U.S. DEPARTMENT OF ENERGY | Office of Science

Future Computing in High Energy Physics

22nd International Conference on Computing in High Energy and Nuclear Physics, CHEP 2016

Jim Siegrist
Associate Director for High Energy Physics
Office of Science, U.S. Department of Energy

The Science Drivers of Particle Physics

The P5 report identified five intertwined science drivers, compelling lines of inquiry that show great promise for discovery:

- Use the Higgs boson as a new tool for discovery **2013*
- Pursue the physics associated with neutrino mass **2015*
- Identify the new physics of dark matter
- Understand cosmic acceleration: dark energy and inflation **2011*
- Explore the unknown: new particles, interactions, and physical principles

** Since 2011, three of the five science drivers have been lines of inquiry recognized with Nobel Prizes*

P5 Science and Construction Timeline (abridged)

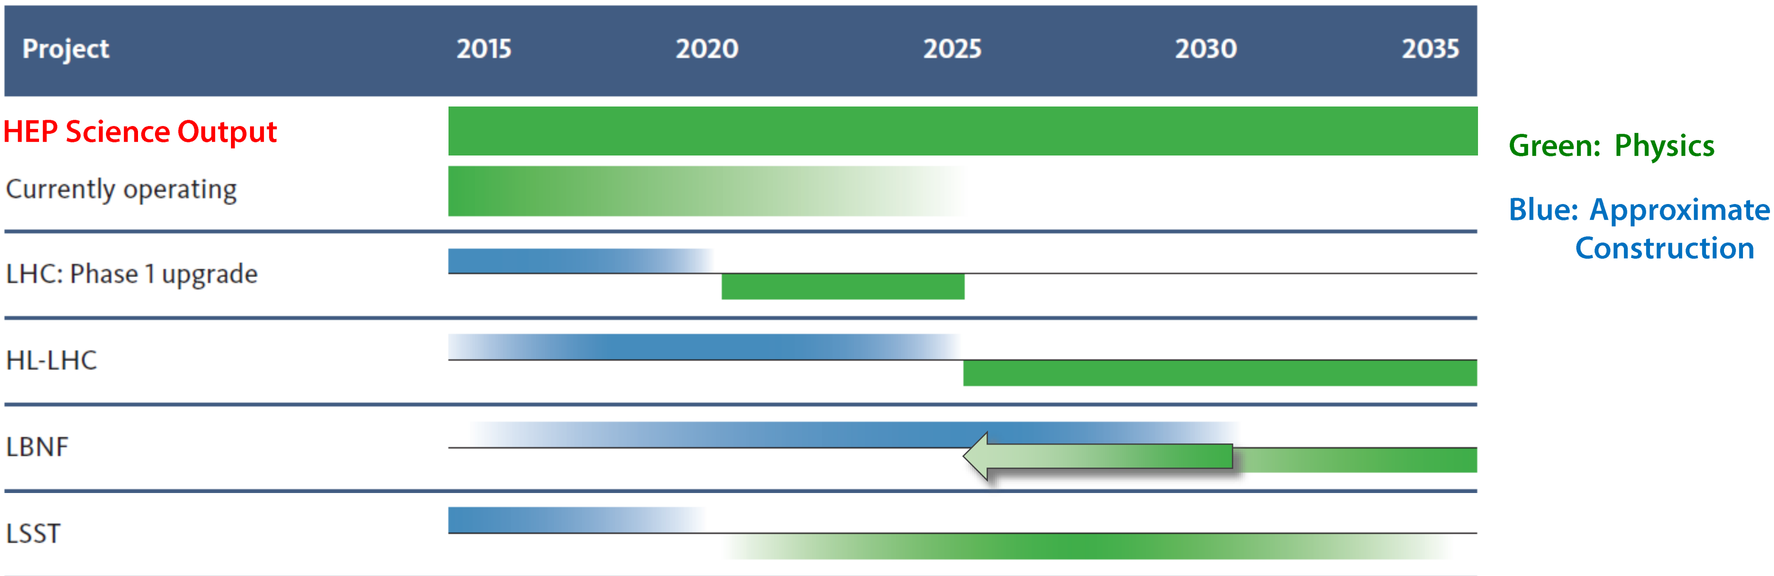

- The P5 report recommends a **limited, prioritized and time-ordered list of** experiments to optimally address the science drivers
  - Covers the **small, medium and large investment scales**
  - Will **produce results continuously** throughout a 20-year timeframe
- HEP is implementing the discovery-driven strategic plan set within a global vision for particle physics as presented in the P5 report
- **Realizing this vision will require a shift in approaching the networking and computing challenges the data from these experiments will present!**

P5 Vision for Computing

- **P5 report recognized the importance of computing:**
  - "Rapidly evolving computer architectures and increasing data volumes require effective crosscutting solutions"
  - "[Need] investments to exploit next-generation hardware and computing models"
  - "Close collaboration of national laboratories and universities across the research areas will be needed"
- **P5 Recommendation 29:**
  - Strengthen the global cooperation among laboratories and universities to address computing and scientific software needs, and provide efficient training in next-generation hardware and data-science software relevant to particle physics. Investigate models for the development and maintenance of major software within and across research areas, including long-term data and software preservation.
- **HEP Response to P5 Recommendation 29:**
  - Initiated HEP Center for Computational Excellence (CCE) http://hepfce.org/

Computing Enables HEP

- Computing is an integral part of theory, experiment, technology development
  - Many recent successes only possible because of significant community effort to develop and advance the necessary computing tools!
- Example: US LHCNet Networking
  - October 1986: 1st Annual Workshop on Energy Research Computing:
    - "Just as we expect a computer to perform as if we are the only user, we expect the network to give that same appearance."
  - July 2012: Higgs boson discovery
    - US LHCNet transatlantic networking capabilities enabled U.S. physicists to play important roles in LHC operations and analysis

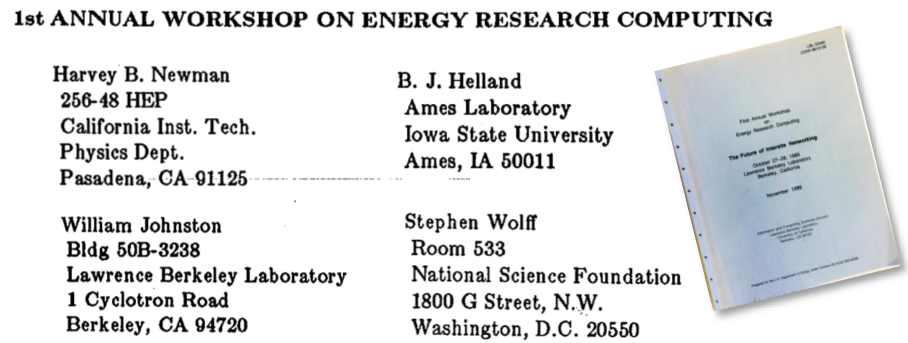

1st ANNUAL WORKSHOP ON ENERGY RESEARCH COMPUTING

Harvey B. Newman
256-48 HEP
California Inst. Tech.
Physics Dept.
Pasadena, CA 91125

B. J. Helland
Ames Laboratory
Iowa State University
Ames, IA 50011

William Johnston
Bldg 50B-3238
Lawrence Berkeley Laboratory
1 Cyclotron Road
Berkeley, CA 94720

Stephen Wolff
Room 533
National Science Foundation
1800 G Street, N.W.
Washington, D.C. 20550

