U.S. DEPARTMENT OF
ENERGY

Big Data and Scientific Discovery

Bill Harrod
Office of Science
William.Harrod@science.doe.gov

February 26, 2014

Big Data and Scientific Discovery

- **Next generation scientific breakthroughs require:**
  - **Major new advances in computing technology:** energy-efficient hardware, algorithms, applications, system software
  - **Dealing with "big data":** advances in data management, data analytics, visualization and machine learning algorithms
- **Ubiquitous data explosion**
  - Experimental, observational, sensor networks, and simulation

- **Computing / data throughput challenges**
  - Can no longer scale up from or modify existing solutions
  - Problem compounded by need for international collaboration/coordination
- **Complicated workflows**

Extreme Scale Science

Next Generation of Scientific Discovery

Genomics

Data Volume increases to 10 PB in FY21

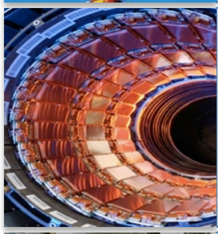

High Energy Physics (Large Hadron Collider)

15 PB of data/year

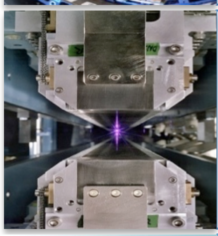

Light Sources

Approximately 300 TB/day

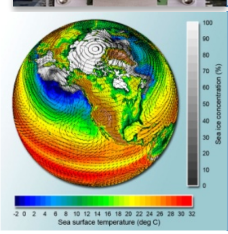

Climate

Data expected to be 100 EB

Data explosion is driven by exponential technology advances

Data sources

- Scientific Instruments
- Scientific Computing Facilities
- Simulation Results
- Observational data

Big Data and Big Compute

- Analyzing Big Data requires processing (e.g., search, transform, analyze, …)
- Extreme scale computing will enable timely and more complex processing of increasingly large data sets

"Very few large scale applications of practical importance are NOT data intensive." – Alok Choudhary, IESP, Kobe, Japan, April 2012

DOE Office of Science Computational Facilities

Titan System Specifications:
- Peak performance of 27.1 Petaflops
  - 24.5 GPU + 2.6 CPU
- 18,688 Hybrid Compute Nodes with:
  - 16-Core AMD Opteron CPU
  - NVIDIA Tesla “K20x” GPU
  - 32 + 6 GB memory
- 200 Cabinets; 710 TB total system memory; 8.9 MW peak power

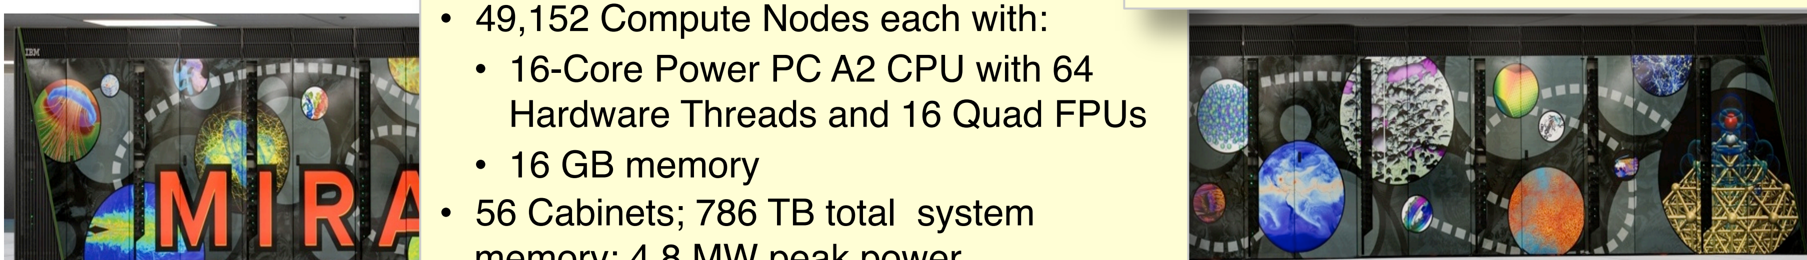

Mira System Specifications:
- Peak performance of 10 Petaflops
- 49,152 Compute Nodes each with:
  - 16-Core Power PC A2 CPU with 64 Hardware Threads and 16 Quad FPUs
  - 16 GB memory
- 56 Cabinets; 786 TB total system memory; 4.8 MW peak power

NERSC System Specifications:
- Hopper XT5 (2010)
  - 1.3PF, 212TB, 2.9 MW peak power
- Edison XC30
  - Based on DARPA/DOE HPCS system
  - 2.4PF, 333TB, 2.1 MW peak power

