CMS Scientific Computing

Robert Harris
Fermilab
Scientific Computing Division

Fermilab User's Meeting
June 11, 2015

Outline

- CMS ready for Run 2
  - Data and Computing Model
  - Multi-Threading
  - Simulations and Software for HEP
- CMS preparing for Higher Luminosity at the LHC
  - Simulation and reconstruction for phase II upgrade
  - Computing on Demand from Amazon Web Services

I will introduce a few of the areas where my colleagues in the Fermilab Scientific Computing Division are actively involved. Many thanks for their work and their help preparing this talk.

Data Model

- CMS now has a global data pool
  - One central condor pool for all types of resources and applications
- Users have access to any data, anytime, anywhere (AAA)
  - **Data Federation** will make CMS datasets available transparently across the Grid
  - In run 1 CMS moved away from MONARC hierarchical data model

Old Data Model

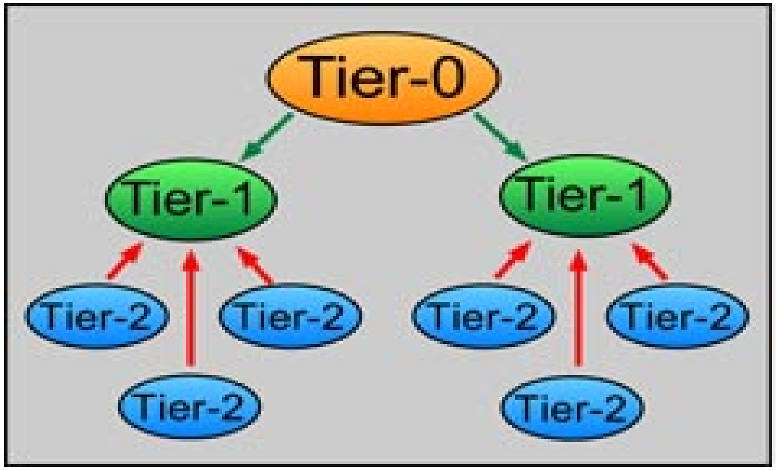

New Data Model

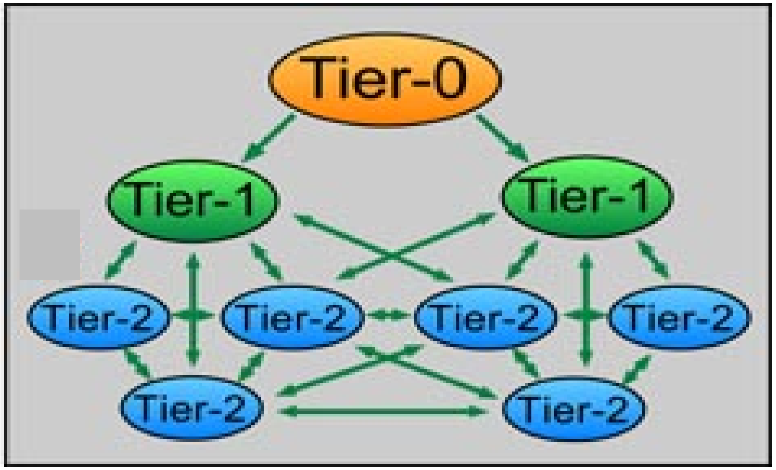

Run 2 Computing Challenge

- Run 2 was projected to require an order of magnitude more computing than Run 1 at similar budget levels
  - HL-LHC will require even more due to pile-up and new detectors

- Strategy to meet the challenge
  - Increase computing efficiency
    - See next slide
  - Improve software algorithms
    - A factor of 3 speedup over run 1 !
  - Exploit parallelism in the software to utilize newly available computer architectures
    - Multithreading to be discussed