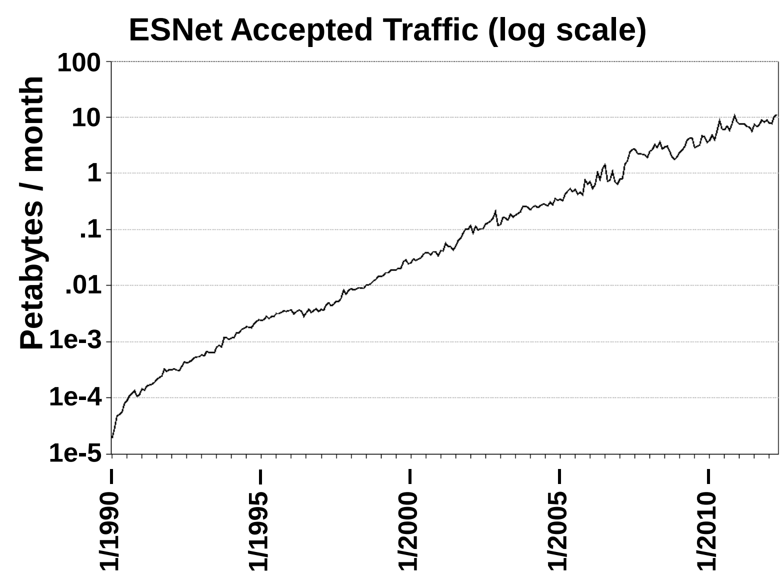

Example: WLGC Successful Global Collaboration

Tier-0
[CERN and Hungary]: data recording, reconstruction and distribution

Tier-1
permanent storage, reprocessing, analysis
U.S. Tier-1 sites (DOE-supported):
CMS: Fermilab
ATLAS: BNL

Tier-2
Simulation, end-user analysis
U.S. Tier-2 sites at local universities (largely supported by NSF)

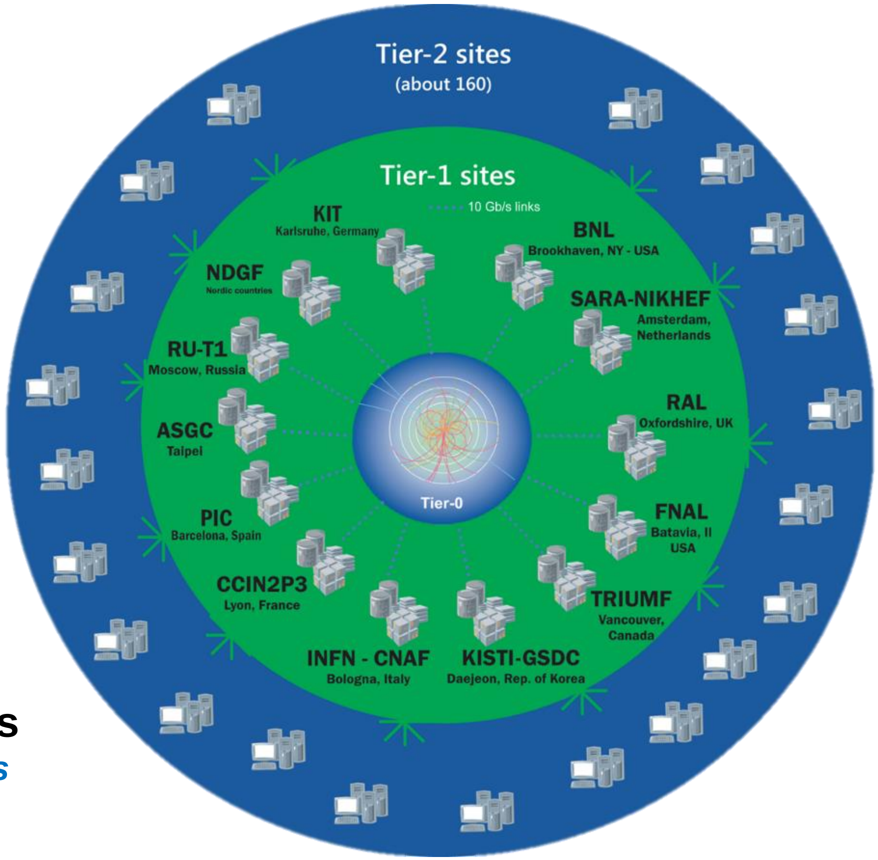

Tiers-1 and 2: nearly 173 sites, 35 countries

~300,000 cores

Average 173 PB of storage

> 2 million jobs/day

10 Gb links

The Worldwide LHC Computing Grid (WLCG):

Integrates computer centers globally (including the U.S.) to provide computing and storage resources into a single infrastructure accessible by all LHC physicists for data analysis

Example: art Software Framework

- art provides common software framework for data reconstruction and analysis
  - Challenge: Evolve art software to meet the needs of experiments and keep pace as commodity hardware platforms evolve
- Leveraging common computing and software infrastructure needed to turn data to scientific results enables experiments to focus on reconstruction and analysis tasks
  - art now used by all of Fermilab’s modern Neutrino and Muon experiments
  - Provides framework for developing common algorithms (LArSoft) for all Liquid Argon experiments
  - MicroBooNE takes advantage of this infrastructure for its first results

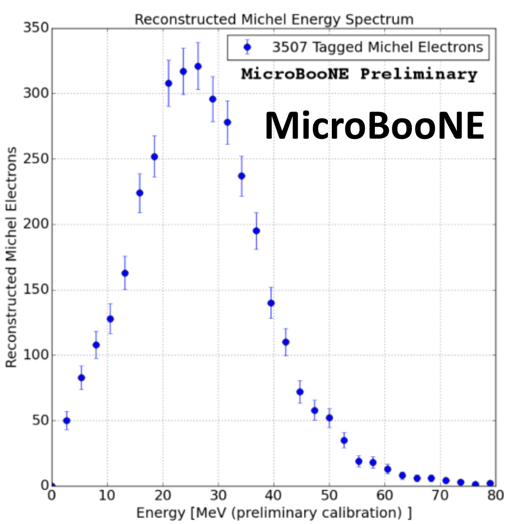

Fully automated reconstruction

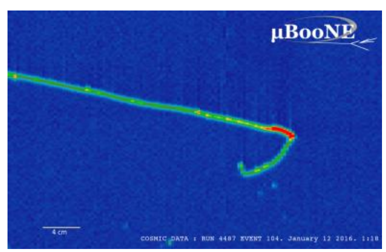

LArSoft toolkit co-developed with all LArTPC experiments

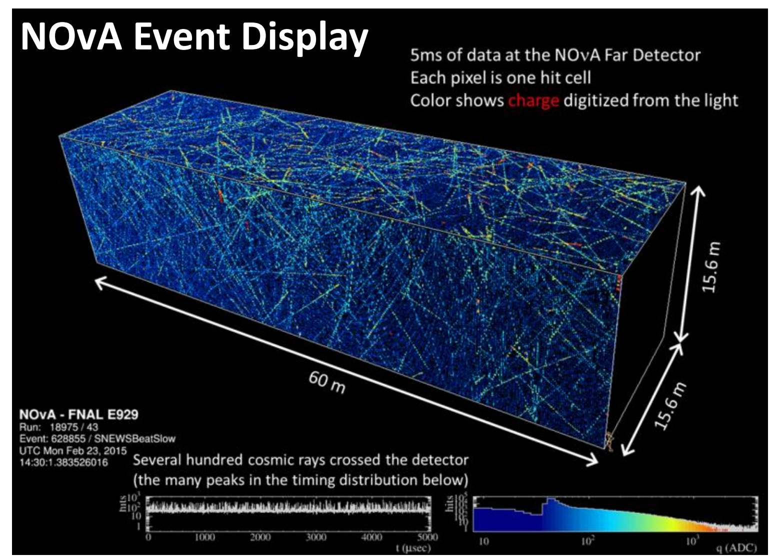

Example: Computational Cosmology

- Computational cosmology is leveraging High Performance Computing
  - HEP & ASCR coordination on multiple efforts including Exascale Computing
    - New ASCR ECP project for cosmological simulations
  - Cosmic simulations, emulators, data-Intensive computing and analysis
  - NERSC allocations for simulations and data analysis for CMB experiments
  - ALCC (LCFs and NERSC) and INCITE projects (LCFs) for large-scale cosmological simulations
    - Underlying recent SPT X DES measurement of the kSZ effect
  - Interactive Data Portals for both experimental and simulation data sets

