

ALICE O2 setup for TPC read-out tests for the LHC run 3

Johannes Lehrbach, Filippo Costa
for the ALICE collaboration

johannes.lehrbach@cern.ch

Introduction

The Time Projection Chamber (TPC) is one of the first detectors to build a small test setup which is comparable to the run3 Online-Offline (O2) computing system. The setup consists of 5 First-Level-Processing servers (FLPs) with two-four Common Readout Units (CRUs) each as input nodes and 2 Event-Processing-Nodes (EPNs). The primary goal is to test the whole readout chain at conditions which are very close to the expected running conditions during run3. It's one of the first small scale setup with the CRU to readout Frontend Cards (FECs), where 4-10 CRUs are used in a single setup. The setup will be extended even further during LS2 to enable the commissioning of the TPC sectors in the clean room. The secondary use case of this setup is to test the O2 software in a realistic scenario with real data as input. Due to a delay in the availability of the CRUs the build up of the test setup was delayed and happened just before CHEP.

Main objectives of the TPC test setup

Primary goals for the TPC group

- TPC readout test of the Frontend Cards (FECs) in realistic conditions
- Testing of a whole TPC sector after the first readout tests
- TPC sector commissioning starting in LS2 2019

Secondary goals for the O2 use case

- Providing a software environment which is as close to the run3 environment as possible
- Testing O2 software and tools with real detector input
- Integrate as much functionality as possible and gain experience with the complete software stack

Test setup

The setup consists of one Login node which is connected to the CERN general purpose network and provides all necessary services. The initial setup for the first testing phase has 5 First Level Processor nodes, which house 2 CRUs each. This setup can be easily extended with more FLPs in case additional servers are available. For the first phase test of the readout the data analysis is not foreseen and only 2 Event Processing Nodes (EPNs) are foreseen which allows already a first small test of the data flow at reduced rate.

