

# Modular Toolsets

*for online integration of  
high performance computing clusters*

Mohammad Al-Turany, Dennis Klein, Anar Manafov, **Alexey Rybalchenko**  
Software Development for Experiments (SDE) group, GSI Helmholtz Centre for Heavy Ion Research

**Jennifer2 Workshop on fast Realtime DAQ and Trigger Systems**


Gießen, April 8, 2024




# HPC Clusters

for particle physics experiments


Today HPC clusters are connected directly to the data acquisition systems and integrated into the online systems of particle and nuclear physics experiments.



3.4  
TB/s





90  
GB/s




## Dedicated Cluster (ALICE)


- Under the experiment control
- Can have a resource management system (Slurm), but must not



1.0  
TB/s




1  
GB/s




## Generic Cluster with a dedicated online partition (GSI)


- Under IT control
- Resource management system (Slurm)



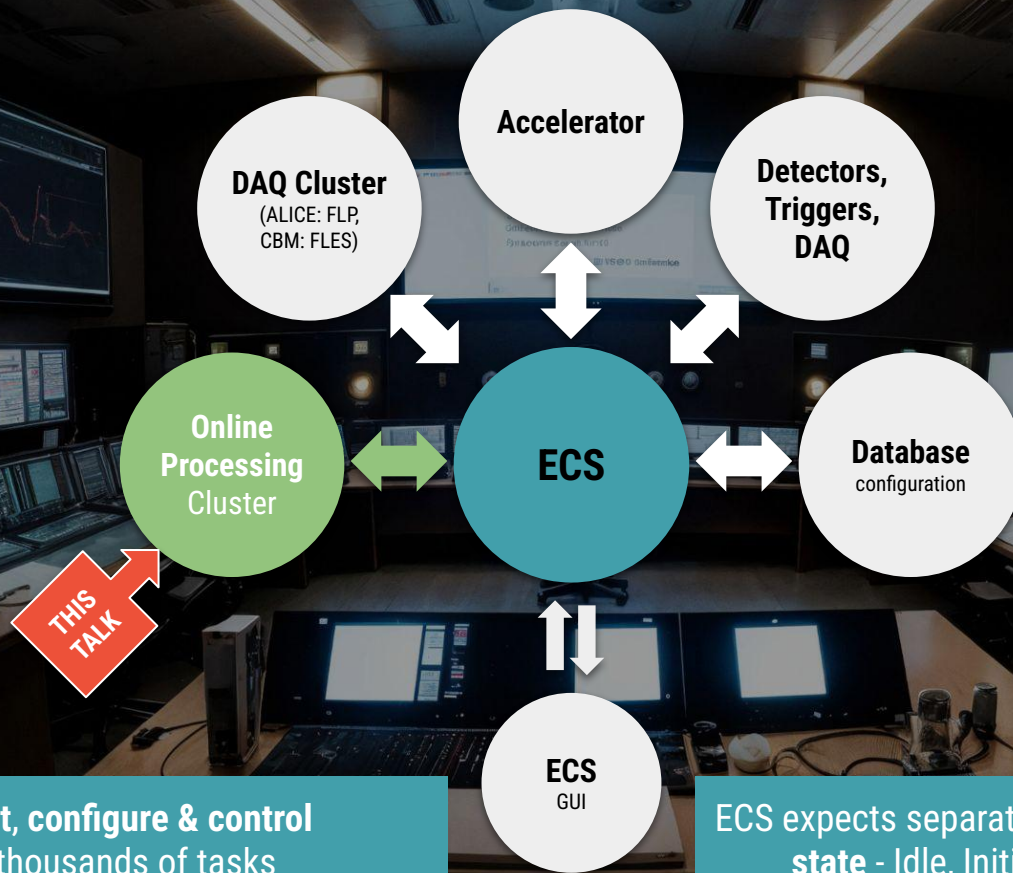
0.3  
TB/s



1  
GB/s



# Experiment Control Systems



Deploy, connect, configure & control  
hundreds or thousands of tasks

ECS expects separate subsystems to be in a **defined state** - Idle, Initializing, Running, Error, etc.