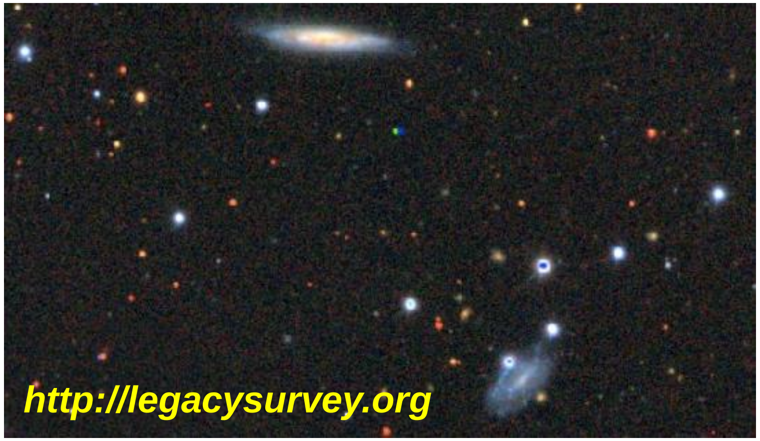

Data Release 3 for DECam Legacy Survey providing target selection for DESI experiment

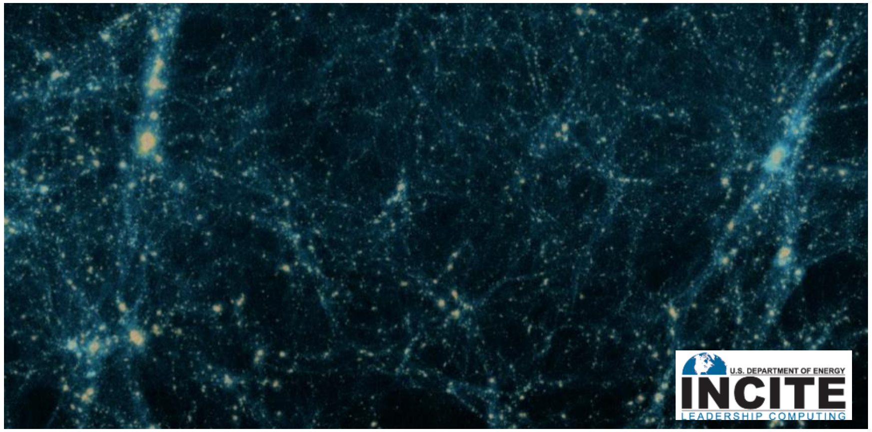

Snapshot of the “Q Continuum” cosmology run with HACC on Titan (about 1/20,000 of the actual volume)

Projected Shortfalls in HEP Computing Resources

- Expected to be 10x – 100x shortfall by 2025
  - Shortfalls
    - Data movement – needing smart networks
    - Hardware for simulations, data analysis, and storage
    - Workforce, expertise & training
  - Computing Ecosystem critical to workflows and results
  - Can this be handled entirely within HEP resources and within HEP subprograms?
  - Need a shift in strategy to best prepare for future while managing current operations and using resources external to HEP
- Sources:
  - SNOWMASS Computing Frontier Reports
  - 2013 Computing Planning Meeting
  - FCE (CCE) Working Group Reports
  - HEP-ASCR Exascale Requirements Report
  - Other studies, talks and input
  - European initiatives
  - Publicly available facts…

Excellent LHC Performance and Near-term Challenges

- LHC continues to set new performance records
  - Unprecedented peak instantaneous luminosity over 1.3 x 10^{34} $cm^{-1}s^{-1}$ exceeds design luminosity by >30%!
  - Data accumulation up to 30-40% beyond the goal of 25 fb^{-1} for 2016
- Congratulations to the CERN accelerator team for the hard work in operating the LHC, and to the experiments for the high performance efficiency in acquiring data!

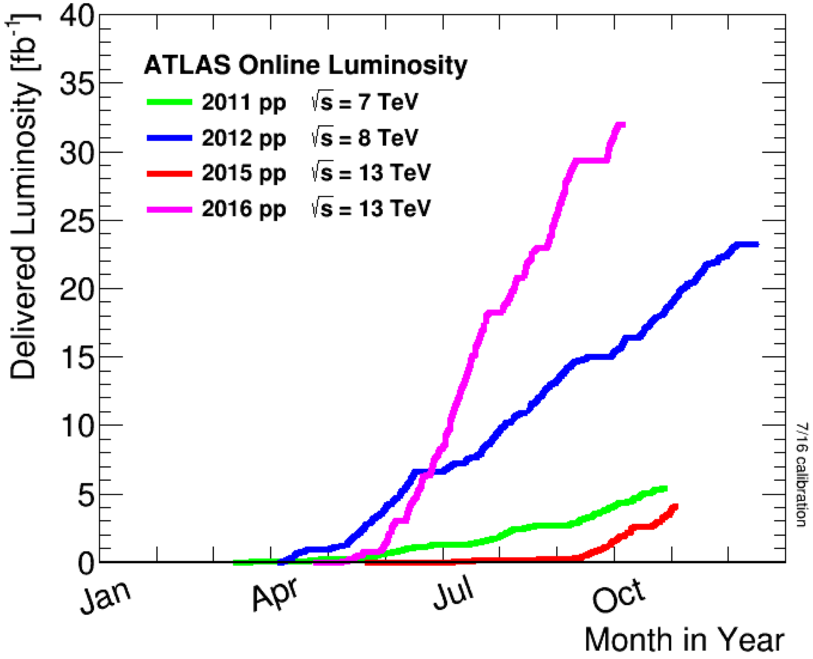

- LHC performance has immediate challenges in computing resources that are needed to support operations and analysis efforts as a result of the increased amount of data generated by ATLAS and CMS by 20-40% in Run 2
  - Increases anticipated for additional CPU, Disk, and Tape resources by ~20% in FY 2017
- DOE is coordinating with the experiments, CERN, and its partners through the international process to address these issues