Energy Science Network (ESnet)
100 Gigabit Upgrade: Project Complete

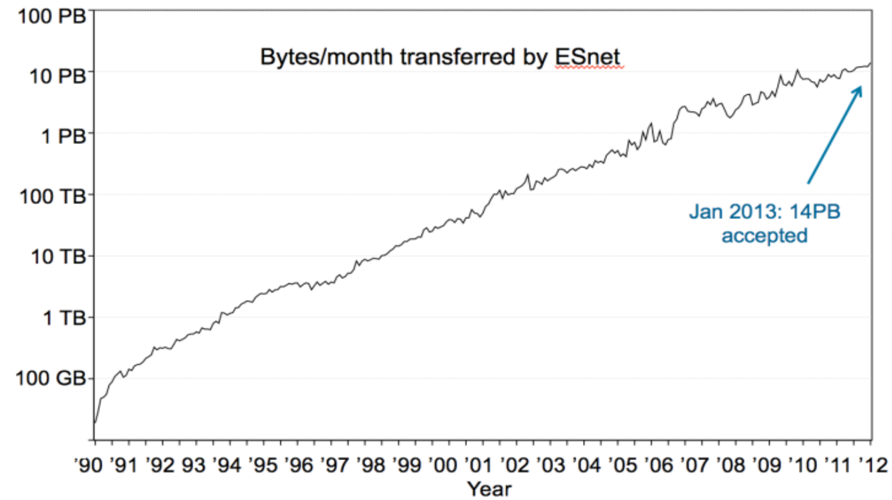

Network Specifications:
- Connects 40 DOE sites to internet
- Growing 2x commercial nets – 50% of traffic is from "big data"
- First 100G cross USA
- Network can grow to 88 independent 100G channels
- Deploying 100G production connections at ANL, BNL, FNAL, LBNL, LLNL, NERSC, ORNL

One Possible View of the Future

Individual University and Lab PIs | National and Int'l collabs | Research + Industry | Industry

Scientific Discovery Engines

Common Services

ESnet

Exascale Simulation Facilities | Massive Throughput Simulations | Data Serving and Archiving Facilities | Analytics & Visualization Systems

ESnet

Local Nodes	Local Nodes	Local Nodes	Local Nodes	Local Nodes
Light sources, etc. (BES)	Colliders (HEP/NP)	Sequencers (BER)	Cosmology (HEP)	"Omics" data (All)

Scientific Discovery Environments

- **Current environment not optimal for:**
  - Geographically distributed data, users, and facilities
  - Dynamic, interactive workflow
  - Computational steering and real-time in-situ analysis

- **Future need for shared, real-time visualization, comparison of simulations with expt'l results, and control of active processes**

- **International collaboration can push the boundaries of scientific discovery by creating unified international teams from distributed users and facilities**

Collaboration is inherently a "big data" issue

Uncertainty Threatens Future of Computing

- **The world has changed – technology is changing at a dramatic rate**
  - Dennard scaling has ended
  - End of Moore's Law looming
- **The IT marketplace is also changing dramatically**
  - PC sales have flattened
  - Handhelds dominate growth, H/W and S/W
  - HPC vendor uncertainty
- **Need to drive innovations at all levels of technology**
  - Processors & Memory (Node)
  - System Designs
  - System Software
  - Algorithms
- **Data volume and variety explosion**
  - Data management
  - Analytics / visualization / machine learning

