HEP Center for Computational Excellence

Presented to JLAB Computing Round Table

Salman Habib for
Kerstin Kleese van Dam, Peter Nugent, and Rob Roser

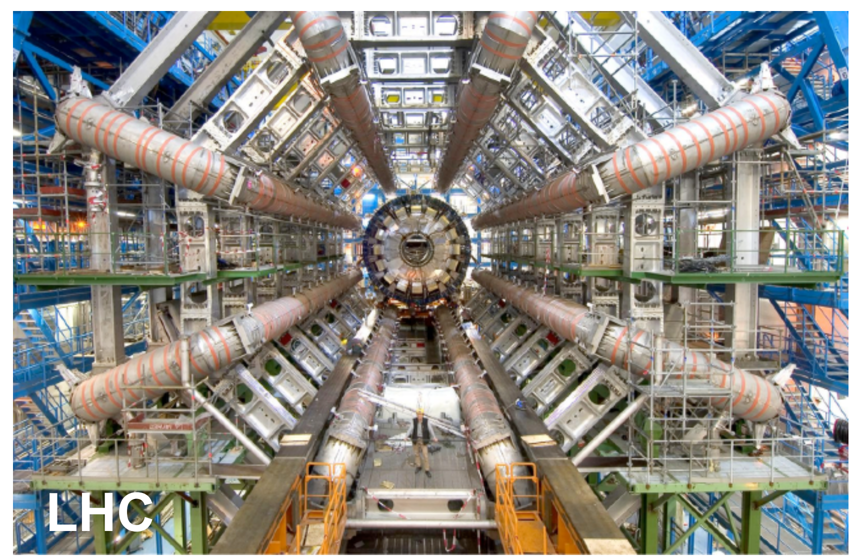

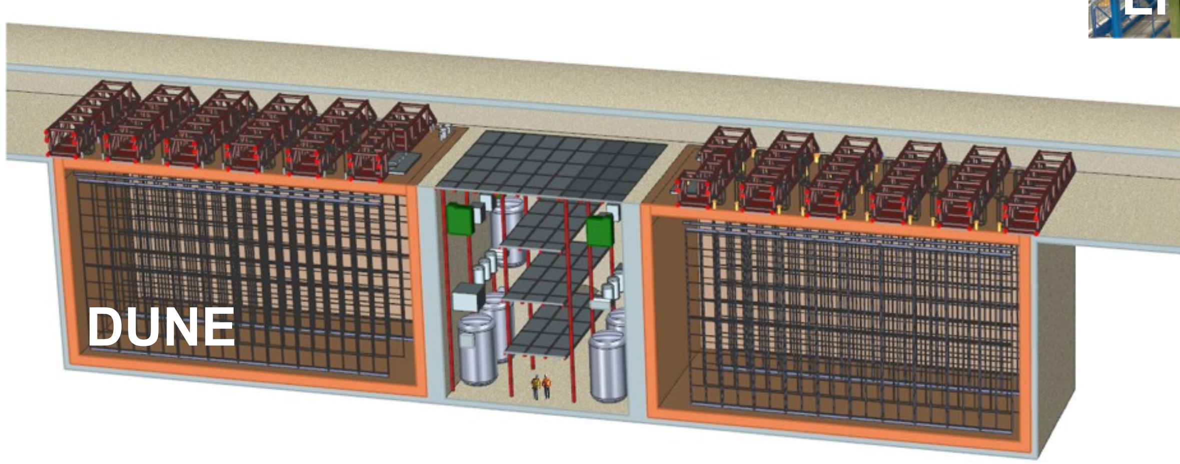

Argonne NATIONAL LABORATORY
Fermilab
BROOKHAVEN NATIONAL LABORATORY
BERKELEY LAB

HEP Computing Frontier

- Computing is an essential enabling and empowering component of almost all aspects of HEP science
- HEP computing has a long history (~60 years) including notable contributions to High Performance Computing (HPC), High Throughput Computing (HTC), and large-scale Data Science
- Substantial resources are devoted to computation and data science as an essential aspect of HEP's scientific enterprise
- Winds of Change: New challenges posed by hardware evolution and rapid increase in data rates/volumes in an era of flat/declining budgets
- Wide recognition that concerted response leveraging HEP community expertise and best practices from the outside world is needed (HEP Topical Panel, Snowmass, P5)

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
- Pursue the physics associated with neutrino mass
- Identify the new physics of dark matter
- Understand cosmic acceleration: dark energy and inflation
- 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*

Argonne NATIONAL LABORATORY · BROOKHAVEN NATIONAL LABORATORY · Fermilab · BERKELEY LAB

From the P5 Report —

- **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 laboratories and universities across the research areas will be needed”
- **P5 recommnedation 29:**
- Strengthen the *global collaboration* 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

What is the HEP–CCE?

- Official DOE HEP response to P5 recommendations, Snowmass, and the HEP Topical Panel on Computing and Simulation (Dec 2013): “Computing Frontier” as key strength of HEP
- Cross-cut exchange to promote excellence in HEP computing, simulation, data transfer/storage/preservation and promote the flow of information and best practices across HEP Frontiers
- Identify and undertake focused R&D tasks addressing frontier challenges posed by rapidly evolving computer architectures, next-generation data-intensive computing, and development of shared software elements
- Undertake information gathering tasks – community-based studies and surveys
- Access point for agency/office partnerships (ASCR, NSF, NASA…) and HEP connections (CERN, DPHEP, G4, HSF, OSG, USQCD...)

Why NOW and WHY HEP-CCE?

Given flat/declining budgets, it is increasingly important to leverage work that has already been done and to use expertise that already exists

Present computing model in HEP handles virtually all computing within experiments/frontiers

- Need to strengthen and expand horizontal component to best utilize HEP resources, **avoid duplication,** optimize technology, **develop external resources & strengthen partnerships**……

Rapidly evolving computer architectures and exponentially increasing data volumes REQUIRE effective responses & cross-cut solutions

- Evolution of commodity hardware and rise of cloud computing
- Enhance opportunities of interfacing with ASCR
- Understanding HEP's core competencies and leveraging that expertise moving forward

Dark Matter

Higgs Boson For Discovery

Pursue physics associated with neutrino mass

Dark Energy and Inflation

Explore the Unknown (new particles and fields)

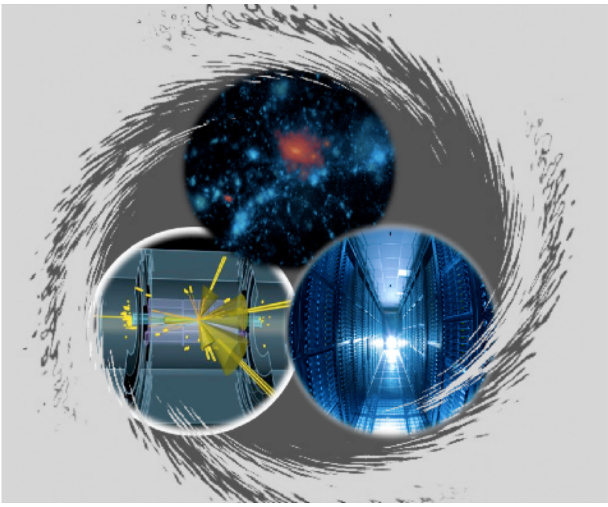

Argonne NATIONAL LABORATORY

BROOKHAVEN NATIONAL LABORATORY

Fermilab

BERKELEY LAB

HEP Center for Computational Excellence Summary

- **Primary Mission**
  - Bring next-generation computational resources to bear on pressing HEP science problems
  - Develop cross-cutting solutions leveraging ASCR expertise and resources
- **Technical Challenges**
  - Hardware and software evolution
  - New algorithms for fine-grained data analysis and I/O on HPC systems
  - HEP workflow management for the exascale ecosystem
- **Engagement Examples**
  - Software management for HPC systems (containers)
  - Edge services for HPC systems
  - Petascale data transfer project with ESnet
  - Distributed large-scale data analytics

