Vision and Strategy for Computing at Fermilab

Elizabeth Sexton-Kennedy

Fermilab PAC

16 Jan 2019

Introduction

- Fermilab Computing Vision of the 10 year scale challenge
- Fermilab Computing Strategy
- Addressing the Charge Question

Fermilab Computing Vision

Instrumentation

- Large scientific achievements in the past decades have been enabled by large advances in instrumentation.
- Large silicon detectors and cameras with high granularity are driving us to large computing and data challenges.
- Large costs of these projects require an international scope for them and their computing.

Fermilab

Host Laboratories

- In line with P5 priorities, the experiments we know for sure will be running in 2025 are CMS @ CERN and DUNE @ FNAL.
- Both are the host labs for major HEP investments and must support large user communities to do their science.
- Today computing is a major cost in the operations and analysis of any experiment
  - Last year the LHC used 750Kcores operating 24/7 across the globe, to purchase that amount of computing on the commercial cloud (0.10$/hr/core) would cost 100M$/year

- In this context it makes sense to make common cause with CERN for the support of computing for the HEP community.

A Data Centric Vision for International HEP

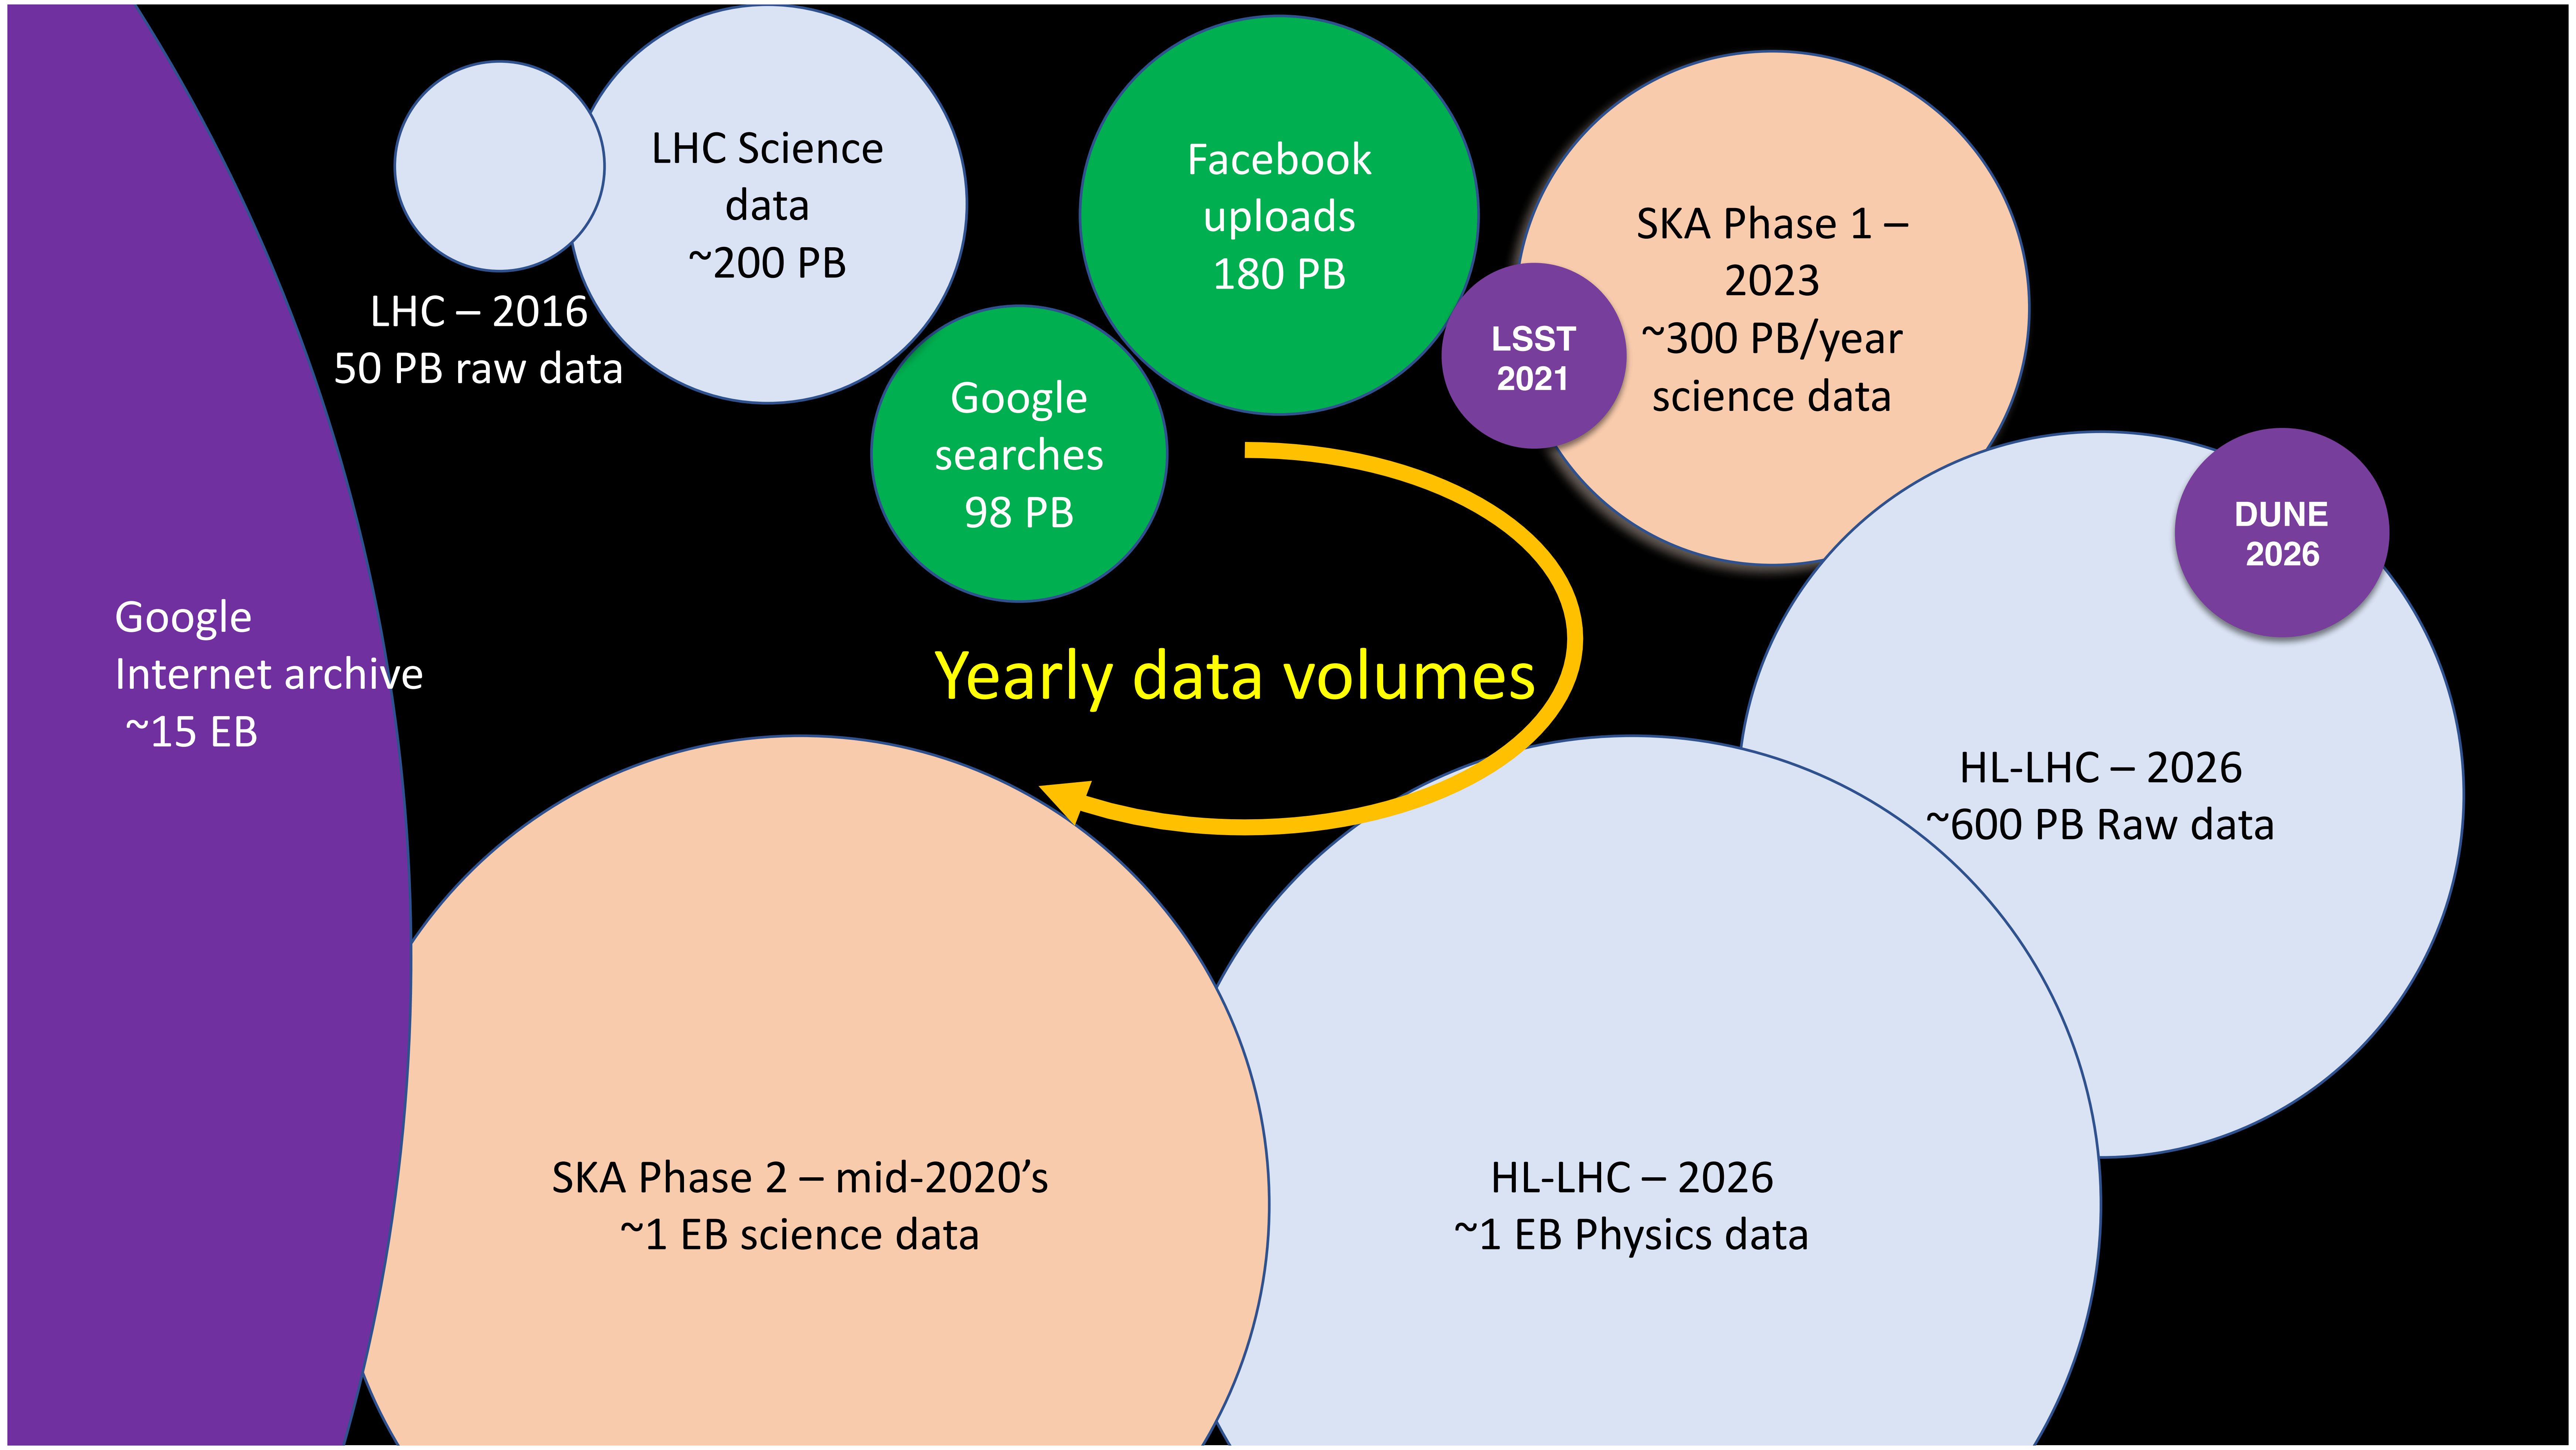

- Yearly raw data collection rates in the next decade

Networking is at the Core

- International science requires international data movement and storage.

- Transatlantic links are expensive so HEP must share them.

