



Couple of random thoughts as a way of introduction

Paolo Calafiura (LBNL)

Projected computing needs vs resources



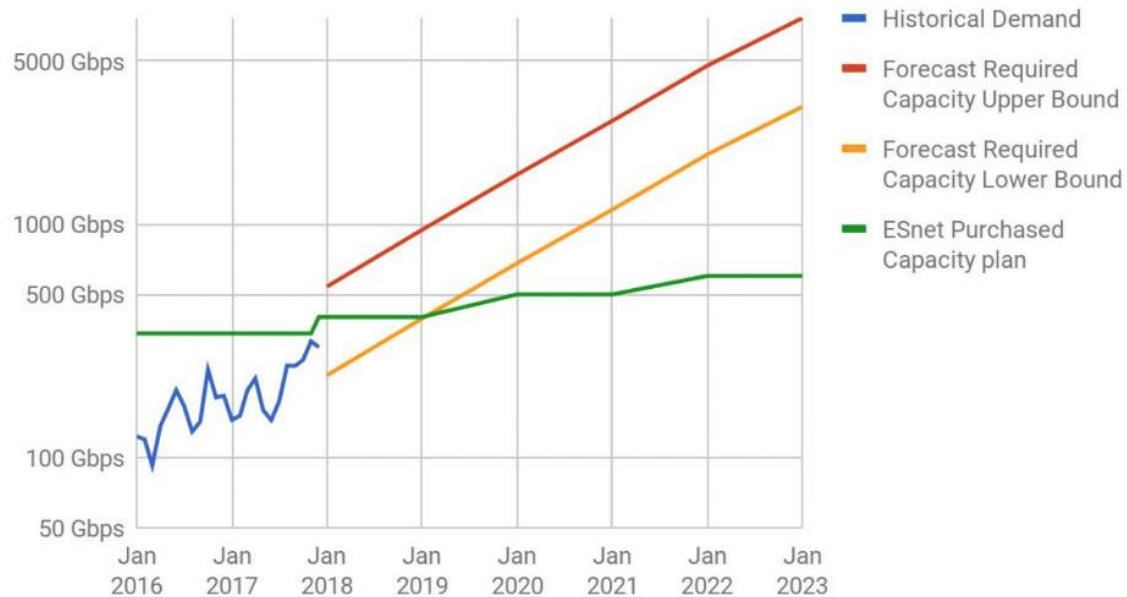
(from the HSF Community White Paper)

What about networking?


ESnet Forecast

(aka we are not the only one good at making scary projections)


European Demand and Capacity Forecasts



HOME · LHCONE
Fermi National Accelerator Laboratory




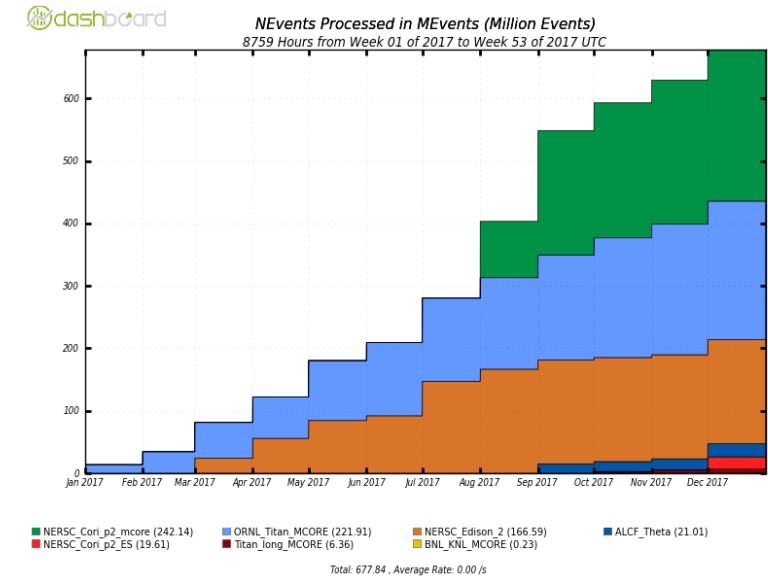
HOME · LHCONE
Brookhaven National Laboratory



- To first order this is coming from CMS, but it will be very instructive to understand the sources as we consider introducing pileup pre-merged overlay, WAN data access

HPC contributions

- Last year DOE HPC produced 0.7B G4 events, 6% of total (11.7B)
 - NERSC was equivalent to ~50kHS06 site
 - We used ~2% of cori, ~20% NERSC HEP allocation
- 2018 HPC request: CPU for 1.3B G4 event
 - Scaling up in out year
- 2019 goals: port full MC chain, derivations



US Facility and the HPC

- Additional HPC CPU cycles free resources for storage and networking
- HPCs riskier than grid sites:
 - Long (by our standards) downtimes (scheduled or not)
 - New, harder to use architectures
 - Changing policies (allocations!)
- US Facility job is getting harder
 - QOS for high-priority workflows (T0 spillover and similar)
 - Remote I/O for HPCs (and perhaps one day clouds?)
 - Data-intensive, fast turnaround analysis workflows, instead of steady-state G4 production