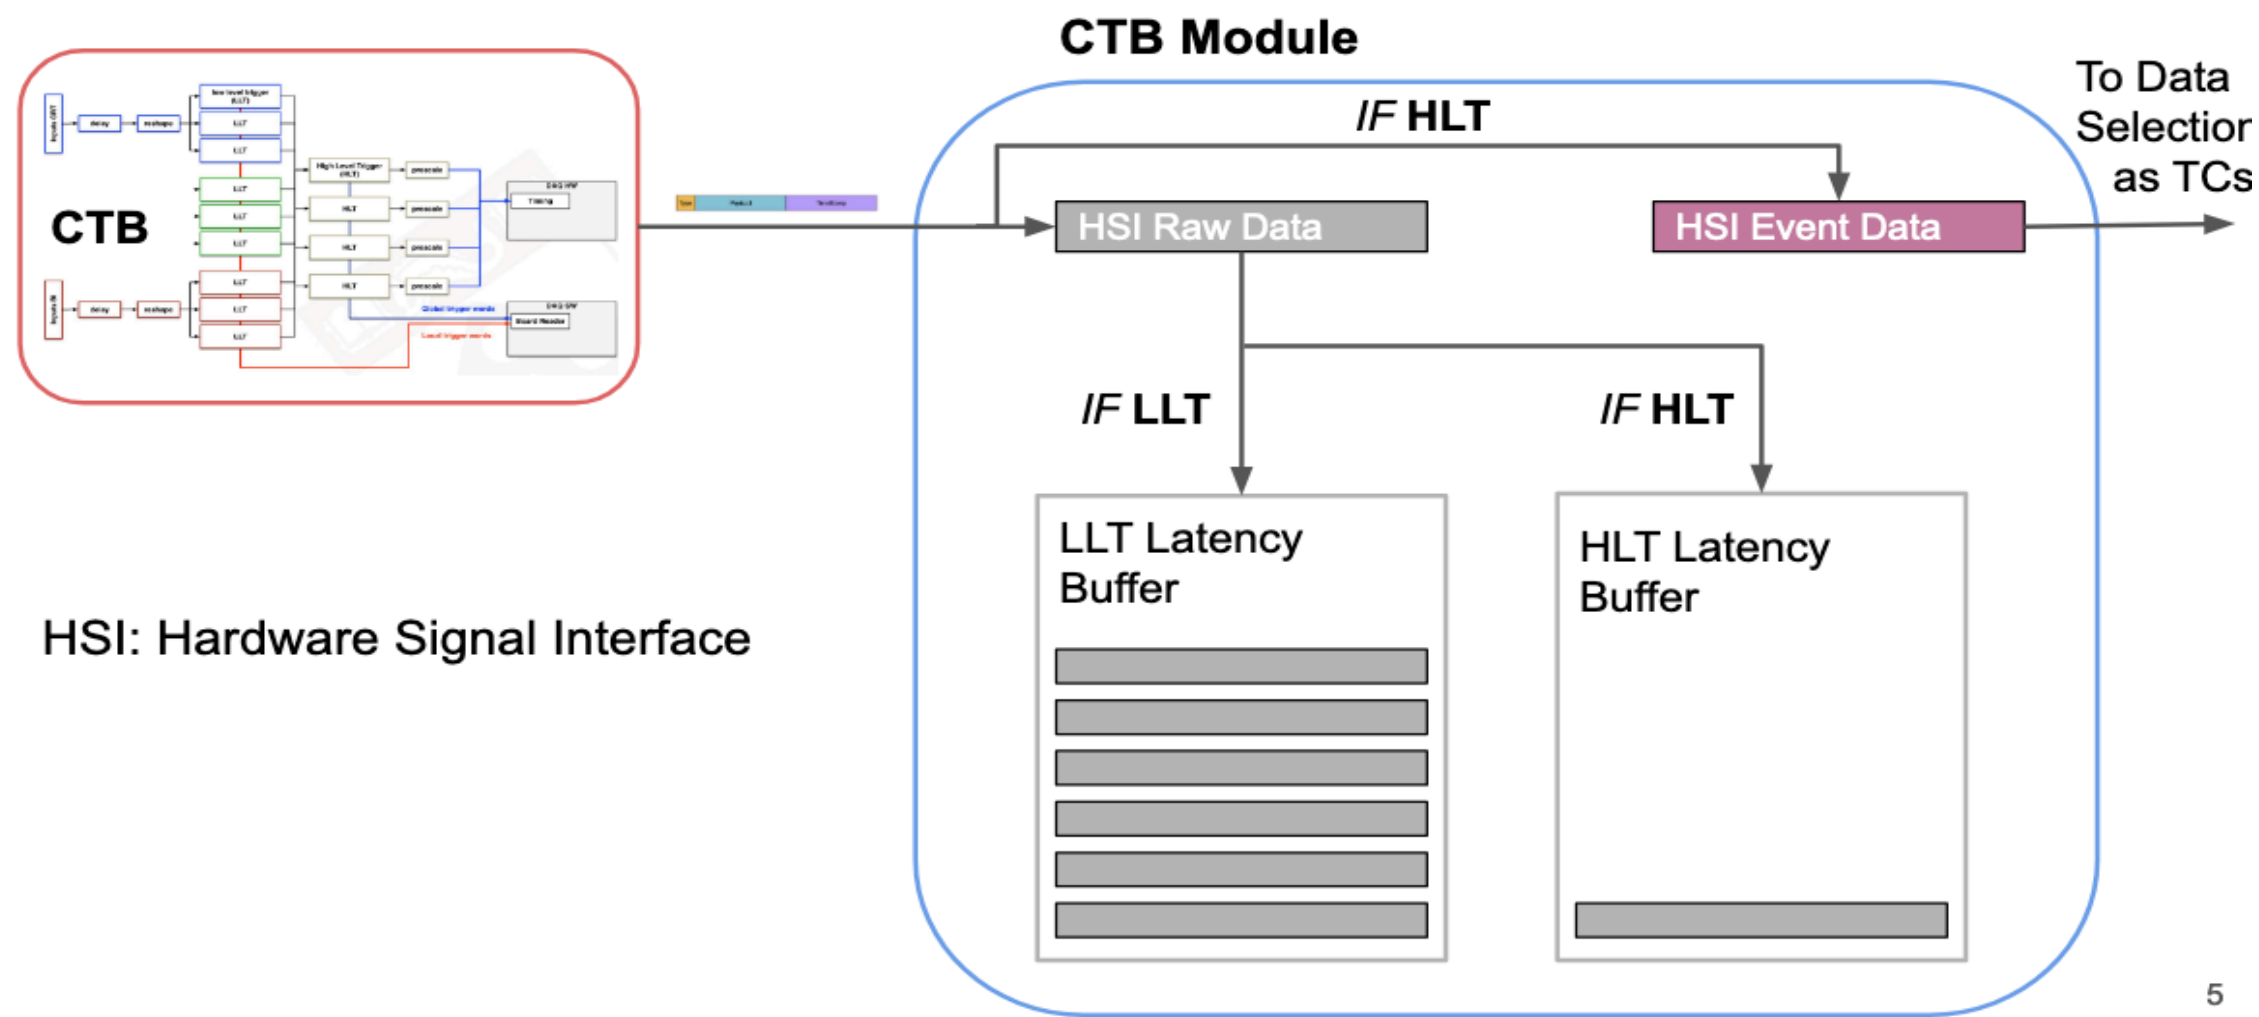





# NP04 Hardware Triggering

- Updates to CTB firmware to accommodate new timing system (Jon S, James S, Ben H)
- Updates to CTB software for dunedaq (was artdaq in last PD-SP run!) (Marco R)
- Smooth running using CTB as HSI source within dunedaq