CMS Reconstruction Time vs. Release

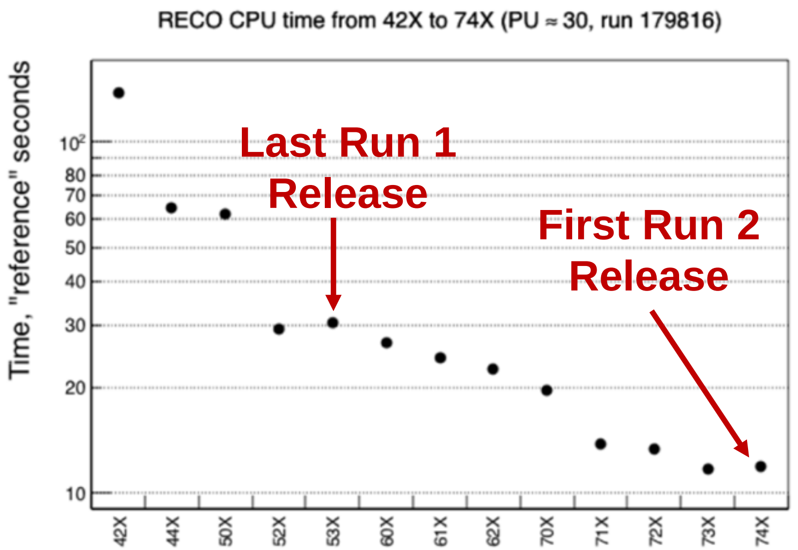

Computing Model

- Run 2 CMS computing will be an even more coherent & efficient system
  - One **central queue** for all resources and all workflows will facilitate prioritization between analysis and production
  - T2's can now do reco and T1's can do analysis.
  - The **HLT farm** is now an integrated resource, outside the LHC running periods
    - Increases T1's by ~50%

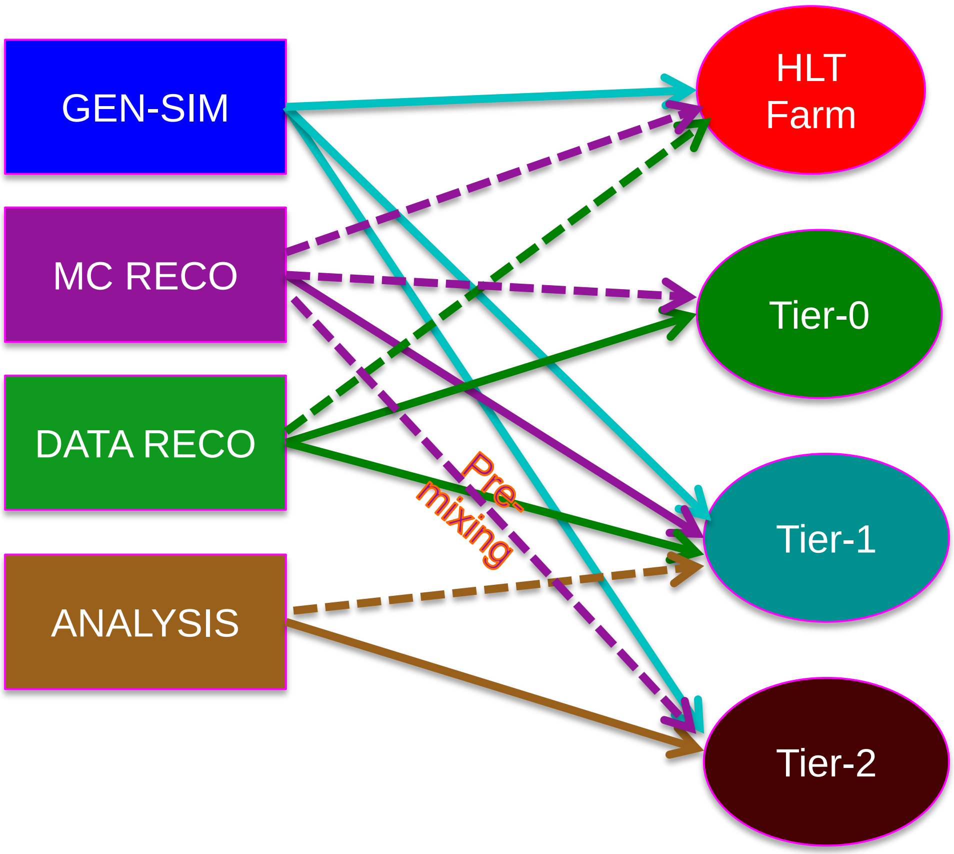

Multi-Threading Software

- Multi-threading allows many calculations to run simultaneously
  - ➔ Essential to take advantage of modern multi-core CPUs and GPUs

- CMS framework now supports multi-threading !
  - ➔ A major technical achievement essential for run 2 requirements

- Existing prompt reco now runs in multithreaded mode
  - ➔ Already a small speedup due to reduction in memory usage without optimizing code to take advantage of multi-threading

