


SWAN Users' Workshop



Contribution ID: 23

Type: **not specified**

Using SWAN to analyse tape server logs.

Friday, 11 October 2019 16:35 (20 minutes)

The LHC data are stored mainly on tape. When the experiments want to analyse a dataset, they issue a recall request, to move the data from tape to disk so that they are available for analysis.

The request propagates through a number of systems with a number of distinct events the process: user issues the request, the request is received by the tape system, the tape is mounted, etc.

In order to better understand the performance of the system in its entirety and identify bottlenecks, I'm using SWAN and SPARK to combine, filter, and process the logs of the different systems in a unified "request timeline". That way we can know in what component we have the largest delays and thus focus our effort in optimizing it.

Primary author: Mr KAKLAMANOS, Georgios (CERN)

Presenter: Mr KAKLAMANOS, Georgios (CERN)

Session Classification: User's use cases