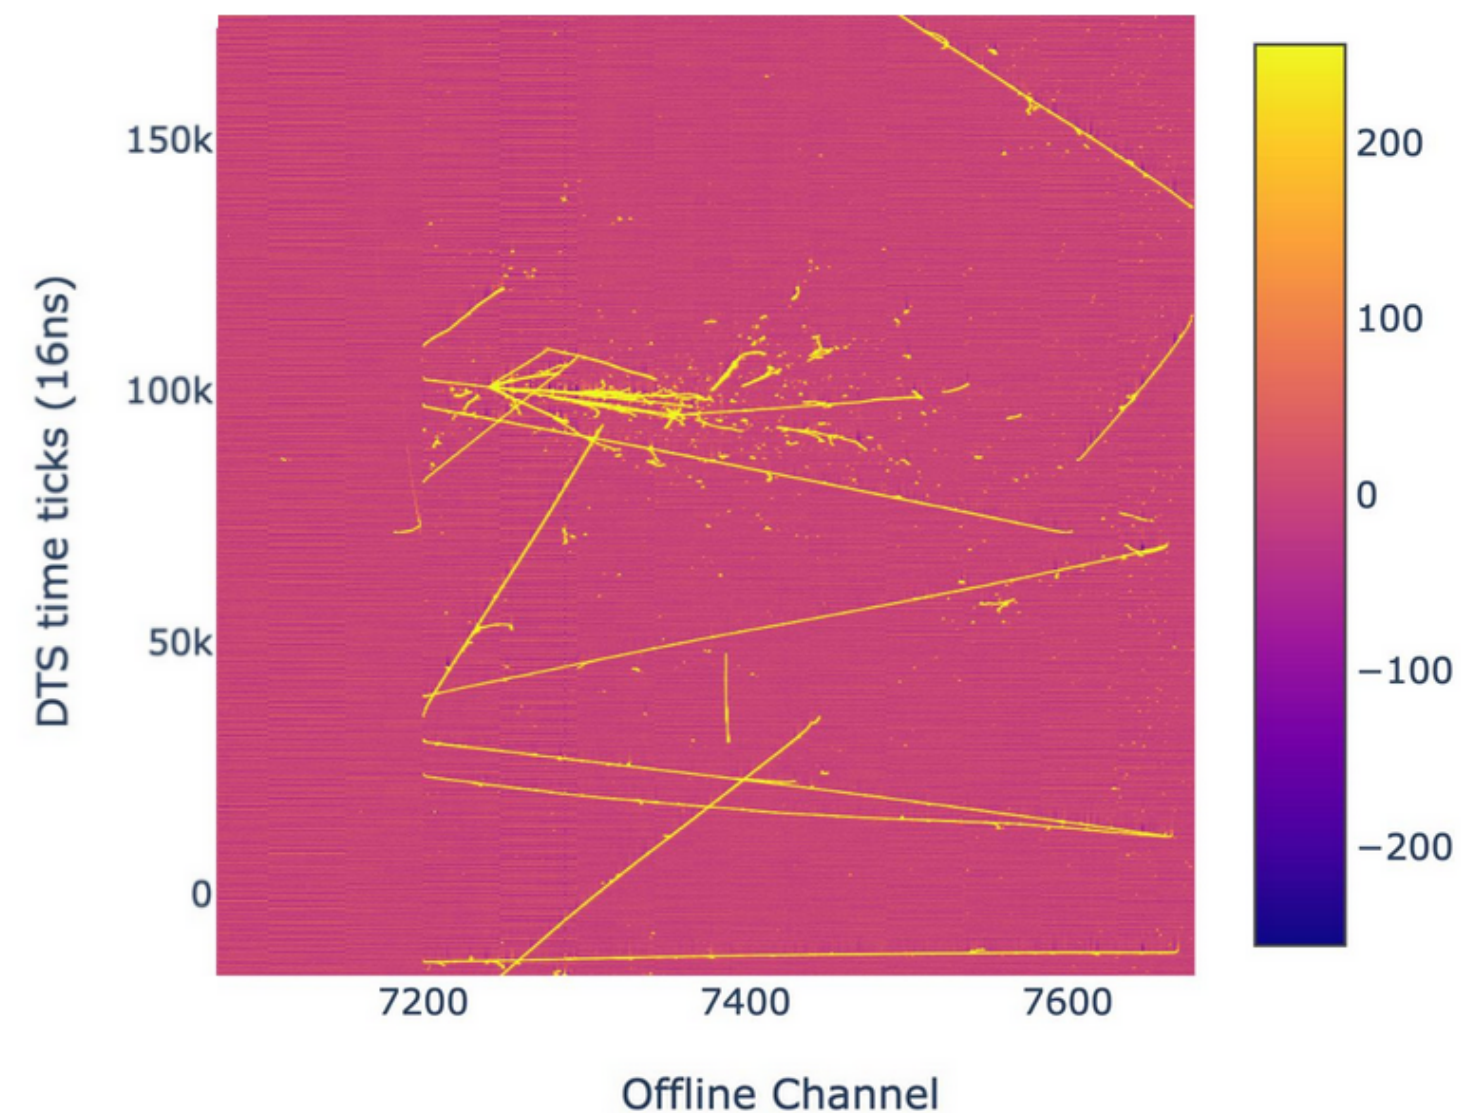
## Data Path



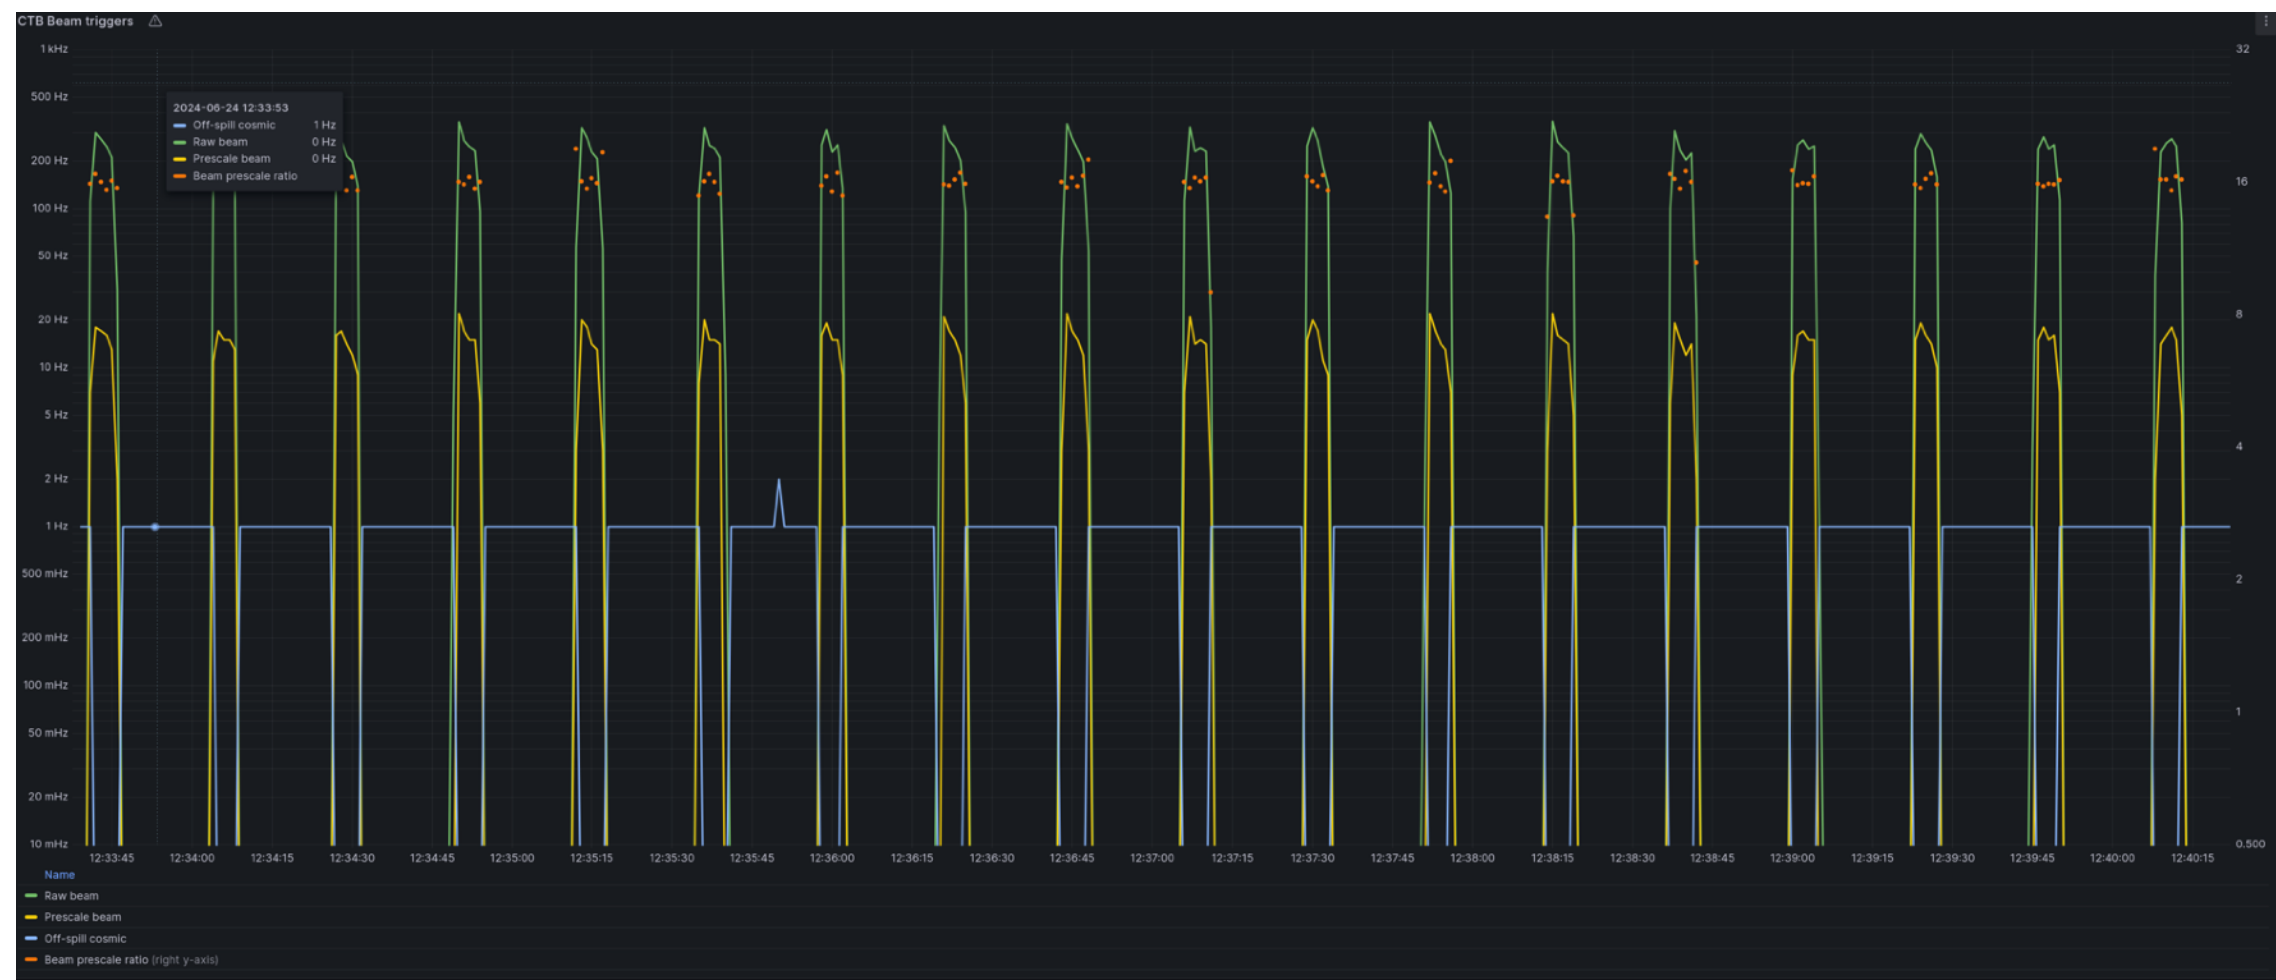
HSI: Hardware Signal Interface



Run 27307, Trigger 2694, APA2 Plane 2  
2024-06-19 20:11:32+02:00 (CERN)




Trigger



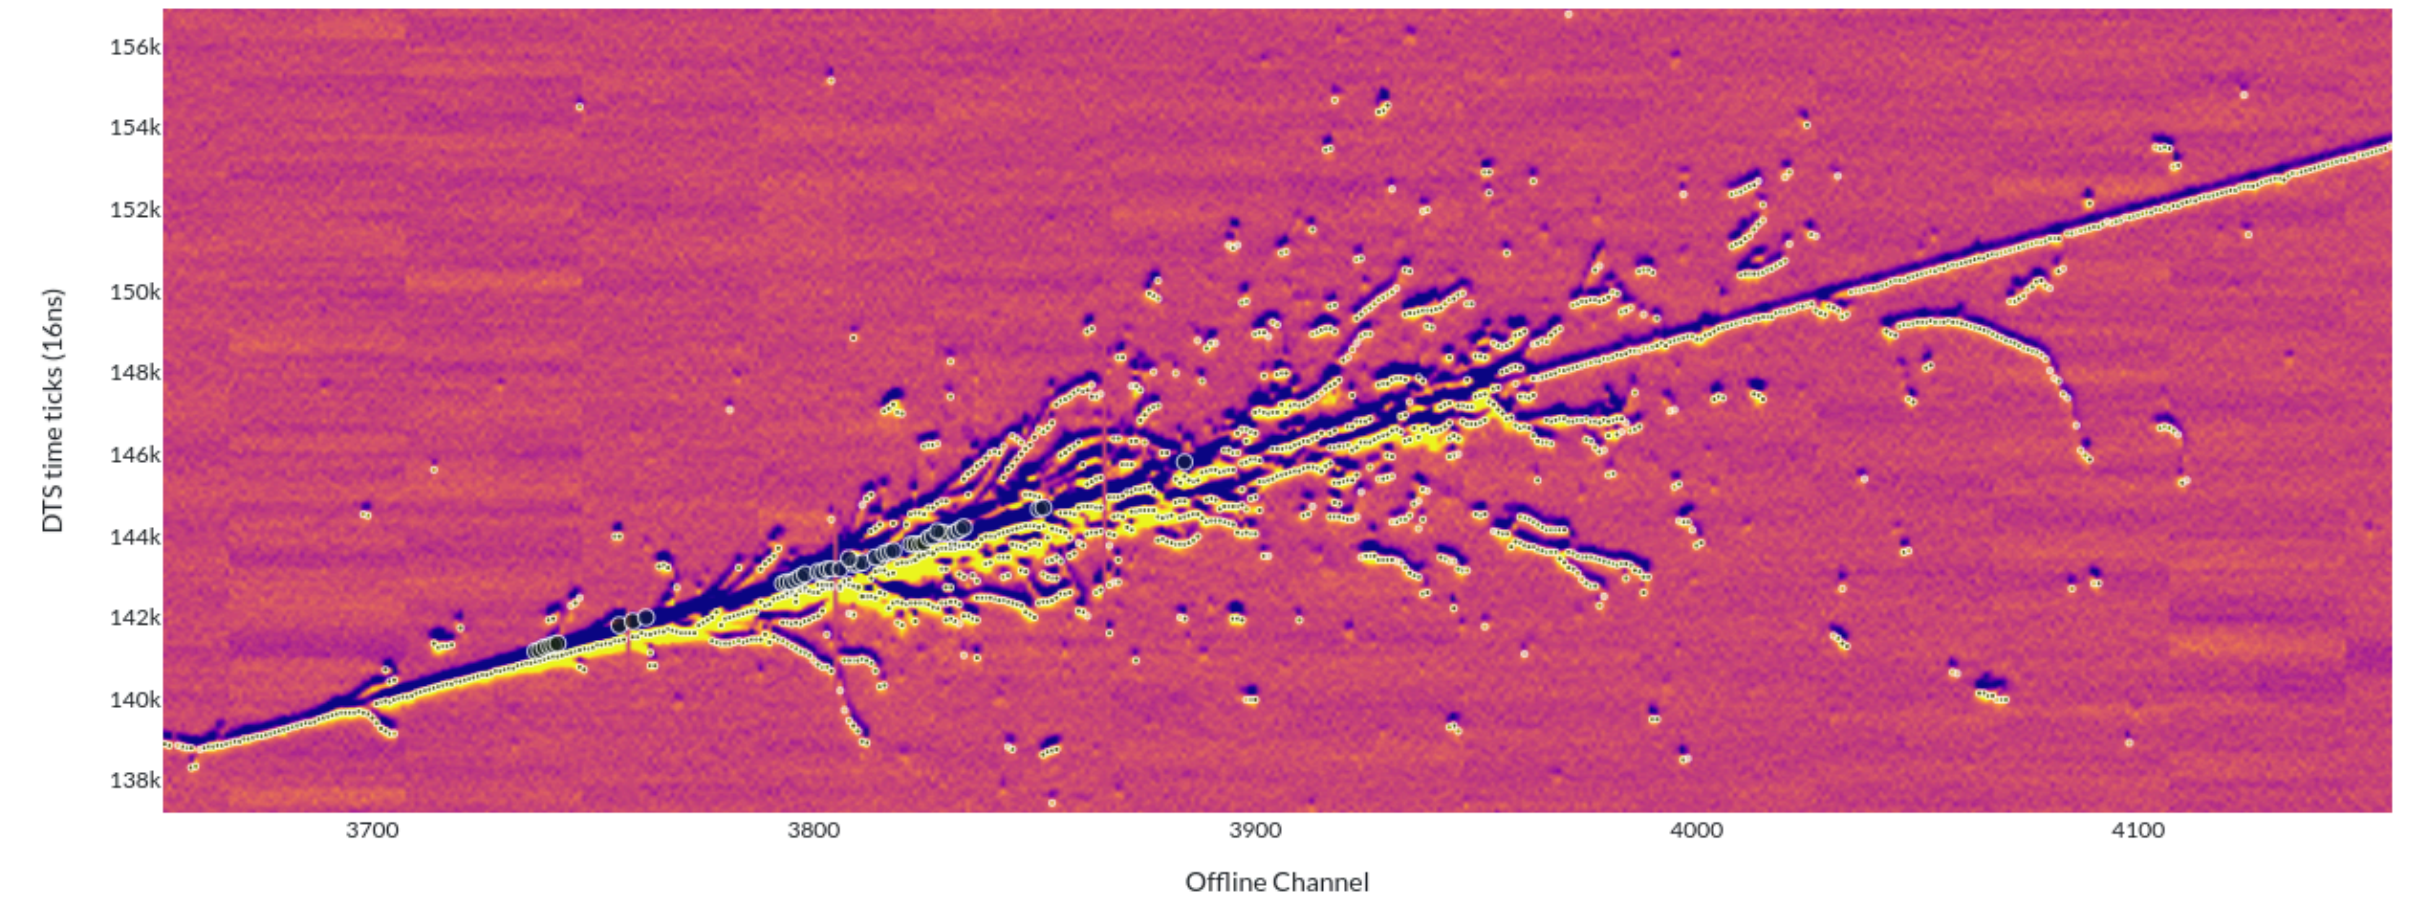


# NP04 Software Triggering

- Ran full software trigger chain for the first time on NP04
  - Included “ground shake” trigger to help understand these events
  - And a simple BSM trigger to catch high energy events
  - Generating Trigger Primitives & Activities on all APAs, all planes