- Substantial speedup now possible in future from multi-threaded code

Moore's Law and Multi-Cores

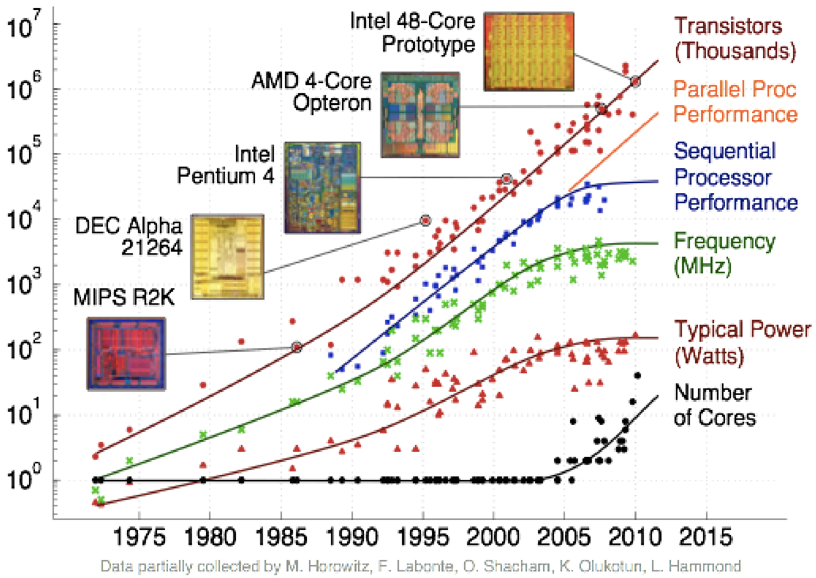

Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

Image from : http://www.csl.cornell.edu/courses/ece5950/handouts/ece5950-overview.pdf

Generation, Simulation and Tools

- PYTHIA generation
  - PYTHIA 8 commissioned for CMS physics during run 1
  - New physics and computing models incorporated
  - Truly the workhorse for physics simulations at the LHC

- GEANT simulation
  - Fermilab contributes to GEANT4 physics & technical improvements
  - Fermilab performing R&D and prototyping on GEANT V
    - A next generation detector simulation toolkit to run multithreaded applications in modern computing architectures.
- Tools
  - Fermilab contributes to the development of ROOT and the vision of HEP software physics tools.
  - Fermilab participates in HEP Software Foundation (HSF).
- Software for CMS that benefits the entire HEP community.

Preparing for HL-LHC

CMS Simulation Development

- CMS simulation was critical to the outstanding physics performance of CMS in run 1
  - Simulation development will be required to sustain the amazing level of agreement between data and Monte Carlo we have come to expect.
- Upgraded CMS detectors require ongoing development in run 2
  - Geometry, simulated hits, digitization, pileup tools, etc.
- Phase II simulation and software will be a particular challenge
  - Re-engineering to exploit benefits of parallelism on multi-core machines
    - e.g. **Particle level parallelization !**
  - State of the art detectors with unique capabilities