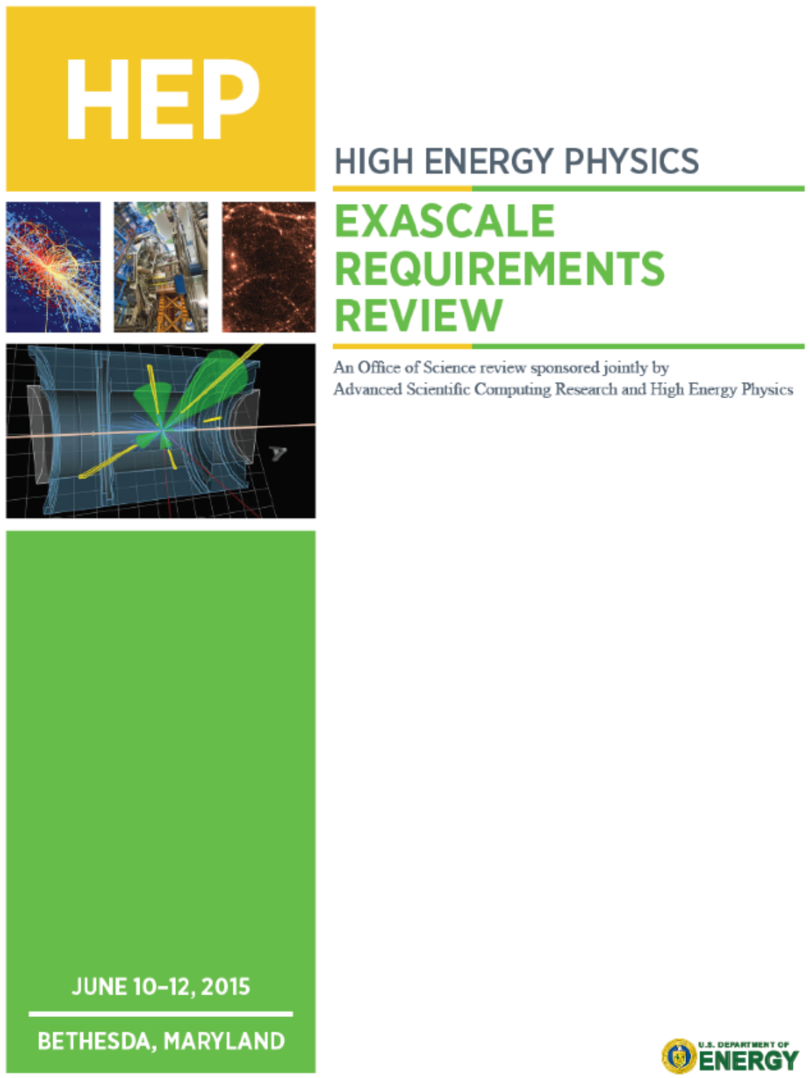

Joint ASCR/HEP Exascale Requirements Review co-organized by HEP-CCE

Argonne NATIONAL LABORATORY | BROOKHAVEN NATIONAL LABORATORY | Fermilab | BERKELEY LAB

Computational HEP and Overall HEP Computing

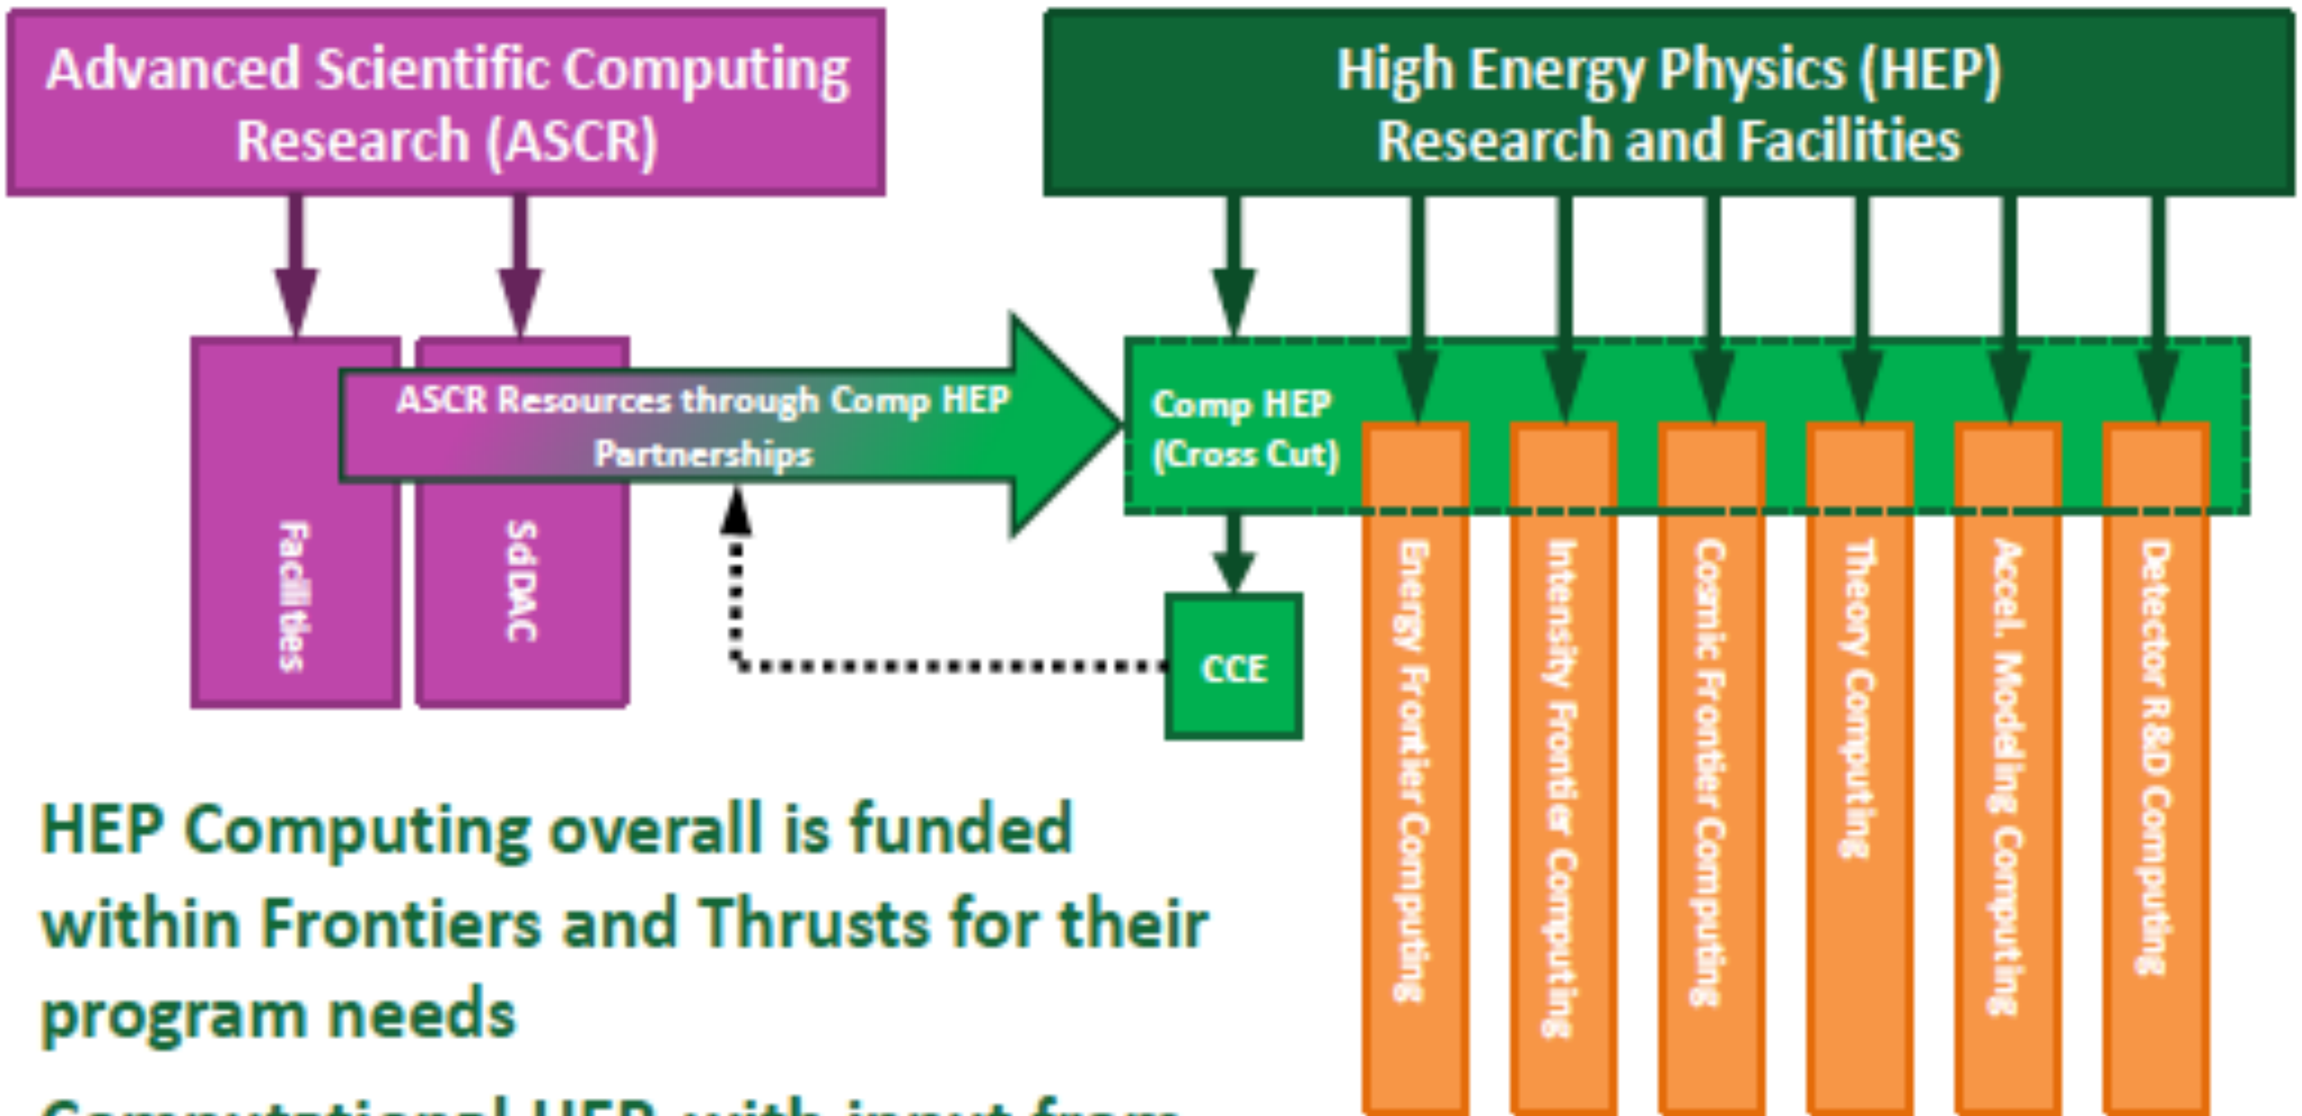

- HEP Computing overall is funded within Frontiers and Thrusts for their program needs
- Computational HEP, with input from CCE, identifies where external partnerships & cross cuts are possible and fosters them

Some CCE Highlights

- Led an ASCR/HEP exascale computing review
- Published our 3 working group reports **arXiv:1510.08545**
- Led the HEP portion of an ASCR workshop entitled Data Management, Visualization, and Analysis of Experimental and Observational Data (EOD) Workshop
- Established a web presence
  - **http://hepcce.org/**
- Coordinated submission of 5 white papers for ASCR Exascale application call
- Executed on a number of target areas
  - Mini-apps
  - Data Transfer Project
  - Software containers (and related work)
  - Summer Student Program
  - Enabling Commercial Clouds to handle Peak Loads

“Production” Example: Large-Scale Data Movement

- **Offline Data Flows:** Cosmological simulation data flows already require ~PB/week capability, next-generation streaming data will require similar bandwidth
- **ESnet Project:** Aim to achieve a production capability of 1 PB/week (FS to FS, also HPSS to HPSS) across major compute sites
- **Status:** Success achieved! numbers from a simulation dataset “transfer test package” (4 TB)
