ALICE Day
UPB, Oct 2024

Alice Suiu - asuiu@cern.ch

About me…

- ACS PhD student (3rd year)
- Bachelor's degree and master's degree at ACS
- Started the CERN collaboration in 2020 during my master's degree
- Real-time data management of a high-throughput data processing system - PhD Thesis
- EPN2EOS - transfers data from ALICE experiment to a remote persistent storage

EPN2EOS in the data transfer path

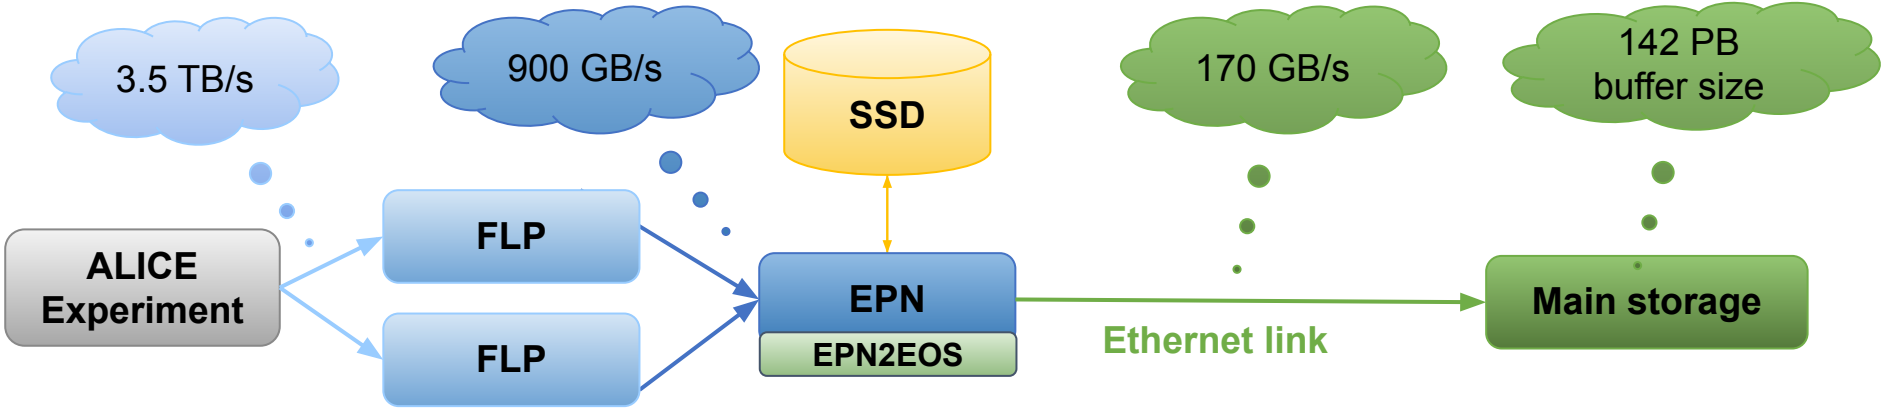

FLP - First Level Processor

- 200 computers
- 8000 *optical links*

EPN - Event Processing Node

- 350 computers, each equipped with one 4TB SSD
- ***Infiniband*** transporting time frame
- **SATA link** between EPN and SSD