Silicon Sampling Endcap Calorimeter

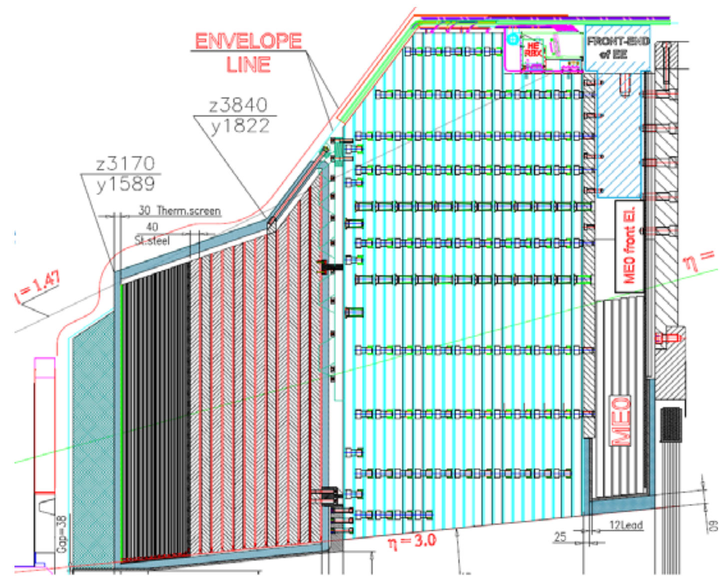

EE HEF
Silicon Sampling

HEB
Scintillator Sampling

High Granularity Endcap Calorimeter

- Silicon pixels allow tracking in this "particle flow" calorimeter

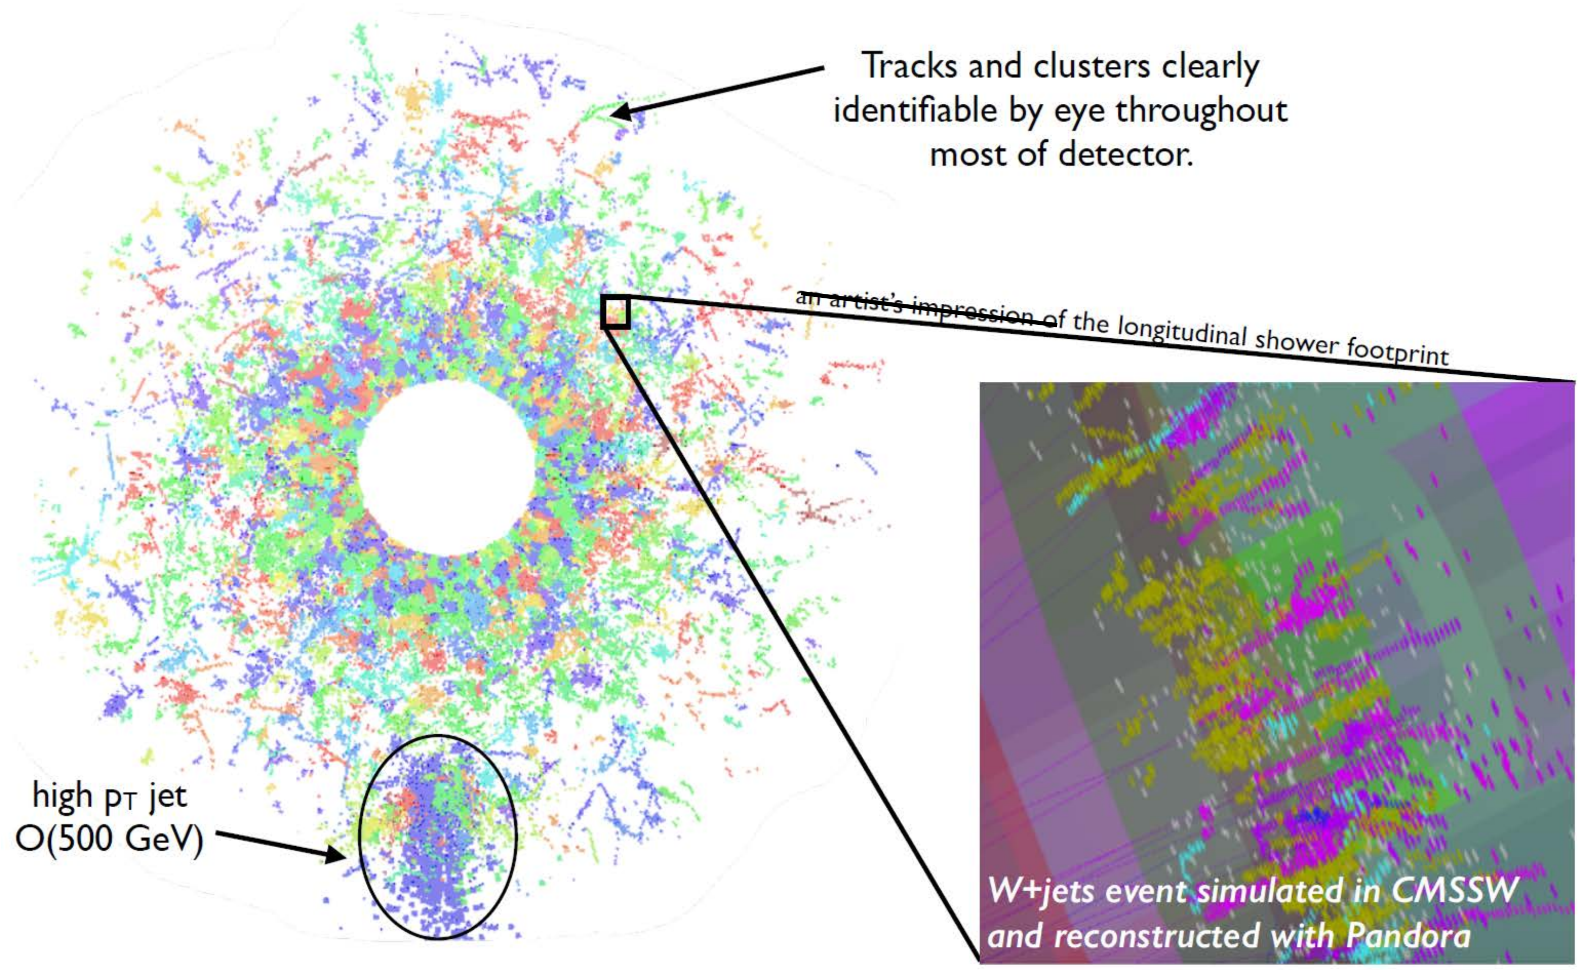

HGCAL Reconstruction

- A significant effort to develop HGCAL reconstruction has begun
  - Take full advantage of unique detector capabilities in particle flow algorithm
  - Reconstruction within a challenging environment of ~140 pileup events
- Solve problem of speed of reconstruction in this pixel calorimeter
  - A challenge with 6.5 million channels and 200 thousand hits per event.
  - Computational geometry and graph theory used so far.
