



PRINCETON
UNIVERSITY

Resource modeling group

David Lange

March 12, 2019


Context

- WLCG/HSF Naples meeting ideas: define a **common FW for modeling the computing requirements** of LHC (and HEP in general) experiments
- We have since understood that experiments are different enough that reaching this goal might mean a tool that is not useful (eg, too low fidelity) to capture requirements well enough to be input to cost modeling.
- However, being more open with code and ideas means we have made good progress this year



Implementation approaches

- Experiments each have active efforts to keep short and long term resource estimates up to date
 - Driven partly by requests from C-RSG / LHCC / national reviews
- There are numerous models in use. These vary in complexity, but have a common underlying approach
 - Start from model current processing activities (data, MC processing), disk replication factors, and tape usage
 - Extrapolate current practice into the future accounting for LHC conditions, trigger rates, expected luminosity, etc
 - Adjust based on R&D expectations (eg, application software improvements, analysis data tier evolution)

The idea behind resource models is quite simple



Example projections



Outputs and capturing R&D

- Model outputs are yearly estimated resource needs
 - We see little need to monitor results at timescales shorter than yearly. Resources are always full. Any peaks/valleys in planned processing requirements get smoothed out or filled in by opportunistic processing.
- Difficult to anticipate impact of R&D or behavioral changes before they happen
 - Models based on current practice
 - Conversely : We can anticipate by setting “goals” for how R&D may change input parameters (examples: future data tier size reductions, application software speedup)

State of affairs

- Collaboration around approaches and implementations
 - No common code base, but increased openness in codes
- Parameters behind estimates (which are numerous!) are generally kept internal
 - Experiments are different enough that the methodology of deriving these parameters is specific to each experiment
 - Can be/have been validated against actual grid job performance attributes
 - Or not?
 - Should Atlas / CMS use common representative benchmarks to derive application parameters (keeping in mind that there are important differences in our approach to generators for Monte Carlo simulation)
 - Infrastructure parameters more in common (but mostly discussed/agreed long ago)

Interfacing to cost group - model outputs

Metrics currently reported

1. CPU (eg, HS06 seconds of compute needed)
 2. Disk (PB of storage)
 3. Tape (PB of storage)
- We presume there is no need to separate by Tier0/Tier-1/Tier-2 in the context of HL-LHC. In any case the boundaries between these are blurring

Where should this work go:

Priorities for New metrics in models

Potential new metrics: Tape I/O

- More interesting (which is to say worrisome)
 - HL-LHC brings large increase in data volume – eg, 50x increase in RAW data volume for CMS
 - I/O capabilities are increasing less quickly than storage capabilities
- Aggregate I/O estimates are straightforward to include
 - Beyond just computing the total PB to tape per year, its important to estimate average need during data taking periods
 - Safety factors needed in order to account for bursts and finite disk buffer sizes are likely best understood by studying monitoring data

Potential new metrics: Networks

- Local:
 - Probably too hard and too variable to include in models
 - Data center setup and configuration
 - Job stage-in / stage-out
- Regional (from dataset transfers and remote file reads)
 - Initially a lower-bound estimate
 - Eventually a way to evaluate DOMA R&D
- Inter-regional (from dataset transfers and remote file reads):
 - Initially a lower-bound estimate
 - Eventually a way to evaluate DOMA R&D

Potential new metrics: Accelerators

- We can hide behind HS06 or its successors so that we can think of accelerators in the same way as CPUs
- Otherwise too early to be able to include in models explicitly

Questions for discussion

- Is this status sufficient for answering questions in this WG? So far it appears to be, and experts from each experiment are engaged to answer specific studies when they arise.
- Any particularly important model parameters, methods or assumptions that should be revisited?
- Additional metrics that are useful for us to extract from models (for this WG or others)?