


Metrics? Efficiency? Cost?

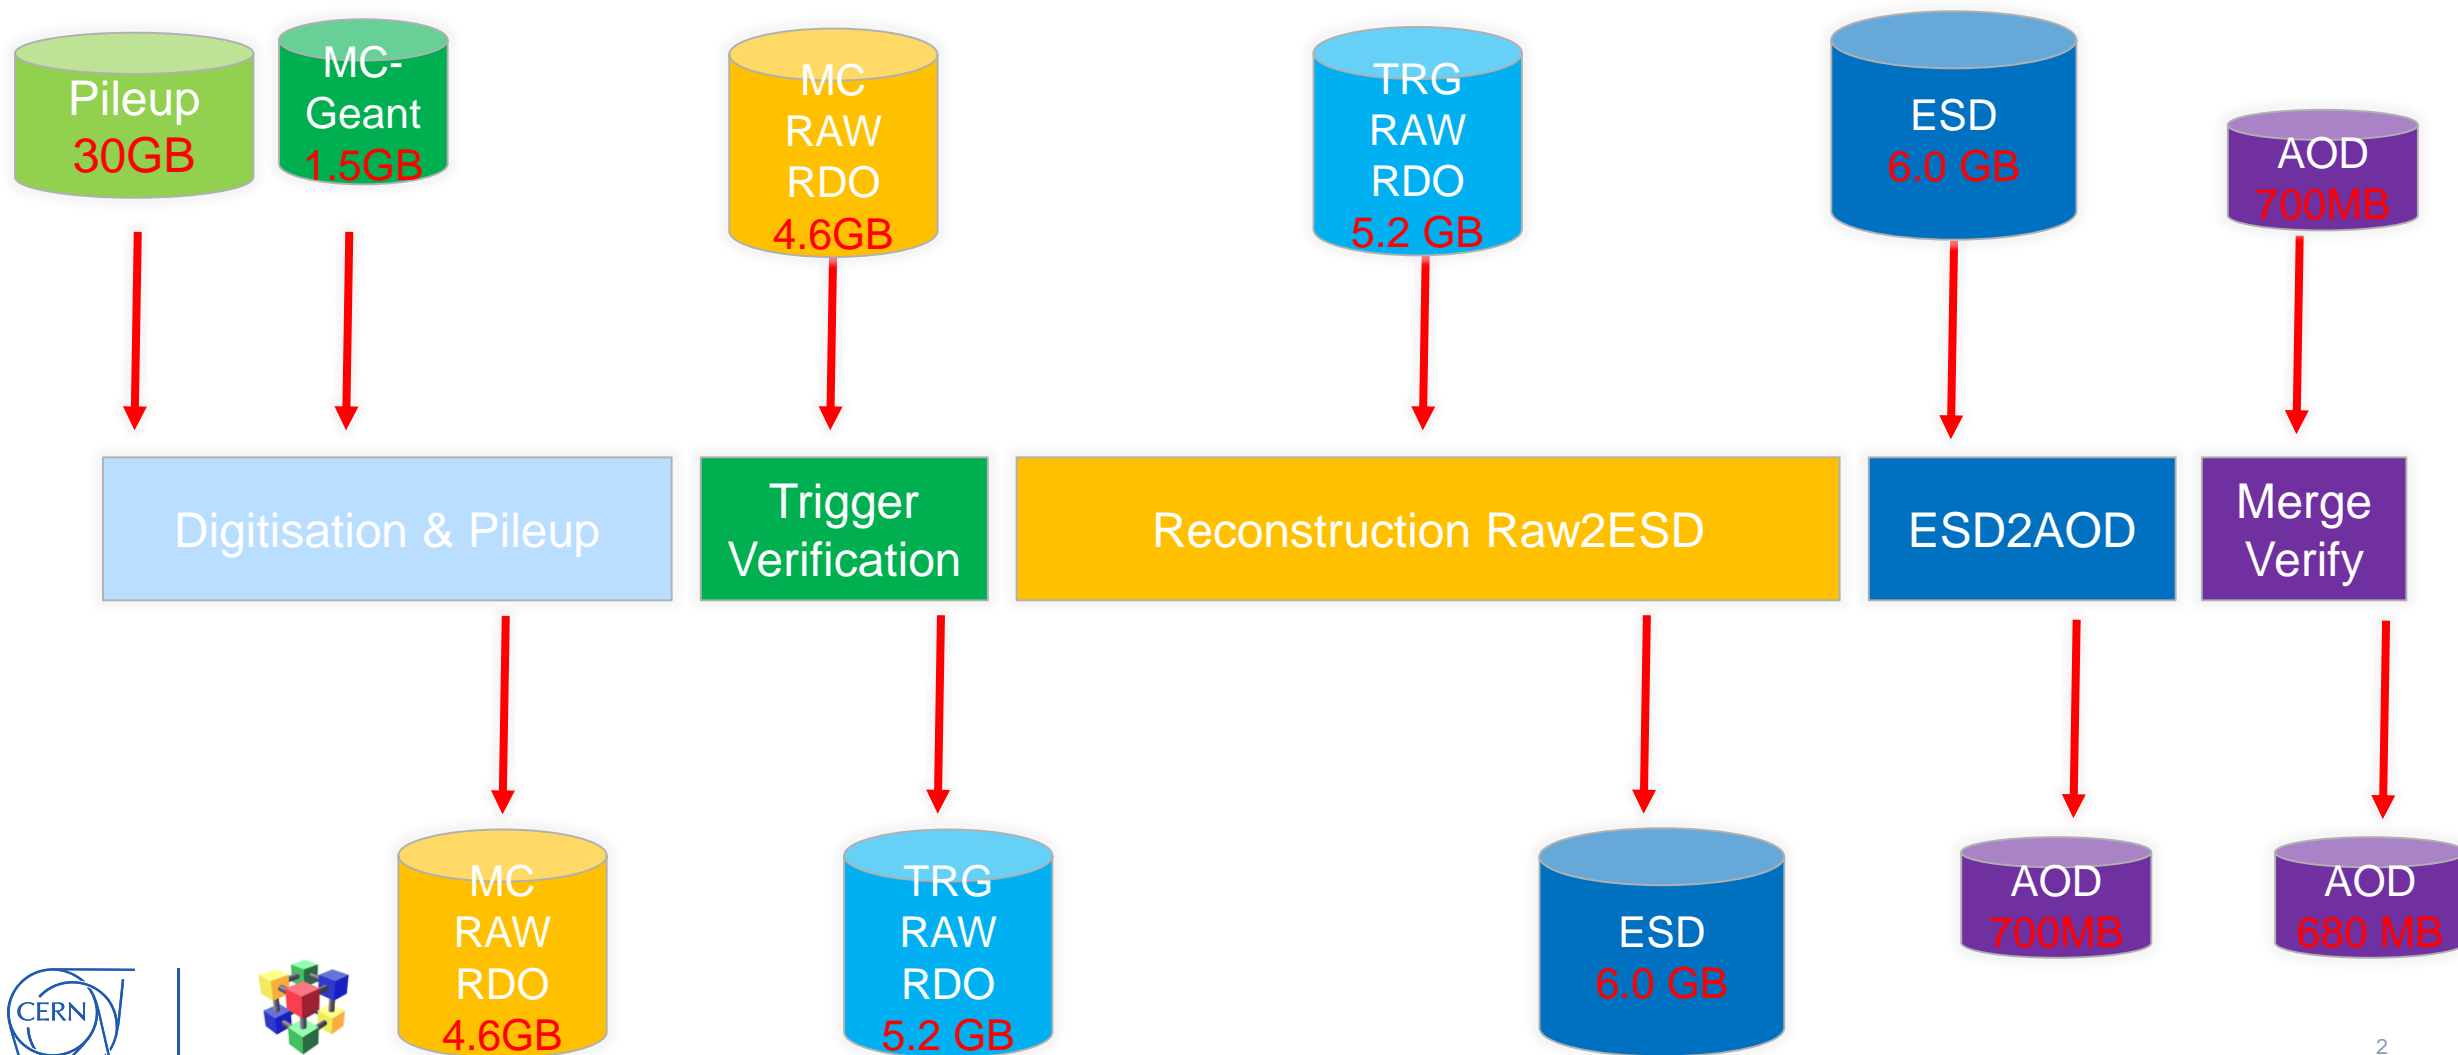
Markus Schulz, Andrea Sciaba



The problem

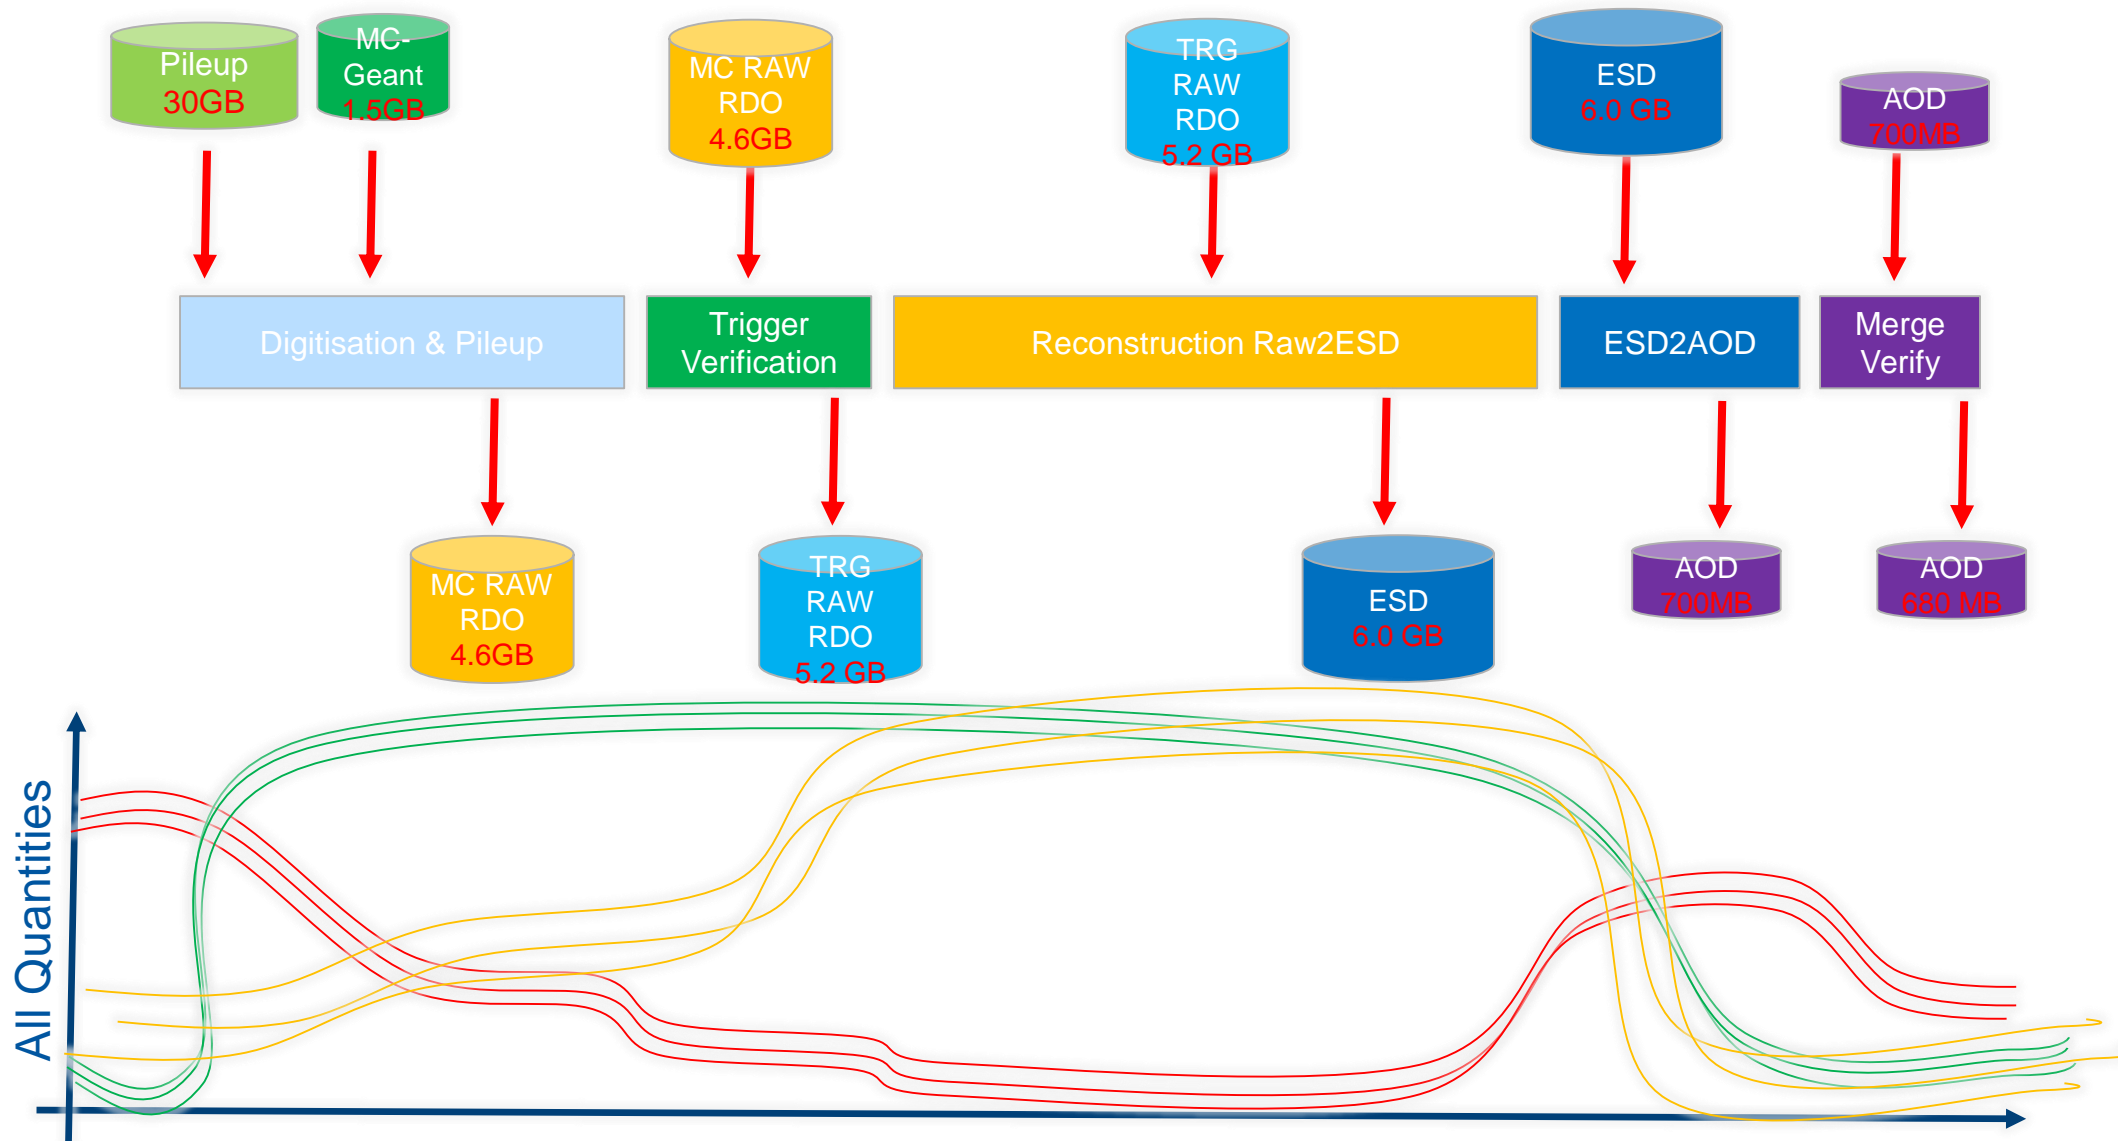
- Our workloads are not simply structured
 - They often are workflows on their own
 - Mix of different activities
 - Different usage at different times
- Example: Atlas pileup + digitization + reco jobs
 - 8 concurrent processes
 - Following pictures inspired by measurements, but not measurements. For illustration only!

A Workload: 2000 MC Events, 8 cores



We have to describe the jobs resource usage in many dimensions

- CPU/Core
 - User, system, idle, Nice, I/Owait.....
 - Instruction level parallelism
 - Cache utilization
 - Active Cores
 - Performance counters.....
- Memory
 - Virtual, physical, swap, throughput
- Storage
 - Volume, access, read/write speeds, physical I/O vs logical I/O
 - ...
- Network
 - Bandwidth, latency



We Monitor most of this, but:

- ❑ What is limiting the jobs throughput?
- ❑ What fraction of the maximum performance do we utilize?
- ❑ What is the throughput/resource unit?
- ❑ What are the dependencies?
- ❑ What phases should be separated?
 - And how?
- ❑

Performance and Efficiency and Cost?

- What does this mean for us?

- Events processed for a workload in a given time?
 - Using how many resources?

- Fraction of the theoretical capability of the system utilized?
 - Doing what?

High level:

- ❑ Defining **P**erformance
- ❑ Defining **E**fficiency
- ❑ Defining **C**ost
- ❑ Classify Job Phases
- ❑ Define Metrics
- ❑ Build Differential equations describing the dependency of **PEC** from the metrics