  - Parallelization being explored and has clear benefits.
- Synergies with linear collider and LArTPC communities (PandoraPFA, etc.)

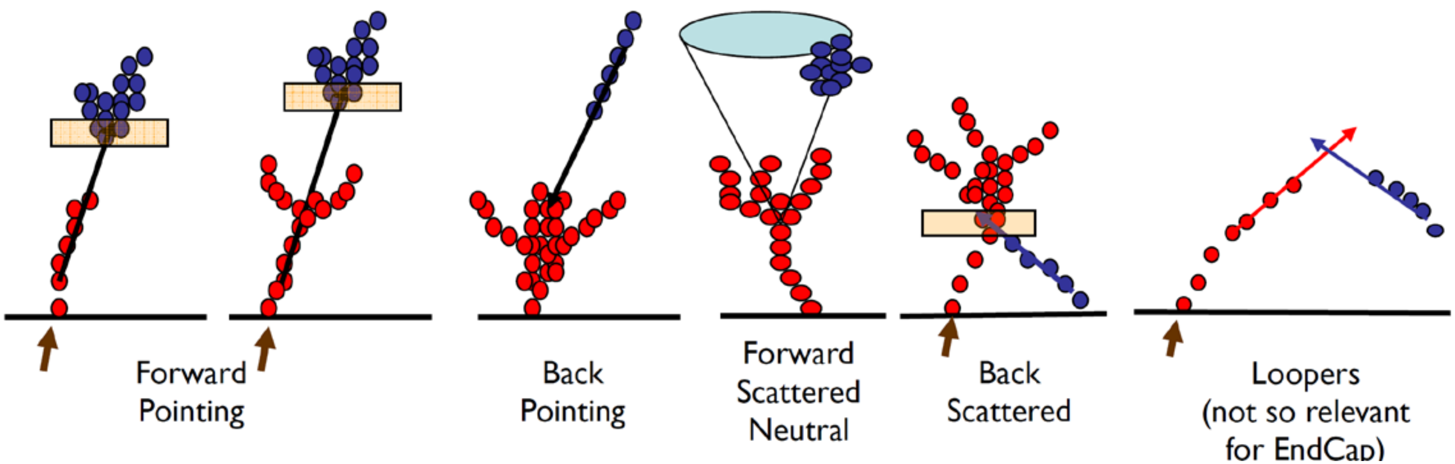

Computing on Demand

- Reprocessing and MC production currently take a long time limiting physics capability
  - ➔ Owned computing resources are limited
- Only need large processing resources during “bursts” of time after code is ready
  - ➔ Remainder of time is a relatively steady usage for prompt reco and analysis
- Computing on demand “rents” resources when needed from a BIG source
  - ➔ Amazon Web Services (AWS) has many millions of processor cores to rent !
  - ➔ Our computing needs are just a small drop in their very big bucket.