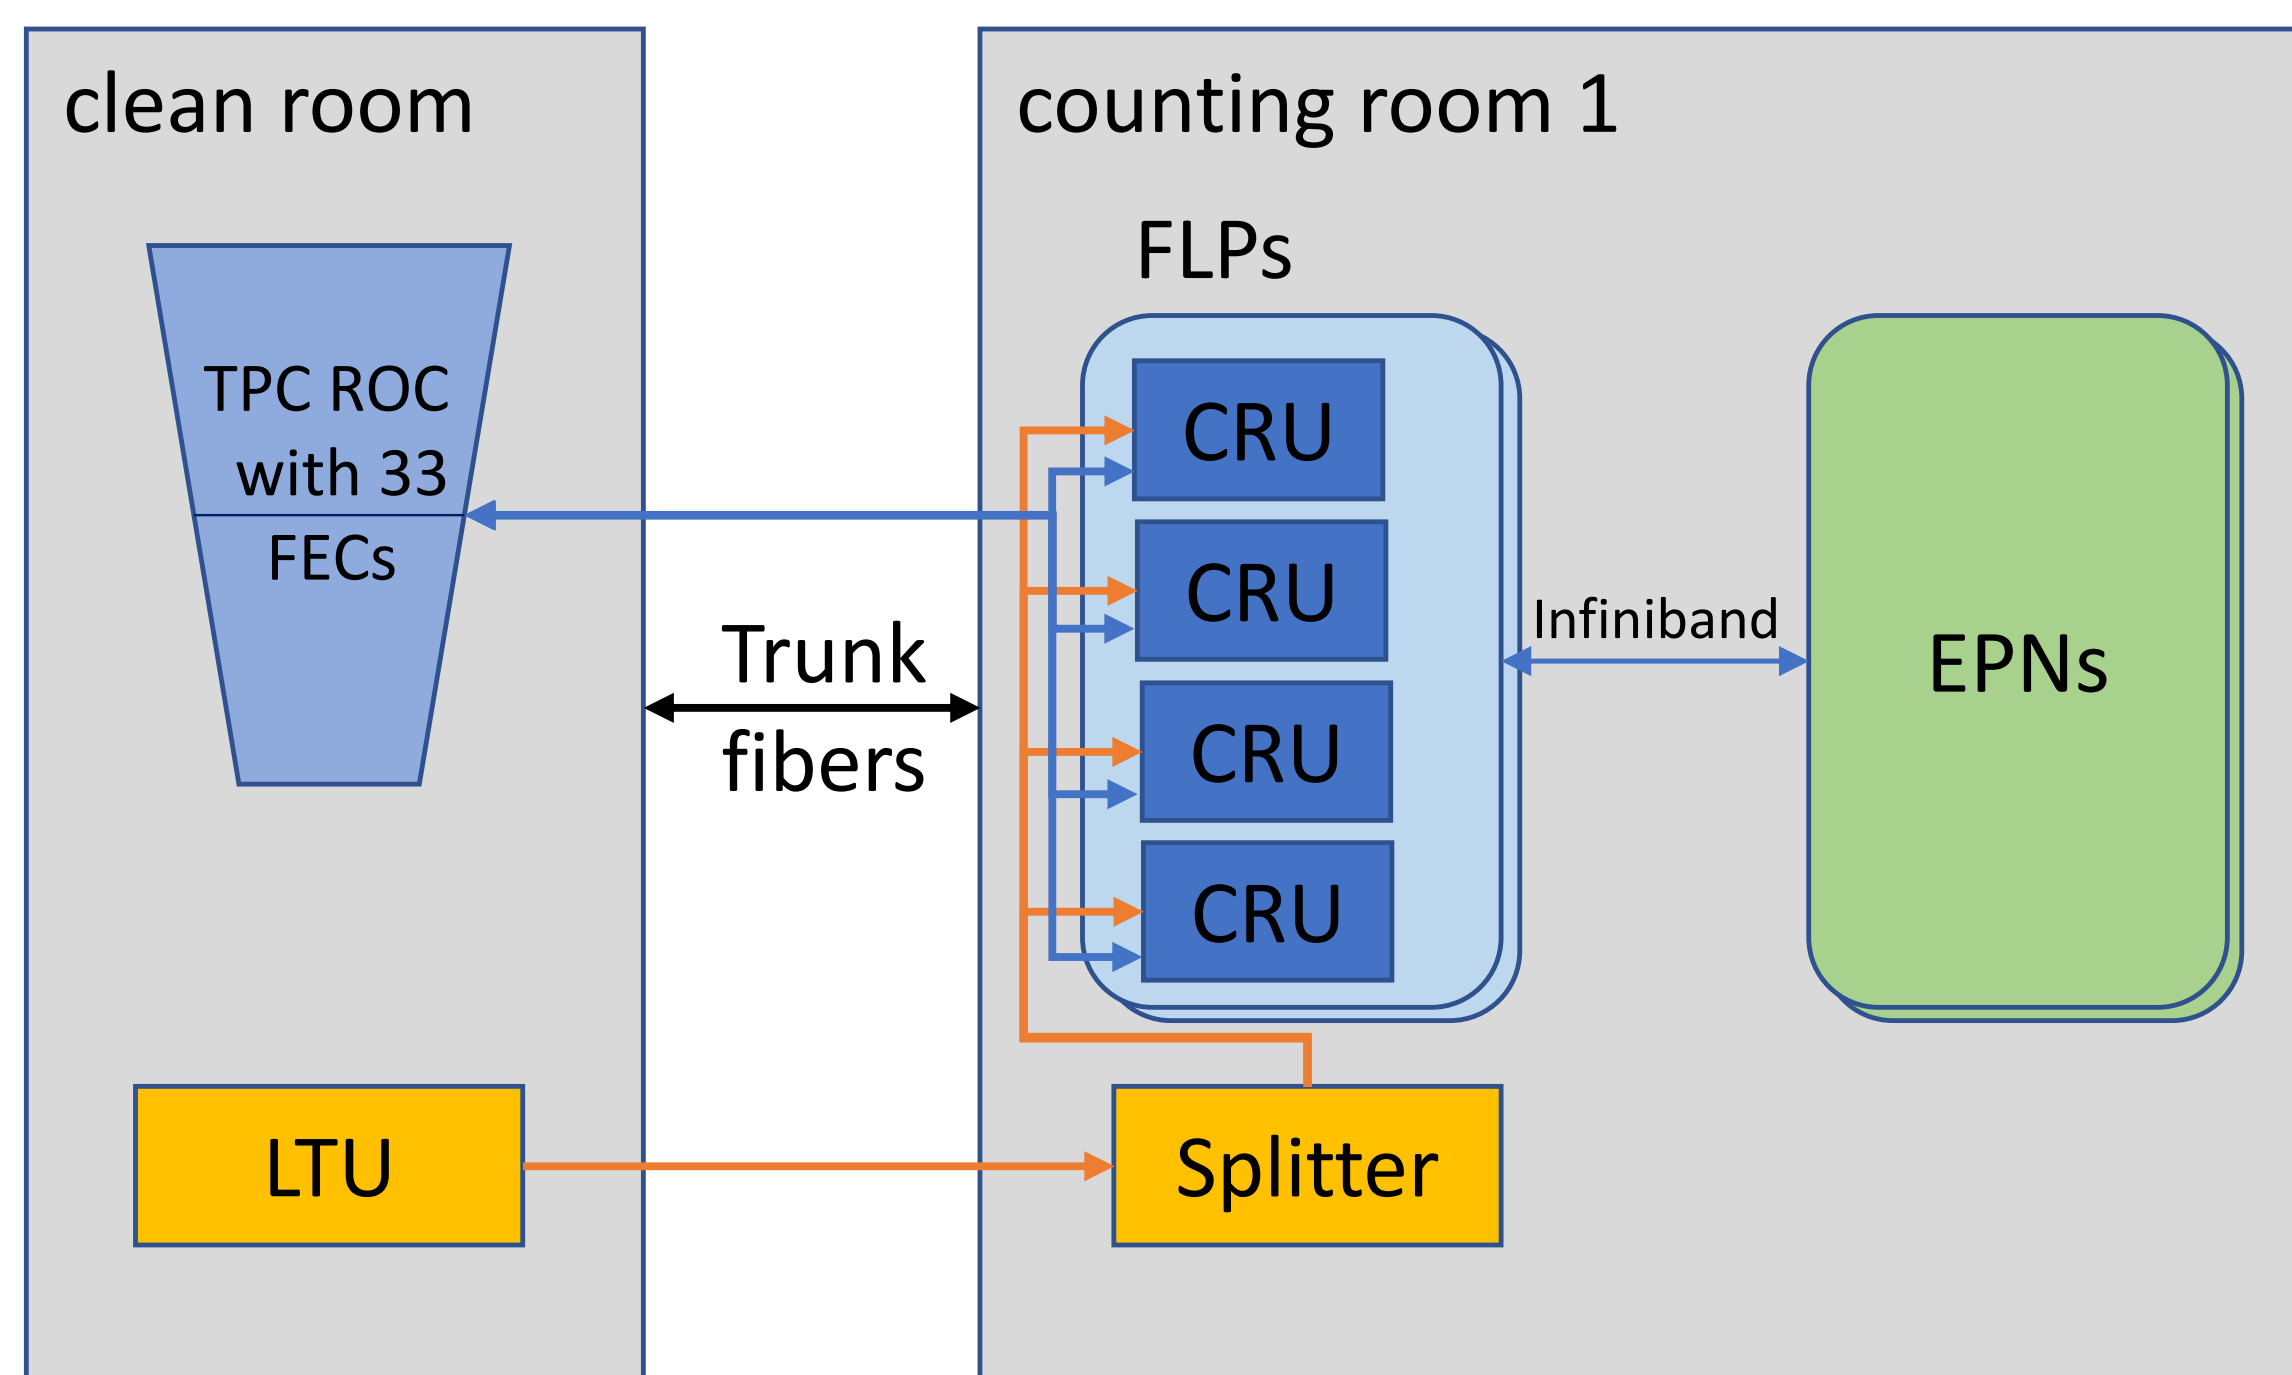


Figure 1: Test setup at Point 2 in the first phase: TPC Inner Readout Chamber (IROC) with 33 FECs and a Local Trigger Unit (LTU) in the clean room which are connected to the FLP servers with 2-4 CRUs in the Counting Room 1 (CR1). Additional EPNs to have the possibility to test the time frame building

O2 integration cluster

The ALICE HLT development cluster which consists of 60 run1 servers became a new purpose as an O2 integration cluster. The goal of this cluster is to provide an realistic test setup for developers to test their specific part of software as well as integrate this into the common O2 software. As an example, the run3 monitoring solution with Collectd, Influxdb and Grafana was integrated there and is now available on this cluster to assist developers and collect additional needs to improve the whole system toward run3. As soon as the software is sufficiently tested we will also provide this to dedicated detector test setups like the TPC test setup. One of the goals providing the tools in an early development phase is to gather feedback and feature requests towards run3 to accelerate the O2 software development. Currently the time frame building and load balancing between FLPs and EPNs is tested and will be made available to the TPC test setup afterwards.