“Ground Shake” event



HighE event



# Core Software for operations

Aided operations of detector apparatuses

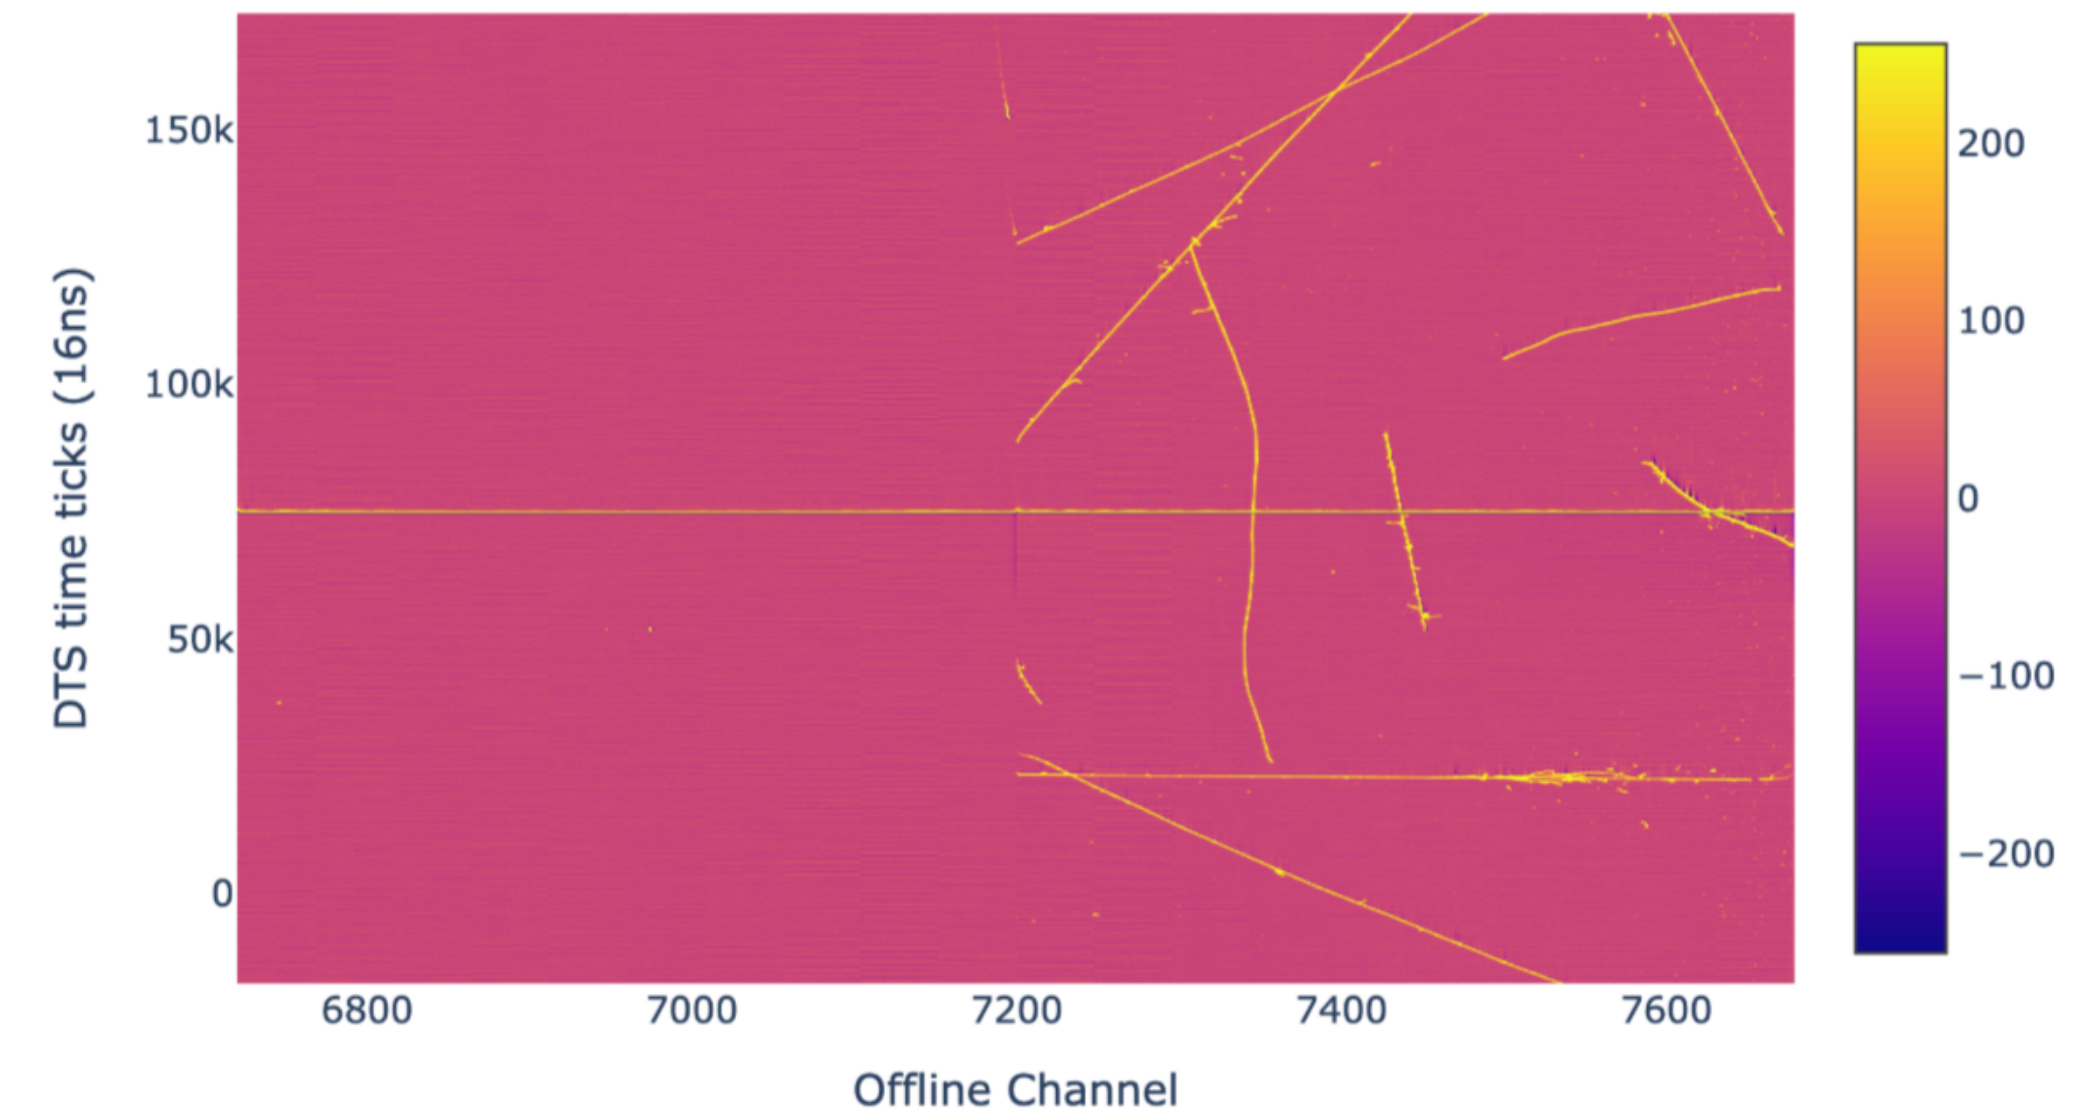
NP04:

- Data move callbacks vs. message passing (Major resource utilization reduction)
- DPDK software stack optimization (Optimal NIC and buffer configurations)
- 10G downstream network tuning
- Improvements to the handling of the TPs in the TPStreamWriter, and to its reporting
- Incremental improvements to the file-transfer metadata that we send to offline

ICEBERG:

- Significant amount of work was done to support calibration runs, etc.

Run 28536, Trigger 3392, APA2 Plane 2  
Trigger Type (CTBOffSpillCosmic), 2024-08-02 15:37:40+02:00 (CERN)





# Readout hardware performance optimisation

The tuning and optimization stages enabled the stable operations of NP04 with several generations of dual socket Intel and AMD readout servers.

- Skylake (np04-srv-021/022) from 2017
- Cascade Lake (np04-srv-028/029) from 2019
- Ice Lake (np02-srv-003/002) from 2021
- Sapphire Rapids (np04-srv-031) from 2023
- EPYC Zen3 (np04-srv-001/004) from 2021

**Confining ourselves to older generation, low-mid range CPU configurations drove substantial effort in low-level tuning of both software and hardware.**

**This resulted with major resource utilization overhead reduction of the readout system!**

# EHN1 Activities

---

## A true consortium-wide effort to pull this off

Commissioning and operations, experts at the ready both on- and off-site (on-call or not), and many data-taking shifts from DAQ consortium

*Thank you!*





```
for object to mirror  
mirror_mod.mirror_object  
operation == "MIRROR_X":  
mirror_mod.use_x = True  
mirror_mod.use_y = False  
mirror_mod.use_z = False  
operation == "MIRROR_Y":  
mirror_mod.use_x = False  
mirror_mod.use_y = True  
mirror_mod.use_z = False  
operation == "MIRROR_Z":  
mirror_mod.use_x = False  
mirror_mod.use_y = False  
mirror_mod.use_z = True
```

```
selection at the end -add  
mirror_ob.select= 1  
mirror_ob.select=1  
context.scene.objects.active  
("Selected" + str(modifier  
mirror_ob.select = 0  
= bpy.context.selected_object  
data.objects[one.name].select  
print("please select exactly
```

# *Highlights from WGs and Activities*

```
operator):  
X mirror to the selected  
object.mirror_mirror_x"  
mirror X"
```





***Facility & Integration***



# F&I Working Group 2024

Made progress on detailed network and infrastructure specifications.

Advanced the software services environment toward a fully containerized and automated deployment model with daq-kube and Pocket and maintained operations for data taking at EHN1 and ICEBERG.

## → Kubernetes and Microservices:

- ◆ Transitioned to AL9 clusters
- ◆ Riehecky completed & refined [daq-kube](#), major improvements in services deployment and Pocket usability

## → EHN1/ProtoDUNE Operations:

- ◆ AL9 upgrades
- ◆ Supported performance testing by on servers by tuning network parameters, CPU pinning, and examining RAID configurations
- ◆ Addressed disk space management with automated pruning and data cleanup policies.

## → Network Specifications

- ◆ refining technical details, addressing feedback from stakeholders, and ensuring alignment with the evolving requirements, incorporated differences in FD1 and FD2 topologies

## → ICEBERG Integration:

- ◆ Achieved ethernet readout from WIBs and FEMBs
- ◆ Supported CE integrating DAPHNE v3





# ***Software Coordination***



# Software Coordination (I)

---

- In February, introduced support for alternating v4 “datataking now” and v5 “longer term development (OKS, ...)” lines of development
  - v4 support now tapering (one “nightly” per week, Saturday night)
- Published several frozen releases from both lines to /cvmfs, including:
  - `fddaq-v4.3.0` (Feb.), `fddaq-v4.4.0` (Apr.), `fddaq-v5.0.0` (May), `fddaq-v5.1.0` (Jul.), `fddaq-v4.4.8` (Sep.), `fddaq-v5.2.0` (Nov.),
  - Documentation at `https://dune-daq-sw.readthedocs.io/en/<release name>`
- Dropped support for SL7 in July, switched from C++17 to C++20 as our standard in August





# Software Coordination (II)

- **daq-release GitHub Workflows:** Along with nightly **integration tests**, also run unit tests, clang-format (whitespace) check, a suite of tests for daq-buildtools, and weekly code linting
- **Automated Slack Messages** relaying status of builds and tests
- **Nightly Doxygen** publishing for continuous, up-to-date documentation

Summary

Jobs

- create nightly tag
- Run unit tests
- send\_slack\_message
- cleanup\_test\_area

Run details

- Usage
- Workflow file

utilities

Test	Status
NamedObject_test	Passed
TimestampEstimator_test	Passed
ReusableThread_test	Passed
WorkerThread_test	Passed
Resolver_test	Passed

cmdlib

Test	Status
Unit tests have not been written for cmdlib	

serialization

Test	Status
Serialization_test	Passed

Part of the output of the Dec-11-2024  
“Nightly code check workflow”





# Software Coordination (III)

---

- Mass code changes/merges when needed (e.g., on more than one occasion OKS-related package and class names have been overhauled)
- Improvements and bug fixes for various **daq-buildtools** commands, though it's a mature product and is less of a focus than previous years (likewise **daq-cmake**)
- Ongoing work on updating externals for the fddaq-v5.3.0 push (bumping package versions, e.g., updating gcc 12.1.0 to 13.2.0, and dropping obsolete packages)
- Investigating separate environments for building and running
- **Open to ideas, and Software Coordination will continue to hold periodic meetings**





# Control, Configuration and Monitoring



THIS HELICOPTER MUST BE OPERATED IN COMPLIANCE WITH THE OPERATING LIMITATIONS SPECIFIED IN THE APPROVED HELICOPTER FLIGHT MANUAL. MINIMUM COCKPIT WEIGHT 170 LBS.

SELECTIVE PASSENGER LOADING

WHEN BOTH CREW SEATS ARE OCCUPIED ONLY ONE (1) MID PASSENGER IS PERMITTED UNLESS THERE ARE TWO (2) AFT PASSENGERS.

WHEN ONLY ONE (1) CREW SEAT IS OCCUPIED NO MORE THAN TWO (2) AFT PASSENGERS ARE PERMITTED UNLESS THERE IS ONE (1) MID PASSENGER.

ABOVE 4,150 LB. GVW ALTERNATE PASSENGER LOADING FROM SIDE TO SIDE REFER TO RPM WEIGHT AND BALANCE FOR ADDITIONAL LOADING INFORMATION.



# Run control

## New run control system (drunc)

Distributed system replaces monolithic prototype – multiple processes spread over different hosts

## Communication via gRPC

Integrated with OKS for configuration

Designed for multiple users and role-based permissions e.g. shifter, DAQ expert...

Poised to be default for NP02 running

```
drunc-unified-shell > boot
[10:55:52] INFO process_manager_driver.py:149 drunc.process_manager_driver: Booting session local-2x3-config
Reading database
Included schemas: ['schema/appmodel/wiec.schema.xml', 'schema/appmodel/application.schema.xml', 'schema/appmodel/trigger.schema.xml', 'schema/confmodel/dunedaq.schema.xml', 'schema/appmodel/fdmodules.schema.xml']
Creating new database
Reading dal objects from old db
Copying objects to new db
Saving database
DONE
[10:55:52] INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'local-connection-server' from session 'local-2x3-config' with UUID 3fcd9975-703a-446c-b6a3-7a56ac85463c
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'root-controller' from session 'local-2x3-config' with UUID 41c40825-999e-44d0-b02e-7175f9b2a4b2
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'ru-controller' from session 'local-2x3-config' with UUID bee45fea-9b4d-46ba-9932-fc90c16c6ae7
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'ru-01' from session 'local-2x3-config' with UUID 1280b3fd-c88f-42b2-8d95-10b6c66b4472
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'ru-02' from session 'local-2x3-config' with UUID fc4aa628-a545-4710-bf9f-088a7e044944
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'df-controller' from session 'local-2x3-config' with UUID c43c64f9-376b-4312-999a-84ff7e8b76c6
[10:55:53] INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'tp-stream-writer' from session 'local-2x3-config' with UUID f2ea7c76-77d0-46ce-8d2d-57f38c983d85
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'dfo-01' from session 'local-2x3-config' with UUID 604774e2-dfb4-41be-bb9f-65ba96ea72d7
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'df-01' from session 'local-2x3-config' with UUID 1a82f2ac-e000-4c67-97a0-ce66dbfd6871
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'df-02' from session 'local-2x3-config' with UUID f609dd29-75fe-4eb4-aadb-20bd20add7a8
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'df-03' from session 'local-2x3-config' with UUID c37bc69d-a6b4-4c22-baf8-1674f39ac65d
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'trg-controller' from session 'local-2x3-config' with UUID bcebb5a5-9c1c-49fa-ae9f-513cc99ff3df
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'tc-maker-1' from session 'local-2x3-config' with UUID d3e81ee5-51a5-4912-8a00-287a2fc63fa0
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'mlt' from session 'local-2x3-config' with UUID 94f50e3f-6d43-4da3-b2a5-7b5ecd6a2f81
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'hsi-fake-controller' from session 'local-2x3-config' with UUID 7552f768-6b84-44fd-98ab-a8f596f248b3
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'hsi-fake-01' from session 'local-2x3-config' with UUID 135ca9ba-08b9-4830-aafa-8ef090f6532e
INFO ssh_process_manager.py:300 drunc.process_manager: Booted 'hsi-fake-to-tc-app' from session 'local-2x3-config' with UUID b2d73280-36b8-47ed-95c8-28d7dc8b10f6
* Looking for root-controller on the connectivity service... --:-- 0:00:07
[10:56:00] INFO commands.py:54 unified_shell_interface: Controller endpoint is '10.73.136.38:37297'
INFO commands.py:55 unified_shell_interface: Connecting this shell to it... --:-- 0:00:00
* Trying to talk to the top controller... --:-- 0:00:00
INFO shell_utils.py:153 controller_shell_utils: 10.73.136.38:37297 is 'root-controller.local-2x3-config' (name.session), starting listening...
drunc-unified-shell > status
local-2x3-config status


| Name                | Info      | State   | Substate | In error | Included | Endpoint                  |
|---------------------|-----------|---------|----------|----------|----------|---------------------------|
| root-controller     |           | initial | initial  | No       | Yes      | grpc://10.73.136.38:37297 |
| ru-controller       |           | initial | initial  | No       | Yes      | grpc://10.73.136.38:36081 |
| ru-01               | conn apa1 | initial | idle     | No       | Yes      | rest://10.73.136.38:57101 |
| ru-02               | conn apa2 | initial | idle     | No       | Yes      | rest://10.73.136.38:46591 |
| df-controller       |           | initial | initial  | No       | Yes      | grpc://10.73.136.38:34939 |
| tp-stream-writer    |           | initial | idle     | No       | Yes      | rest://10.73.136.38:46211 |
| dfo-01              |           | initial | idle     | No       | Yes      | rest://10.73.136.38:36263 |
| df-01               |           | initial | idle     | No       | Yes      | rest://10.73.136.38:57679 |
| df-02               |           | initial | idle     | No       | Yes      | rest://10.73.136.38:53041 |
| df-03               |           | initial | idle     | No       | Yes      | rest://10.73.136.38:52033 |
| trg-controller      |           | initial | initial  | No       | Yes      | grpc://10.73.136.38:42993 |
| tc-maker-1          |           | initial | idle     | No       | Yes      | rest://10.73.136.38:49259 |
| mlt                 |           | initial | idle     | No       | Yes      | rest://10.73.136.38:45547 |
| hsi-fake-controller |           | initial | initial  | No       | Yes      | grpc://10.73.136.38:44259 |
| hsi-fake-01         |           | initial | idle     | No       | Yes      | rest://10.73.136.38:52433 |
| hsi-fake-to-tc-app  |           | initial | idle     | No       | Yes      | rest://10.73.136.38:51509 |


Current FSM status is initial. Available transitions are conf.
drunc-unified-shell > |
```



# Run control

User friendly web-based tools being developed on top of new run control

The screenshot shows the Run Control web interface. At the top, there are navigation links for 'Run Control', 'PROCESS MANAGER', and 'CONTROLLER', along with a user greeting 'WELCOME, DIEGO' and a 'HELP' link. The main content area is divided into two panels:

- Process Control:** This panel features a top bar with buttons for 'ROOT', 'RESTART', 'FLUSH', and 'KILL'. Below this is a search bar and a table listing running processes. The table columns are: UUID, Process Name, User, Session, Status, Exit Code, and Logs. All processes listed are in a 'RUNNING' state.
- Messages:** This panel displays a list of system messages. It includes a search bar and a dropdown for 'All Severities'. The message list has columns for 'Timestamp', 'Severity', and 'Message'. The messages shown are all 'DEBUG' level and indicate successful execution or attempts to execute 'ps'.