Cumulative data transfer

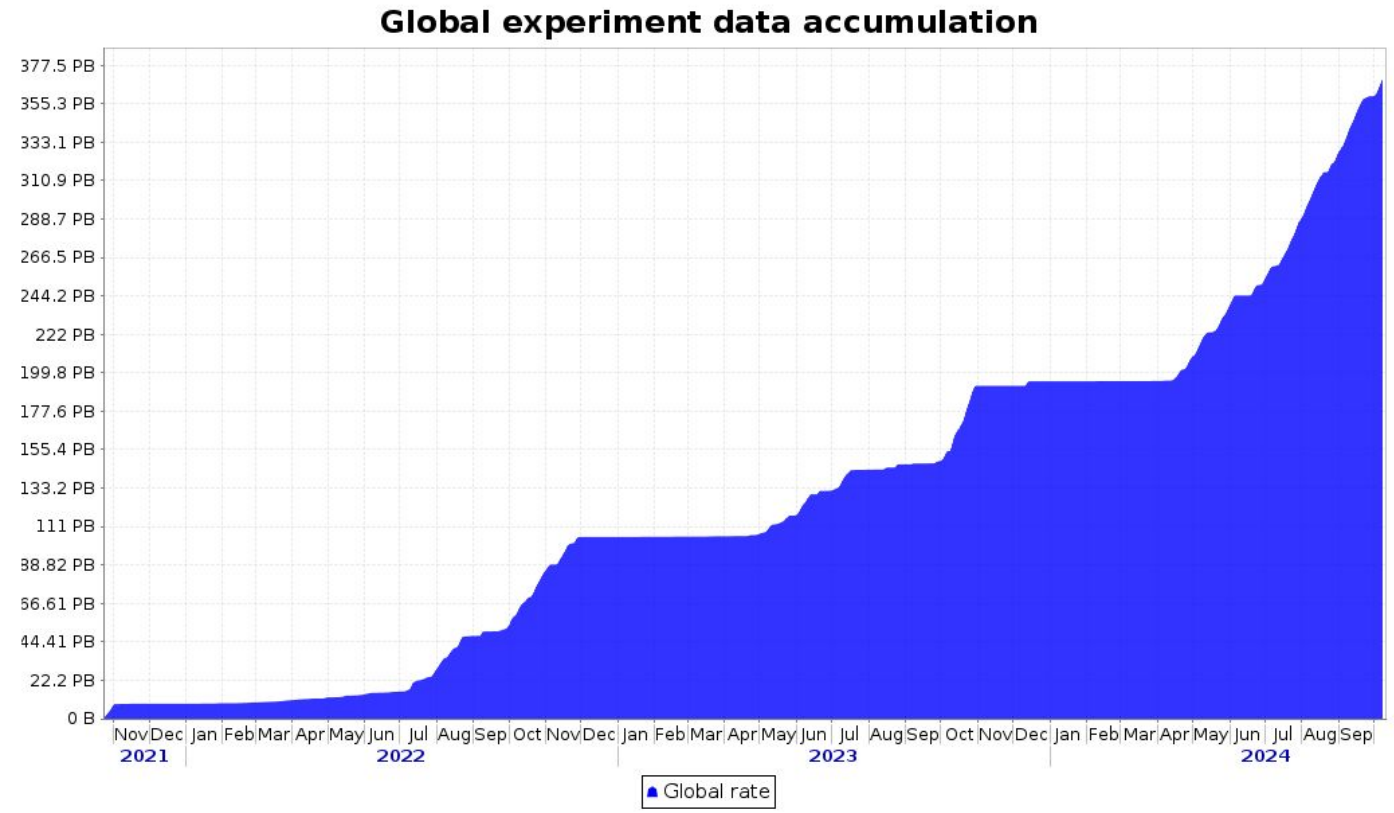

written in Java

running as a daemon service on EPNs from November 2021

transferred **380 PB** of data

total number of transferred files: **100 M**

average file size: **3 GB**

Error-handling system, monitoring and messaging

Monitoring System

- Log messages and monitor the system
  - Number of active transfers/registrations (*Ongoing*)
  - Maximum number of transfers/registrations that could run in parallel (*Slots*)
  - Number of files in the transfer/registration queue (*Queued*)
  - Data transmission rate (*Copy rate*)
  - Success rate, in files/second
  - Error rate, in files/second
- Send alerts to list of recipients with details about the error condition

	Services	Data file transfers						Catalogue registration				
Location	reporting	Ongoing	Slots	Queued	Copy rate	Success rate	Failure rate	Ongoing	Slots	Queued	Success rate	Failure rate
P2	354	1715	2601	3117	229.9 GB/s	61.52/s	0.167/s	3	330	0	60.63/s	0

Monitoring Page

- Other metrics:
  - EPN node that reports the metrics (*Machine*)
  - Time that has passed since the last restart of the EPN2EOS (*Uptime*)
  - Version of the EPN2EOS that runs on the EPN (*Version*)
  - Different types of errors (Transfer errors, Invalid metadata files, Registration Errors)
  - Local disk status (*Write status*)

			Data file transfers						
Machine	Uptime	Version	Ongoing	Slots	Queued	Queued size	Copy rate	Success rate	Failure rate
1. epn000	17d 0:56	v.1.29	8	8	40	76.57 GB	525.7 MB/s	0.133/s	0

	Catalogue registration					
Machine	Ongoing	Slots	Queued	Success rate	Failure rate	Rejected
1. epn000	0	1	0	0.15/s	0	0

	Accumulated error files			Target SE
Machine	Transfer	Missing source	Invalid meta	Write status
1. epn000	0	0	0	

Disk Error Alerts

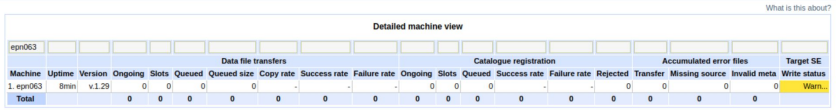

- If the EPN local disk usage reaches **50%** of its capacity, a corresponding message is shown on the monitoring page showing the disk space status.