Network

The O2 integration cluster provides a 1Gbit Ethernet network for provisioning, control and monitoring as well as a faster 40Gbit Infiniband network for data transport. O2 software tests to integrate RDMA functionality of the data transport were successfully performed on this test cluster. The TPC read-out test has already a newer generation of Infiniband with 56Gbit. IP over Infiniband (IPoIB) is configured everywhere to provide a normal IP NIC interface via Infiniband.

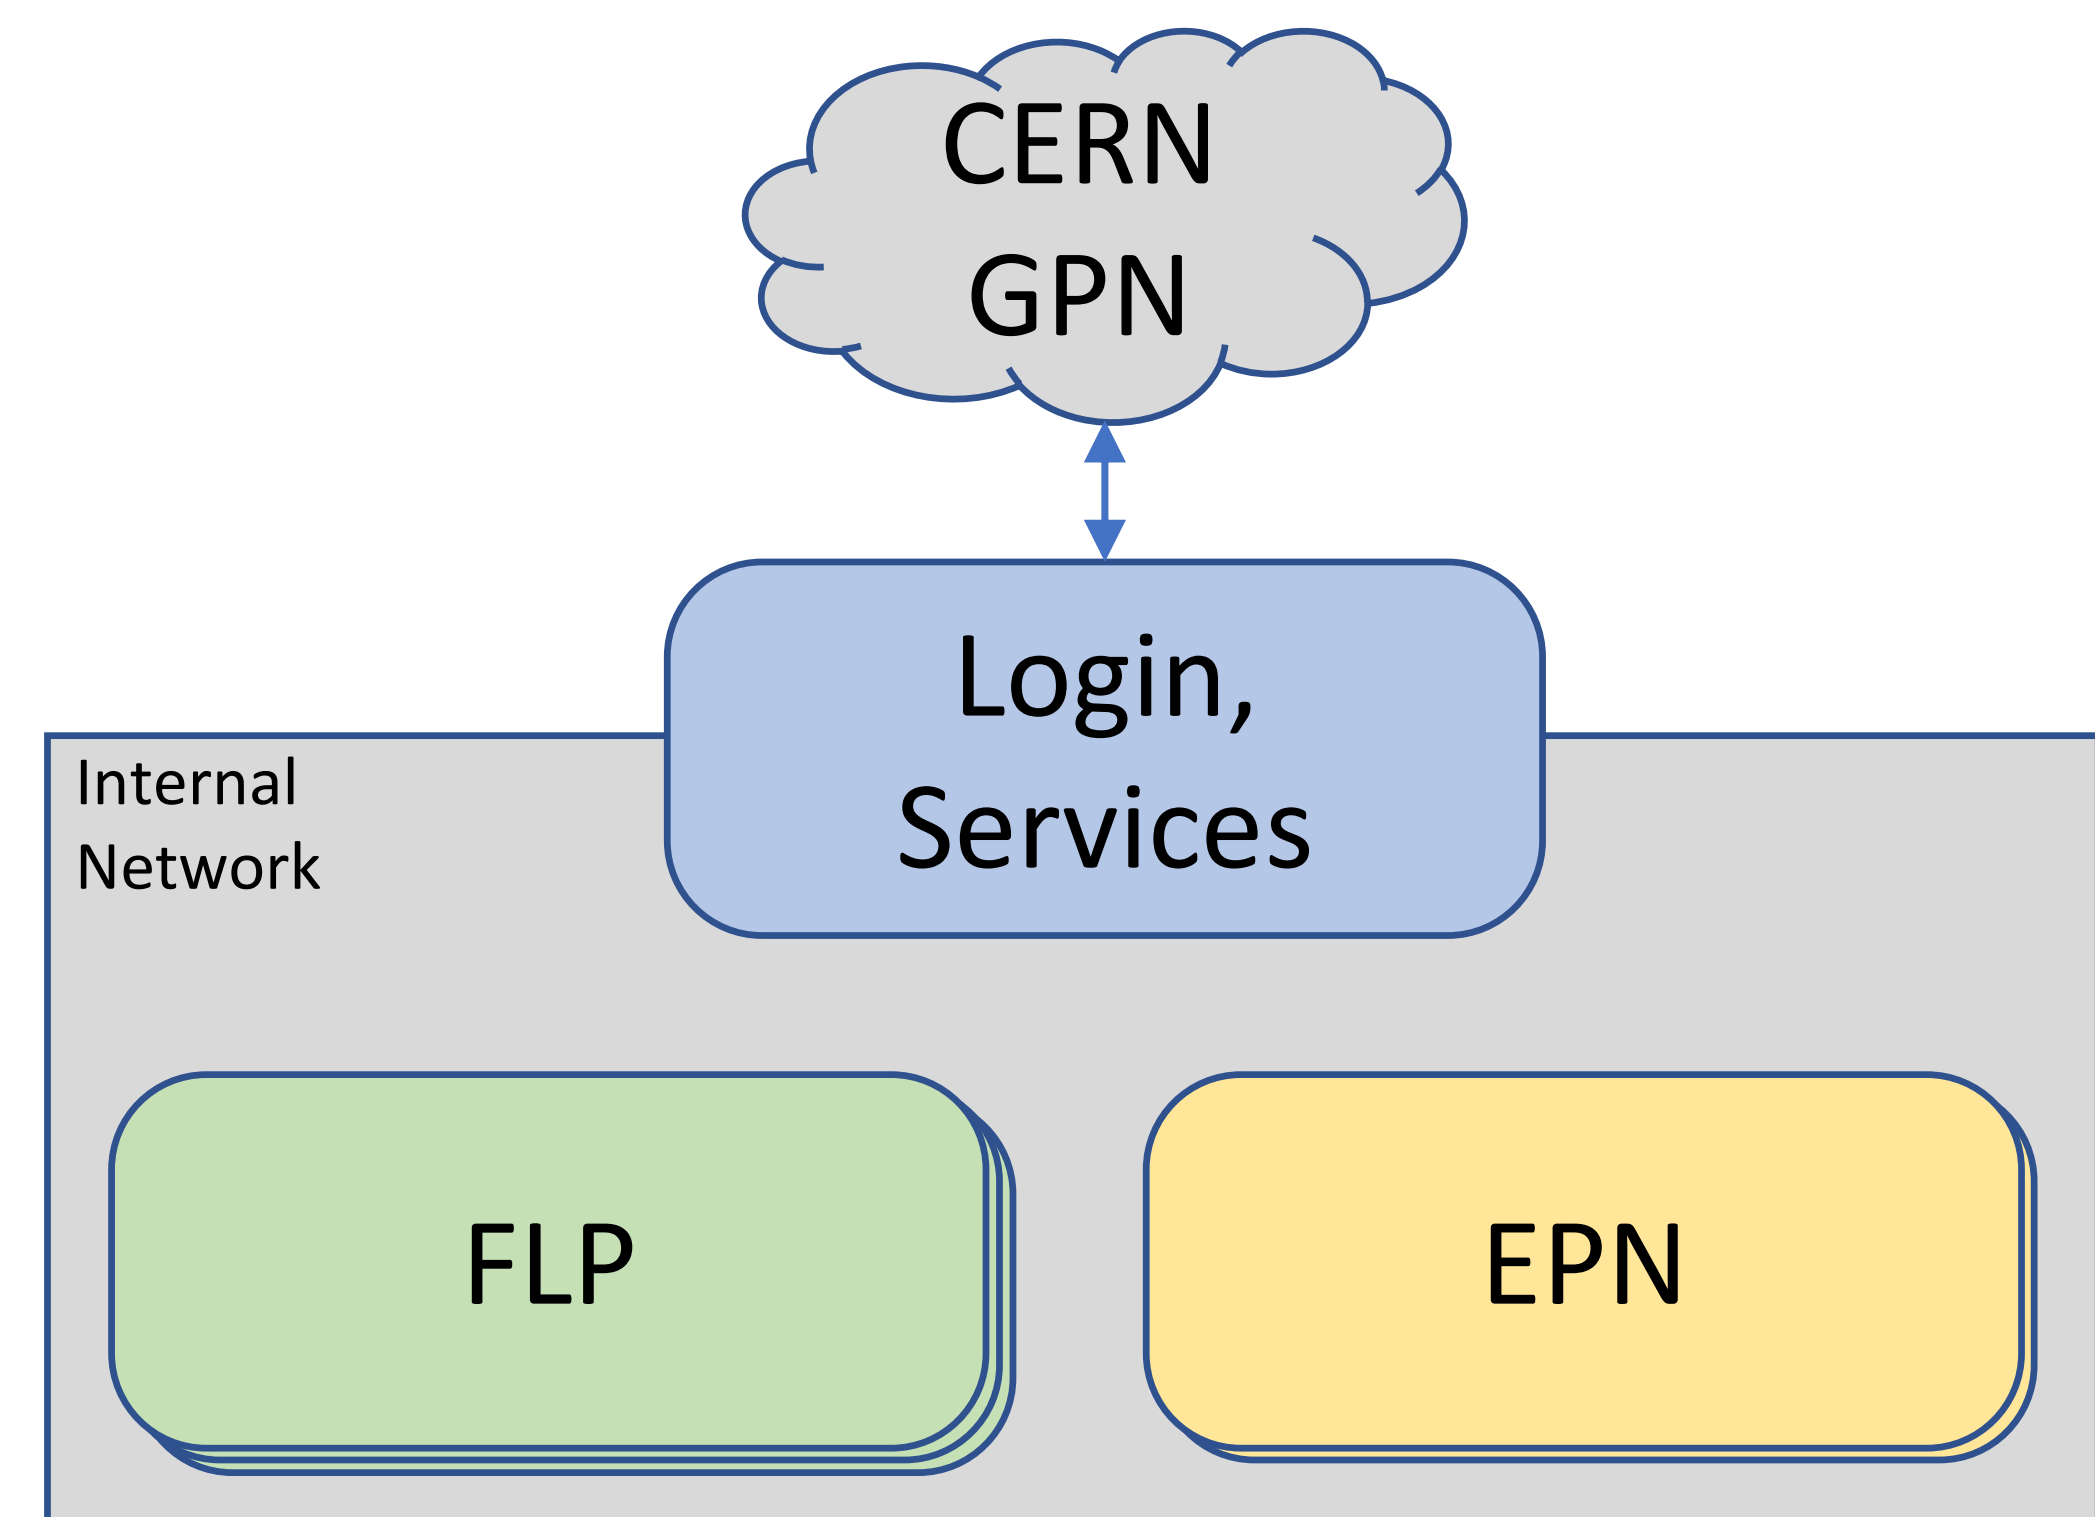


Figure 2: Schematic of the network setup for the test: GPN connectivity for remote login via a gateway/login node. Internal 1Gbit Ethernet network for control via IPMI, provisioning and monitoring as well as Infiniband FDR 56Gbit for data transport

Tools

To make the test setups as reliable as possible from the software side we try to automatize as much as possible. This way we can guarantee a defined setup all the time by simply re-applying the configuration again or re-building everything in the worst case. The current baseline tools for run 3 for provisioning and configuration is Foreman and Ansible. With the newest releases of Foreman we can use a plugin to easily integrate the Ansible playbooks and assign