- **Future:** Automate entire process within the data workflow including retrieval from archival storage (HPSS); add more compute/data hubs (BNL underway, just solved Globus-dCache handshake problem)

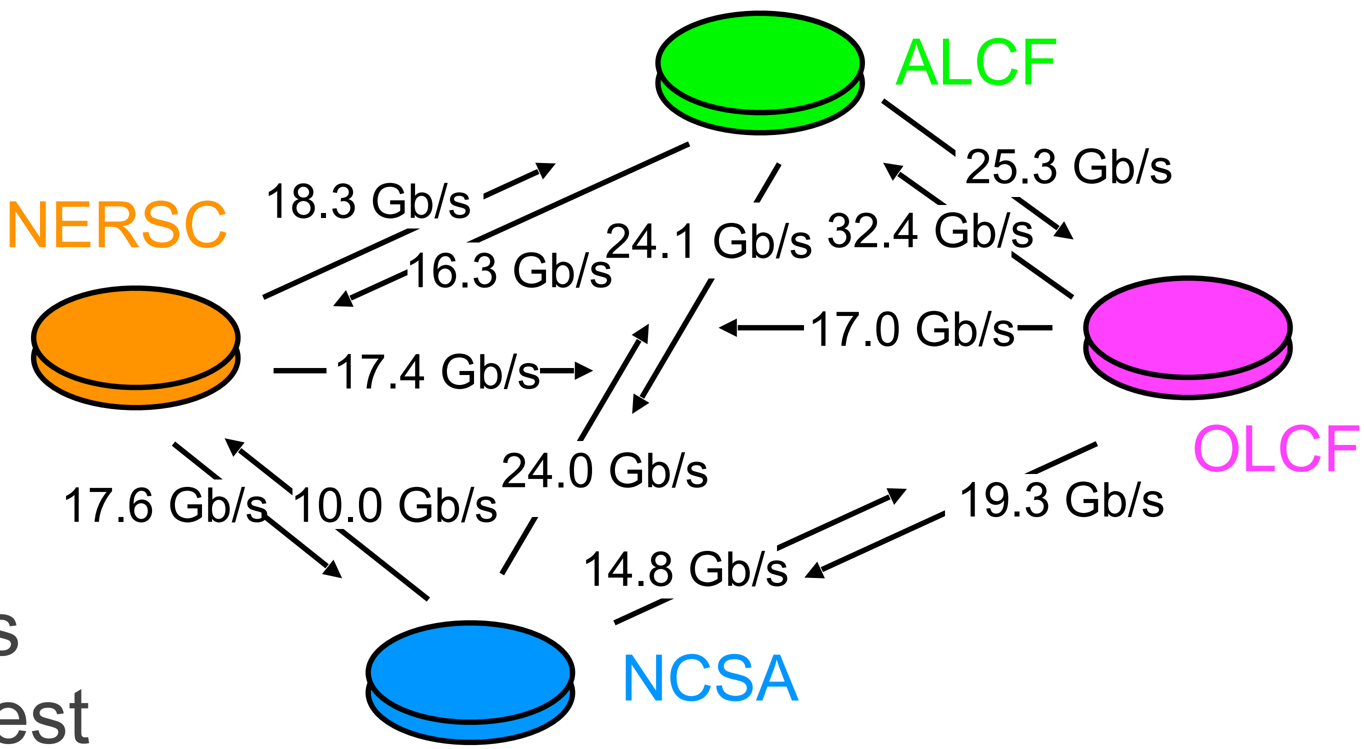

Petascale DTN project, courtesy Eli Dart, HEP-CCE/ESnet supported joint project

globus online

ESnet
ENERGY SCIENCES NETWORK

HEP-CCE

The Realities of HEP Computing Moving Forward

- Due to funding constraints, we will have to optimize computing resources for “average” demand and not peak
- Need to find creative solutions for those instances where we need more than we have and need it fast (“pledged” vs. “non-pledged” resources)
- Importance of leveraging HPC facilities – already demonstrated at level of several percent of ATLAS production computing, equivalent to all of Spain’s contribution (LeCompte et al.)
- Computing in the Cloud will be an important player in future
- Helping each other across Frontiers/projects – sharing core competencies essential
- Finding ways to not just develop key software base – but to maintain and evolve it going forward

Projected Shortfalls in HEP Computing Resources

Expectation of 10 X — 100 X shortfalls in compute by 2025 including:

- **Storage and Data movement** – need smart networks, optimization of compute, data movement, and storage
- **Hardware** for simulation, data analysis, and storage
- **Workforce**, highlighting need for expertise and training

Entire computing ecosystem critical to workflows and results; rapid hardware and software evolution is a critical concern

Can these challenges be handled entirely within HEP resources and programs?