This screenshot shows the Run Control web interface with a different set of panels:

- Finite State Machine:** A table showing the state transitions for the system. The states are INITIAL, CONFIGURED, READY, RUNNING, and DATAFLOW\_DRAINED. Transitions include 'CONF', 'SCRAP', 'START', 'ENABLE\_TRIGGERS', 'DRAIN\_DATAFLOW', 'DISABLE\_TRIGGERS', 'STOP\_TRIGGER\_SOURCES', and 'STOP'.
- Application Tree:** A tree view showing the hierarchy of applications. The root is 'root-controller', which includes 'ru-controller' (with sub-items 'ru-01' and 'ru-02'), 'df-controller' (with sub-items 'tp-stream-writer', 'dfo-01', 'df-01', 'df-02', 'df-03'), 'trg-controller' (with sub-items 'tc-maker-1' and 'mlt'), and 'hsi-fake-controller' (with sub-items 'hsi-fake-01' and 'hsi-fake-to-tc-app').
- Messages:** A panel for viewing system messages, similar to the one in the first screenshot.

An 'Arguments to run transition: start' dialog box is overlaid on the Finite State Machine panel, showing a 'Run number\*' field, a 'Disable data storage' checkbox, a 'Trigger rate' field, a 'Run type' dropdown (set to 'TEST'), and a 'File logbook post' field. There are 'CANCEL' and 'CONFIRM TRANSITION' buttons at the bottom.

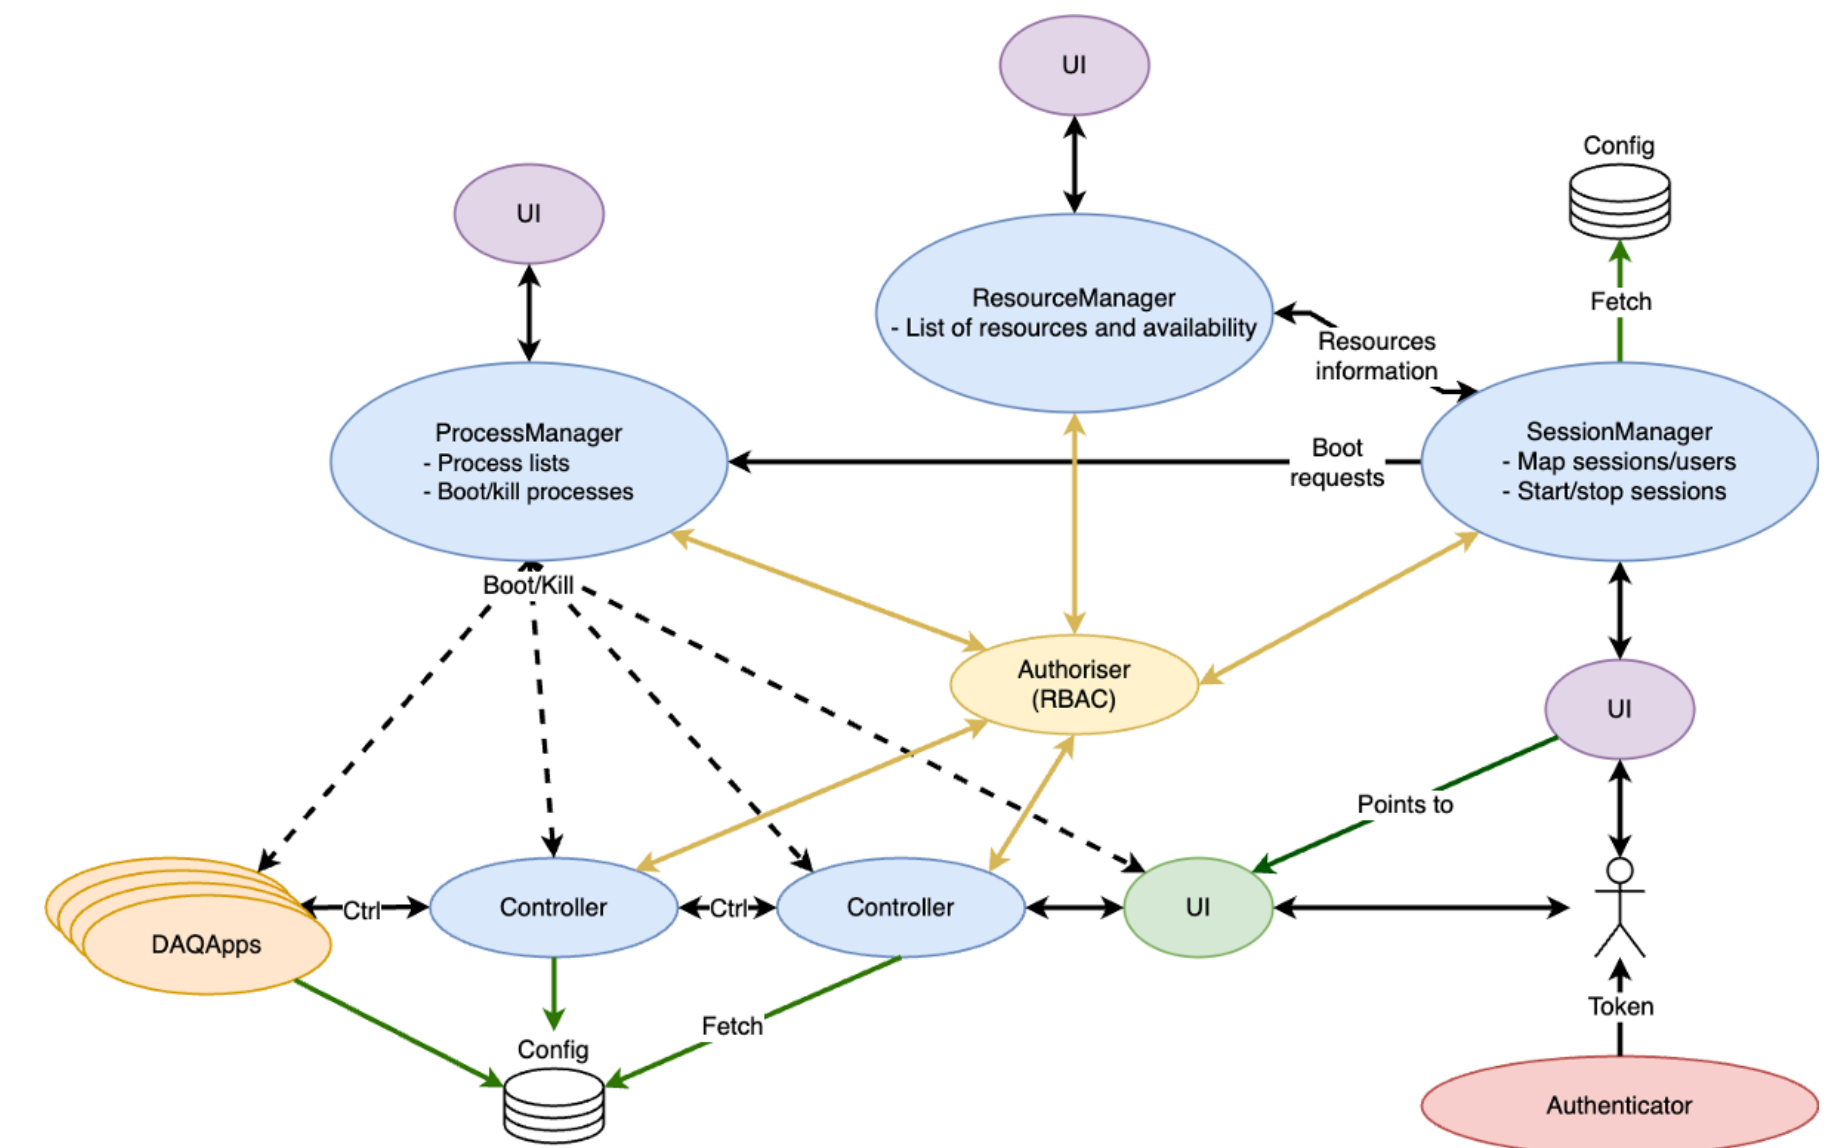
This screenshot shows the 'Application tree' panel in the Run Control web interface. It displays a table with the following columns: 'Application Name', 'Host', and 'Detector'. The table lists various applications and their configurations:

Application Name	Host	Detector
root-controller	localhost	-
- ru-controller	localhost	-
- - ru-01	localhost	-
- - ru-02	localhost	-
- df-controller	localhost	-
- - tp-stream-writer	localhost	-
- - dfo-01	localhost	-
- - df-01	localhost	-
- - df-02	localhost	-
- - df-03	localhost	-
- trg-controller	localhost	-
- - tc-maker-1	localhost	-
- - mlt	localhost	-
- hsi-fake-controller	localhost	-
- - hsi-fake-01	localhost	-
- - hsi-fake-to-tc-app	localhost	-



# Run control

What's next?



- New run control has feature parity with older nanorc
  - Working on stability, bug fixes and improving usability – coming very soon
  - TUI and GUI demos in progress – testing and development in 2025
- Still missing
  - Authentication and roles to authorise shift, expert etc. actions
  - Top level session manager – entry point for users
  - Resource manager – properly handle multiple sessions, clashes etc.
  - Kubernetes support – separate process manager for K8s



# Configuration

Transition to OKS framework for describing configuration data

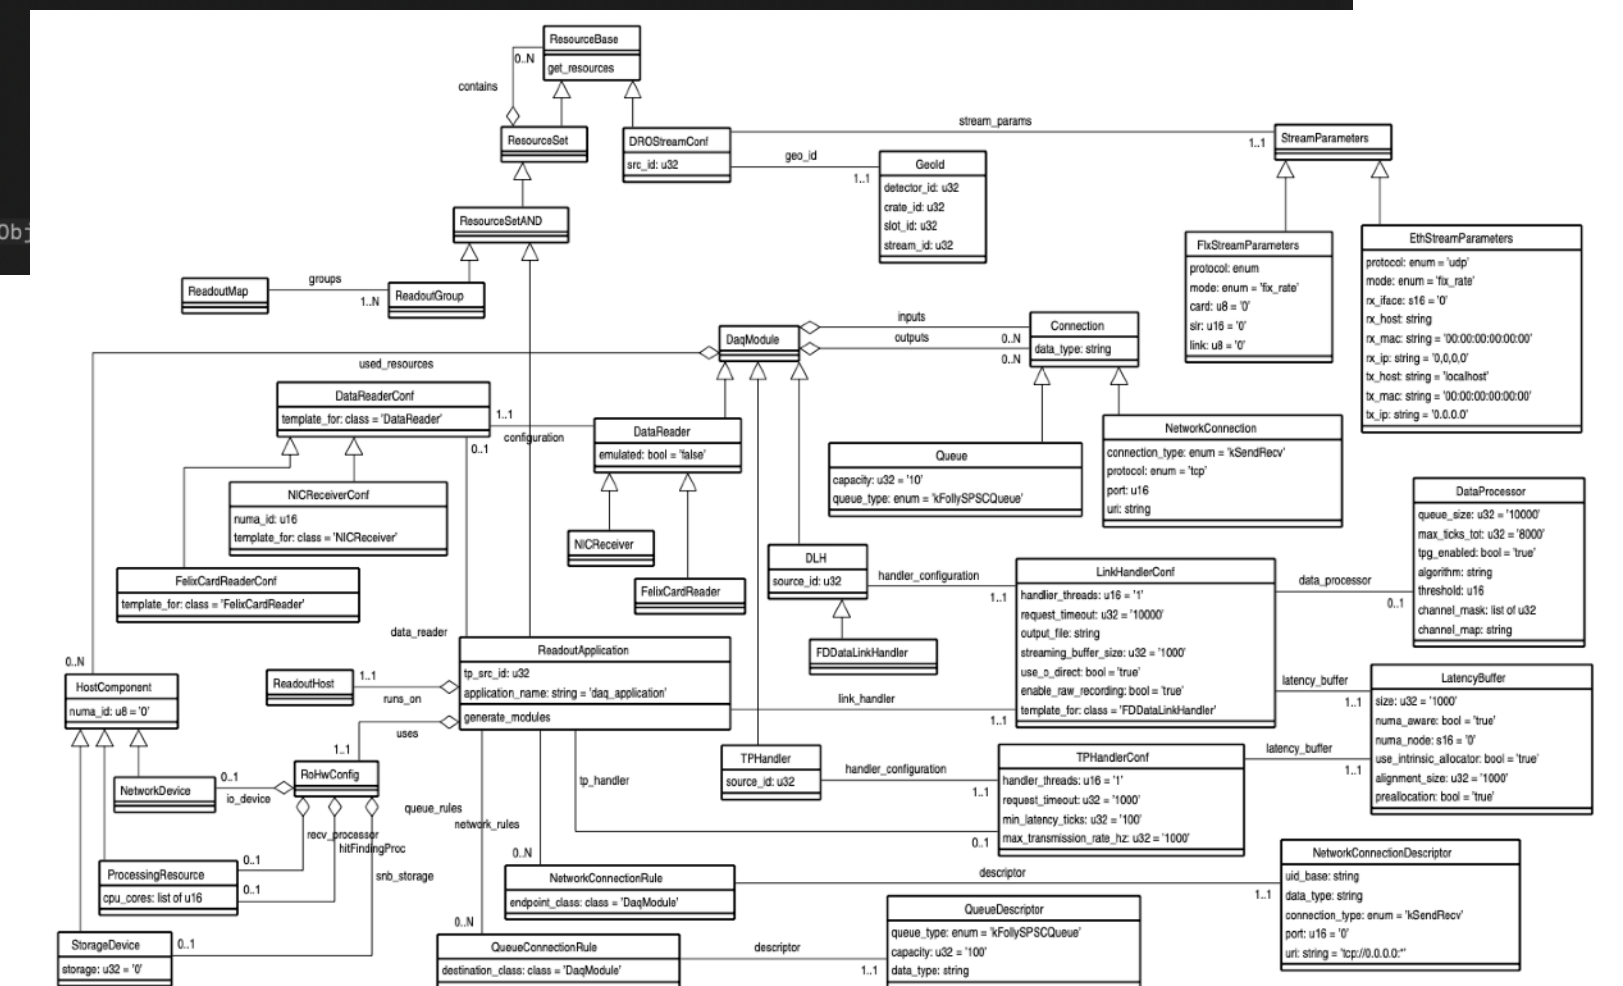
Schema and data defined and stored in xml

Flexible and reusable schema elements to describe system

Data Access Layer with interface via C++ and Python

The screenshot shows the CIDER configuration tool interface. On the left is a 'Class View' tree showing a hierarchical structure of configuration objects. The tree starts with 'Configuration' and branches into 'Sessions', 'np02-session@Session', 'opmon\_uri', 'segment', 'root-segment@Segment', 'controller', 'segments', 'crp4-segment@Segment', 'crp5-segment@Segment', 'df-segment@Segment', 'trg-segment@Segment', 'mlt@MLTApplication', 'hsi-segment@Segment', and 'applications'. The 'def-tc-handler@DataHandlerConf' class is highlighted in yellow. On the right is a 'Session View' showing a 'Welcome to CIDER!' message and a list of opened configuration files and connected databases. Below this is a table with columns 'Attribute', 'Value', and 'Type'.