Computing with Owned Resources

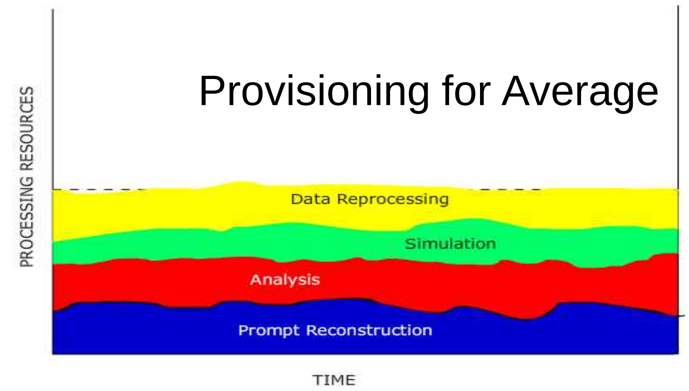

Computing on Demand

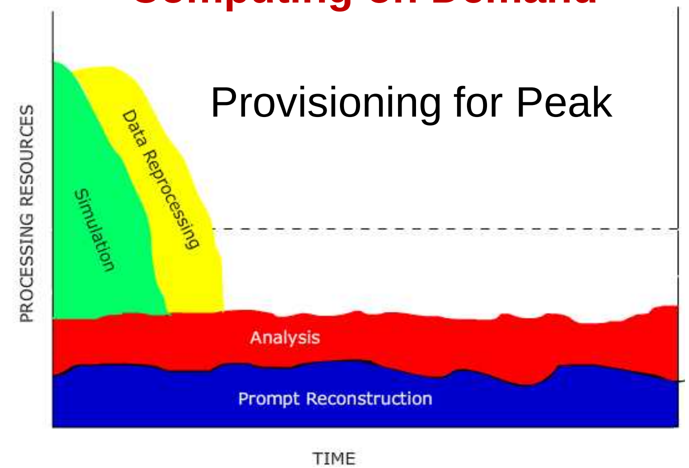

Computing from the HEP Cloud

- Pilot project to integrate computing on demand has begun at Fermilab
  - Has the **highest priority** within the Scientific Computing Division
- CMS Monte Carlo reconstruction using AWS
  - First use case and a critical part of the project
- CMS awarded research grant by Amazon
  - Study feasibility of large scale use of AWS by CMS
  - Use ~50K cores for one month, which is 50% of the current CMS worldwide resources.
  - Simulation and reconstruction of billions of events which will be used by CMS early in run 2.
  - Data will originate and be stored at Fermilab
- Tests new method of provisioning resources
  - Opens up a **new class of HEP computing**

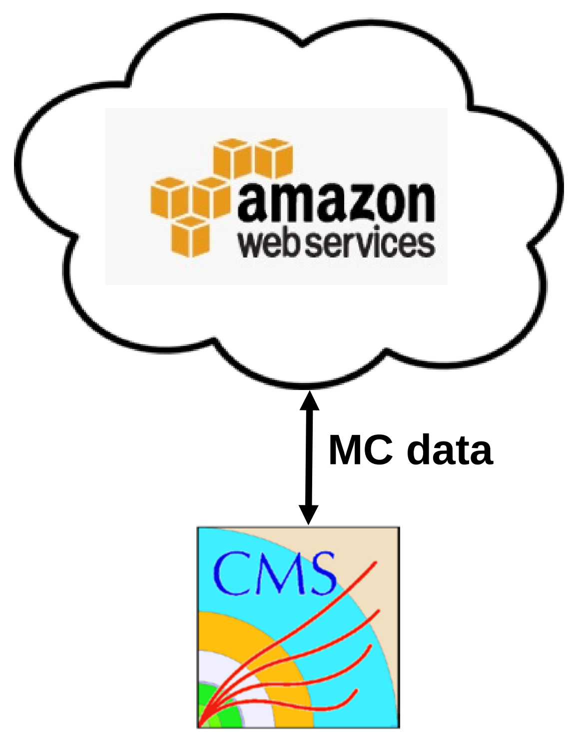

Conclusions

- CMS scientific computing is ready for Run 2
  - Our global data pool provides users with access to any data, anytime, anywhere within a coherent and efficient computing system.
  - Multi-threading software framework in place will allow us to use modern processor architectures efficiently and handle Run 2 data rates
  - Generation, simulation and data analysis tools are solid.
- CMS scientific computing is preparing for HL-LHC
  - Simulations with particle level parallelization utilizing multi-threading
  - Simulation and reconstruction of phase II detector upgrades
  - Computing on demand
- World class scientific computing supporting the energy frontier