



Post LS3 infrastructures for DAQ

What will we need in 11 years... ?
Some projections...or...



CMS TIG March 17, 2016
attila.racz@cern.ch for CMS DAQ group



Different extrapolations...

- Today, we process the events with ~ 0.5 MHS06
 - Our latest compute node (Q1'16) gives 660 HS06 and burns 350W
- By looking at Pileup (140/200), trigger rate (500/750kHz) after LS3
 - Guestimation of post LS3 processing power is 5/11 MHS06
 - Targeting infrastructures for 11 MHS06 even if not needed from the beginning of run4 (Q1'27), ~ 11 years from now (Q1'16)
- Based on our purchases since Q4'2007
 - If exponential, CPU $\rightarrow 31\%$ per year, TPD $\rightarrow 5\%$ per year
 - Compute node of Q1'27 is 12.86 kHS06@600W
 - 855 nodes, 513kW, 14 racks@38kW (64 nodes per rack)
 - If linear, CPU $\rightarrow 73$ HS06 per year, TPD $\rightarrow 5$ W per year
 - Compute node of Q1'27 is 1.4 kHS06@405W
 - 7860 nodes and 3.2 MW and 123 racks@26kW
- IT considers **20-25%** compute power increase per year at constant price (it is not our case but we have a limit...) and let's add a 2% per year for TPD
 - Compute node of Q1'27 is 4.9 kHS06@435W
 - 2244 nodes, 876kW, 36 racks@28kW
- All above is only for HLT. We need to add the event builder, storage and services
 - 800 links at 100Gb/s (~ 500 nodes, ~ 500 kW, 25+5 racks@20kW)
 - Maybe keep a "Cloud" for CMS : 200-300 kW more and 10-15 racks



Predictions...

- Nobody can foresee what will happen in computing in 10 years...☺
 - Usage of GPUs... ARM architecture... Requirements on HLT... Cloud
- Different extrapolations give a power consumption of post LS3 data center ranging from ~1MW to 2-3MW, from 40 to 150 racks and power per rack from 20 to 40kW
 - We don't believe to the 1MW picture...and no element to choose between 2/3 MW...
 - Even 20-25 % per year on CPUs seems very optimistic to us
- Let's assume 2-3MW for the post LS3 CMS Data center
 - In line with the guestimates of the other LHC experiments (ouf!)
 - ~80 racks with 20 to 40 kW per rack
 - **Good news ! We have the real estate and the building can support the load**
- Is there an additional 1-2MW available on-site ?
 - If yes, fine, we take it !
 - If no, build another tower ? Add chillers on the parking lot ?
 - Find synergies with HL-LHC infrastructures ?



Pipes and racks

- With a ΔT of 6 deg, water flow into SCX5 must be:
 - 285 m³/h (2MW) or
 - 430 m³/h (3MW)
 - Today we have a max of 150 m³/h (limited by pipes diameter I think)
- Racks of 40kW requires a flow of ~5-6m³/h per rack
 - Considering racks of 120x80cm²+door, instead of 90/100x60 today
 - Today, too much of rack space is used for powering and networking
 - Deeper racks will allow us to use side walls for power/networks
- Piping in SCX5 must be modified or changed
 - Either partition existing circuit into 2-3 sub-circuits of 1MW each and modify last stage of distribution
 - Or redo everything with adequate dimensioning
 - Water flow into racks must be tunable on an individual basis (20 -> 40 kW)
 - **Studies to be asked to EN-CV**



Electrical distribution

- 5x 250A canalis feeding 8 racks gen2 (16kW of heat dissipation)
 - Max power is ~190 kW to be shared : ~22-24 kW per racks
- 14x 160A canalis feeding 8 racks gen1 (10 kW of heat dissipation)
 - Max power is ~123kW to be shared : ~15 kW per rack
- Projections are giving 20 to 40 kW per rack
 - Next tick after 250A is 400A... Is our actual cabinet compatible ? Is the transformer big enough ?
 - With new rack width, 6 racks instead of 8, 30kW per racks on 250A canalis (could be OK for non HLT racks)
 - For HLT racks, go to 400A Canalis (~50 kW per HLT rack)
 - **Again studies to be asked to EN-EL**
- DAQ project to be under full UPS coverage... in LS2, LS3 ?
 - Last year studies ended up in a cost around ~800kCHF... and 2 big containers (9x3m²) on the parking lot...