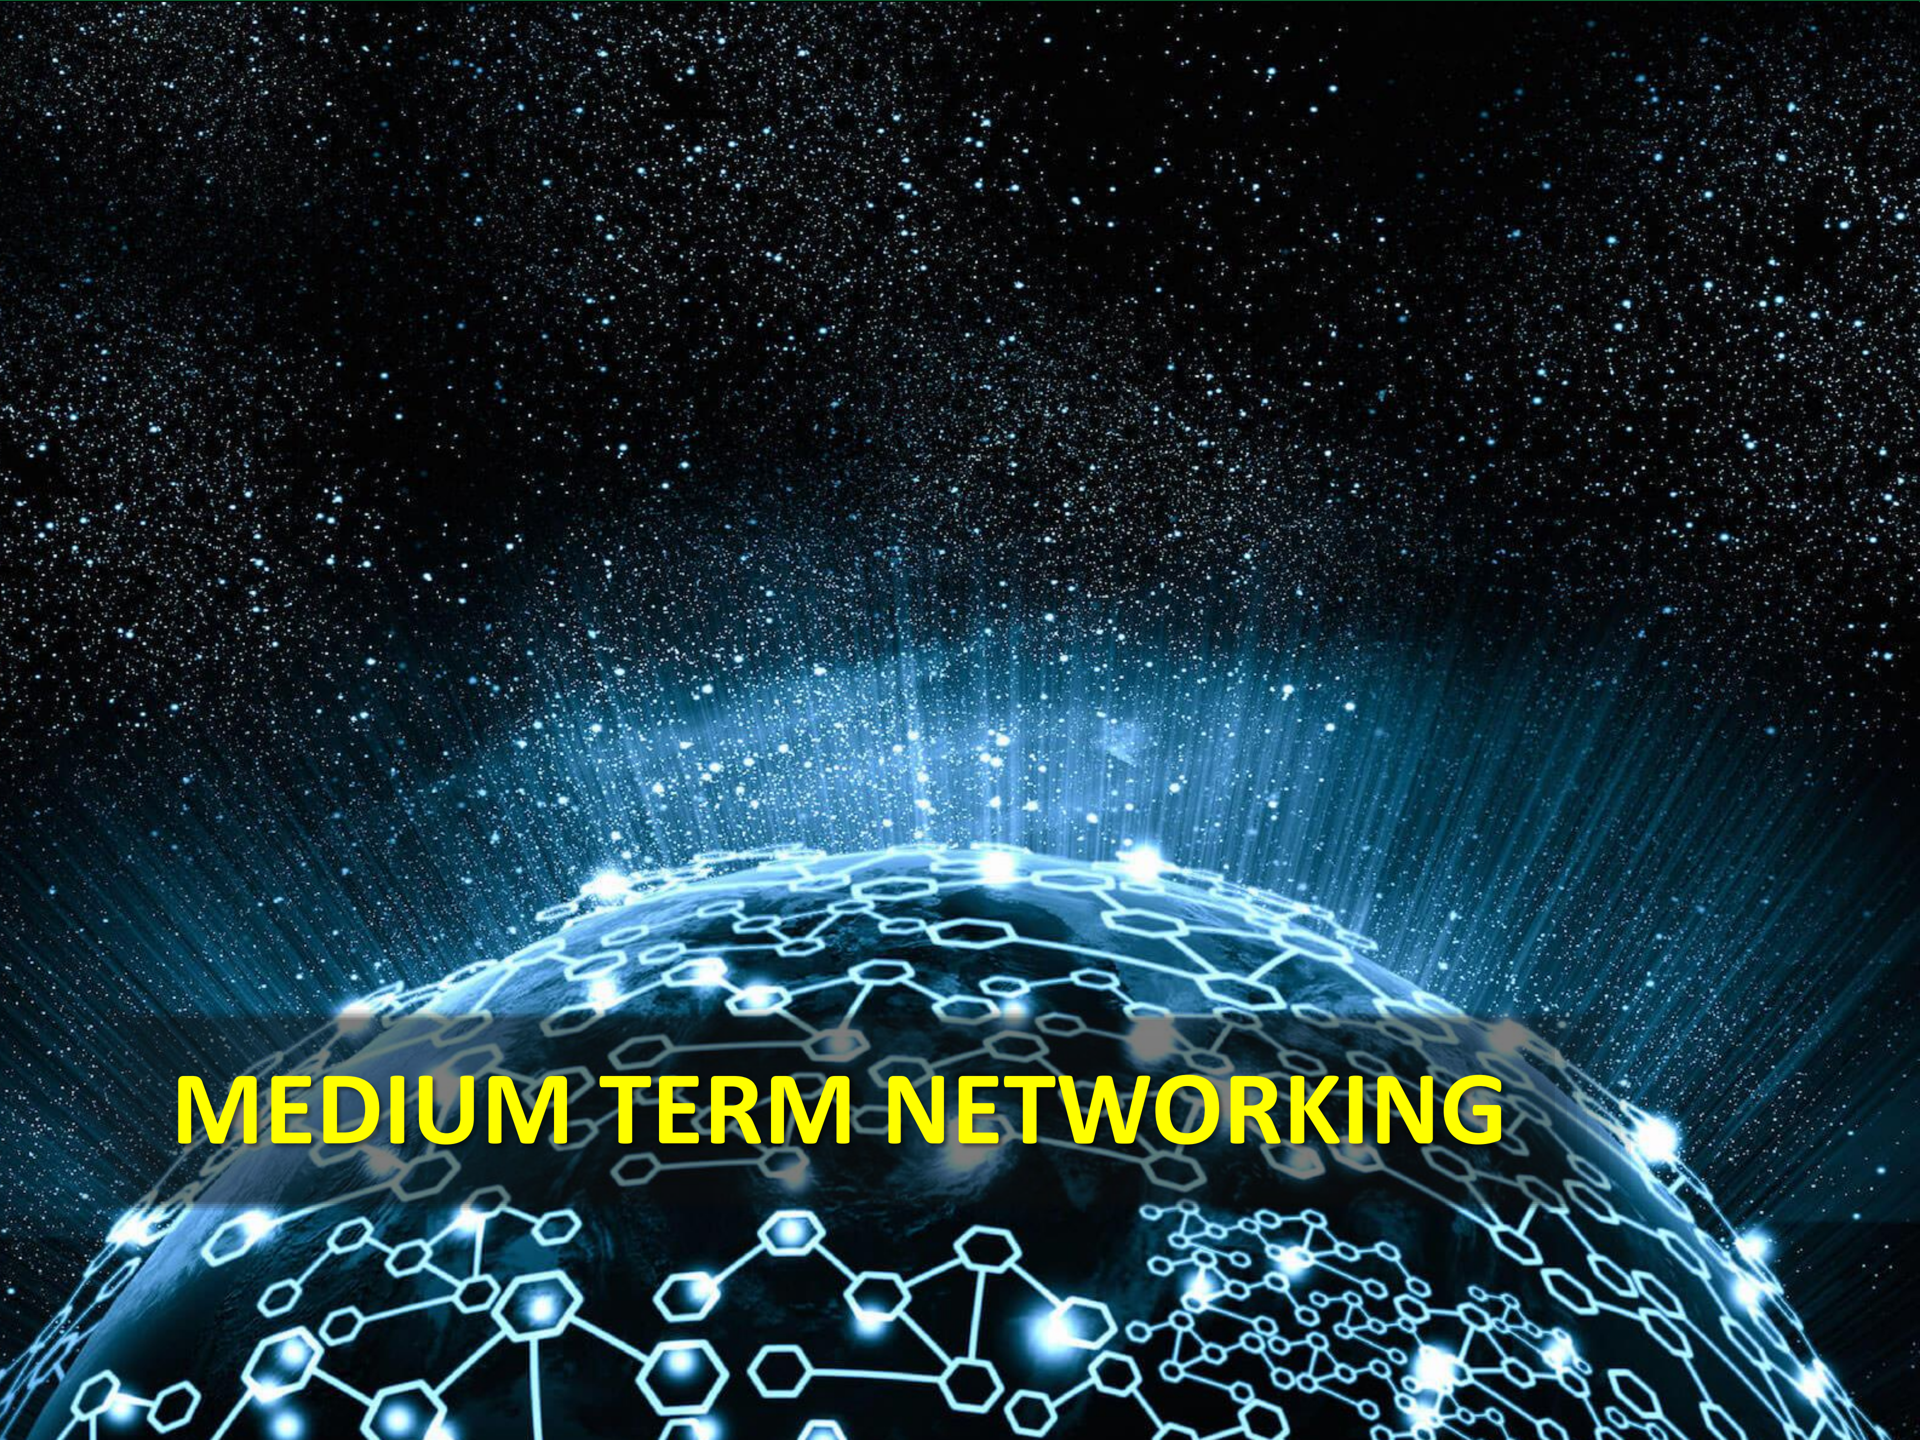
MEDIUM TERM NETWORKING

ESNet Projected Traffic

- ESNet provides world-class support for scientific discovery for Office of Science researchers and their collaborators
- Rate of increase follows or exceeds historical trend of 10x / 4 yrs.
- HEP traffic will compete with other Office of Science programs
- Current system will not keep pace with projected traffic
  - Projection reaches 1 Exabyte/month by ~2020, 10 EB/mo. by ~2024