Need a shift in strategy to best prepare for the future while managing current operations and using resources external to HEP

Argonne NATIONAL LABORATORY BROOKHAVEN NATIONAL LABORATORY Fermilab BERKELEY LAB

What to Do? Many White Papers and Reports —

HEP

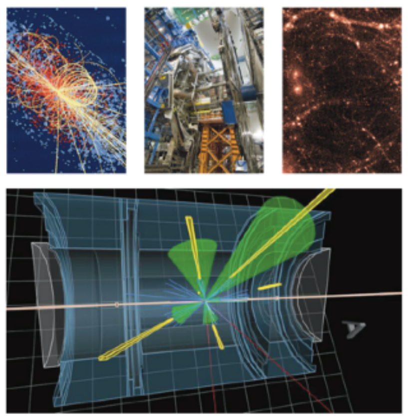

HIGH ENERGY PHYSICS

EXASCALE REQUIREMENTS REVIEW

An Office of Science review sponsored jointly by Advanced Scientific Computing Research and High Energy Physics

Lead Authors

HEP

Salman Habib[1] and Robert Roser[2]

ASCR

Richard Gerber,[3] Katie Antypas,[3] Katherine Riley,[1] and Tjerk Straatsma[4]

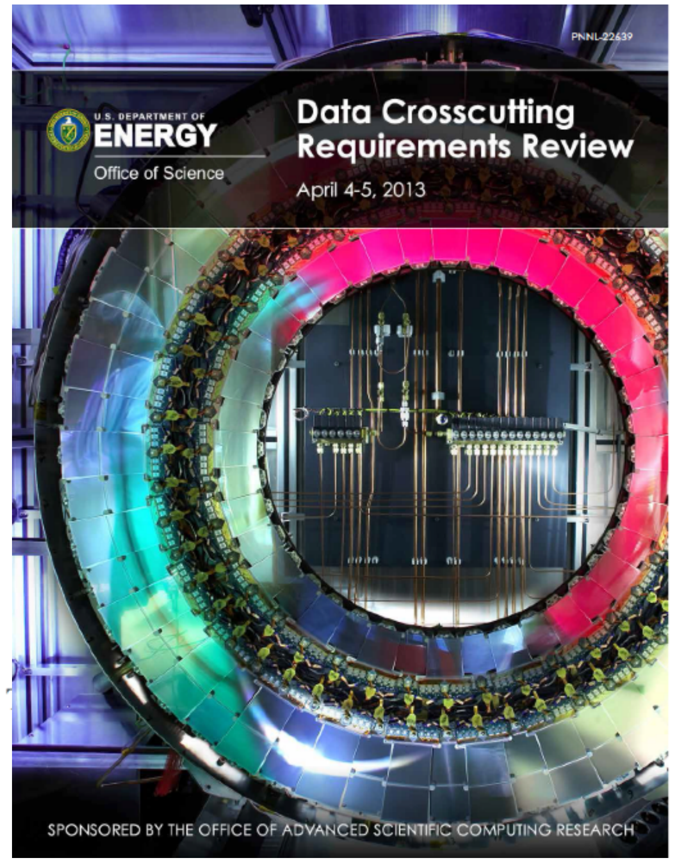

HIGH ENERGY PHYSICS FORUM FOR COMPUTATIONAL EXCELLENCE: WORKING GROUP REPORTS

I. APPLICATIONS SOFTWARE
II. SOFTWARE LIBRARIES AND TOOLS
III. SYSTEMS

Lead Editors: Salman Habib[1] and Robert Roser[2] (HEP-FCE Co-Directors)

Applications Software Leads: Tom LeCompte[1], Zach Marshall[3]
Software Libraries and Tools Leads: Anders Borgland[4], Brett Viren[5]
Systems Lead: Peter Nugent[3]

Applications Software Team:
Makoto Asai[4], Lothar Bauerdick[2], Hal Finkel[1], Steve Gottlieb[6], Stefan Hoeche[4], Tom LeCompte[1], Zach Marshall[3], Paul Sheldon[7], Jean-Luc Vay[3]

Software Libraries and Tools Team:
Anders Borgland[4], Peter Elmer[8], Michael Kirby[2], Simon Patton[3], Maxim Potekhin[3], Brett Viren[3], Brian Yanny[2]

Systems Team:
Paolo Calafiura[3], Eli Dart[3], Oliver Gutsche[2], Taku Izubuchi[5], Adam Lyon[2], Peter Nugent[3], Don Petravick[9]

Report from the Topical Panel Meeting on Computing and Simulations in High Energy Physics

Sponsored by the U.S. Department of Energy, Office of Science, High Energy Physics
December 9-11, 2013 Rockville Hilton Hotel, Rockvil

Steering Committee

Paul Avery (co-Chair)	U Florida
Salman Habib (co-Chair)	Argonne
Amber Boehnlein	SLAC
Robert Roser	Fermilab
Stephen Sharpe	U Washington
Heidi Schellman	Northwestern
Craig Tull	LBNL
Torre Wenaus	BNL

Planning the Future of U.S. Particle Physics

Report of the 2013 Community Summer Study

L. A. T. Bauerdick, S. Gottlieb, G. Bell, K. Bloom, T. Blum, D. Brown, M. Butle E. Cormier, P. Elmer, M. Ernst, I. Fisk, G. Fuller, R. Gerber, S. Habib, M. Hildreth, S. Hoeche C. Joshi, A. Mezzacappa, R. Mount, R. Pordes, R. Rebel, L. Reina, M. C. Sanchez, J. Shank, A. Szalay, R. Van de Water, M. Wobisch, S. Wolbers

Chapter 9: Computing

High Energy Physics and Nuclear Physics Network Requirements

HEP and NP Network Requirements Review
Final Report

Conducted August 20-22, 2013

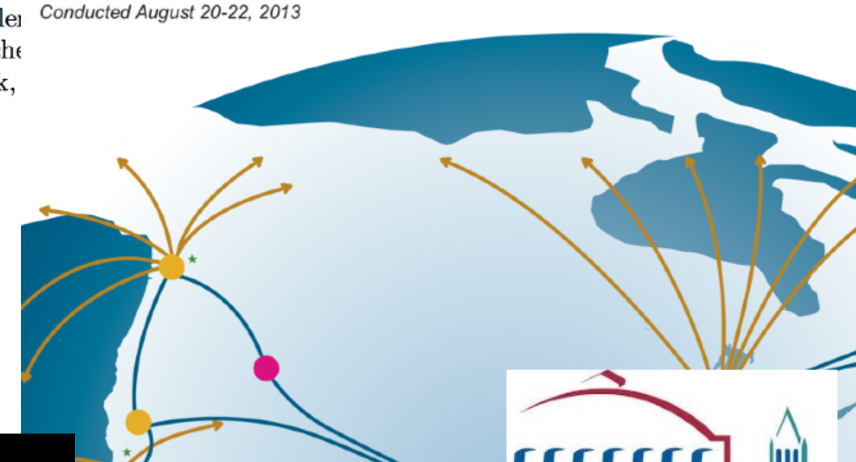

Argonne NATIONAL LABORATORY

BROOKHAVEN NATIONAL LABORATORY

Fermilab

BERKELEY LAB

Connection to Exascale Computing Project

Major DOE SC and NNSA joint project to arrive at a scientifically usable architecture for exascale computing in the early 2020's — *largest science project within DOE*