Attribute	Value	Type
post_processing_delay_t	625000	u64
input_data_type	TriggerCandidate	string
generate_timesync	False	bool
template_for	TriggerDataHandlerModule	class





# Configuration

OKS integrated into DAQ applications and run control – prototype schemas available

What's next?

- Need to learn how to design flexible, factorised schemas for different aspects of the DAQ system
- Need to develop tools to perform configuration management – allow shifters and experts to interact with the configuration system in suitable way, ensure robust bookkeeping, checking consistency of config etc. – prototyping underway
- Still missing many things for final system – databases, interfaces to offline ...









***Core Software***



# Application Framework review

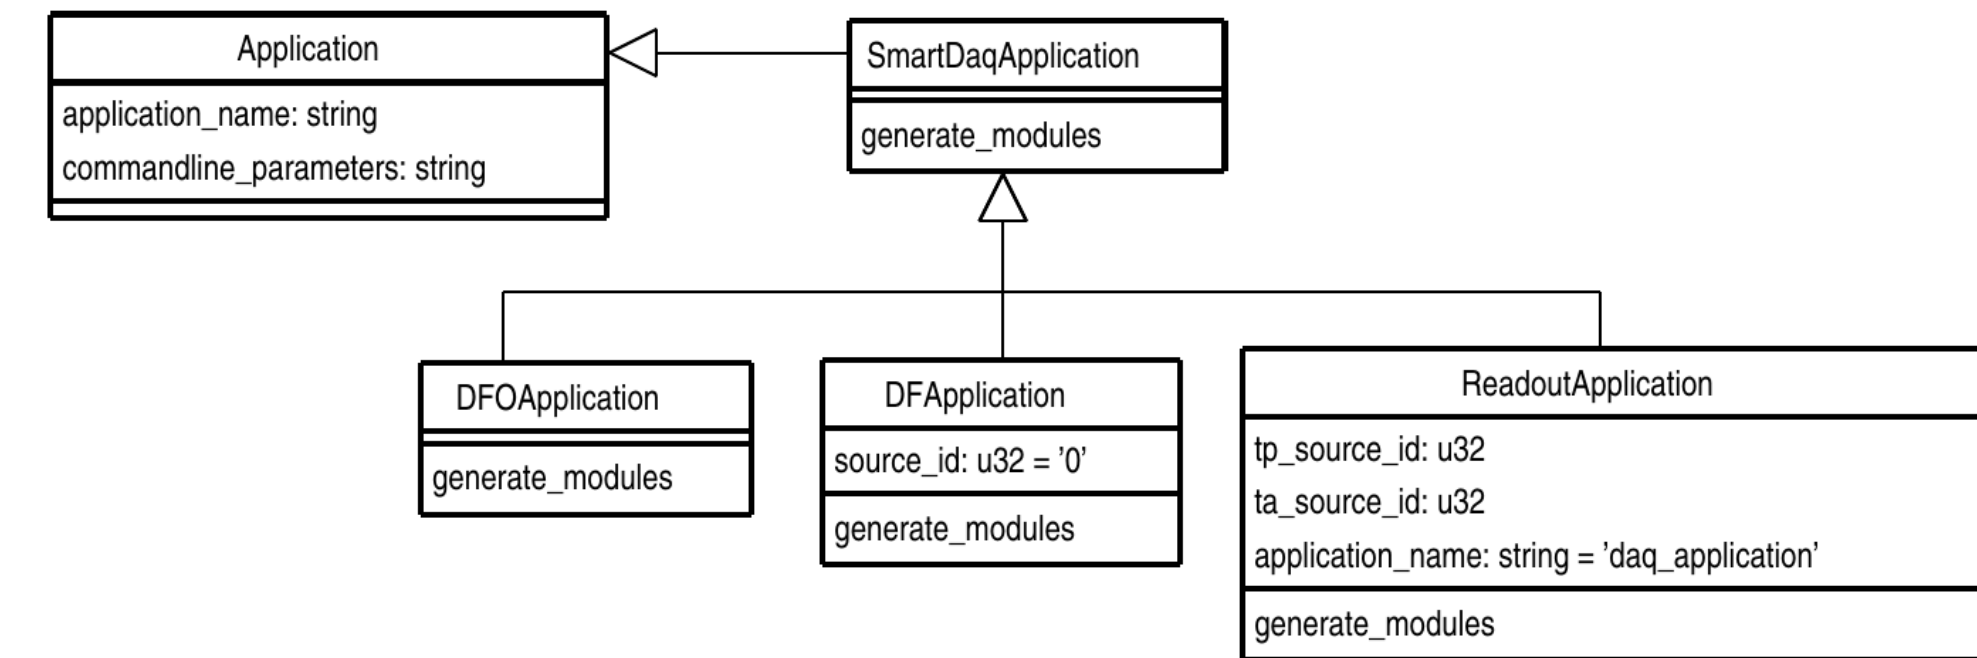
- Held workshop at CERN Jan 29-Feb 2
- Completed “[Functional Review](#)” of framework and identified several needed changes
  - **Action Plans** for coordinating command processing
  - **Configuration and Control interfaces** with full OpMon facility overhaul
  - Review of connectivity patterns
  - Cleanup of threading utilities, provide a better API to DAQModules
  - Improve the DFO protocol to avoid having a single-point-of-failure
  - “Runtime parameters”, handling of configuration domains of validity
- Started working on requirements and interfaces document for the appfwk
- Presented Application Framework at RT2024



# OKS based configuration

Schema driven subsystem configurations:

- Branched “production/v4” from develop, started work on DUNE-DAQ v5 using develop branches
  - Produced initial v5.0.0 release May 1
  - Continued development, targeted detector operations for v5.2.0 release in November
  - Configuration parameters for DAQ system topology and DAQ module settings were copied to the new system (and improved, as time permitted)
- OKS branch is now primary development, working on reconciling changes made on production/v4 branches to reunify codebase





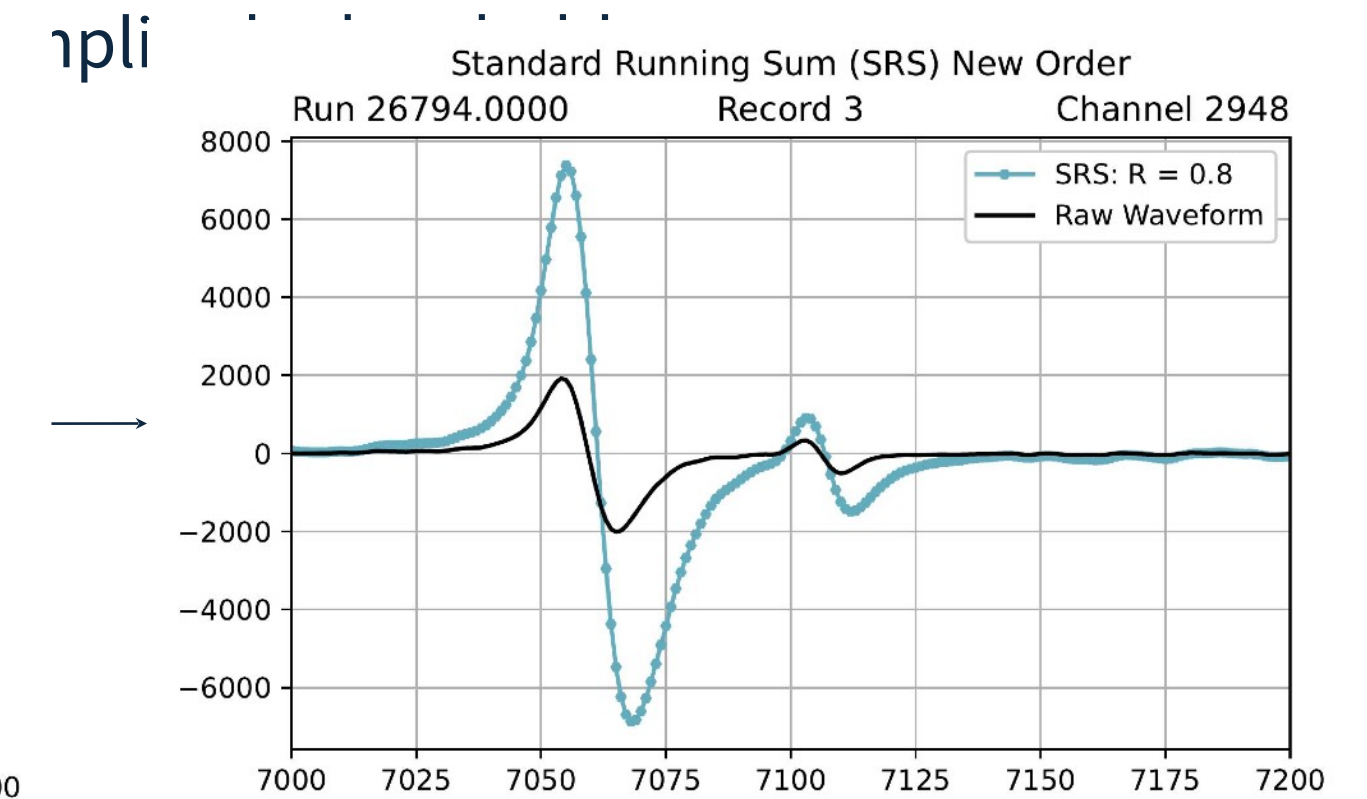

# Trigger Primitive Generation

## Architectural developments:

- **TPG Library:**
  - Decoupled from readout
    - Improved evaluation of online TPG performance.
  - Modularity
    - Easier algorithmic developments.
  - Minimal dependencies
    - Easier detailed, offline performance testing, e.g., frame expansion testing and improvements.
- **TPG By Plane:** TP configuration and sending by readout planes.
  - Improved data sending performance and stability.

## Algorithmic improvements and changes:

- **Absolute Running Sum & Running Sum Verification**
  - Removed overflow cases.
  - Improved time complexity.
  - Consistent performance between both.
- **32 Bit ADC Integration**
  - No longer maxing out at 16 bit limit.
- **Signal Width Filtering**