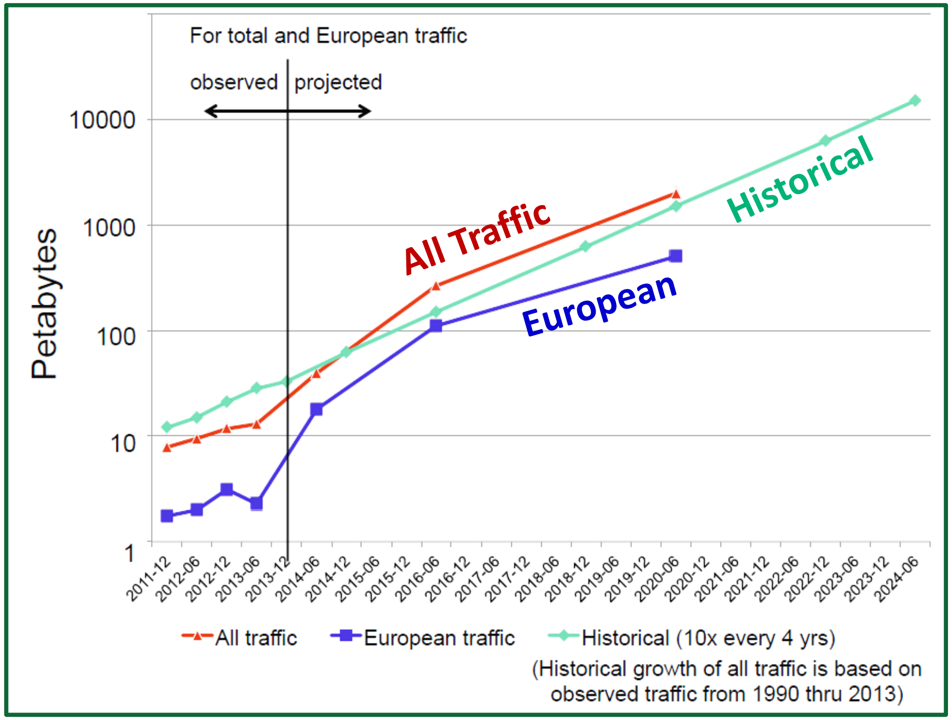

43 PB/month in April 2016

LHC Run 2 is accelerating the growth

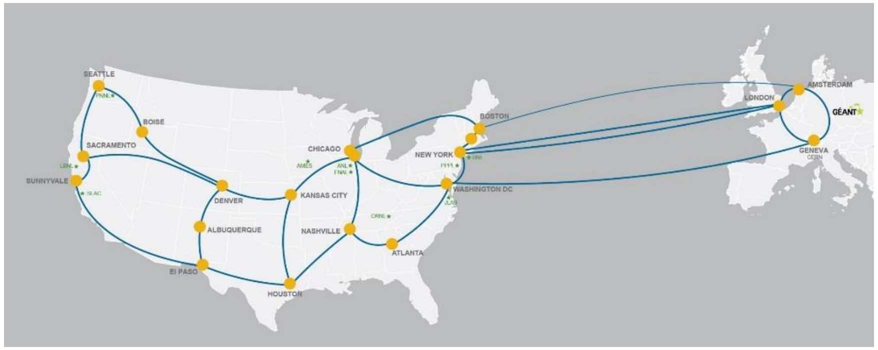

Next Generation of Networking

- Vision: Distributed environments where resources can be deployed flexibly to meet the demands
- Software-Defined Networking (SDN) is a natural path to this vision:
  - Separating the functions that control the flow of traffic, from the switching infra-structure that forwards the traffic through open deeply programmable "controllers"
- With many benefits:
  - Replacing stovepiped vendor HW/SW solutions by open platform-independent software services
  - Virtualizing services and networks: lowering cost and energy, with greater simplicity
  - Adding intelligent dynamics to system operations
  - A major direction of research networks and industry
  - A sea change that is still emerging and maturing

