



“Top concerns”: introduction to day-2 discussions

CMS

WLCG ws - Feb 2nd, 2016

Caveats

- ◆ not at all an exhaustive report on CMS preparations for HL-LHC
- ◆ focus on stimulating WLCG-related discussions
- ◆ on all areas of discussions in the next 5 sessions in agenda today, a concise and specific CMS input is attached

Phase-II

Phase-II will require software/computing systems capable to collect, process, archive and serve for analysis 5 to 7.5 kHz of complex events, and to create a corresponding simulation sample



Processing

Effective growth in CPU power at fixed cost is 25%/yr

- ◆ Not enough to meet the HL-LHC scale for CMS (trigger output rate x pileup)
- ◆ A new approach (and R&D) is needed

What will WLCG facilities look like in 10 years?

- ◆ CMS will need to evolve following industry trends to be cost efficient -> Expect more heterogeneous resources
- ◆ Infrastructure for non-x86 architectures, specialised compute centres
- ◆ Compute centres for large scale tests when experiments have the software? (building on prototype work)
- ◆ Mechanism for optimising resources for use cases

Storage

Data volume increase in Phase-II

- ◆ ~ by a factor 5 to 7.5 in # evts, ~ by a factor 4 to 6 in event size

Effective growth in disk space at fixed cost is 20%/yr

- ◆ basically we need factor of 2-3 beyond growth

CMS is moving towards a much smaller event data tier for analysis

- ◆ we need to follow its impact on replications, versions needed, etc

To make further progress, CMS needs new techniques to need even fewer replicas than today (already really few..), or just novel approaches

- ◆ so far, strong investment with xrootd-based data fed on optimising ability of processing and analysis application to read data over high-latency links while maintaining reasonable CPU efficiency
- ◆ Next-generation data federation complementary to local access + dynamic replication + clean-up on-top + more predictive placement tactics + data locality less relevant also in a cloud env + ..
- ◆ Can we build a predictive model including disk/tape/network access?

Networking

Networking capacity is evolving (not thanks to us)

- ◆ global traffic growing at a rate of 30%/yr
- ◆ Not expecting disruptive technology change - as optical networks - on the timescale of Phase-II, though

CDN systems offer a huge potential for flexibility in data access, but do we have sufficient networking for any model of evolved data federations?

- ◆ a ~10k processing farm (average of a "T2" center in 10 yrs?) would be served data with a O(10) Gbps link; for analysis, with primary data source remote, a 10k jobs task may need O(100) Gbps
- ◆ such distribution systems appropriate for some of the largest NRENs users?

Network somehow less in the LHC/WLCG radar?

- ◆ more focussed efforts as a community could be a value



R+D on application sw and computing techniques

Difficult to predict precise direction of technologies on the Phase-II timescale

- ◆ focus on adapting to generic trends, rather than exploiting specific products available on the market today
- ◆ CMS does a lot of specific product testing - but targeted as prototyping for future architectures

CMS has a multi-threaded capable event-processing framework

- ◆ reduction of burden on the WM system, dramatic improvement in memory use
- ◆ crucial for most future processors technologies with more cores per CPU
- ◆ a concern is teaching physicists enough programming to get to truly many-core applications for analysis by HL-LHC

Analysis software

- ◆ Will analysis workflows in ~2025 look like those today? Not likely
- ◆ How can we take advantage of other data analysis and data reduction techniques and approaches
 - e.g. reduction of the actively accessed and processed sample would significantly decrease the scale of computing resource needs