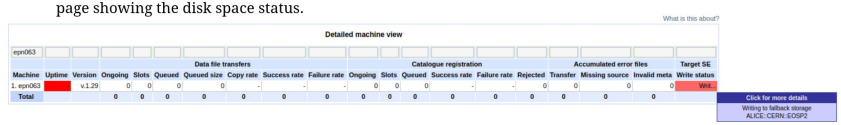

- If the disk usage reaches **90%** of its capacity, the EPN2EOS starts transferring the files to the fallback storage and sends **an email alert** with a corresponding message.
- If the disk usage reaches **95%** of its capacity, the EPN2EOS stops running and sends **an email alert**. Also, the tool will keep restarting as long as the disk space situation is not solved.

Transfer speed

Maximum aggregated transfer speed: **240 GB/s**

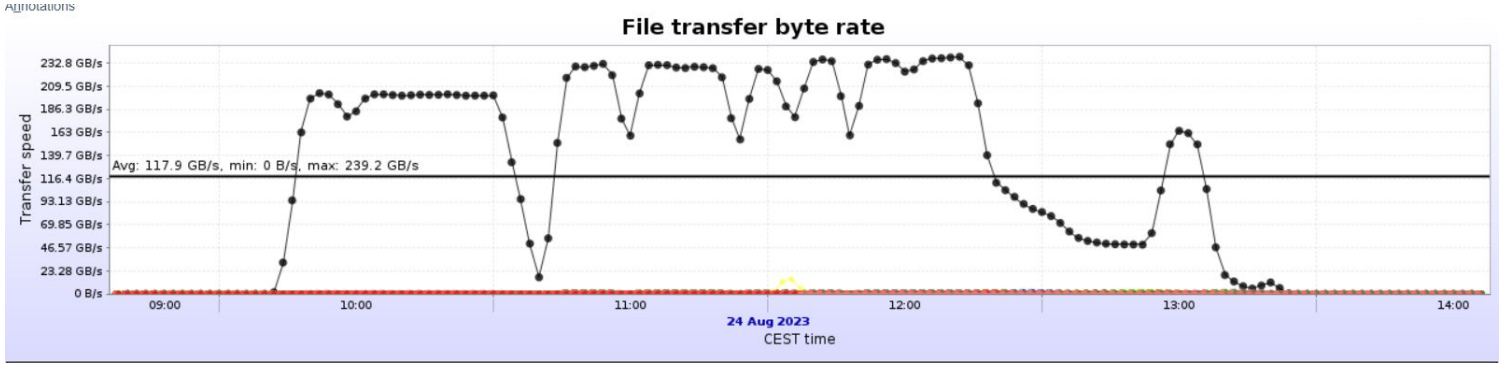

Maximum transfer speed per EPN: **2.5 GB/s**

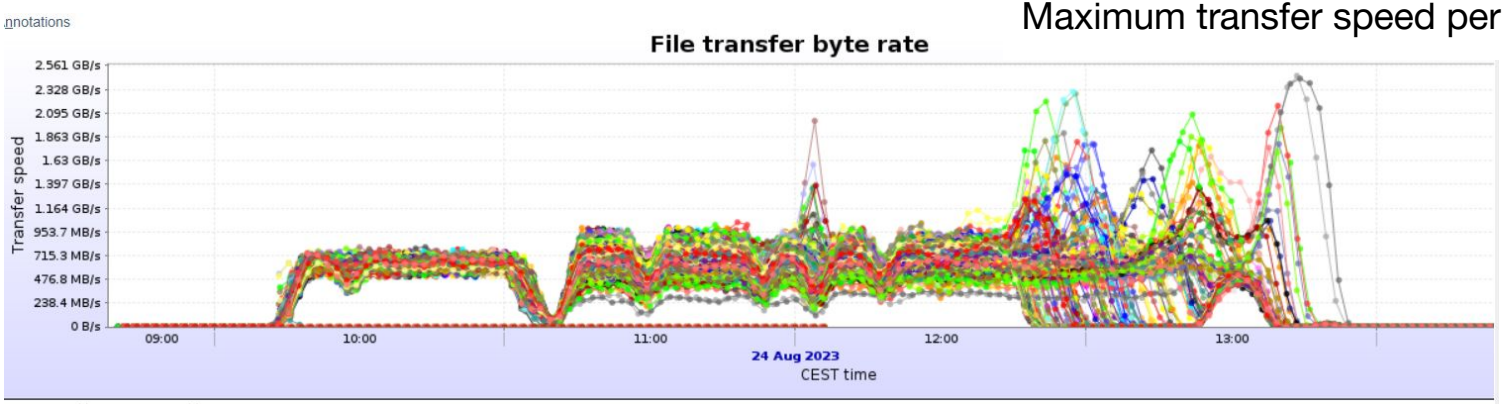

Why CERN?

- open people, eager to help
- project you are working on matters
- work at CERN for several months
- development and accumulation of knowledge from several domains

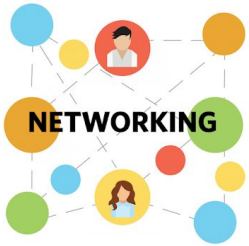

- existence of a connection between physics and computing