***Trigger and  
Data Selection***







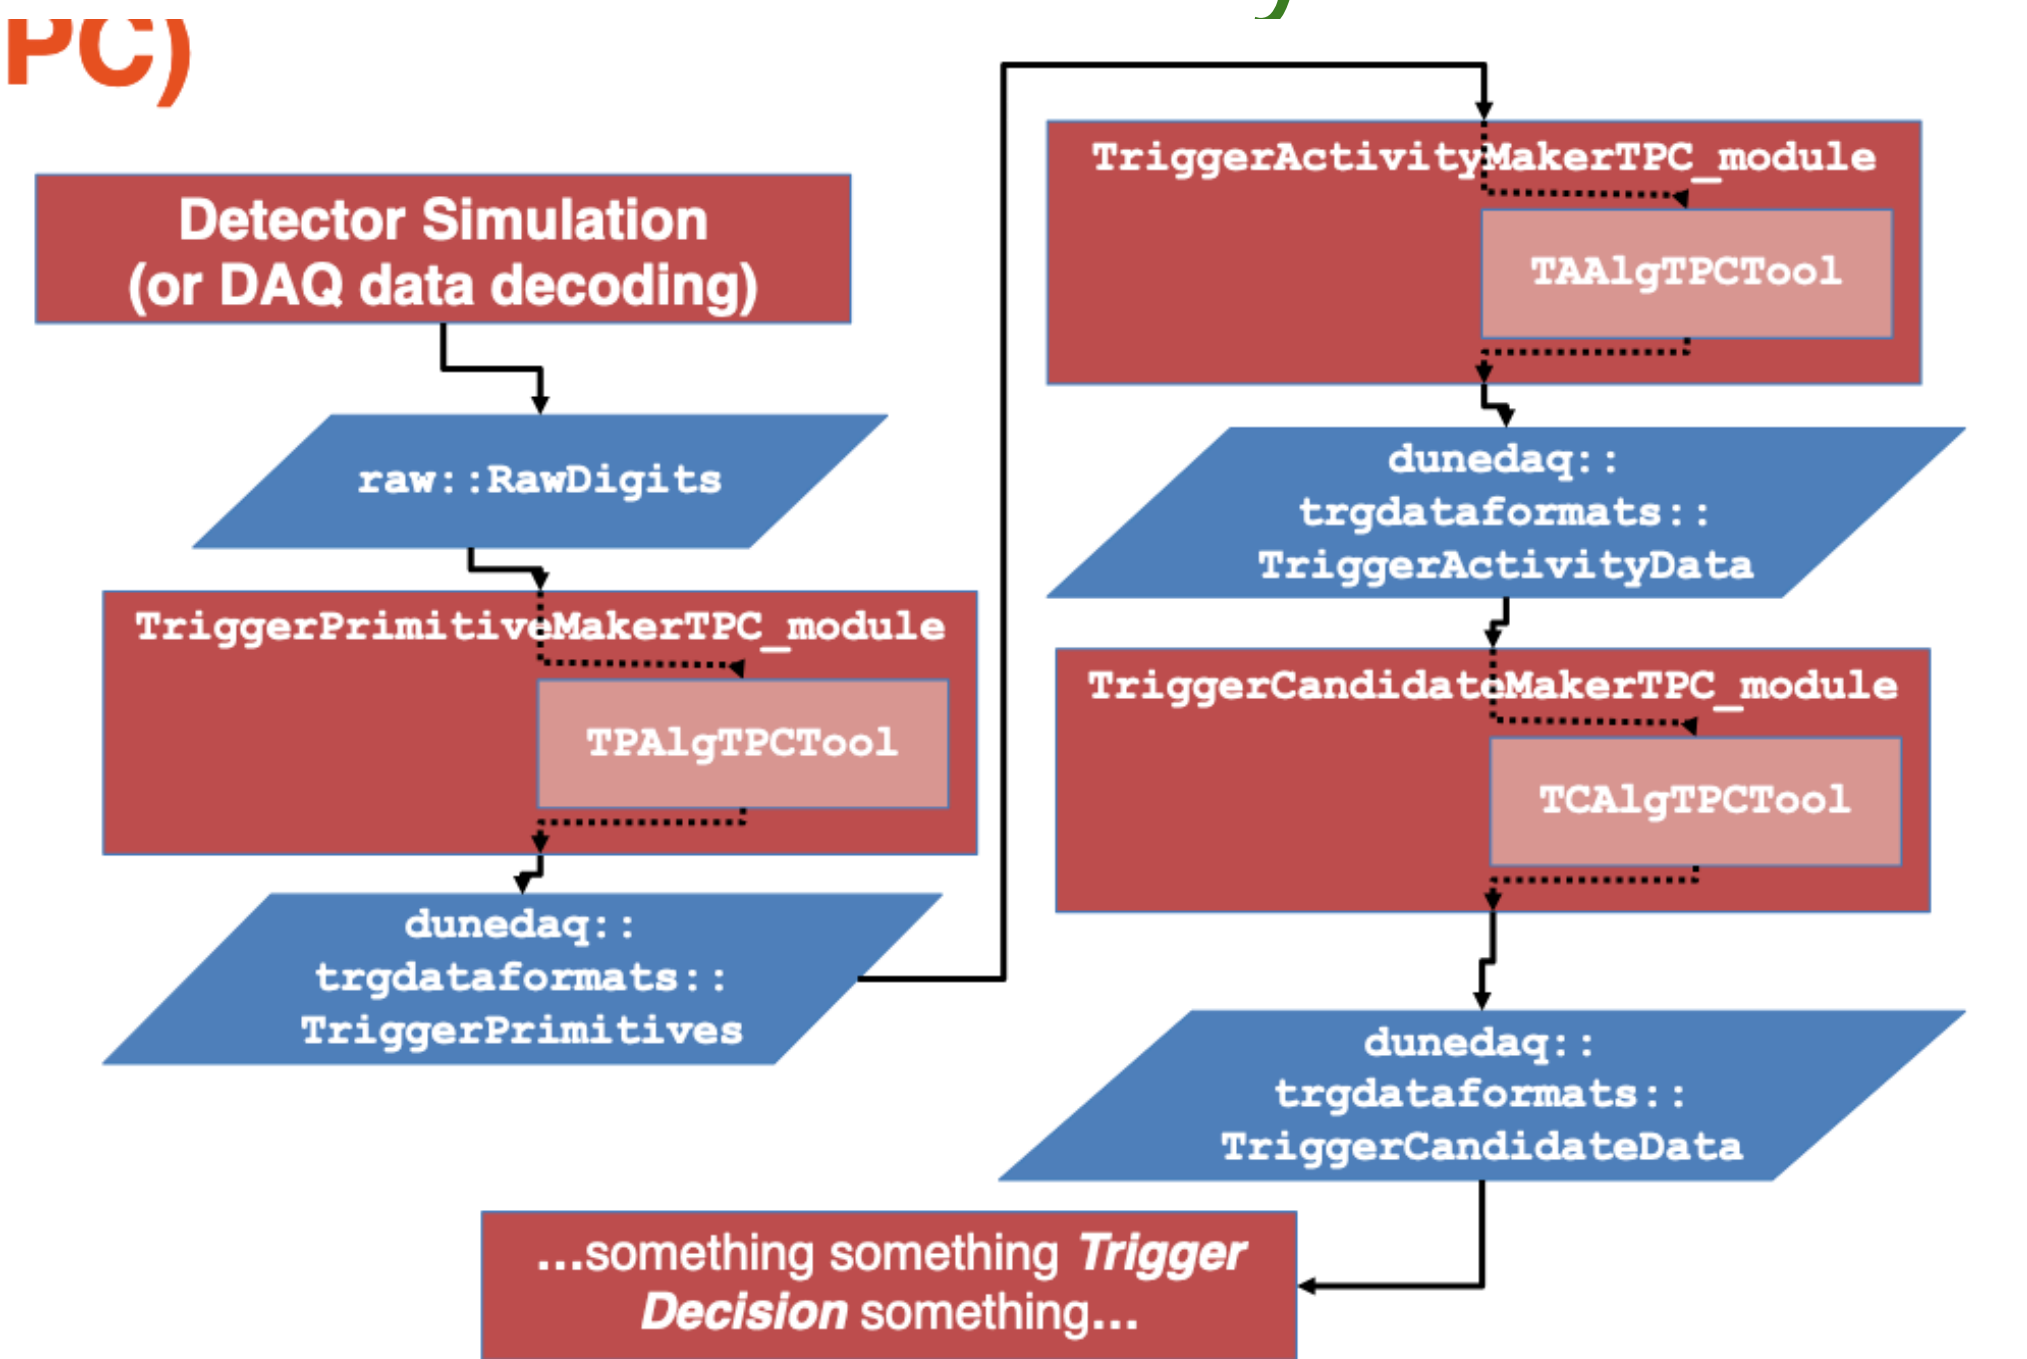
# Other New Trigger Algorithm Work

- Looked at DBSCAN again (Alejandro O.)
  - *May be too slow for our hardware*
- New ChannelAdjacency algorithm (Simranjit)
- New ChannelDistance algorithm (Alejandro O.)
- ADCSimpleWindow is still our “workhorse”
- Extensive TPG tests at NP04
- New lean definition for TPs developed (Alejandro O., Giovanna)




# Trigger Simulation Model (LAr-Trigger)

- Have not had a trigger model distributed to DUNE collaboration
  - Has led to a lot of confusion amongst collaborators on how trigger works
  - And a lot of work on impossible analyses (2-hits...) or complicated trigger algorithms
- Nor anything standard for us
- Work by several people (Wes, David D, James S, Simranjit, Klaudia) has resulted in ability to use much of our DAQ code within LArSoft!




run028508; ALL events

- Online: 82 TAs
- Offline: 82 TAs



• ADCSimpleWindow algorithm





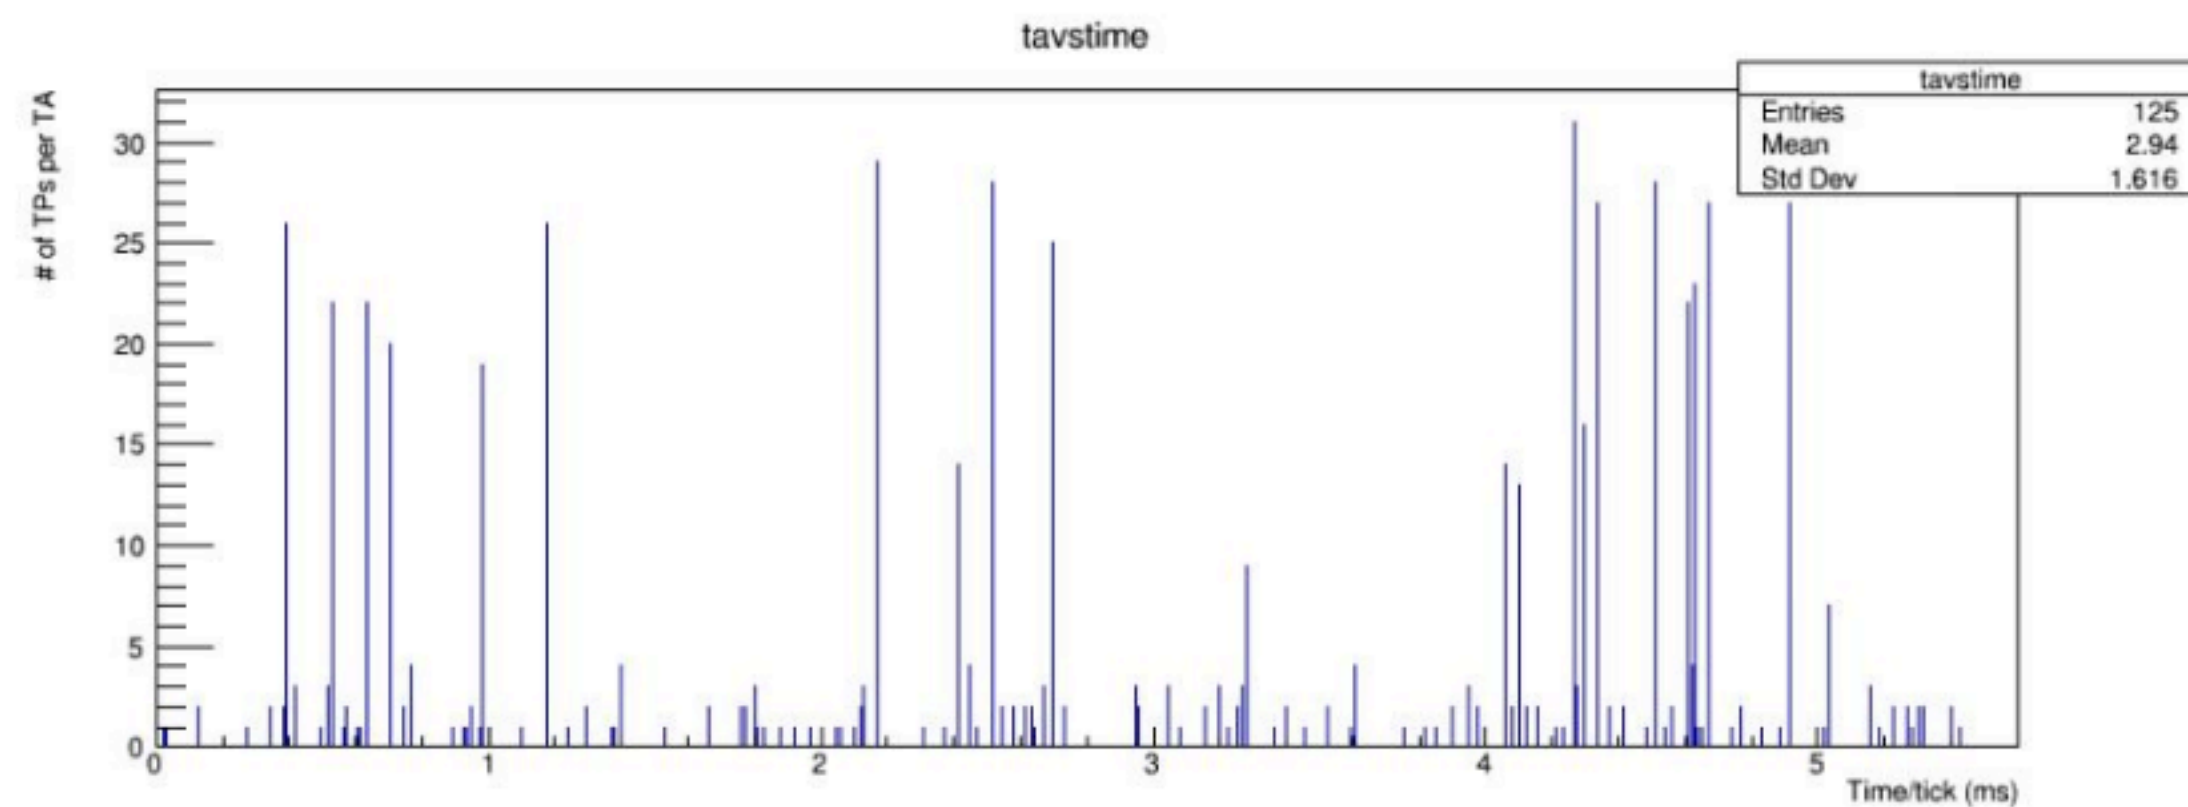



# Trigger Simulation Model (LAr-Trigger)

- Have not had a trigger model distributed to DUNE collaboration
  - Has led to a lot of confusion amongst collaborators on how trigger works
  - And a lot of work on impossible analyses (2-hits...) or complicated trigger algorithms
- Nor anything standard for us
- Work by several people (Wes, David D, James S, Simranjit, Klaudia) has resulted in ability to use much of our DAQ code within LArSoft!
- Gearing up for “official” DUNE MC data sets
- Distribution of model within standard LArSoft package
- Will re-do 6+ year-old studies of efficiencies and rates



# PDS TPs and TAs

- Major progress using DAPHNE TPs to generate TAs
- To date just an ADCSimpleWindow-like approach
- Will need to get this integrated into the simulation like TPC



M. Arroyave

J. Soto


Trigger



# Trigger Workshop!




OVERVIEW OF SYSTEM FOR NEW  
PEOPLE



UPDATES FROM PHYSICS/  
BACKGROUND GROUPS ON  
REQUIREMENTS



ALGORITHM DISCUSSIONS AND  
OVERALL SYSTEM DESIGN



“HACKATHON” DAYS TO  
IMPLEMENT CHANGES

**Week of April 7 at CERN**



---

+

•

○

# ***Timing System***



# DTS firmware and software in 2024

---

- **New firmware features**

- ▶ IRIG decoder
- ▶ DTS timestamp initialisation from IRIG (GIB)
- ▶ DTS timestamp initialisation from upstream DTS source (MIB)
- ▶ Sub-cycle (fine phase) measurement firmware


- **New software features**

- ▶ Support for new firmware features and hardware
- ▶ Timing control, configuration and monitoring ported to v5.x.x DAQ
  - ◆ Integration with drunc and configuration framework
- ▶ Timing “service” software requirements under development






# DTS hardware development in 2024



Pre-production FIB (v2) manufactured and commissioned  
FIB carrier FPGA board (AFC) passed production readiness review



- Pre-production AFC board in manufacture, expected mid-Feb



Pre-production GIB (v2) submitted for manufacture, expected end-Jan

Prototype MIB (v2) commissioned,

- Pre-production MIB (v3) in production, expected mid-January








***Readout Network  
Performance Evaluation***



# Readout Network – Switch Testing Progress


January 2024

Front (left) and rear (right) of rack in January 2024.



December 2024

Front (left) and rear (right) of rack in December 2024.



## Highlights

1. **10 GbE emulator** scaled to **96 links**, corresponding to **768 Gb/s**.
2. **40 GbE emulator** delivered, **28 links**, 672 Gb/s output.
3. Rx firmware delivered.
4. New industrial server to host eight PCIe hardware accelerator cards for Rx system.
5. **Power monitoring** via PDU.