Fermilab

Computing Challenges

- The LHC experiments, Belle II and DUNE face the same challenges
  - HEP software must evolve to meet these challenges
  - Need to exploit all the expertise available, inside and outside our community, for parallelization
  - New approaches needed to overcome limitations in today's code
- Cannot afford any more duplicated efforts, it is unsustainable
  - Currently each experiment has its own solution for almost everything (framework, reconstruction algorithms, …)
  - It is not clear how much can be shared however today's situation means there are lots of opportunities.
- Even if we had the money for all of this duplication we do not have the requisite expertise to do all of the work necessary for all experiments!
  - A number of us recognized this a few years ago and formed the HEP Software Foundation

The Increasing Complexity of Computing

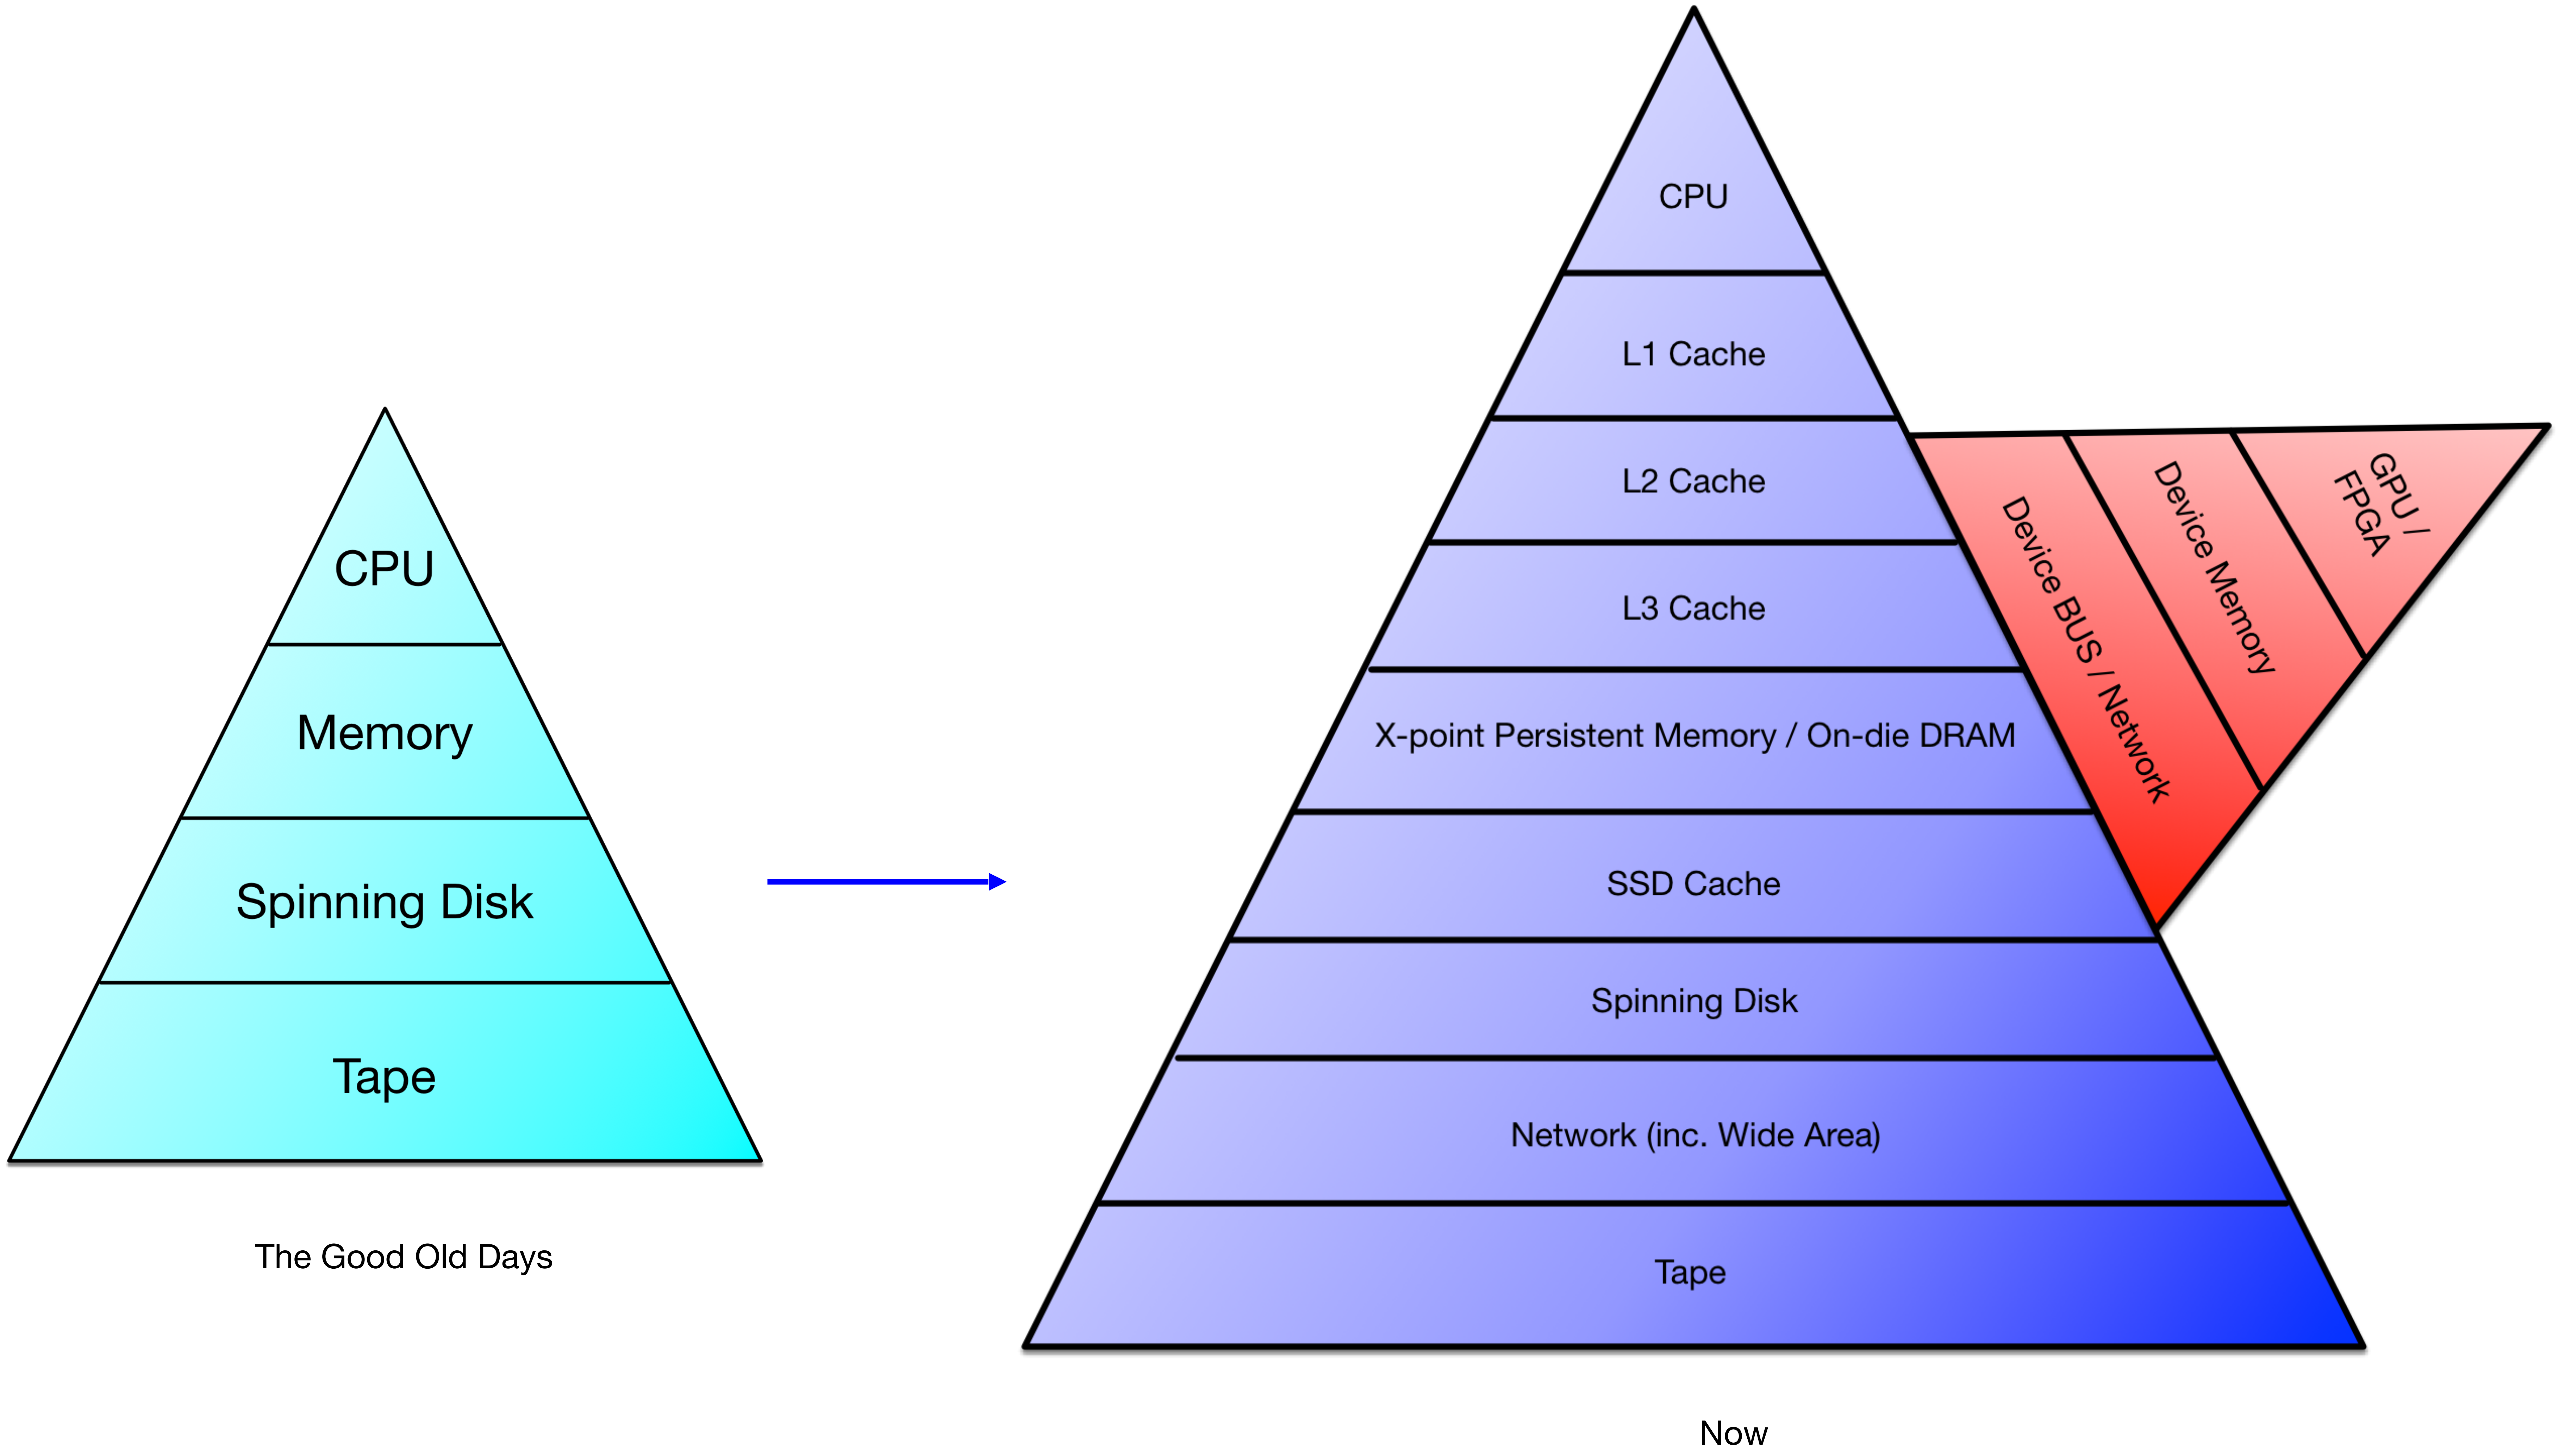

- These changes make the current GRID technical models obsolete

HEP Computing Challenge

- Event or Readout window complexity is accelerating.
- With a flat budget, Moore's law like improvements are the real maximum we can expect on the HW side.
  - HEP software typically executes one instruction at a time (per thread)
  - CPU (core) performance increase are due to more internal parallelism
  - x10 with the same HW only achievable if using the full potential of processors
- Major SW re-engineering required (but rewriting everything is not an option)
  - Co-processors like GPUs, TPUs, FPGAs are difficult to use efficiently and R&D topics
- Increased amount of data requires us to revise/evolve our computing and data management approaches
  - We must be able to feed our applications with data efficiently
- Salvation will come from software improvements, not from hardware
  - We can not buy our way out of this challenge

Fermilab Computing Strategy

WLCG++ -> World Scientific Computing Consortium

- LHC computing is successful both technically and organizationally for over a decade.   CERN proposed an evolution of the model that FNAL supports.