Building on the Caltech/ESnet/Fermilab Pilot Experience

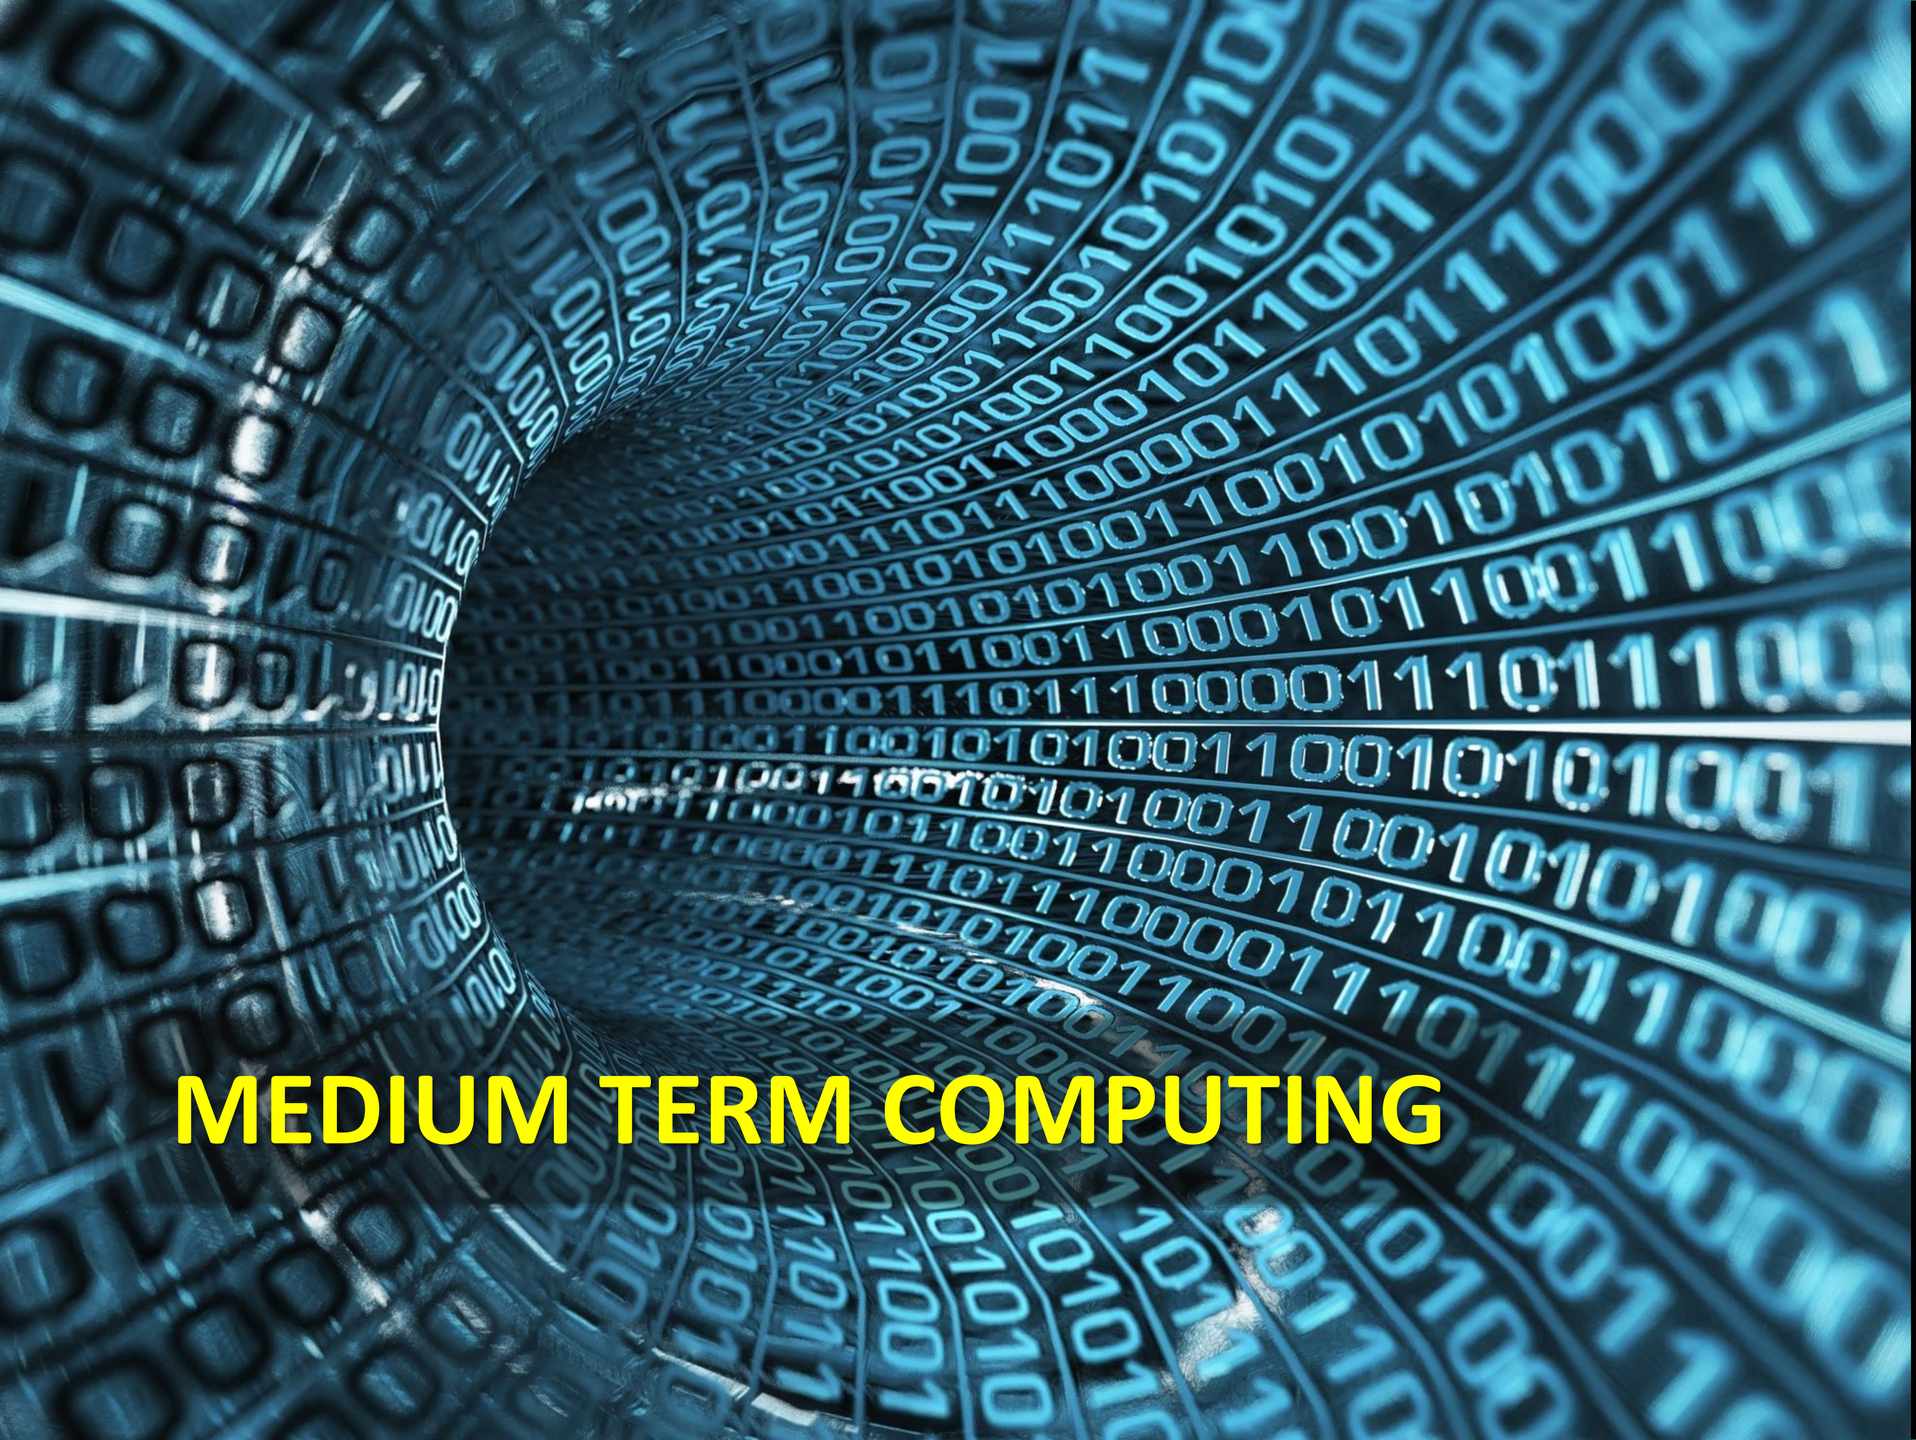

MEDIUM TERM COMPUTING

Options for Future Computing

- **Example: LHC Run 4 (2026 and beyond) will start the exabyte era for HEP!**
  - How will the data be processed and analyzed?

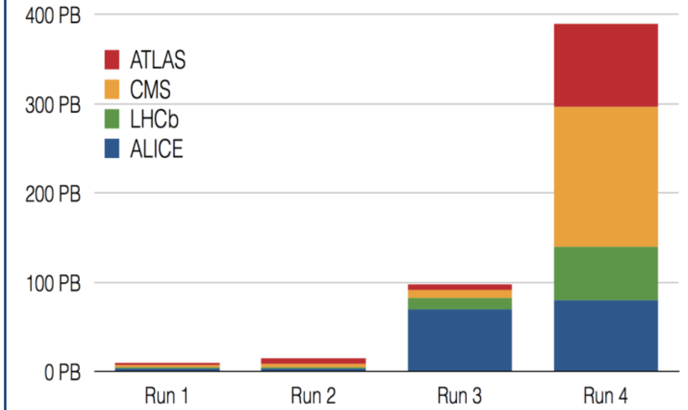

Ref: ECFA HL-LHC Computing Meeting, Oct 2014

☐ **Buy facilities**
- ✓ Own it! No impediment to running at full capacity when needed
- ✗ Must invest for peak utilization even if not used

☐ **Use services from other providers**
- ✓ Others make capital investments
- ✗ Will usage be available/affordable when needed?

- **Evolution of HEP networking provides a promising example for pursuing computing as part of infrastructure not owned by HEP**
  - Like ESNet; not necessary for HEP to purchase all hardware

☑ **Hybrid model**
- ✓ Own baseline resources that will be used at full capacity
  - Reliable cycles available for reconstruction, MC generation, etc.
- ✓ Use service providers for peak cycles when needed
  - Conference analysis season, special collaboration needs, etc.