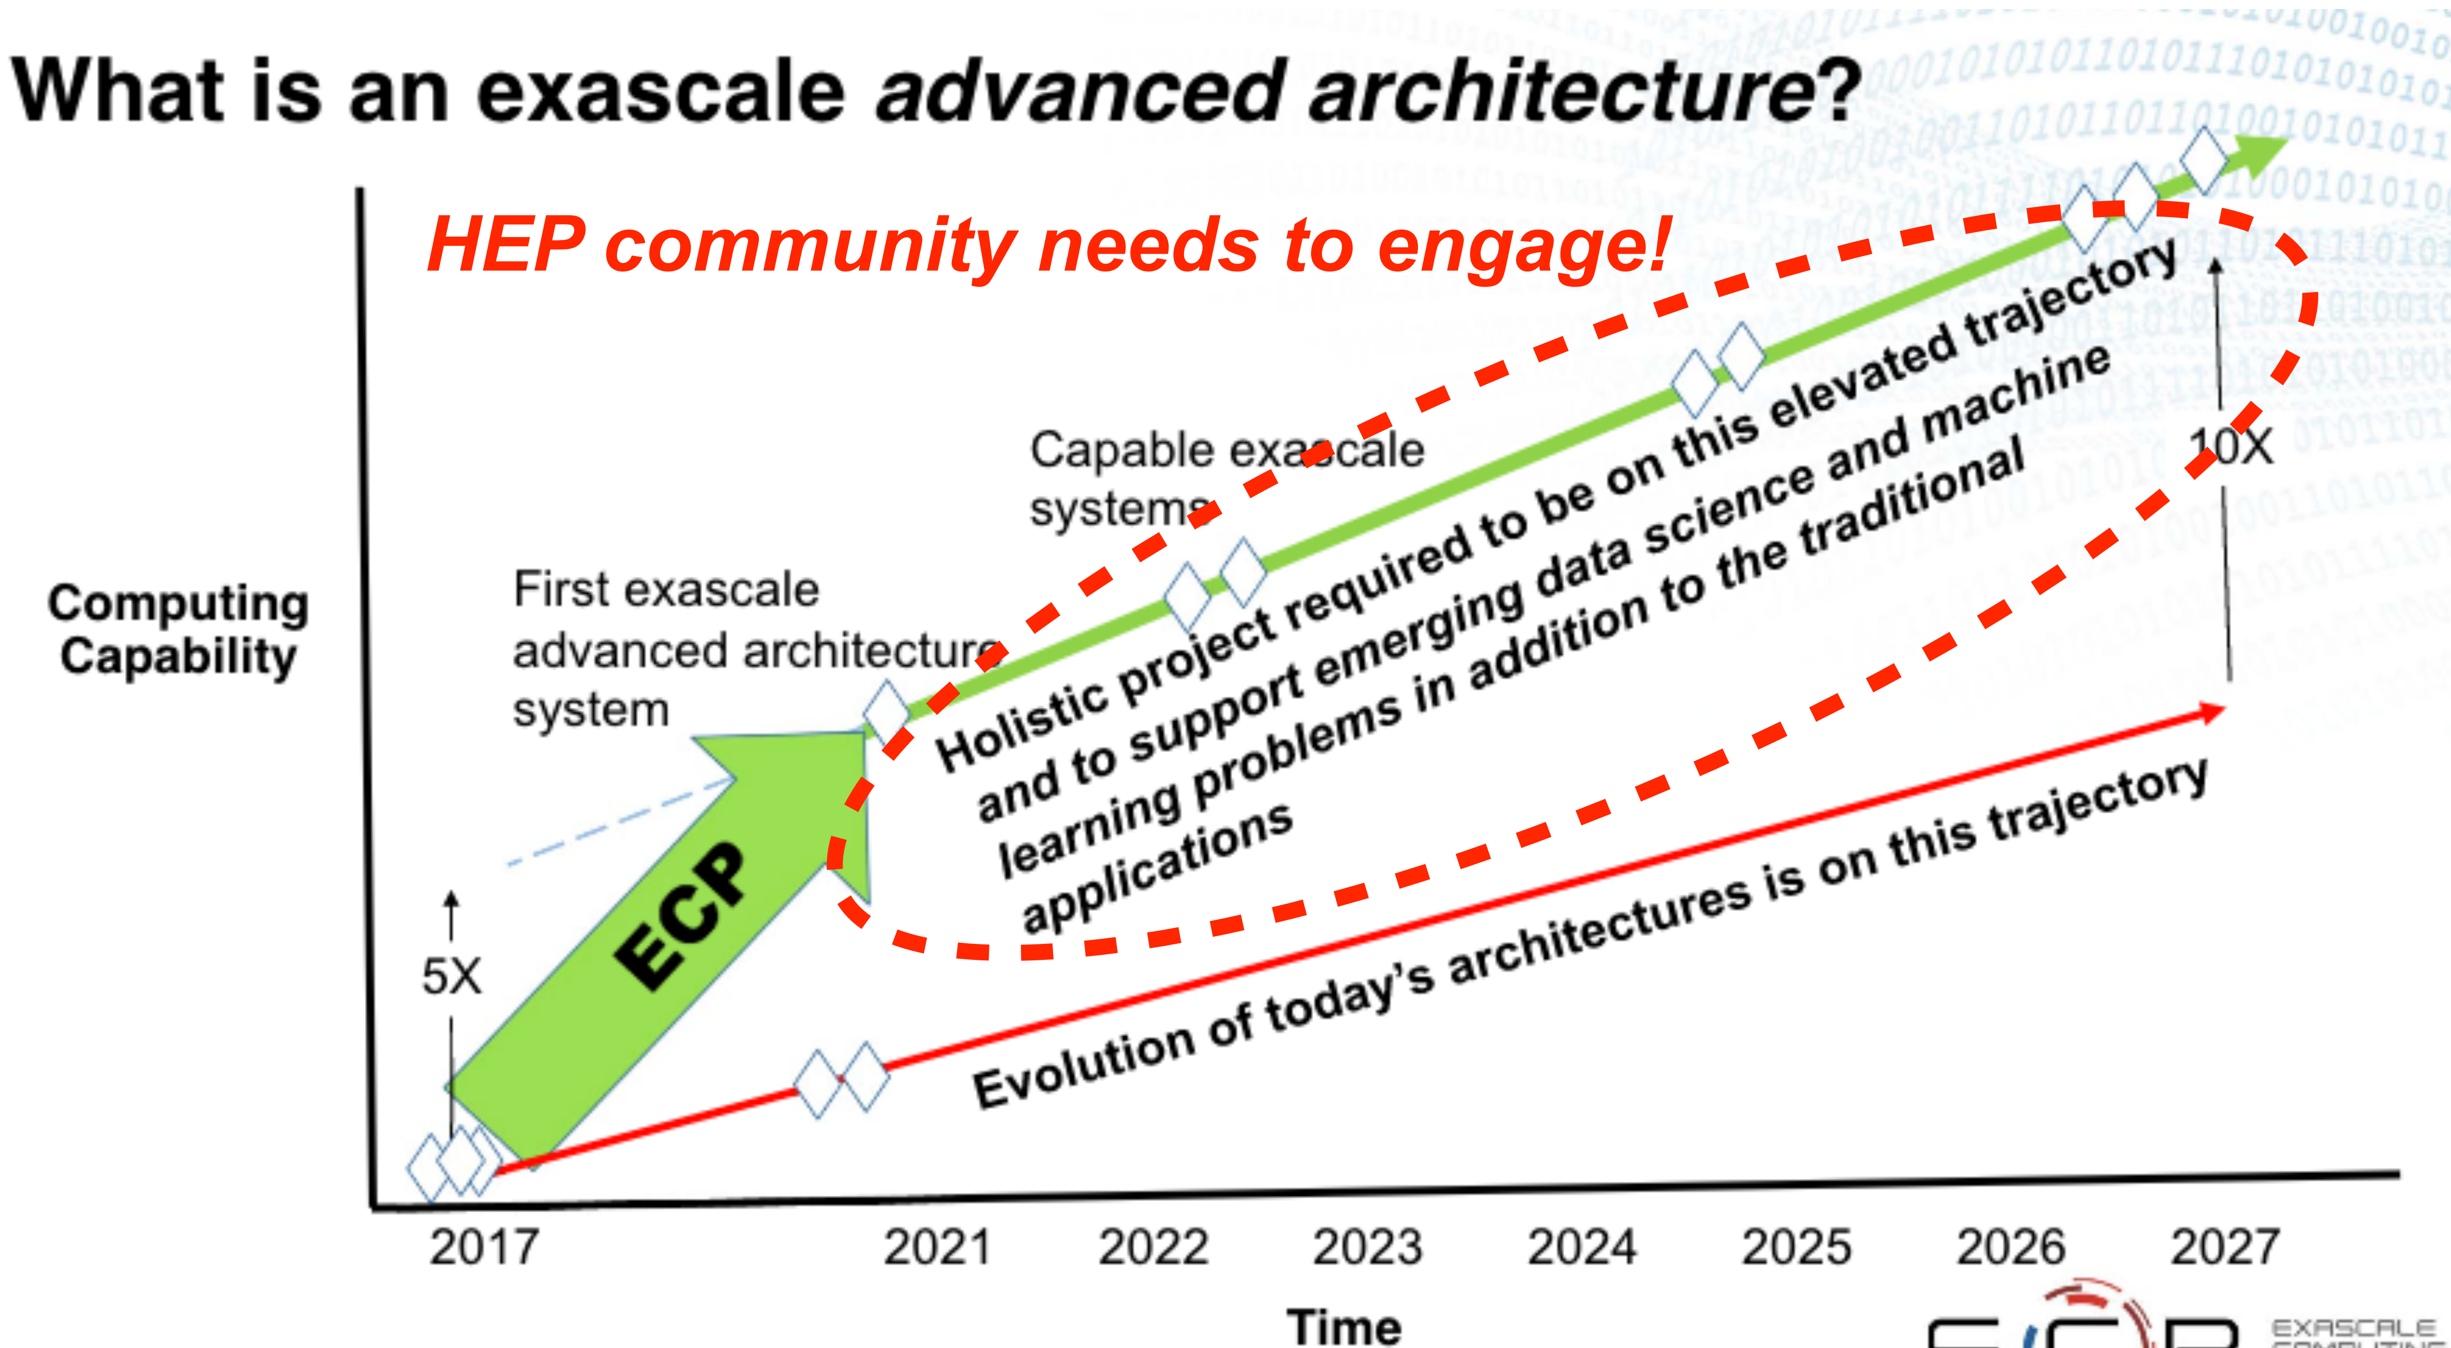

Computing Requirements – Energy Frontier

- **HEP Requirements in computing/storage will scale up by ~50X over 5-10 years**
  - Flat funding scenario fails — must look for alternatives!