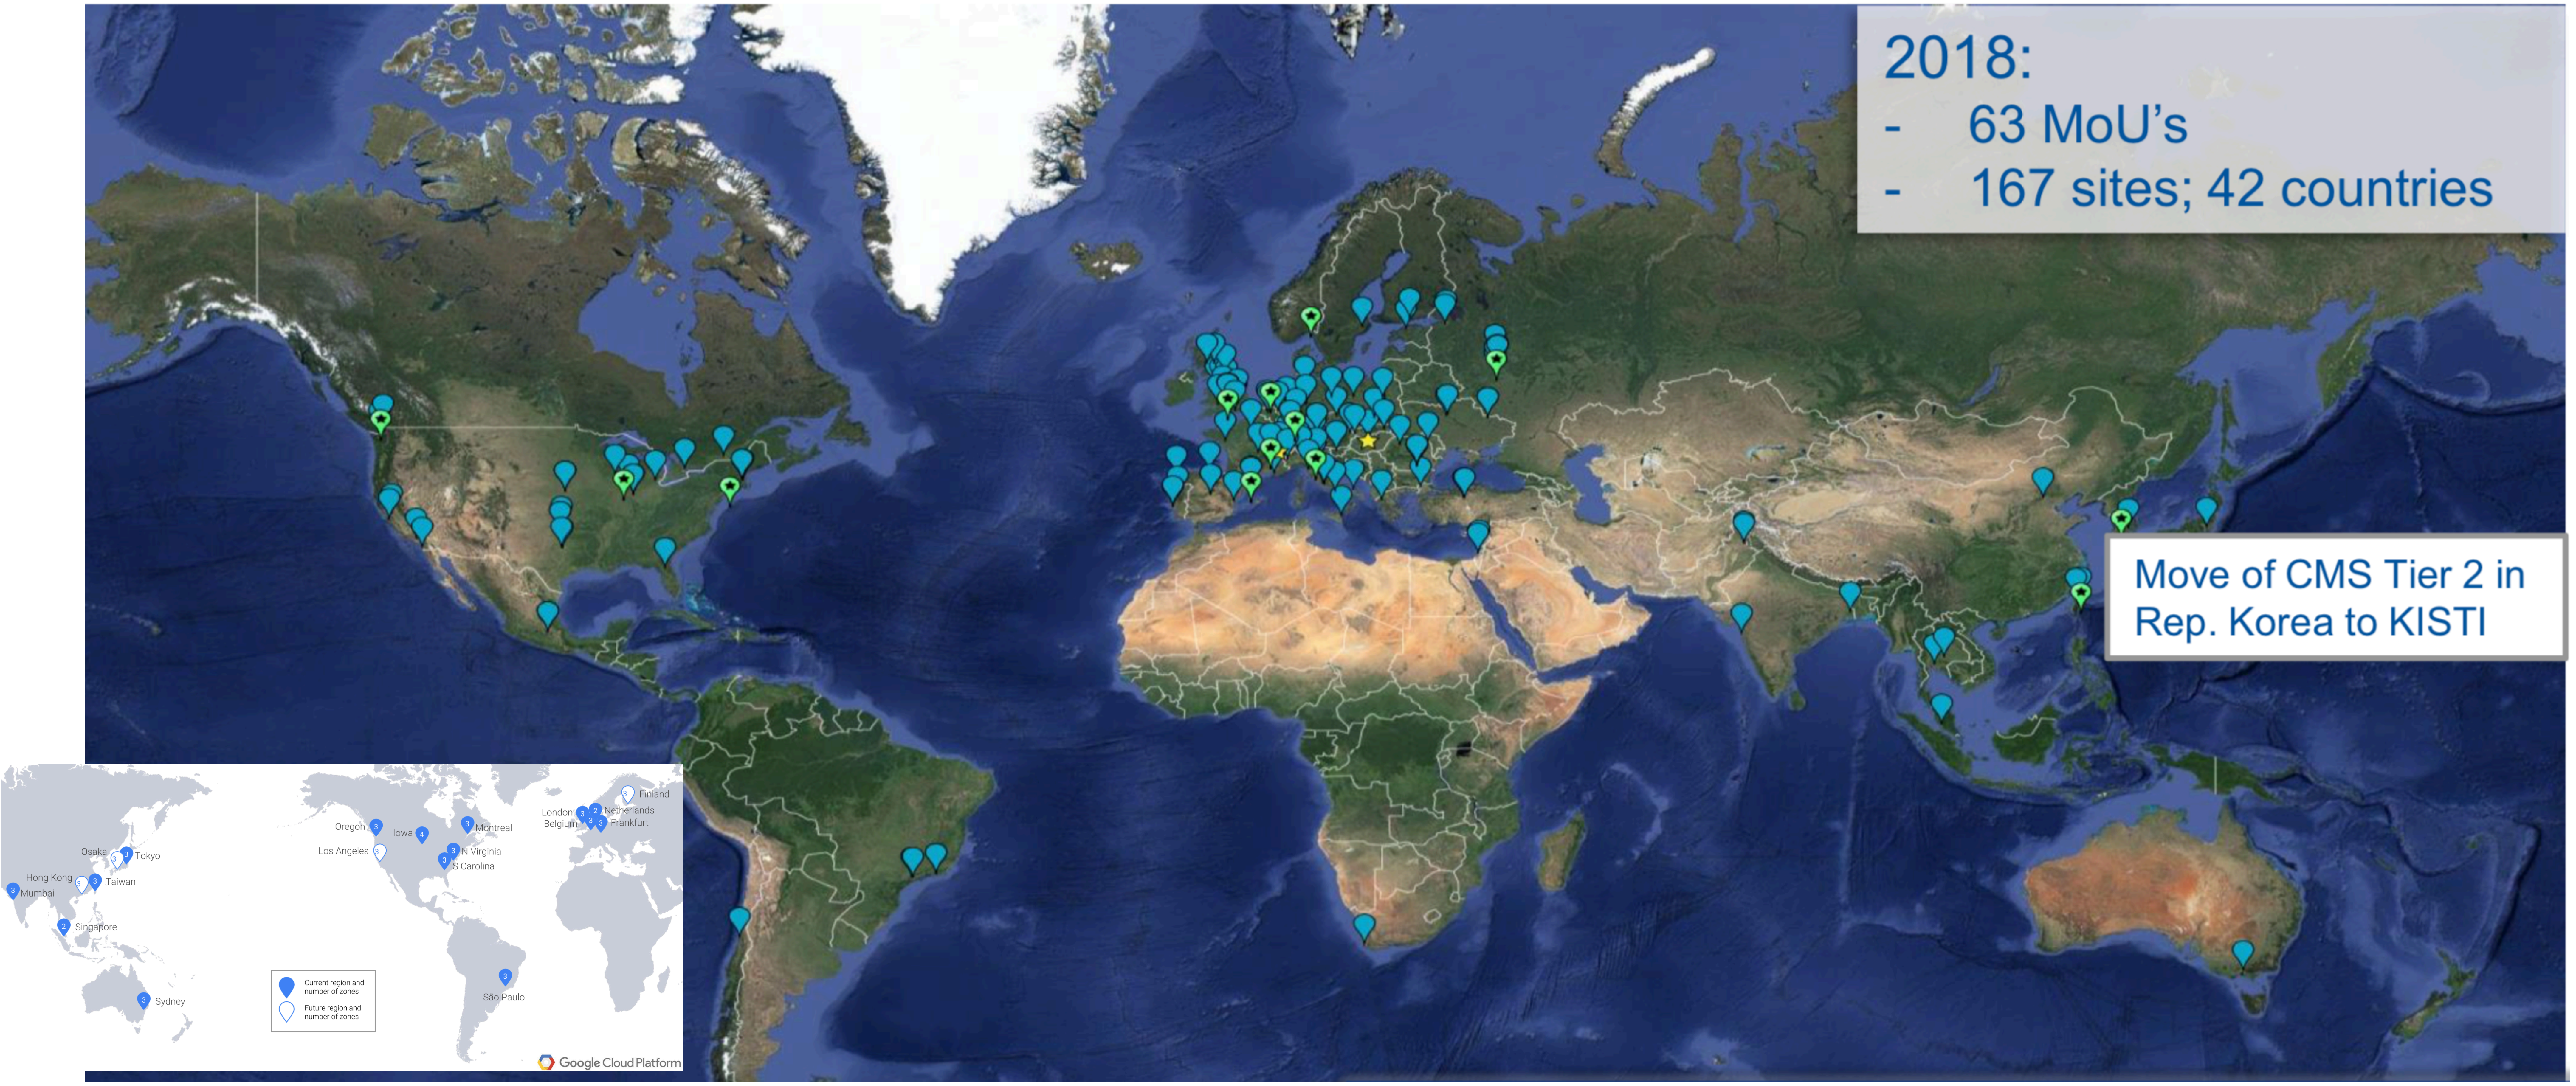

Fermilab

Explaining and Evolving the Model

- Members and funding agencies of the WLCG agree to jointly contribute to the needs of their experiments based on the physics interests of their PhD scientists.
- Computing centers are categorized by their different qualities of service in both the scale and availability of the computing they offer.
  - Tier 1 centers have > 98% uptime and provide tape in addition to CPU and Disk
  - Tier 2 centers originally had a > 80% availability but in practice do much better
  - This has opened the door to clustering into regional centers
  - Resource accounting is done in terms of Disk, Tape, and CPU
- Most national scale computing centers support multiple experiments already with the same technical infrastructure. Fermilab has to catch up.

Joint Projects with Joint People

- Joint projects across frontiers is the only way to get enough expertise working on the many future challenges.
  - Even with an infinite budget the type of scientific software & computing we do is grown not hired off the street… and we don't have unlimited budget for people.
  - Education and training, talent retention, and career path for scientist that work in these areas is imperative. Labs are essential in providing for this.
- A joint project strategy works best with stakeholders that are equals
  - DUNE needs "pre-comissioning" funding for computing in order to takes its proper place on the world stage.
  - protoDUNE is already a running experiment that needs computing and has attracted effort and resources from abroad.
  - Physics deliverables are tied to the availability of computing. If the U.S. wants to control the success of it's physics program, it has to control it's computing.

Addressing the Charge Question

Fermilab's Role

- We ask the committee to review the plans for the experiments and to comment on the role of Fermilab in the global coordination effort initiated by CERN.

- I could not get an answer to the question which CERN initiative so I will talk about all of them.
  - WLCG++
  - Openlab
  - HSF
- By supporting and participating in the above, Fermilab can play a role in getting the buy-in from the experiments needed for these initiatives to succeed.

WLCG ++

- The idea is to separate the resource allocation and RRB policy elements from the technical coordination which would then have light HSF style governance.