these to Groups or single hosts. For run2 we mostly used Puppet for basic configuration, this means a change for the configuration software and a development effort to provide all the necessary Ansible playbooks for run3 which already started now also for existing parts of the configuration. The operating system used for all developments is Cern Centos 7 which is also the current baseline for run3.



Results

So far, starting automatizing already in an early phase for development has proven to be a good approach to provide a comparable testing environment for everybody. Minimizing manual steps also decreases configuration errors even though there are scenarios where a generic solution seems to be not good enough. So far, no conceptual problems were discovered with the chosen tools. The current software and tools work well on this small to medium tests. This has to be scaled up for run3 of course. So far, there were no indications that the scaling could be a problem.

Conclusions

The TPC read-out tests are a nice opportunity for the O2 project to test the software in parallel with the whole read-out chain. The needs for these tests accelerated the developments to meet the detector requirements. The expected feedback from the detectors as users of the O2 software is important input for further developments. The tools for provisioning and configuration work fine and guarantee a reliable basic setup for everyone.

Next steps

After the initial deployment of the test setup at Point 2 further testing is needed to ensure everything is working. In the beginning only 4 CRUs are available for the IROC readout test. New parts of the O2 software like load balancing will be added step by step as soon as they are available and tested. Later this year the setup will be extended to 10 CRUs to enable the sector tests when the additional CRUs become available.

During LS2 this setup will be extended to 20 CRUs to test two sectors in parallel and enable the surface commissioning of the TPC in the clean room at Point 2.

Acknowledgements

Thanks to the TPC group who opened their read-out test setup for the O2 project to build a realistic testbed for further software and integration tests. An early possibility to test everything even in a small scale is very helpful to improve the software and verify the current developments. In addition, thanks to everyone in the O2 project contributing to the setup as well as everyone testing and integrating their software and tools.