Support framing

- Frame supporting racks have been designed for 90x60cm racks
 - Locally modified in 2008 to welcome 100x60cm racks
- Need to see how we can fit ~70-80 racks of 145x80 with minimal work
 - 120 cm internal depth + 25 cm cooling door
 - Options could depend on what is done for piping
 - **Studies to be asked to ?**



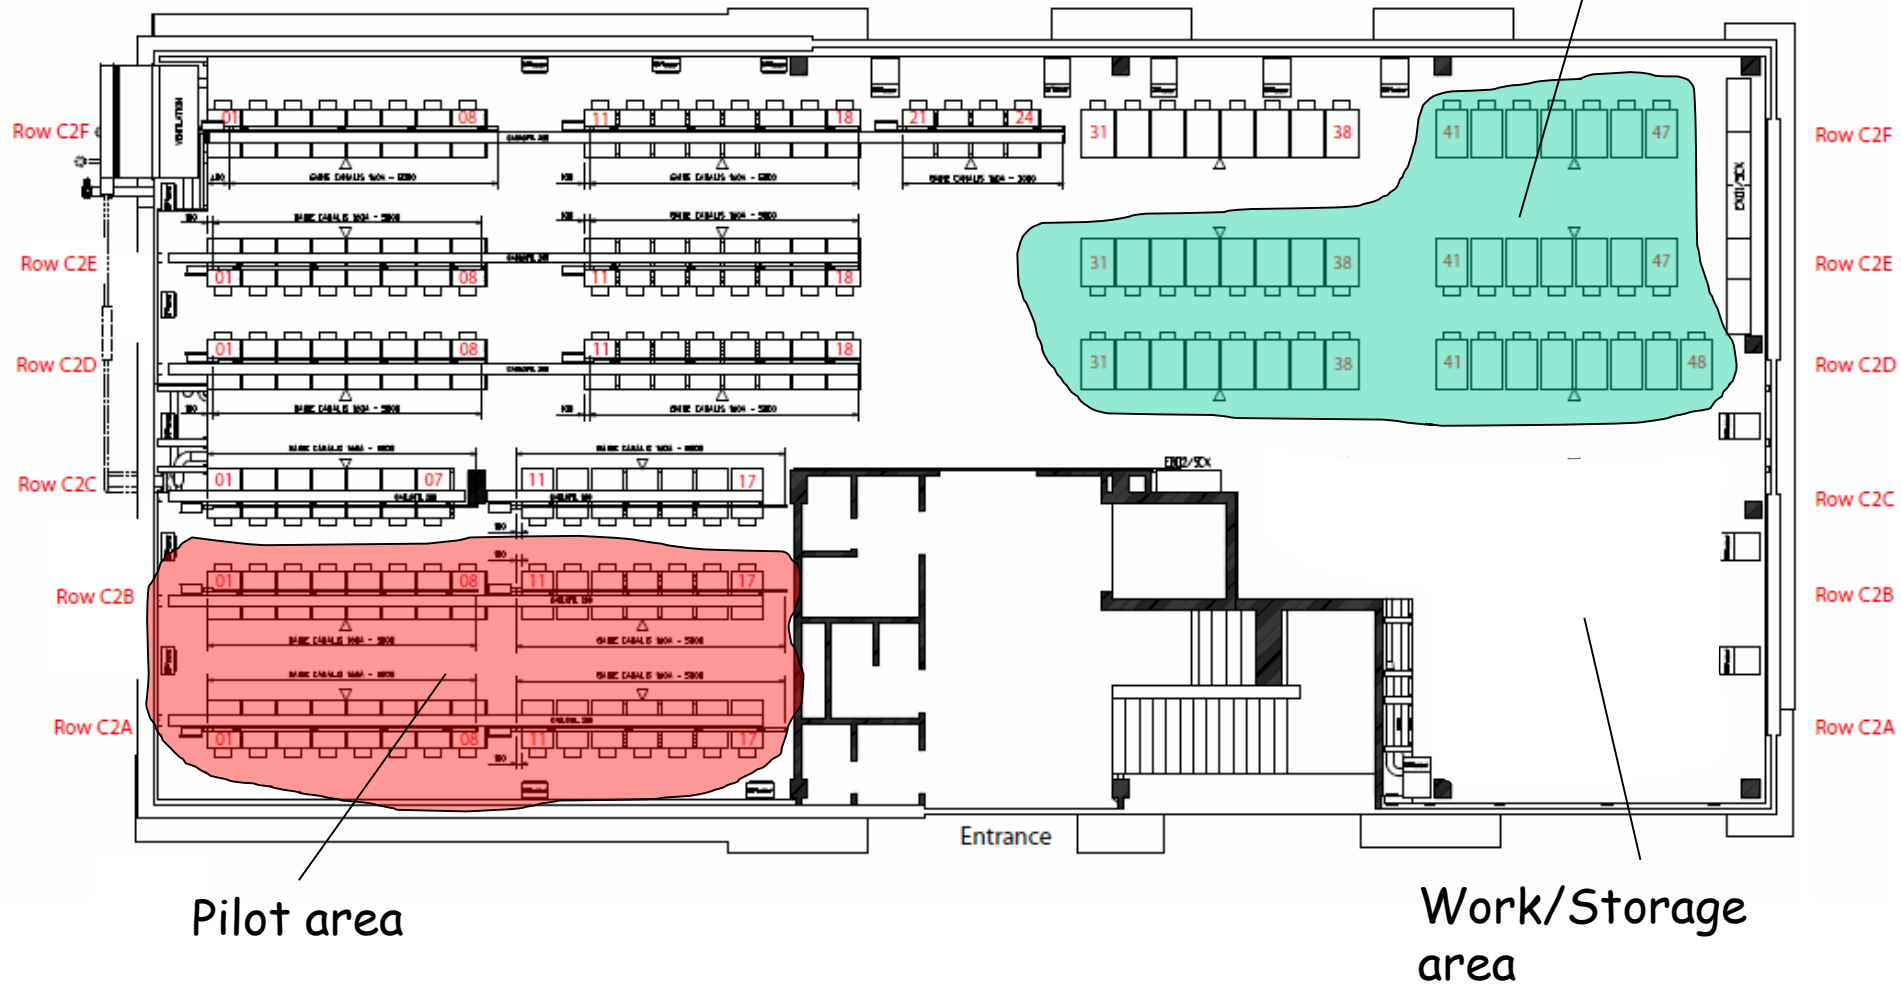
Schedule

- LS3 is ~2.5 years long (Jan'24 ->Jun'26)
 - 1.5 years to dismantle and refurbish the room to be ready to receive servers
 - 1 year to load racks with servers, cable up every thing, commissioning...
 - Super tight...
 - Need to prepare well ahead everything
 - CERN industrial resources may be hard to get for LS3
- DAQ group plan to make a little pilot to check feasibility
 - New racks, plumbing, electrical power, framing...
 - Area [A..B][1..17] of our room is now empty and unused, containing gen1 racks (10kW, 90 cm deep) not suited for todays and future servers
 - Target is mid-end of 2016, as soon as we are in routine operation



Data Center layout today

Gen2 racks 16kW, 100cm





Typical rack (10-70kW...)



March 17, 2016

CMS DAQ Group



Conclusion

- For DAQ project, LS3 is a major upgrade, not only for computers etc but also for infrastructures
 - To summarize, everything is to be redone in SCX5 except the floor !
 - We need to get studies/quotes from EN-EL/CV, mechanics for costs and duration
 - Also, a lot of work in USC area : cable removal campaign for all sub-systems
 - Pilot project to be launched asap after YETS
- Enormous uncertainties for 10 years ahead projections
 - A few% difference in CPU power increase gives MW of difference
 - Not possible to choose between 2 or 3 MW
- DAQ project to be counted in central CMS UPS
 - 700kW max up to LS3
 - 2-3 MW post LS3
- Even if LS3 is far away, what is the next step ?



Extras



Our procurements