- **To achieve P5 global vision, all partners need to bring in their available resources!**

DOE Computing Resources

- DOE's major computing resource is the Advanced Scientific Computing Research (ASCR) program
- ESNET provides valuable network resources for science
  - Initiation of the ESNet TransAtlantic networking is a successful example of ASCR-HEP partnership
- Through National Energy Research Scientific Computing Center (NERSC), ASCR provides reliable HPC resources
  - ✓ Will exist for "free"
  - ✓ Hardware will become available for HEP to buy
  - ✗ Must port code!

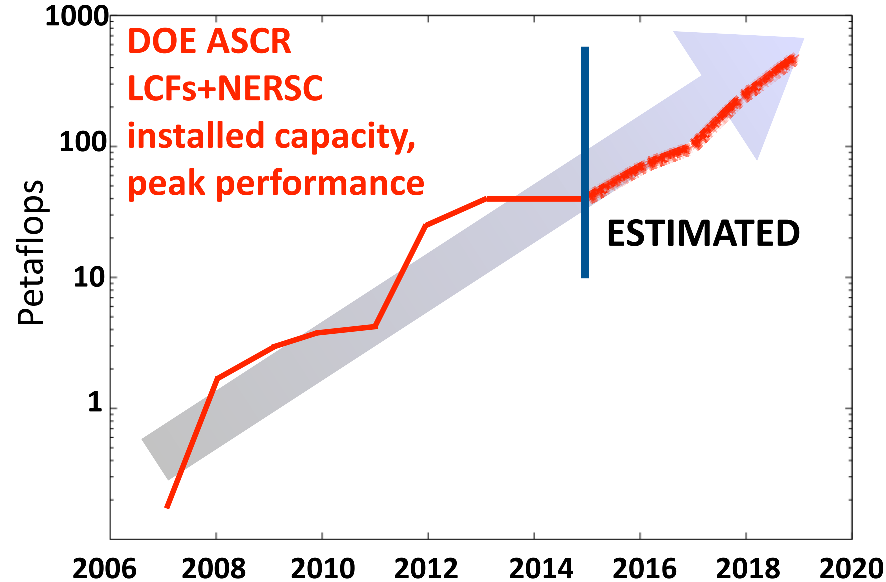

Peak performance as a function of time: approach to exascale by LCF/NERSC systems

- HEP is working collaboratively across the program to optimize use of DOE resources, including HEP cloud, through CCE and other efforts

Computational HEP and Overall HEP Computing

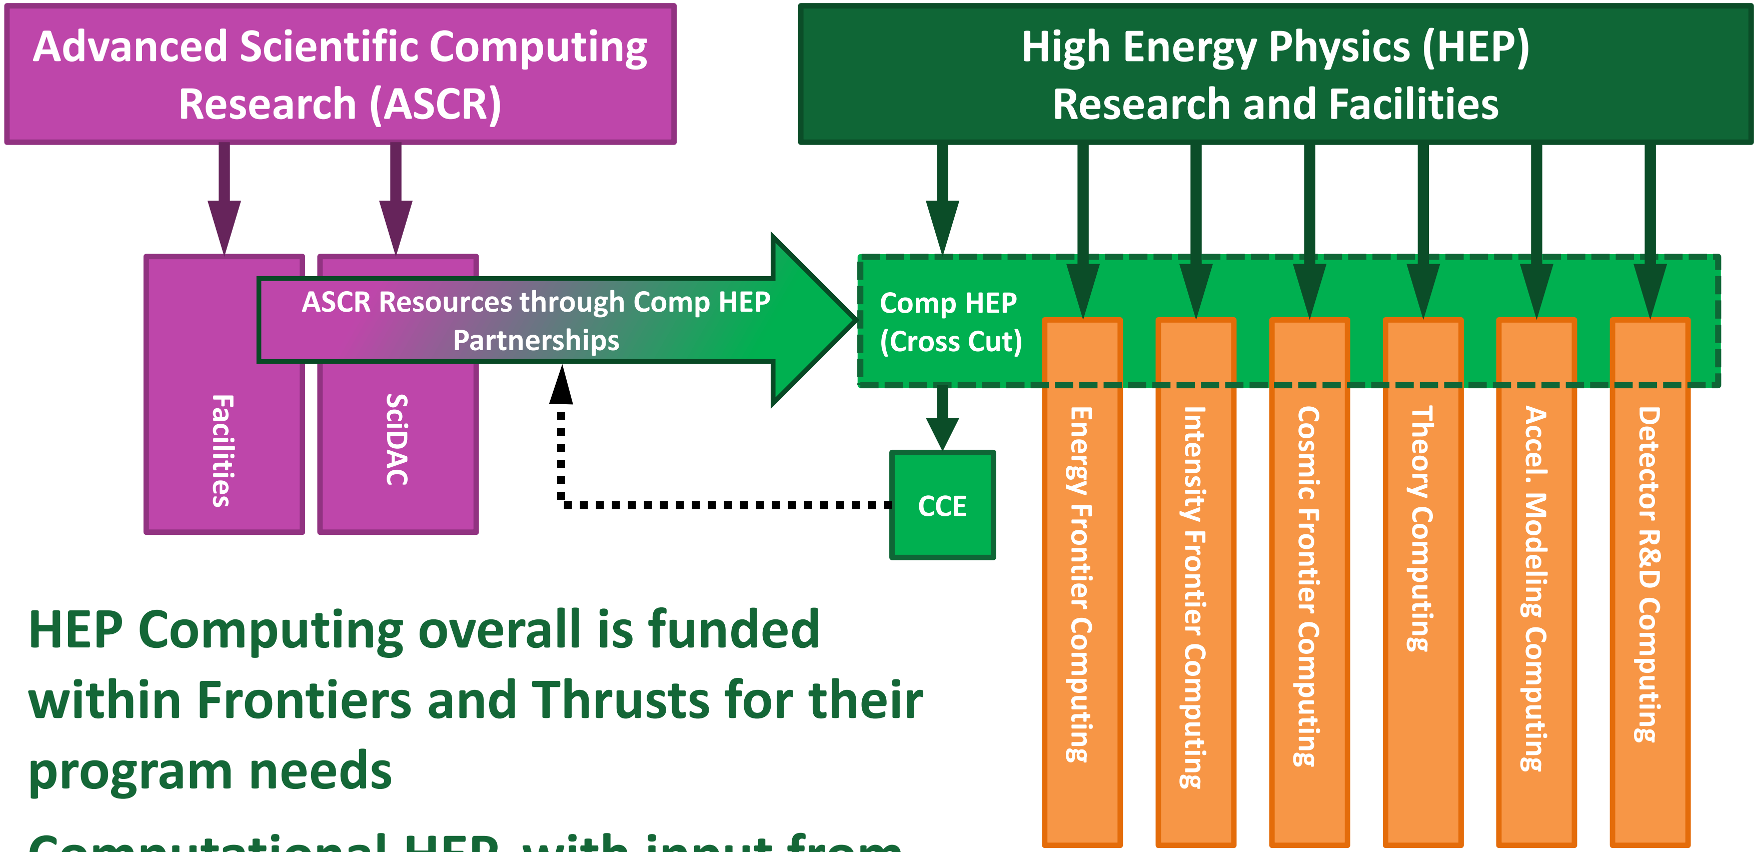

- HEP Computing overall is funded within Frontiers and Thrusts for their program needs
- Computational HEP, with input from CCE, identifies where external partnerships & cross cuts are possible and fosters them