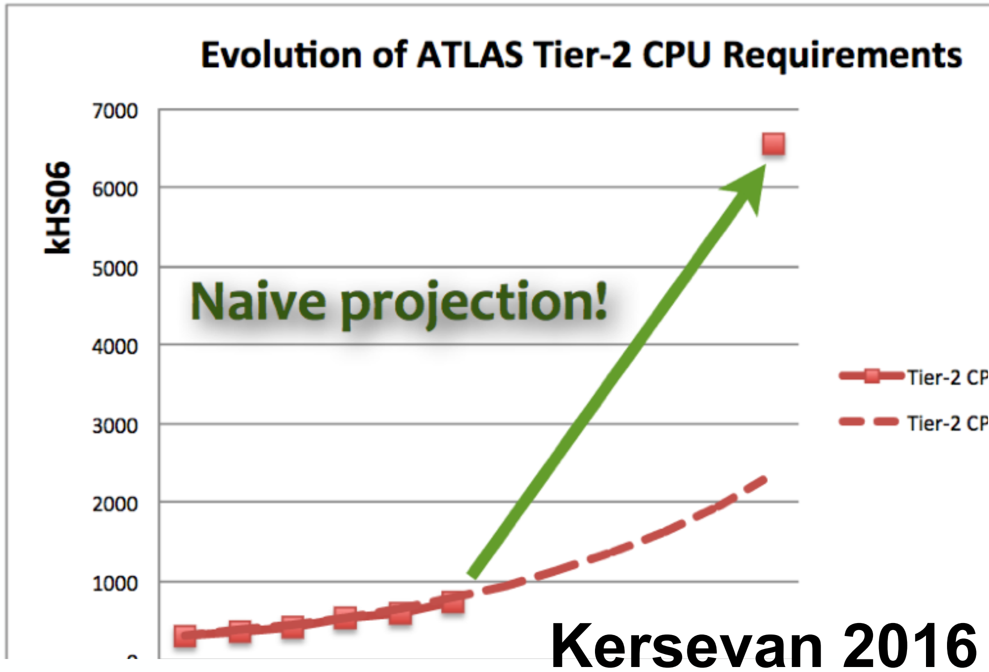

Kersevan 2016

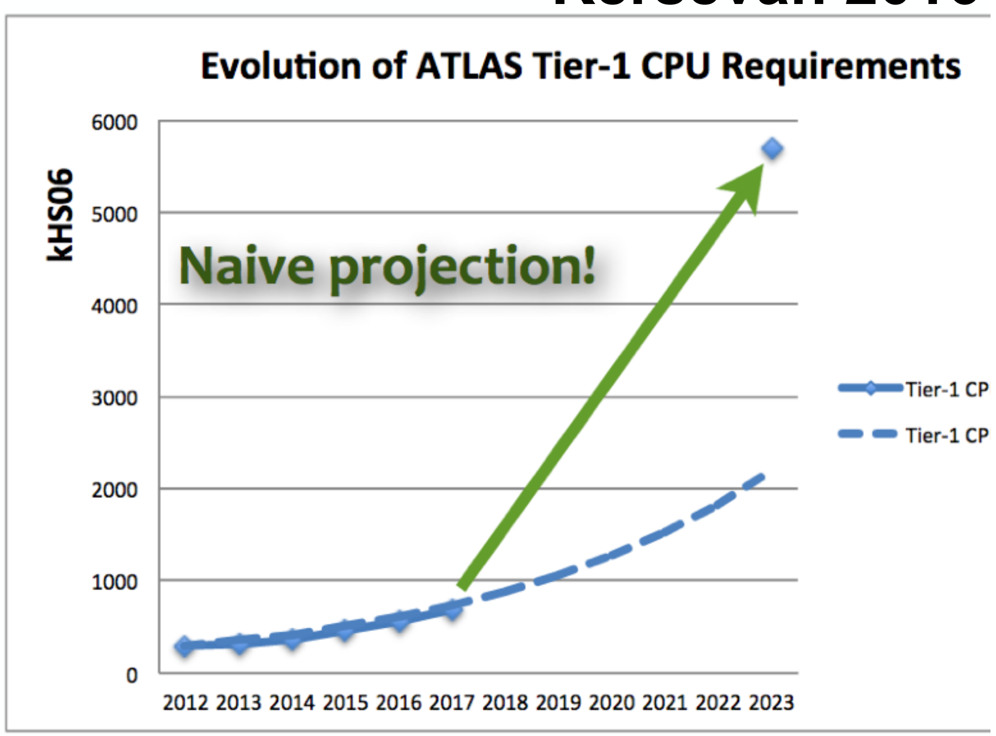

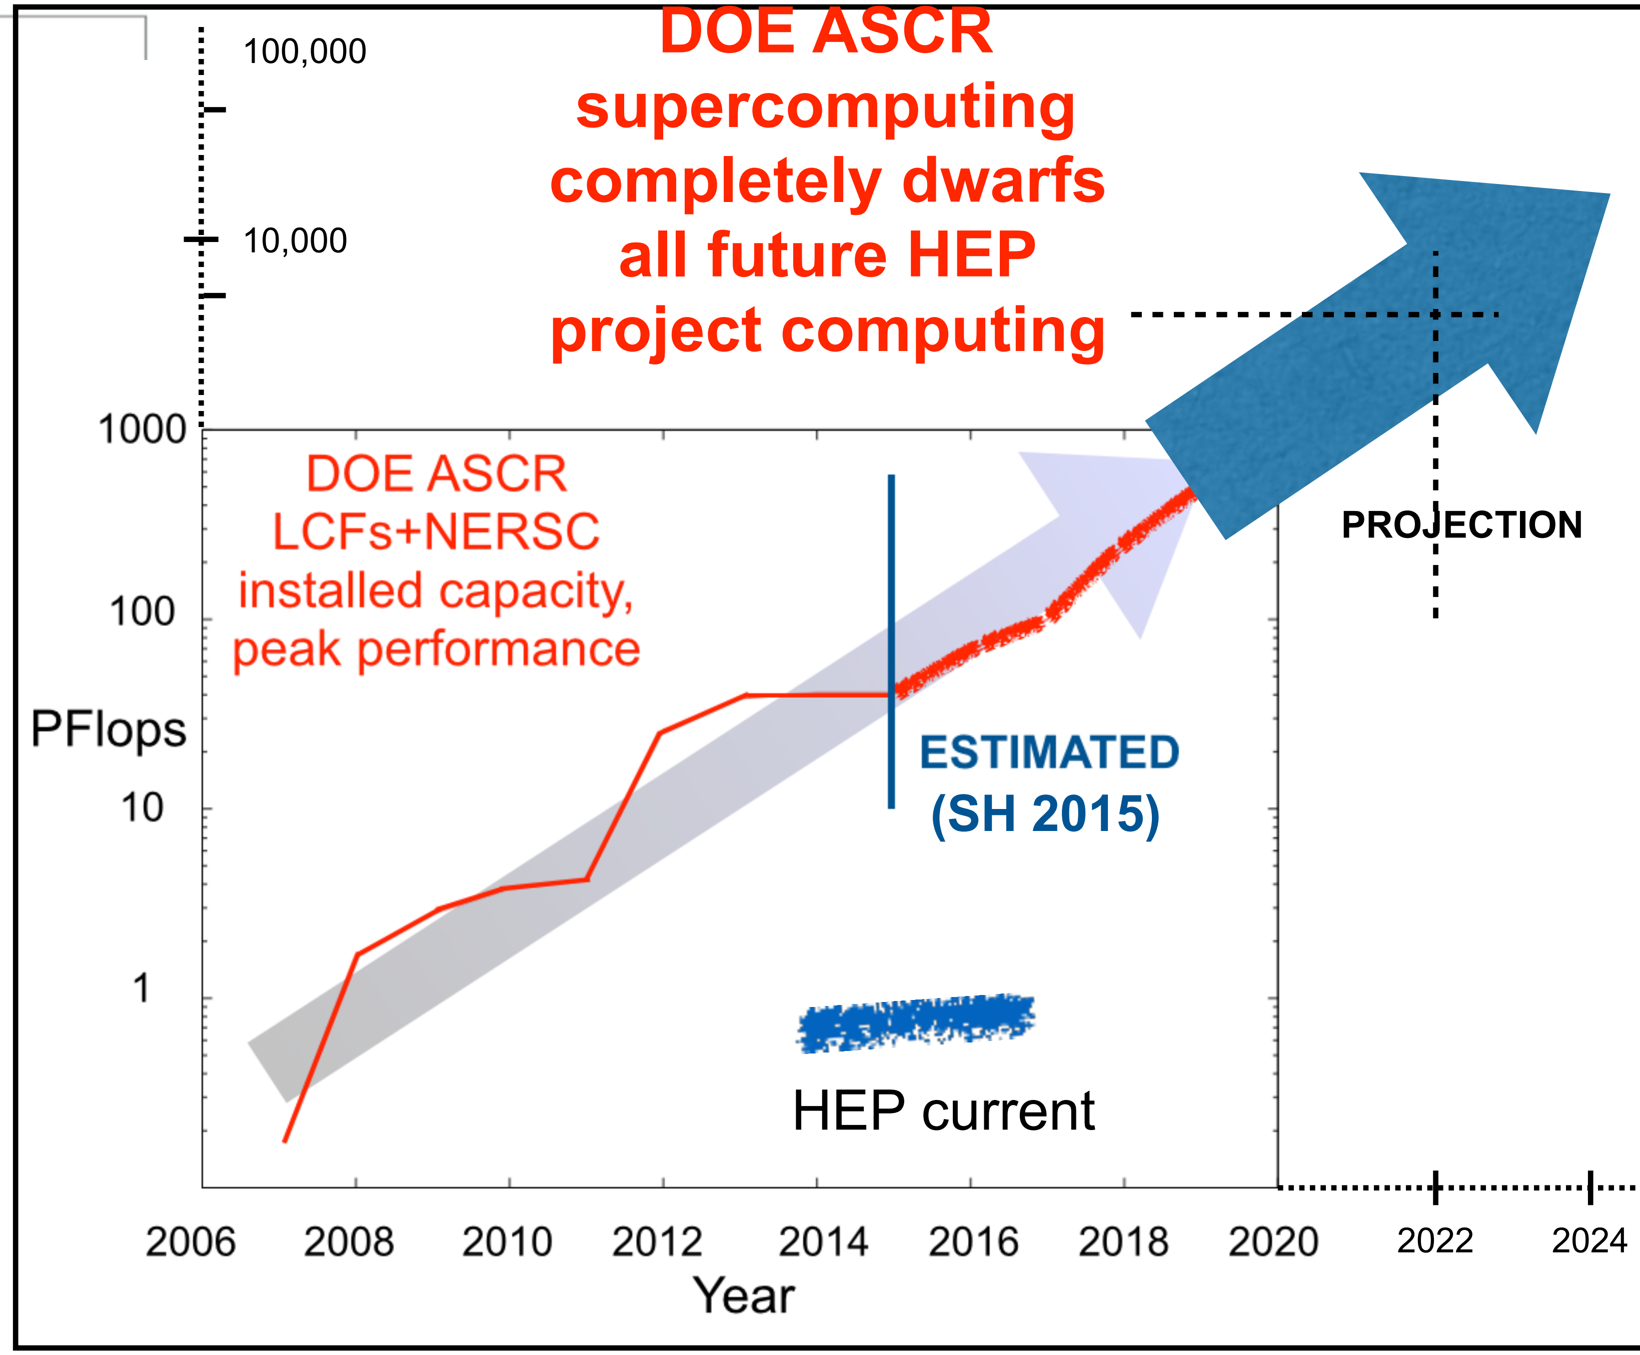

The Impact of HPC on HEP

- Not all problems can be solved using HPC systems, but many can (accelerators, cosmology, event generation/simulation, QCD…)
- Next generation of ASCR HPC machines (staging begins 2016/17, ends in 2018/19) will sum to ~200 petaflops of compute capability
- If HEP experiments use just 10% of that, i.e. 20 petaflops, it is ~20 times what the Grid will provide
- Learning how to leverage these resources to seamlessly supplement/ enhance current capability is important
  - Currently, dealing with HPC systems is painful! Not for the faint of heart!
- New possibilities opened up by HPC platforms will offer unique computational opportunities — but we need to invest to harness them – doesn't come automatically
- Major task is to use HPC resources to act on incoming data, not just on data created in situ; this will require investment in edge services