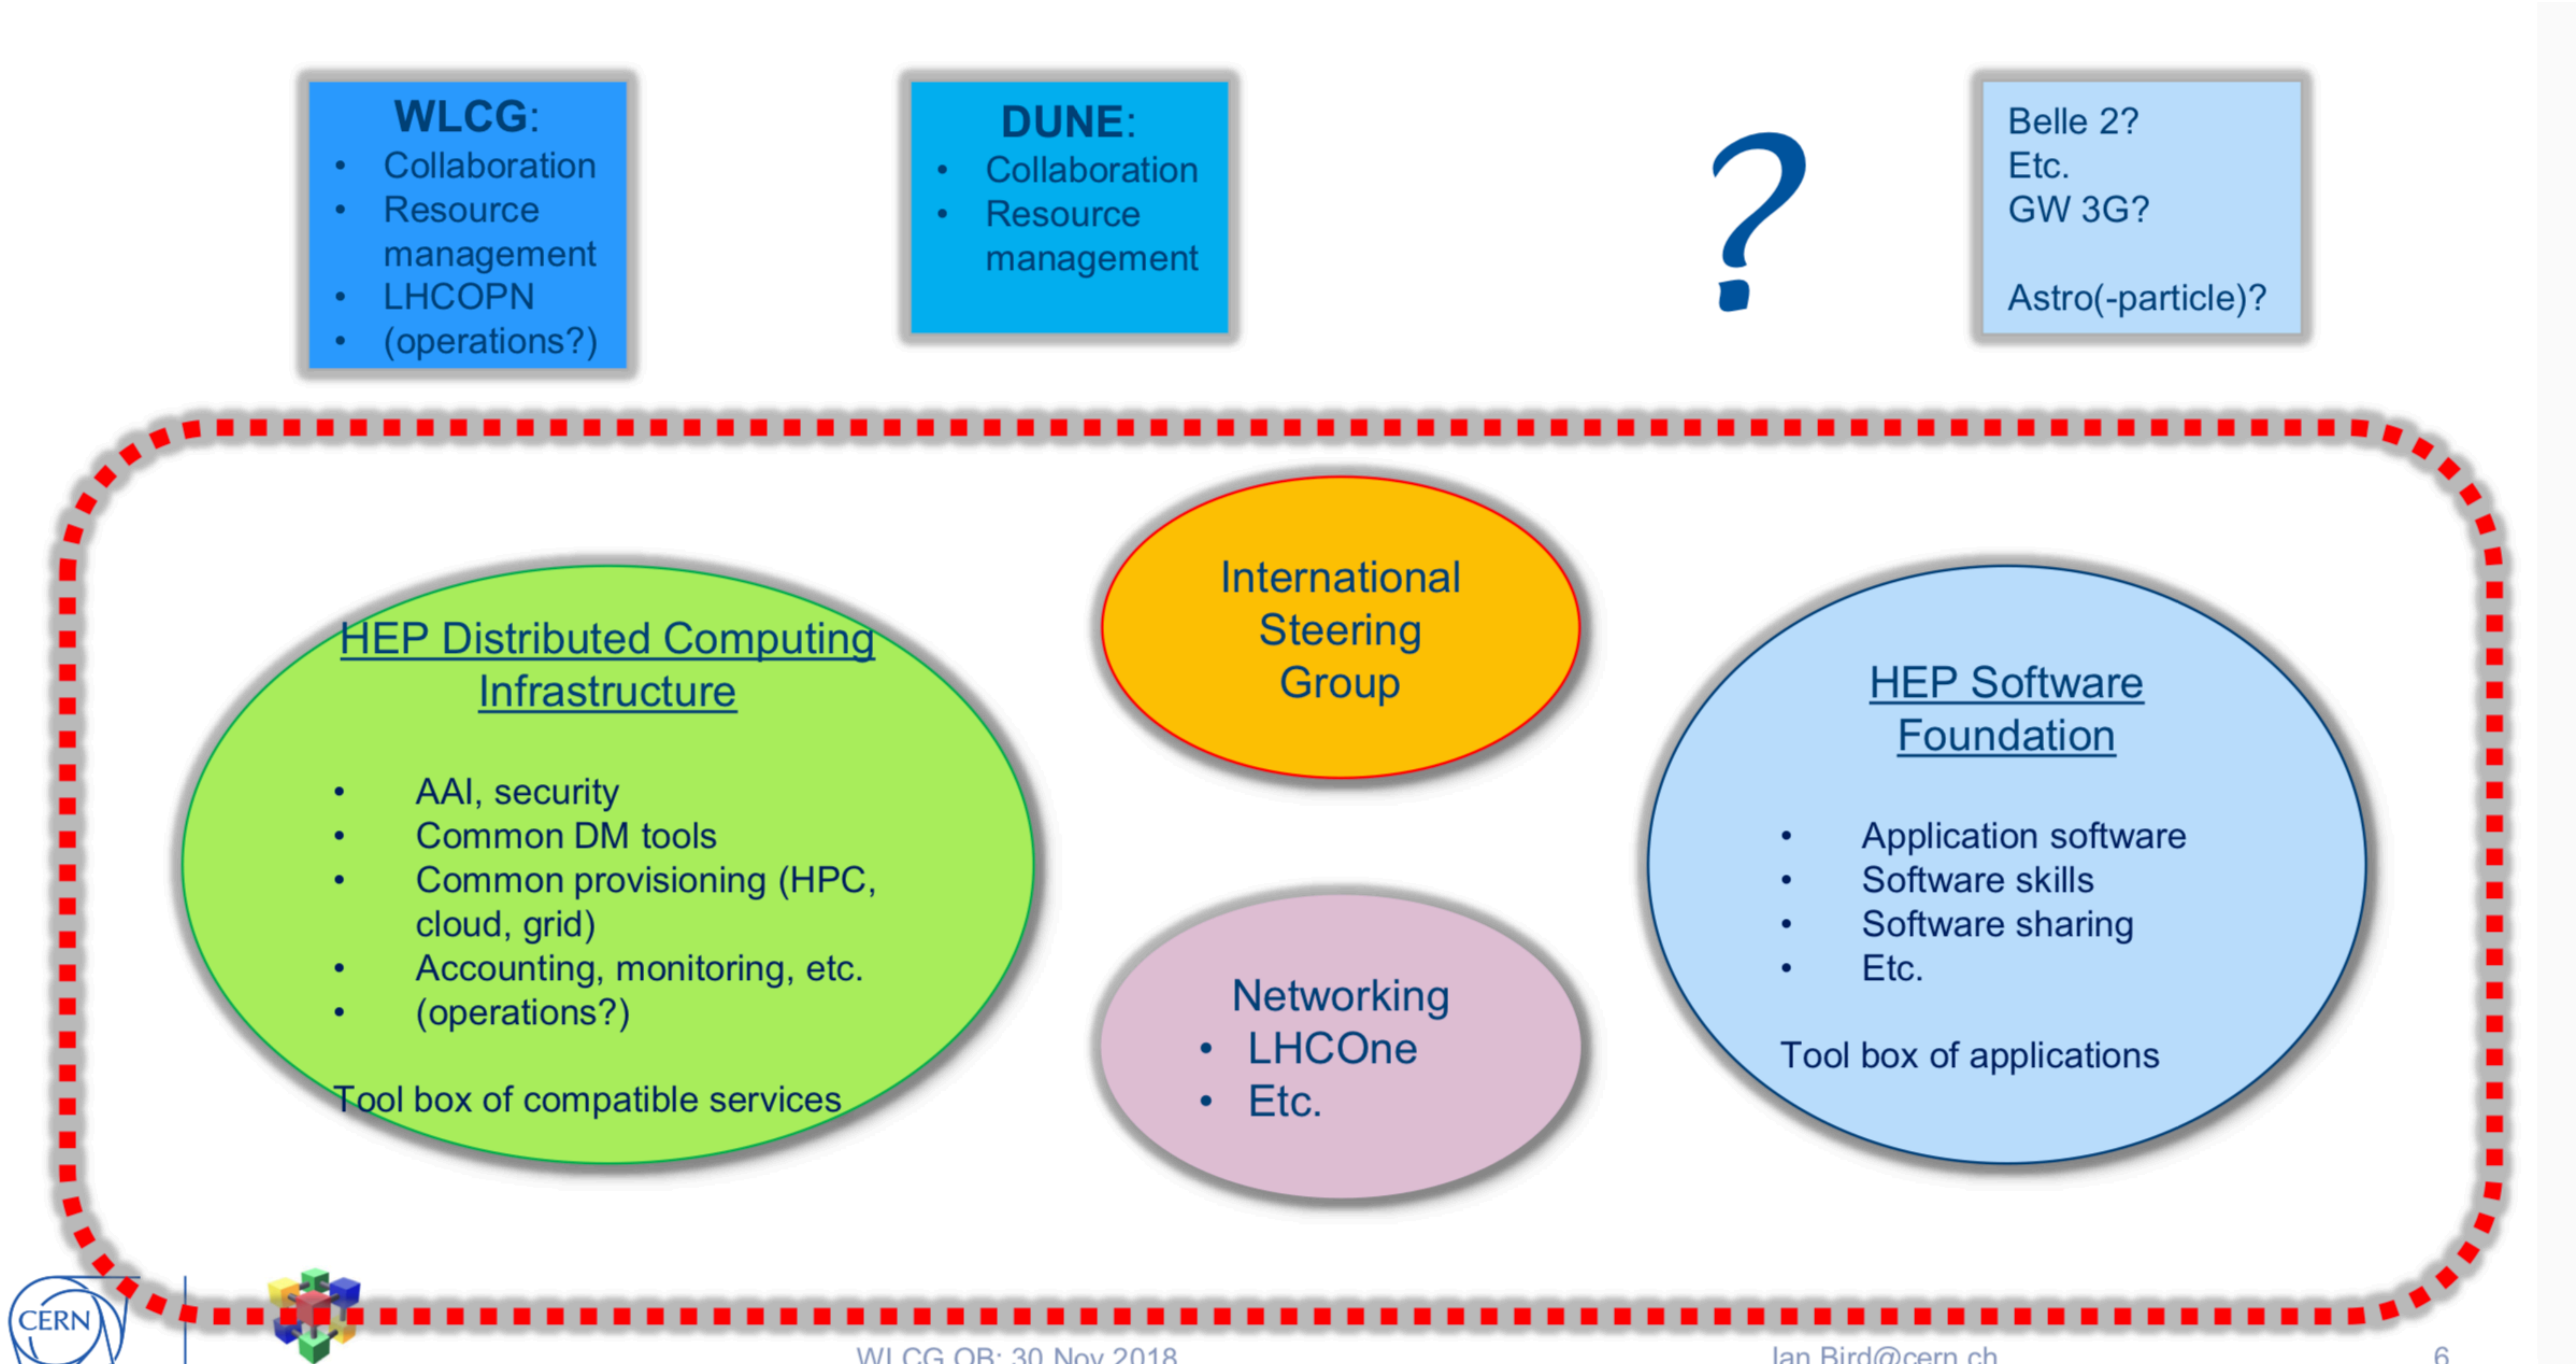

From the WLCG strategy document.

It is a down-selection of topics from the CWP

It is somewhat ATLAS centric, Fermilab could help fix that.

Item	M1 (2019)	M2 (2020)	Long Term
A.1. Distributed Storage Infrastructure	Deploy and operate a prototype across 5 T1s	Offer the prototype as pre-production service to the experiments for beta testing	Upgrade the prototype to a full scale production service
A.2. Caching	Prototype different caching solutions with different protocols	Deploy the valuable solutions to complement the distributed storage pre-production service	Deploy and operate a fully scaled content delivery network
A.3. Storage Hierarchies	Evolve the computing systems of the experiments to fully leverage a high latency and low cost multi-tiered storage architecture	Stress tests the experiment services and the facilities to measure the effectiveness of the multi-tiered model	Run an increasing number of workflows accessing data from high latency storage
A.4. Workflows	Define the relevant metrics based on the cost model. Setup a system allowing to measure the impact of different technologies/decisions on workflows	Evolve the existing workflows to leverage the benefits introduced by the new technologies introduced in M1 and M2	Introduce new workflows, tailored to the new infrastructure and services introduced in M1, M2 and beyond
B.1. Physics Validation	Setup a system automating the Physics Validation process. The validation should require close to no human effort at least when refactoring the code (n algorithmic change)		
B.2. Software Performance	Document coding best practices according to the criteria in Section 4	Organize training based on those best practices and start applying them to the newly written code	Refactor the existing code, based on the best practices.
B.3. I/O	Review the Event Data Model to benefit of new technologies and adapt to new data access pattern (latency hiding). Define the optimal granularity for data management and data access.		Evolve the I/O layer based on the criteria in 4.3
B.4. Algorithms	Identify potential algorithms and compile for the different tasks their characteristics	Evaluate the impact of the most promising algorithms with realistic tests use cases	Implement new algorithms(focus on reconstruction) for HL-LHC leveraging B.2.
B.5 Generators	Regular improvements along the lines of Section 3.1 with yearly checkpoints		
C.1. Fast Simulation	Regular Improvements in Fast Simulation with yearly checkpoints. By 2020, it should be clear what will be the impact of Fast Simulation in managing the HL-LHC cost and the implications to the computing models.		