HEP Computing Future Challenges

- **Future computing will intertwine different computing paradigms**
  - **High Throughput Computing (HTC)**
    - Increasing HEP experiment demand may outpace Grid computing resources
    - **Need**: New hardware/software exploits
  - **High Performance Computing (HPC)**
    - Classic use of HPC resources by theorists
    - **Need**: Event services for simulations and dedicated front-ends for job packaging
  - **Data-Intensive Scalable Computing (DISC)**
    - Analysis of simulations & comparison to observational data without HTC lead times
    - **Need**: True interactive largescale computing
- **Challenge will be to adapt the HEP computing model to optimize operations and analysis workflow to exploit all resources**
  - **Software Stack**: Run arbitrarily complex software stacks
  - **Resilience**: Handle failures of job streams
  - **Resource Flexibility**: Complex workflows with changing computational 'width'
  - **Wide-Area Data Awareness**: Seamlessly move computing to data & vice versa
  - **Automated Workloads**: Run large-scale automated production workflows
  - **End-to-End Simulation-Based Analyses**: Run analysis workflows on simulations using a combination of in situ and offline/co-scheduling approaches

New Architectures Foster Data Intensive Computing

- High bandwidth external connectivity to experimental facilities from compute nodes (Software Defined Networking)
- NVRAM Flash Burst Buffer as I/O accelerator
- More login nodes for managing advanced workflows
- Support for real time and high-throughput queues with SLURM
- Virtualization capabilities with Shifter (docker containers) – CCE

Burst Buffer: Non-volatile storage in HPC system for application I/O acceleration

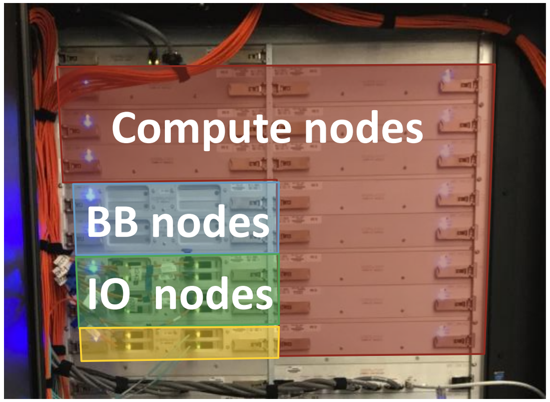

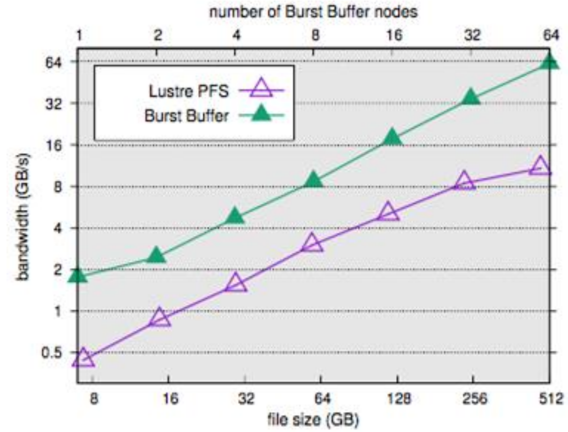

Enhanced Networking for CORI

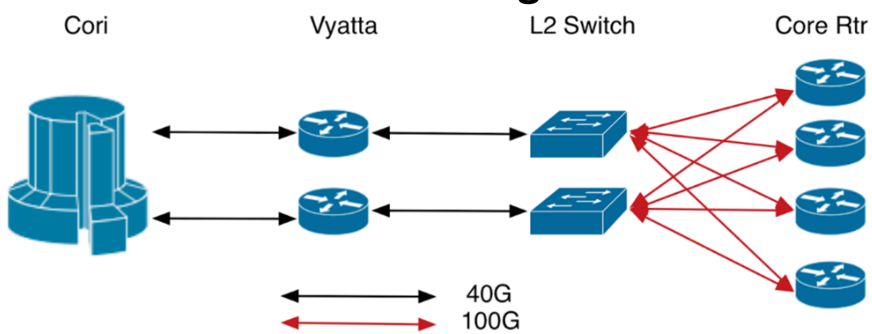

LONG TERM COMPUTING

Relevant National Events External to HEP

- **National Strategic Computing Initiative (NSCI) Executive Order 13702, signed by U.S. President on July 29, 2015**
  - DOE is one of the lead agencies executing the mission – with ASCR and NNSA primarily responsible – Exascale Computing Project
  - Three HEP groups have earned ASCR funded code development projects:
    - Cosmic Frontier codes
    - Plasma acceleration codes
    - Lattice QCD software

- **Many of you participated in HEP-ASCR Exascale Requirements Review**
  - Advantages to HEP:
    - Use of powerful computing capabilities
    - Exploring new hardware for HEP facilities
    - Develop advanced computing ecosystem:
      - Data – data movement – networks – hardware – software
  - Report available at: https://arxiv.org/abs/1603.09303

Quantum information Science (QIS) & HEP

- **The NSCI also pointed to post-Moore computing technologies including Quantum Computing**
  - Quantum Computing is a sub field of the holistic area of QIS that has recently been recognized by a White House Blog and a national Science and technology Committee Report
    - https://www.whitehouse.gov/blog/2016/07/26/realizing-potential-quantum-information-science-and-advancing-high-performance
- **HEP has held two joint meetings with ASCR on QIS:**
  - Grand Challenges at the Interface of QIS, Particle Physics, and Computing
  - Quantum Sensors at the Intersections of Fundamental Science, QIS, & Computing
  - Reports available at: http://science.energy.gov/hep/community-resources/reports/

HEP Quantum Information Science Connections

- **Topic A: HEP Quantum Computational Science**
  1. Quantum computational field theory development and quantum algorithms
  2. Quantum entanglement discovery science
  3. Quantum computational experiments
  4. Innovative computing & data tools for HEP (quantum/post-Moore test beds)
- **Topic B: Quantum Entangled Sensors**
  1. New tools for the dark universe (dark energy & dark matter)
  2. Small experiments for HEP discovery exploiting quantum entangled sensors
- **Topic C: Partnering with Other Communities to use HEP Facilities (Broader Impacts)**
  1. Atom interferometry/fundamental science experiments at HEP Facilities
  2. Experiments using HEP facility resources or providing lab partnership to experiments

- **Value of the HEP QIS Connections**
  - Promotes discovery science & develops cutting edge tools, techniques and technology
  - Works in partnership with other stakeholders reducing costs, gaining & sharing expertise
  - Contributes to National S&T priorities – Main Goal of Connections program
- **Anticipated Partnership Connections**
  - Within SC: ASCR and BES
  - External: Other agencies, Foundations, Industry, Small Businesses

Future of HEP Computing

- Particle physics community has recognized the importance of moving forward in computing through Snowmass, P5, and other community meetings
- Community at CHEP needs to cooperate to develop new computing systems that can be adopted by experiments
  - Not enough resources to rewrite software separately for everyone, therefore…
  - Must work across experiments and laboratories, NOT one-by-one!
- White House decision for Exascale has led to ASCR's work on hardware design
  - HEP community has opportunity to work with ASCR to ensure new HW is useful for future HEP needs
  - Optimized HW then becomes widely available and lower cost to HEP
  - Timescale for this change well-matched to HEP needs (early 2020s)
- New paradigms are on the horizon in computing that HEP should also pay close attention to:
  - Deep learning
  - Smart networking
  - NSCI → Neuromorphic, Quantum computing
- Directed computing investments in partnership with ASCR are needed to meet future needs
  - ASCR welcomes partnerships with HEP through CCE

U.S. DEPARTMENT OF
ENERGY

Office of
Science