# Alice & FAIR Collaboration

modular software toolset

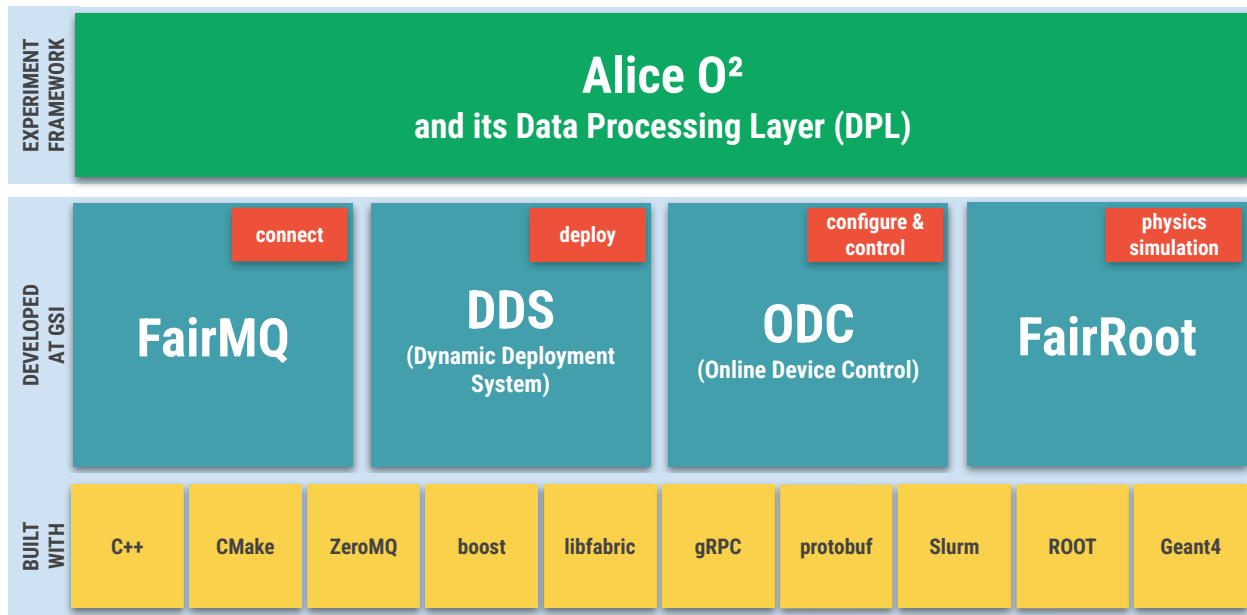


+



ALICE

In a joint collaboration between GSI, FAIR and the ALICE experiment at CERN, a modular set of open source tools for simulation, reconstruction and analysis of particle physics experiments was developed. **It extends the FairRoot framework to provide building blocks for highly parallelized and dataflow-driven processing pipelines** required by the next generation of experiments.




FairRoot

<https://github.com/FairRootGroup/FairRoot>  
<https://doi.org/10.5281/zenodo.3896282>

Alice O<sup>2</sup>

<https://aliceo2group.github.io/>

# Tools developed by the SDE group for online experiments



**FairMQ** is a C++ message queuing framework and library that integrates standard industry data transport technologies and provides building blocks for creation of data flow actors and pipelines (**topologies**).