## Also : 10 GbE Transmitter

1. Released variant of 10 GbE transmitter firmware for DAPHNE v3.





***Readout Servers  
Performance Evaluation***



# Roadmap

Most of the stages are well advanced, thanks to recent critical developments and contributions from:

Alejandro Oranday

Shyam Bhuller

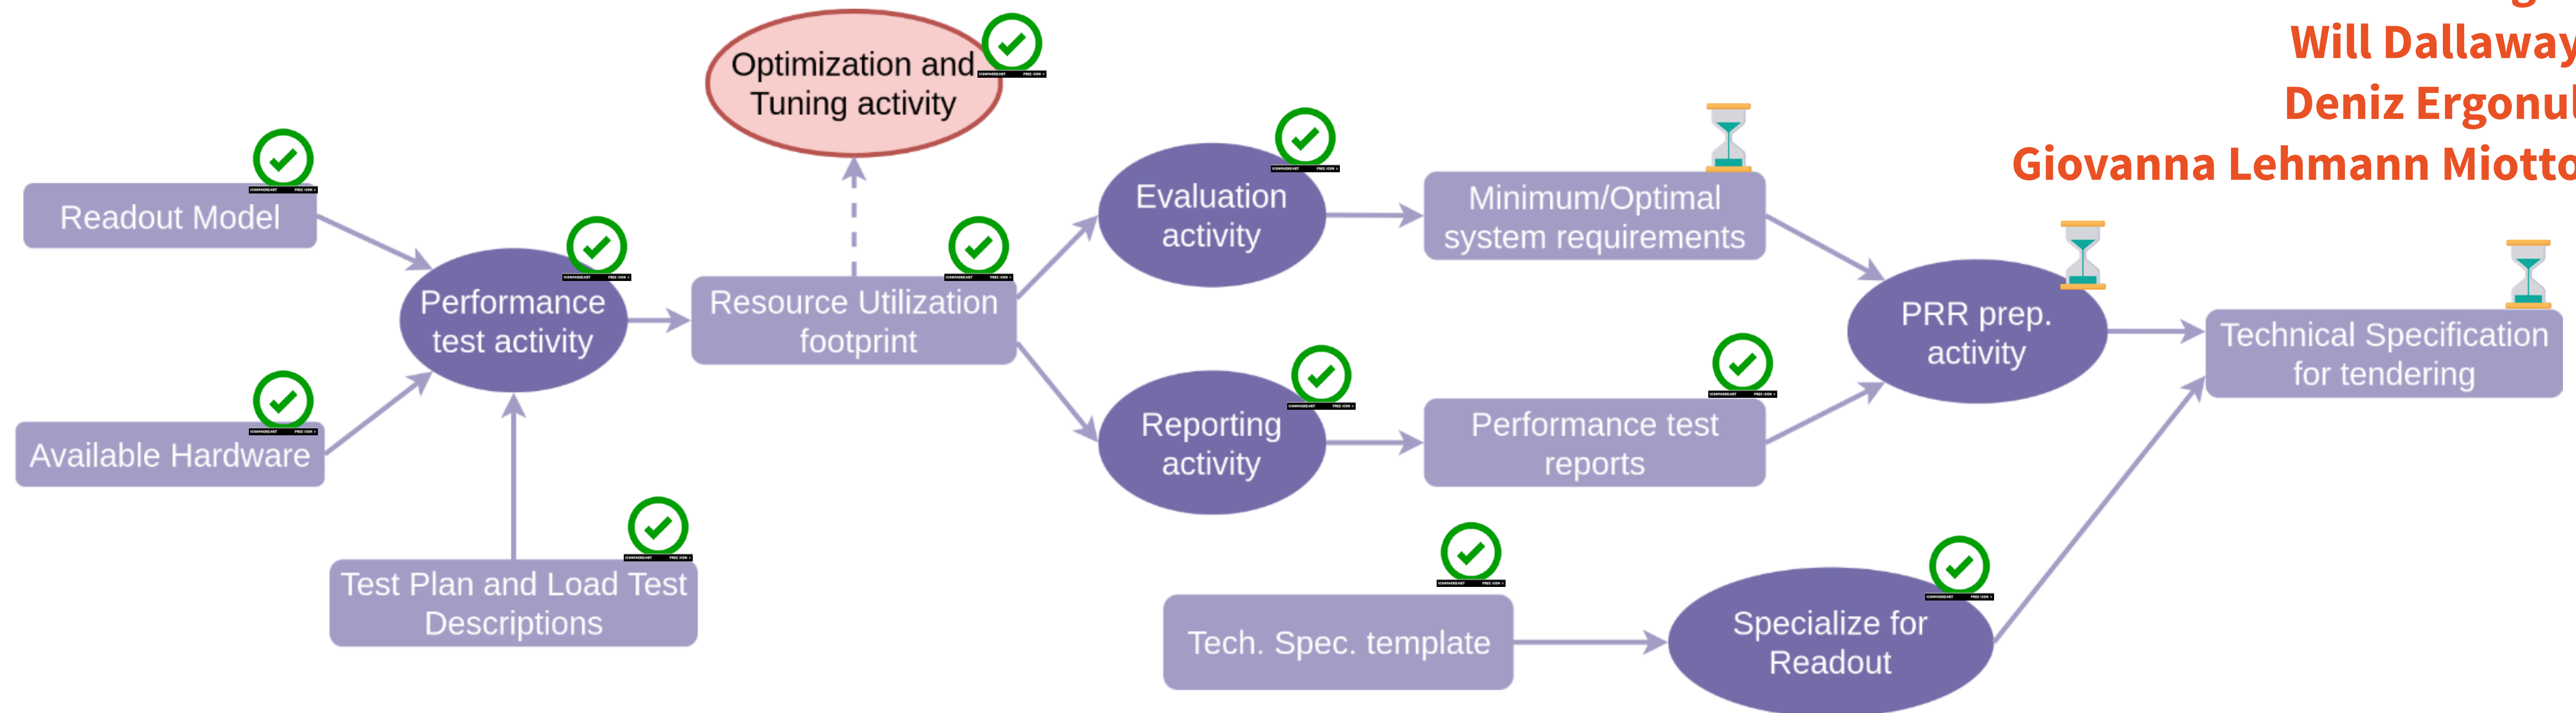
Matthew Man

Danaisis Vargas

Will Dallaway

Deniz Ergonul

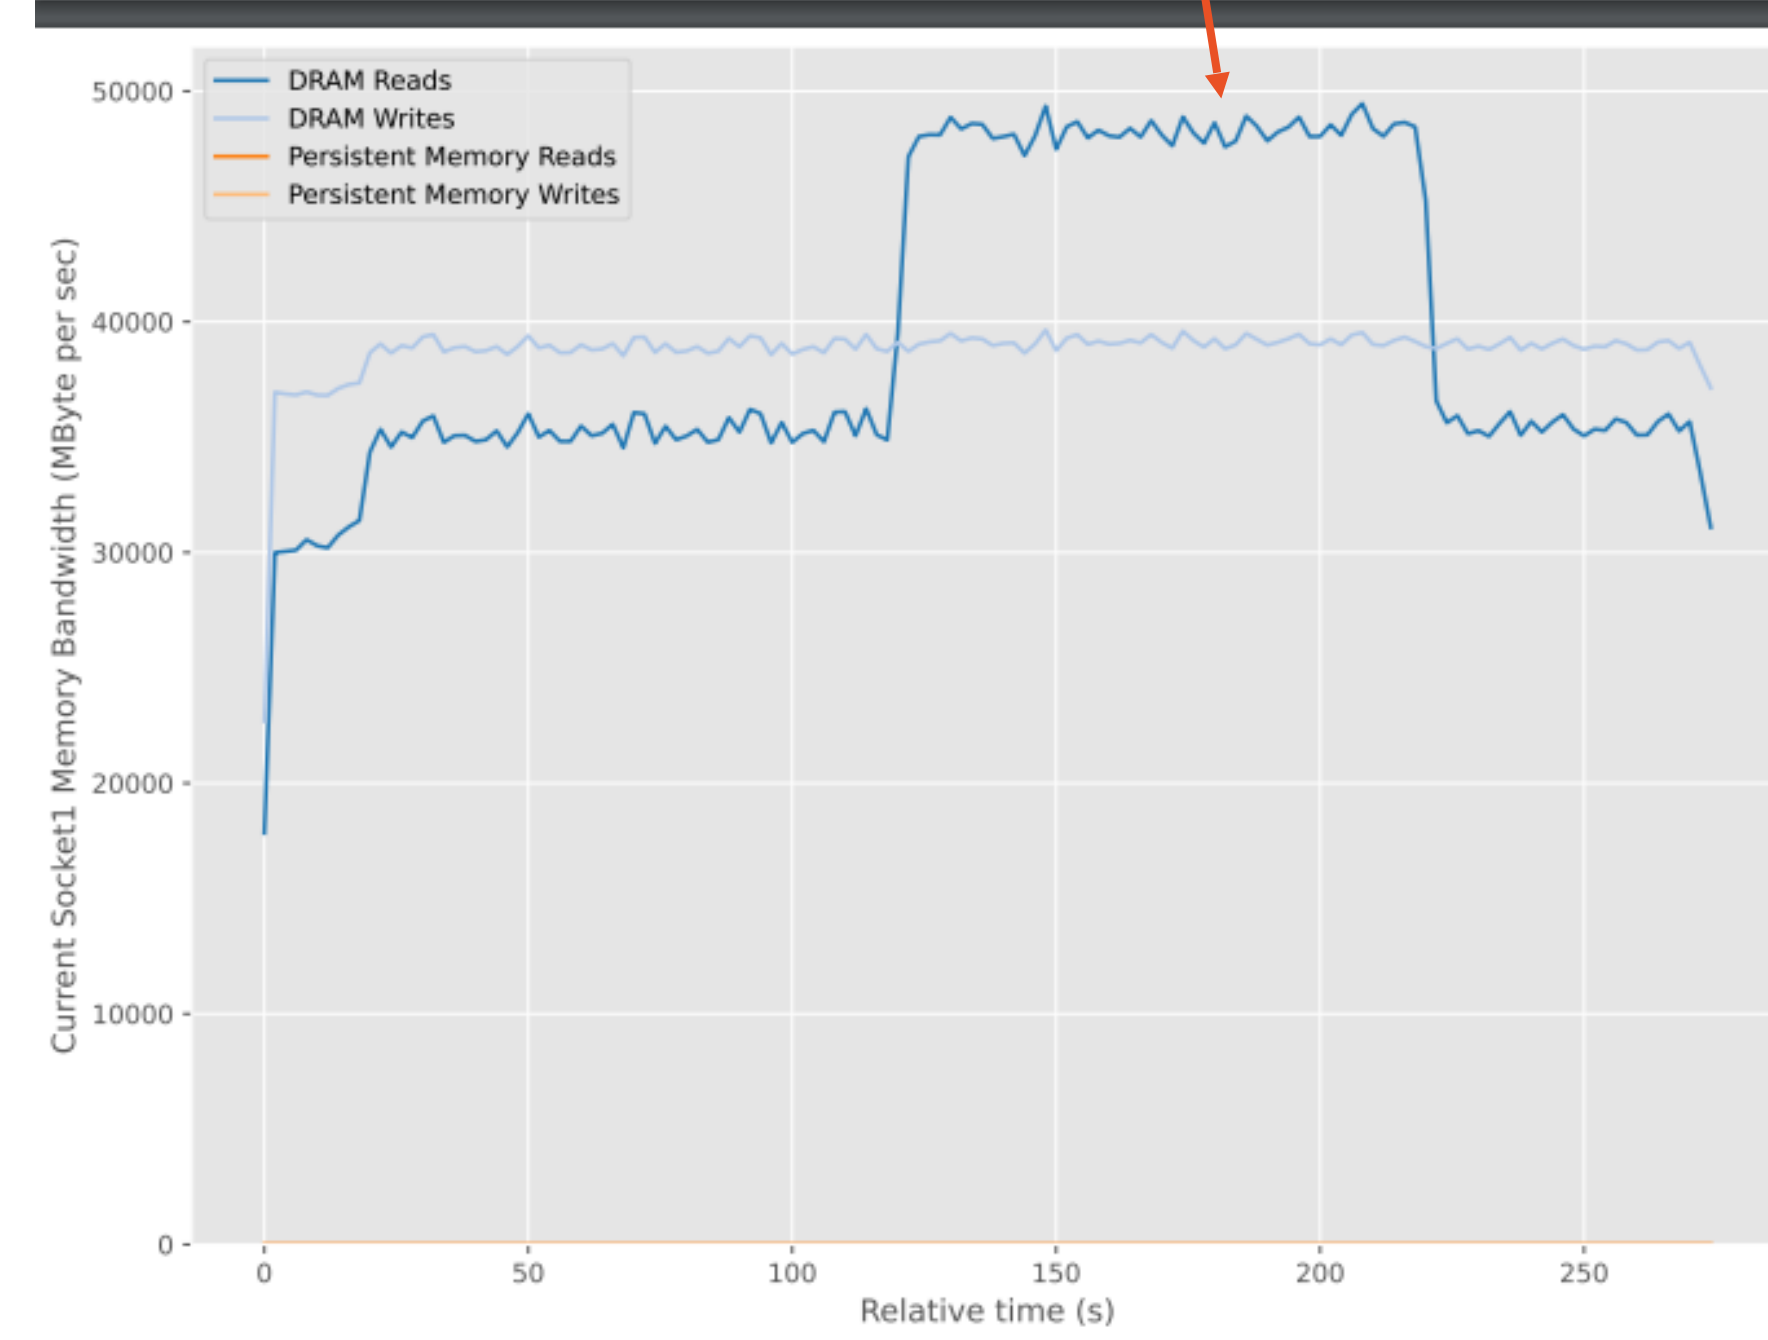
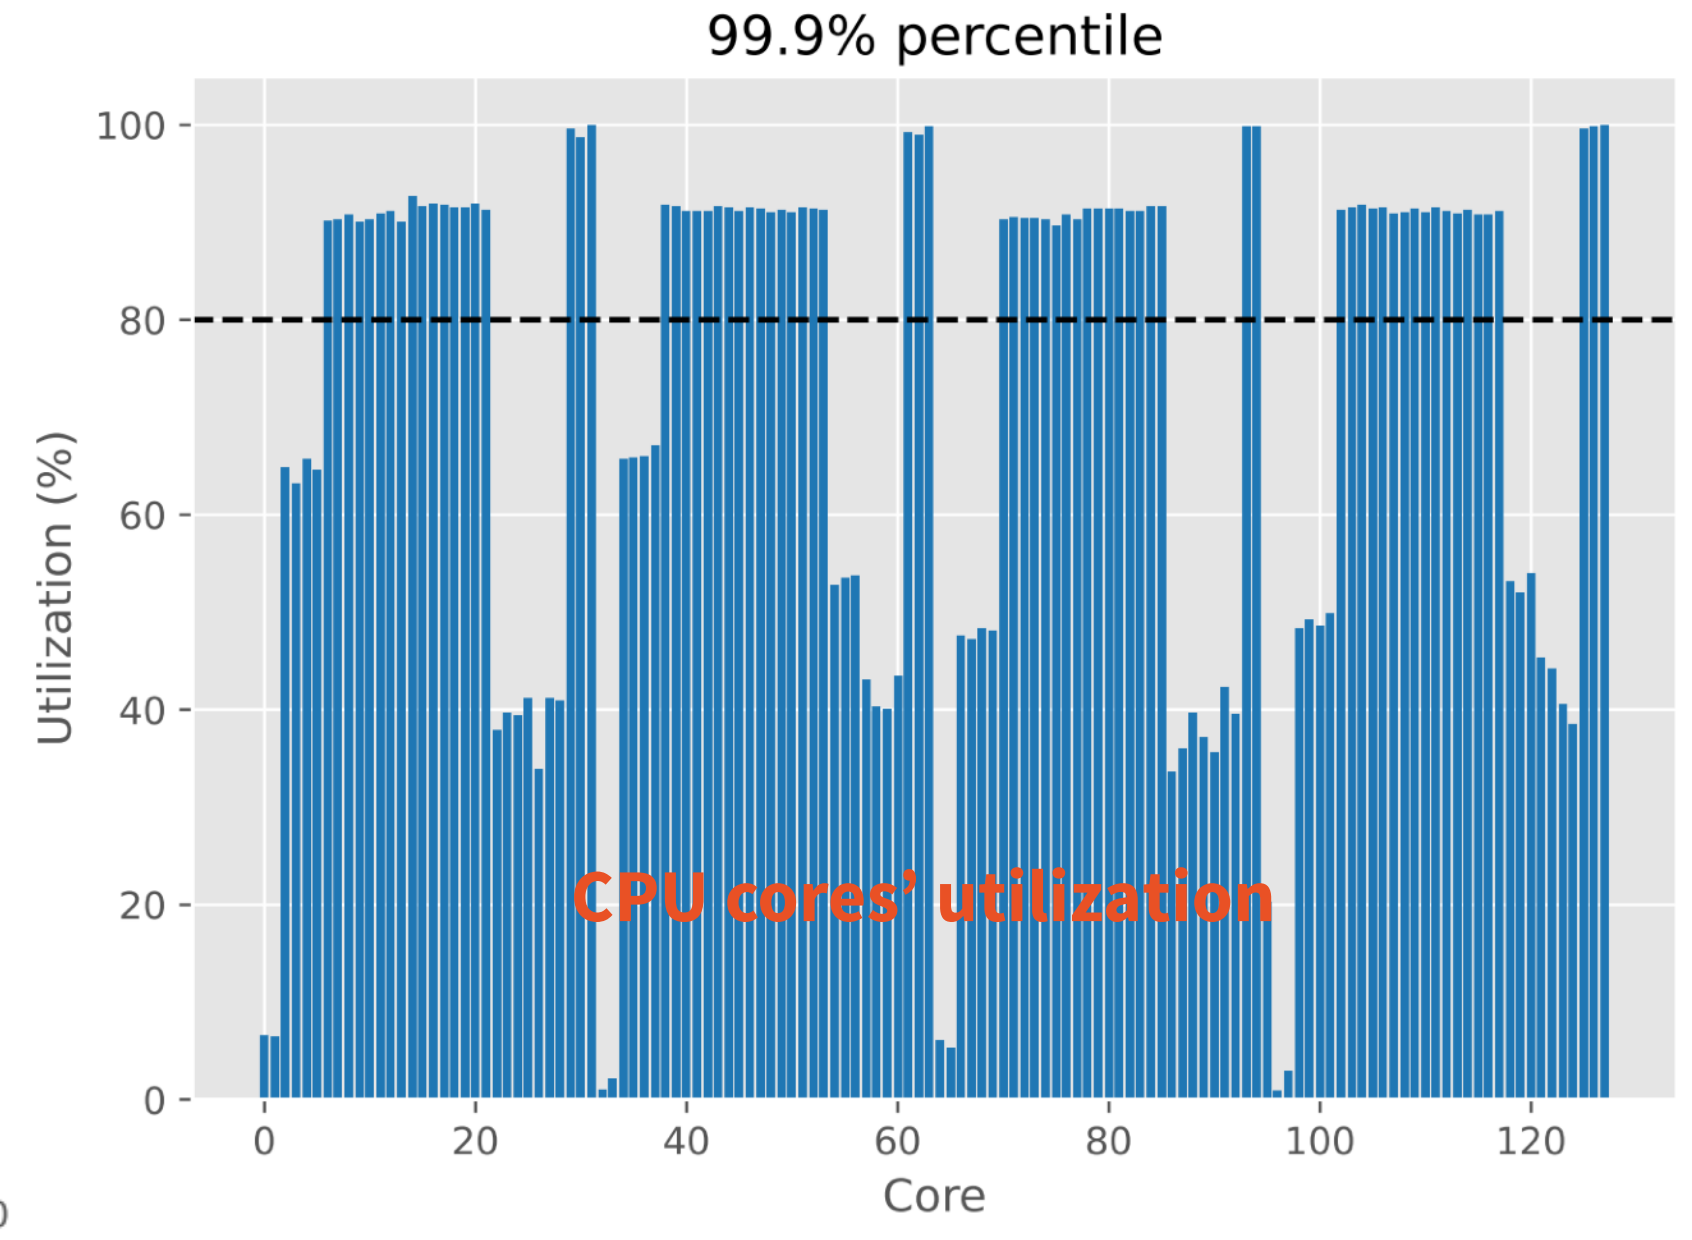