Argonne NATIONAL LABORATORY BROOKHAVEN NATIONAL LABORATORY Fermilab BERKELEY LAB

Connectivity Example: Edge Services

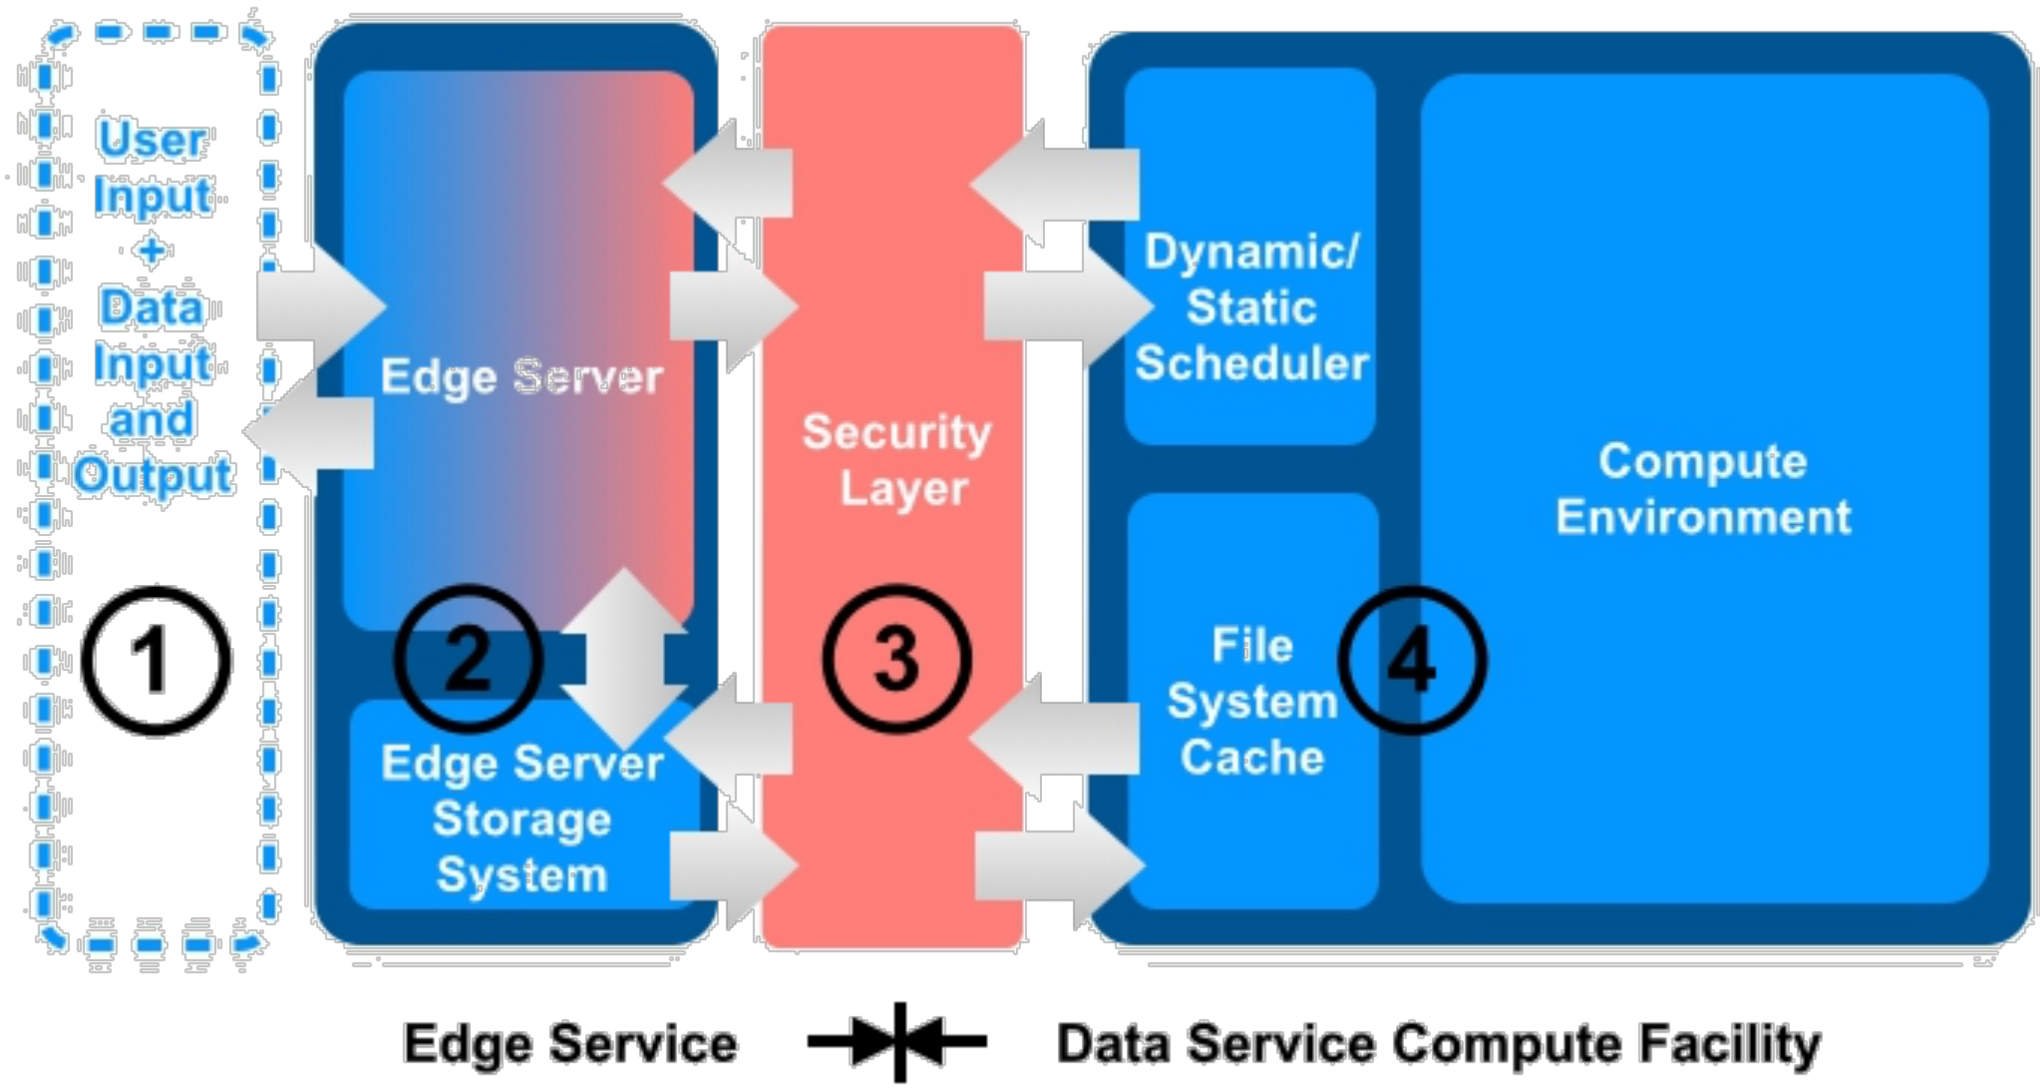

Edge service design must consider a number of factors; security, resource flexibility, interaction with HPC schedulers, external databases, requirements of the user community — modern supercomputers are once again ‘strategic’ resources, not a ‘pile of PCs’!

Argonne NATIONAL LABORATORY BROOKHAVEN NATIONAL LABORATORY Fermilab BERKELEY LAB

Boundary Conditions

- **What's the Problem?**
  - Even if solutions can be designed ***in principle***, the resources needed to implement them are (usually) not available
  - ***Despite all the evidence of its power***, computing still does not get high enough priority compared to building "things"
  - In part this is due to the success of computing — progress in this area is usually much faster than in others, so one can assume that ***computing will just happen (Moore's Law)*** — to what extent is this still true?
- **Large-Scale Computing Available to Scientists**
  - Lots of supercomputing (HPC) available and more on the way
  - Not enough data-intensive scalable computing (DISC) available to users, hopefully this will change over time
  - Publicly funded HTC/Grid computing resources cannot keep pace with demand
  - Commercial space (Cloud) may be a viable option but is not issue-free
  - Storage, networking, and curation are major problems (***sustainability***)

“Data Meets HPC” — Basic Requirements

- **Software Stack:** Ability to run arbitrarily complex software stacks on HPC systems (***software management***)
- **Resilience:** Ability to handle failures of job streams, still rudimentary on HPC systems (***resilience***)
- **Resource Flexibility:** Ability to run complex workflows with changing computational ‘width’, possible but very clunky (***elasticity***)
- **Wide-Area Data Awareness:** Ability to seamlessly move computing to the data (and vice versa where possible); access to remote databases and data consistency via well-designed and secure edge services (***integration***)
- **Automated Workloads:** Ability to run large-scale coordinated automated production workflows including large-scale data motion (***global workflow management***)
- **End-to-End Simulation-Based Analyses:** Ability to run analysis workflows on simulations using a combination of in situ and offline/co-scheduling approaches (***hybrid applications***)

Summary

- As we advance into the next decade, HEP cannot take a "business as usual" approach to computing
  - Combination of upgraded facilities and finer grained detectors will push the envelope of "big data"
- Is HPC the answer?
  - Not clear yet – but it will certainly be part of the equation
- Our Compute and Data model will have to evolve
  - Faster and smarter networks are a given
  - Relationship of compute to data is not obvious
  - Clouds will likely play a role in the future
- HEP compute will look more like HPC centers though perhaps architected more for data processing
- HEP needs ASCR to play a significant role in helping us define a path forward and executing it