connect

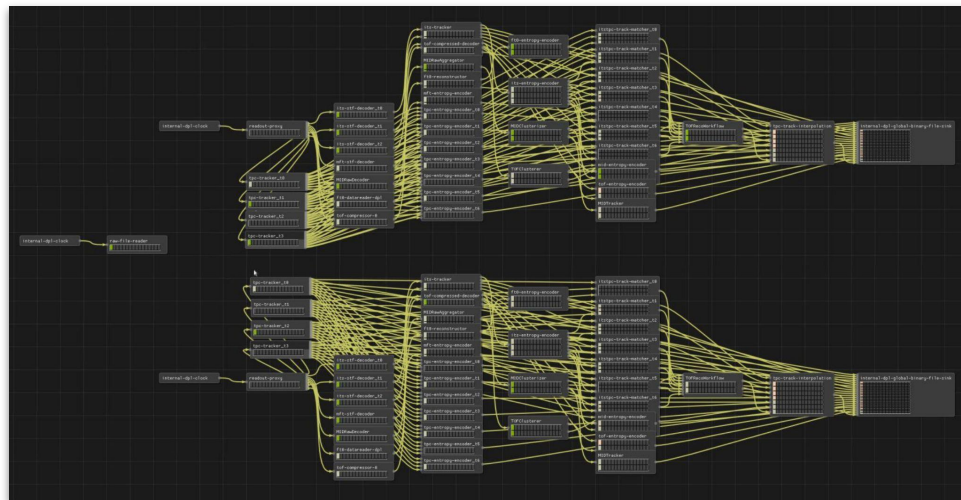
deploy

**DDS** (Dynamic Deployment System) is a toolset that automates and significantly simplifies deployment of user-defined topology of processes and their dependencies on any resource management system.

configure & control

**ODC** (Online Device Control) is a tool to execute topologies of FairMQ devices via DDS, allowing control through ECS commands. It presents a unified topology state to the ECS, while at the same time allowing access to device details.

Usable in an **online setup**, and also in an **"offline" setup** (scheduled batch execution)

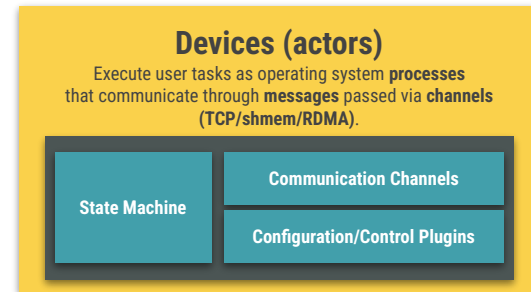


Example of a FairMQ topology on a single node, when used by DPL of the Alice O<sup>2</sup> framework. (reconstruction only, without QA, calibration).

# FairMQ

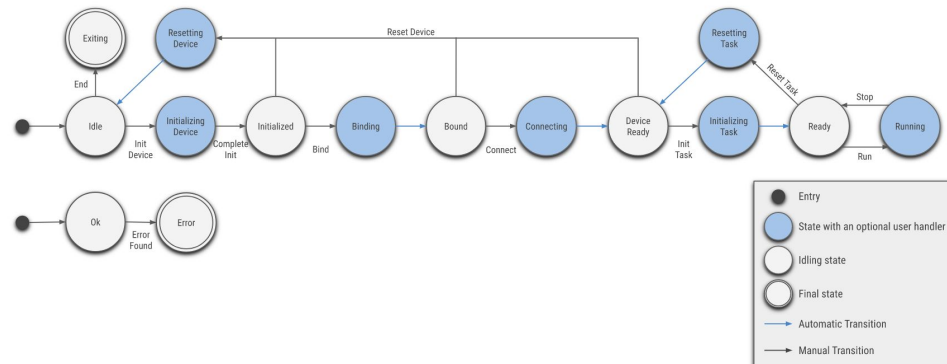
## Library components:

- **Message** and **Channel** provide a message queuing APIs for inter/intra node/process communication.
- Unified API to different data transports (tcp, ipc, shared memory).
- Allow to transparently combine different transport in one device.
- Transport switch via configuration, without modifying the code.



## Framework components:

- Thin and opinionated compositional layer.
- Main framework concept is a **Device** which composes **State Machine**, **Configuration Manager**, and **Command/Configuration Plugins**.
- Devices are composed into **Topologies** that belong to a Session which isolates them.