D.1. Computing Model	Regular Evolution of the computing models, incorporating the findings from all the work plan and setting directions accordingly. Yearly checkpoints		

OpenLab

- Fermilab is the first DOE lab to sign an openlab framework agreement back in the beginning of 2017.
- Industrial partners can provide both people and supplies for agreed upon research projects.
- We have several ongoing involving Fermilab staff:
  - A Big Data project - “CMS Physics Data Reduction” (CRADA FRA-2016-0011-A)
  - Prototyping of a DL-based Particle Identification System for the Dune Neutrino Detector
  - CMS REAL-TIME STREAMING MATCHING INFERENCE ENGINE PROTOTYPE
  - CMS Intel Scaleable Key-Value store for Event Building in CMS and ATLAS

HSF

- Fermilab staff have been on the coordination team of the HSF since it's inception in 2015.
- After the writing of the CWP was sent to archive, there was a meeting in DC to discuss which parts of the work plan would be the focus of NSF and which would be DOE lab based.
- The labs went with their historical strengths in:
  - Full detector simulation
  - Generators
  - Framework
  - Grid middleware
  - Both NSF and DOE want to work on reconstruction, which is fine since there is so much to do there…

Fermilab

Local Partnerships

- In an effort to fully leverage the FRA contract, we are starting a cooperations initiative with Argonne and UChicago in software and computing

- First face to face meeting involving executives on both the business and scientific computing next week.

- Examples
  - opportunities to involve university students in lab projects
  - facilitate CAS opportunities
  - combine purchasing of business systems in order to improve our combined purchasing power

Summary

- The Fermilab Computing strategy is to leverage resources within the broader HEP community to meet the challenging needs of the next decade.
- In partnership with CERN and our local community we will build on the success of the LHC computing program which has been an extraordinary success.

Back up

Community White Paper [1]

HEP Software Foundation

A Roadmap for HEP Software and Computing R&D for the 2020s

- Inspired by the P5 process and guided by its goals
- The Global Community White Paper provides a roadmap to extend commonality to a broader set of software.
  - 70 page document
  - 13 topical sections summarising R&D in a variety of technical areas for HEP Software and Computing
  - Almost all major domains of HEP Software and Computing are covered
  - 1 section on Training and Careers
  - 310 authors (signers) from 124 HEP-related institutions

[1] https://arxiv.org/pdf/1712.06982.pdf

Contents