Giovanna Lehmann Miotto



And we also made a plan that focuses on the most outstanding issues  
Towards a Readout Server PRR



# Highlight plots

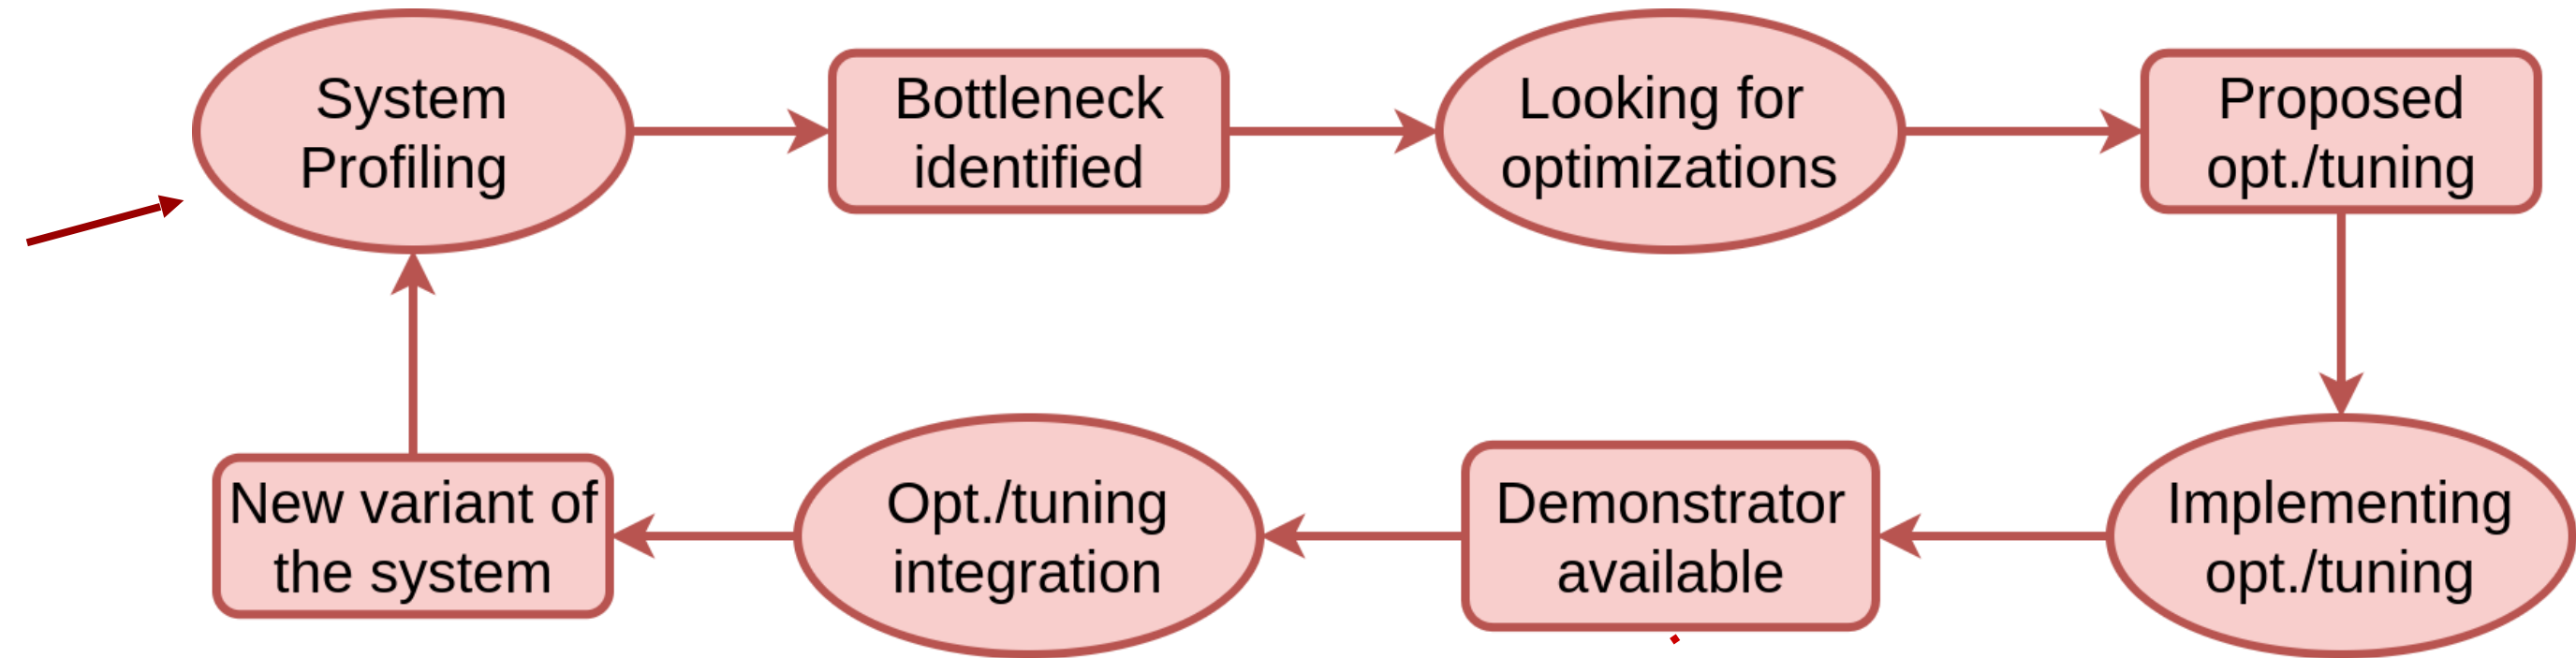


**Performance**



# Optimization and tuning activity

- Steps of a single stage of the optimization pipeline looks like:

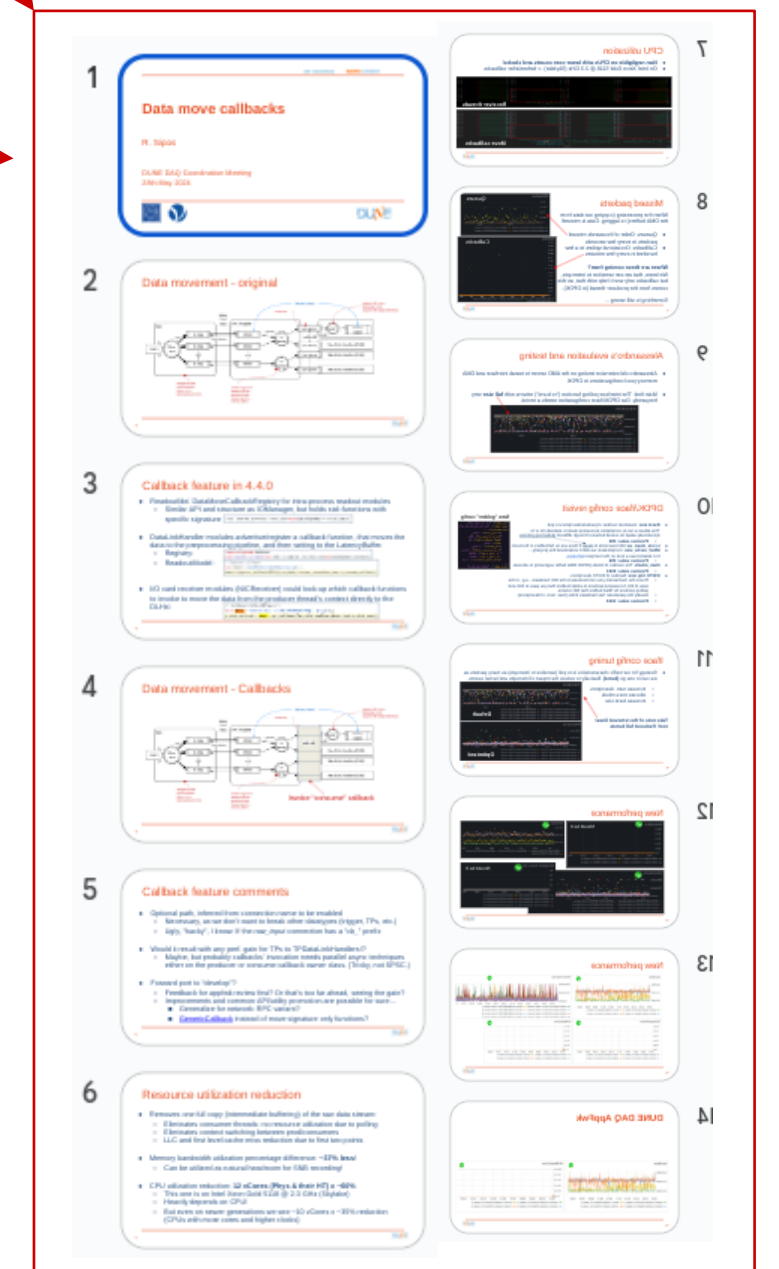


- Few examples (stages) of this work:

- Data insert callbacks for Latency Buffers
  - Received data type conversion to buffered data type
  - SNB store (RAID0, off-kernel, zero-copy mode)
  - Component placement differences
- Kernel isolation of data reception
  - AMD DMA latency
  - 100Gb NIC optimal polling config via DPDK
  - CPU sleep states, AVX mixed workloads, 10G NIC tune, etc.

SW

host



Performance



# Demonstrator for 4x baseline

- **Read out the whole NP04 detector with a single high-performance server**
  - 4 x 100 Gbps Ethernet interfaces,  
Supernova buffer on local high-speed storage elements
  - Goal is to validate every readout component operating on real detector apparatus (important for “Hit finding”)
- Motivation is the **power-draw reduction and price optimization** with this configuration
  - Factor 2 reduction of needed servers with this topology
- **Load balancing and resource isolation techniques are essential** to reach deterministic performance on quasi-real time COTS hardware and software
- Performance testing activity includes **analysis of performance reports that lead to further optimization options** that are continuously fed into the readout system

**Run 32849**

**np04-srv-031**

(technical specification for procurement provided by CERN)



Established DAQ performance test conditions:

- ~ 6 MHz of Trigger Primitives (TPs)
- 5 ms of readout window @ 5 Hz (network limit)
- ~ 7 Gbps + ~1.5 Gbps (TPs) sustained
- Most complex feature extraction algorithm used
- SNB recording sustained at NVMe bw. limit
- No packet loss with stable resource utilization
- **Quasi DUNE FD conditions achieved!**



*How Did It All Turn Out,  
Really?*





# A lot was achieved at the cost of significant delays in some areas

---

- **Planning effort at the start of 2024 helpful, but not sufficient to handle all parallel activities**
- **Operation and support load heavier than initially estimated**
  - ▶ Partly due to our tendency to under-estimate it, partly due to a shifting detector schedule
    - ◆ NP04 beam schedule defined late and then extended, testing, NP04->NP02 LAr transfer uncertainties
- **Complexity of planned changes greater than expected**
  - ▶ Transition from nanorc to drunc proved challenging
    - ◆ and while completed, feature parity has not yet been achieved
  - ▶ Transition to an object-based configuration model exposed our partial understanding of configuration organisation and management
    - ◆ here again the transition is complete, but a significant amount of work remains to understand and use the new system
- **Resource uncertainties required re-planning and in-flight adjustments**
  - ▶ PRRs planned for 2024 have now moved to 2025





# *What Lies Ahead in 2025?*









2024 has been **very, Very, VERY** challenging

We couldn't have done it without the hard work, dedication, and commitment of **every single member of this consortium.**

Despite being a small team, often pulled in many directions, and spread across two continents we've **pushed forward on so many fronts and accomplished lots.**





A sincere **THANK YOU!** to everyone for your energy and perseverance—it's truly made all the difference.

We wish you and your loved ones a **happy, restful holiday season and a peaceful close to 2024.**

Here's to an exciting and successful 2025 ahead!