# DDS

## Dynamic Deployment System

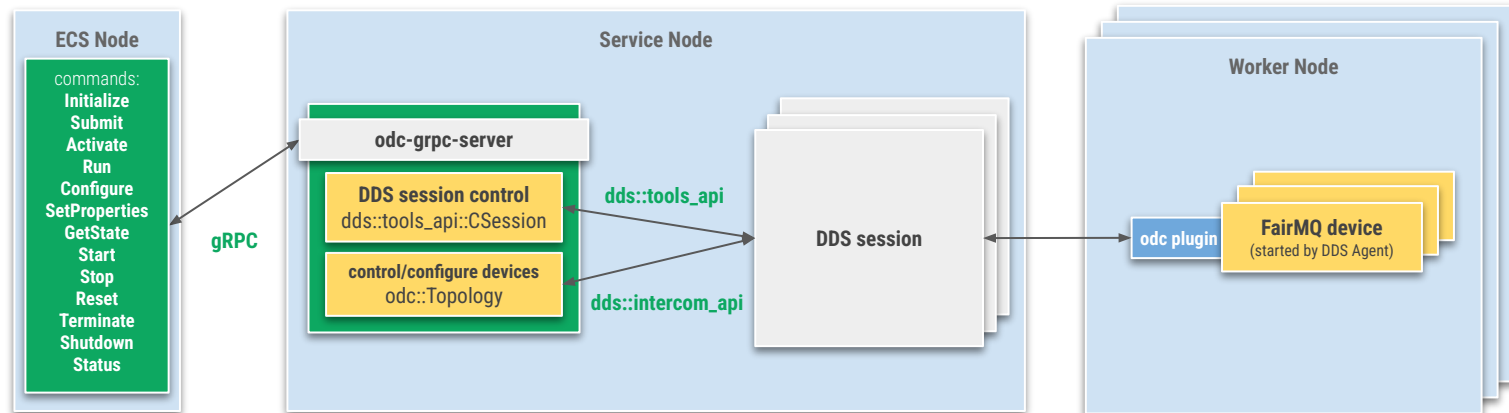
- **Topology-based execution of tasks.**
- **Treats user tasks as black boxes – they can be executables or scripts.**
- **Watchdogging.**
- **Key-value propagation.**
- **Custom commands.**
- **Plugins to abstract from RMS including SSH, localhost, Slurm plugins.**
- **No pre-installation and pre-configuration on the worker node required.**
- **Isolated (and parallel) user sessions.**
- **Lightweight agents on the worker nodes to manage user tasks (single agent can handle hundreds of tasks).**
- **Single responsibility principle command line tool-set and API.**



deploy

# ODC

## Online Device Control



### ODC features:

- Topology state (aggregated device states).
- Fixed set of commands for session control, topology deployment, device state control and configuration.
- Parallel isolated sessions with a single ODC controller process.
- Resource management delegated to DDS which supports **Slurm**, custom SSH and localhost managers (and more).
- **Resilience**: reconnect to running sessions on ODC server restart.
- gRPC server and an example gRPC client.
- Local command-line server for testing.



# ODC

New features **since start of ALICE Run 3**


- **Fault tolerance:**
  - Marking **individual task types as expendable**, so that their failure can be ignored.
  - **Setting a minimum number of error-free task collections**, where certain number of collections can be allowed to fail, without stopping the entire session.
- Support for **multiple heterogeneous Slurm partitions with different configurations**.
- **Dynamic calculation of required resources** from the provided topology file: In the addition to the existing manual resource request, user can now request that the resources are to be calculated based on the tasks in the topology definition file.
- Support for **core-based scheduling** to allow for a more fine-grained task distribution on the available resources. The core requirements are passed to Slurm through DDS.
- **Asynchronous processing of State & Status commands** for better feedback during ongoing commands.

configure  
& control

# LHC Run 3 - ALICE Experiment at CERN

## Deployment in an ongoing experiment

- Deployment & control of large FairMQ Topologies with ODC/DDS on a dedicated online cluster.
- **~130,000** total FairMQ devices (processes) in the online cluster per workflow.
- **~325** processing nodes per workflow. Heterogeneous sets of nodes with different CPU/GPU capabilities.
- **~400** processes per node per workflow.
- Several dozens of workflows executed daily with different setups (different detectors and calibration/QC components).
- intra-node communication with shared memory for processing devices, plus network communication to QC/calibration.



**ALICE**

Run 3 Pb-Pb

$\sqrt{s_{NN}} = 5.36$  TeV

27 September 2023, 04:50

# Additional Resources

FairMQ

<https://github.com/FairRootGroup/FairMQ>  
<https://doi.org/10.5281/zenodo.1689985>  
<https://helmholtz.software/software/fairmq>

DDS

<https://github.com/FairRootGroup/DDS>  
<https://doi.org/10.5281/zenodo.8332289>

ODC

<https://github.com/FairRootGroup/ODC>  
<https://doi.org/10.5281/zenodo.7707968>

FairMQ  
and DPL

<https://aliceo2group.github.io/quickstart/fair-dpl.html>