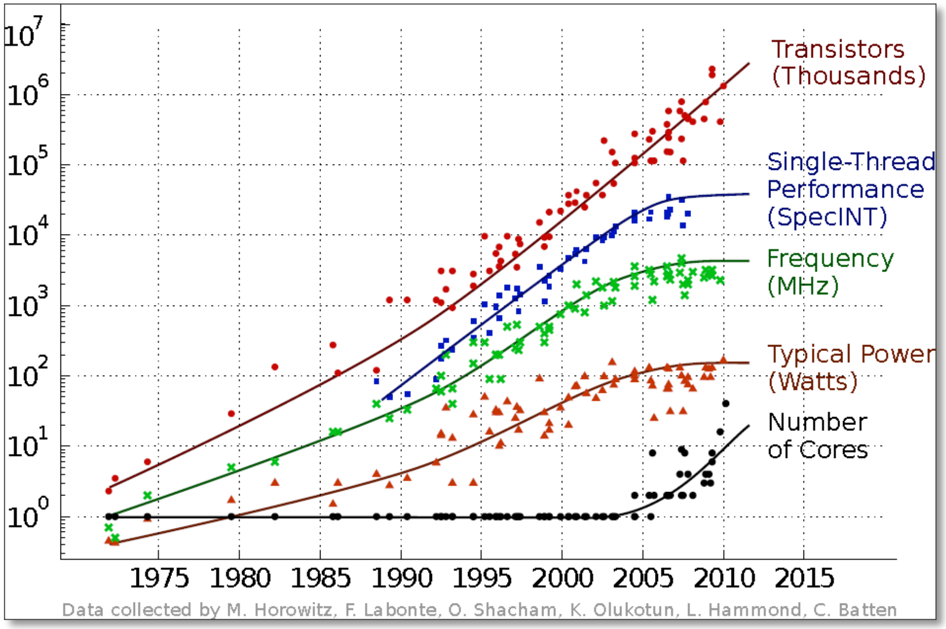

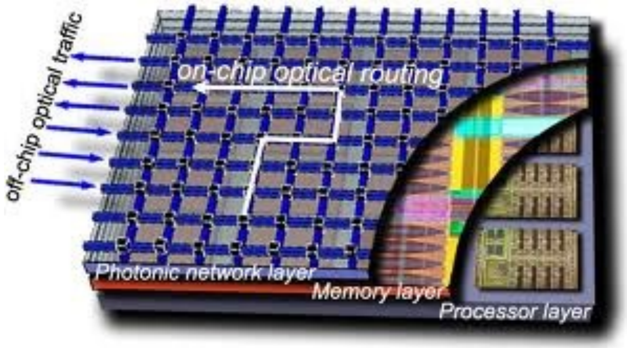

Exascale Computing
The Vision

- **Exascale computing**
  - Achieve order 10^{18} operations per second and order 10^{18} bytes of storage
  - Address the next generation of scientific, engineering, and large-data workflows
  - Enable extreme scale computing: 1,000X capabilities of today's computers with a similar size and power footprint
  - Establish a new trajectory of progress – towards a broad spectrum of computing capabilities over the succeeding decade
- **Productive system**
  - Usable by a wide variety of scientists and engineers
  - "Easier" to develop software & management of the system
- **Based on marketable technology**
  - Not a "one off" system
  - Scalable, sustainable technology, exploiting economies of scale and trickle-bounce effect
- **Deployed in early 2020s**

Advanced Scientific Computing Advisory Committee

Top Ten Technical Approaches for Exascale

1. **Energy efficiency:** Creating more energy efficient circuit, power, and cooling technologies.
2. **Interconnect technology:** Increasing the performance and energy efficiency of data movement.
3. **Memory technology:** Integrating advanced memory technologies to improve both capacity and bandwidth.
4. **Scalable System Software:** Developing scalable system software that is power and resilience aware.
5. **Programming systems:** Inventing new programming environments that express massive parallelism, data locality, and resilience
6. **Data management:** Creating data management software that can handle the volume, velocity and diversity of data that is anticipated.
7. **Exascale algorithms:** Reformulating science problems and refactoring their solution algorithms for exascale systems.
8. **Algorithms for discovery, design, and decision:** Facilitating mathematical optimization and uncertainty quantification for exascale discovery, design, and decision making.
9. **Resilience and correctness:** Ensuring correct scientific computation in face of faults, reproducibility, and algorithm verification challenges.
10. **Scientific productivity:** Increasing the productivity of computational scientists with new software engineering tools and environments.

http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/report.pdf

Scientific Data Challenges

Workflows for computational science must drive fundamental changes in computer architecture for exascale systems

Breaking with the past: *traditional scientific workflow* - simulate or experiment, saving the data to disk for later analysis

Worsening I/O bottleneck and energy cost of data movement combine to make it impossible to save all of the data to disk

***in situ* data analysis, occurring on the supercomputer while the simulation is running.**

Data Management, Analysis and Visualization
Top "10" Technical Approaches

1. Data structures and traversal algorithms that **minimize data movement**
2. Methods for **data reduction/triage** that support validation of results and data repurposing
3. Maintaining the ability to do **exploratory analysis** to discover the unexpected despite severe data reduction
4. **Knowledge representation and machine reasoning** to capture and use data provenance
5. Coordination of resource access among running simulations and data management, analysis and visualization technologies that **run *in situ***
6. Methods of ***in situ* data** analysis that minimize reliance on a priori knowledge
7. **Data analysis algorithms for high-velocity, high-volume multi-sensor, multi-resolution data**
8. Methods for **comparative and/or integrated analysis** of simulation and experimental/ observational data
9. Design of sharable ***in situ* scientific workflows** to support data management, processing, analysis and visualization
10. Maintaining **data integrity** in the face of error-prone systems
11. Methods of **visual analysis** for data sets at scale and metrics for validating them
12. Improved abstractions for data storage that move beyond the concept of files to more richly represent **the scientific semantics of experiments, simulations, and data points**

Lucy Nowell, DOE/ASCR

Use of Co-Design in DOE Exascale Strategy

Approach: Use scientific workflow requirements to guide architecture and system software and use technology capabilities to design algorithms and software

- Three co-design centers focus on specific application domains:
    - **ExaCT**: Combustion simulation (uniform and adaptive mesh)
    - **ExMatEx**: Materials (multiple codes)
    - **CESAR**: Nuclear engineering (structures, fluids, transport)
- Create "proxy apps"
    - Scaled down versions of "full" code
    - Selects parts/patterns from code to drive programming/ architecture

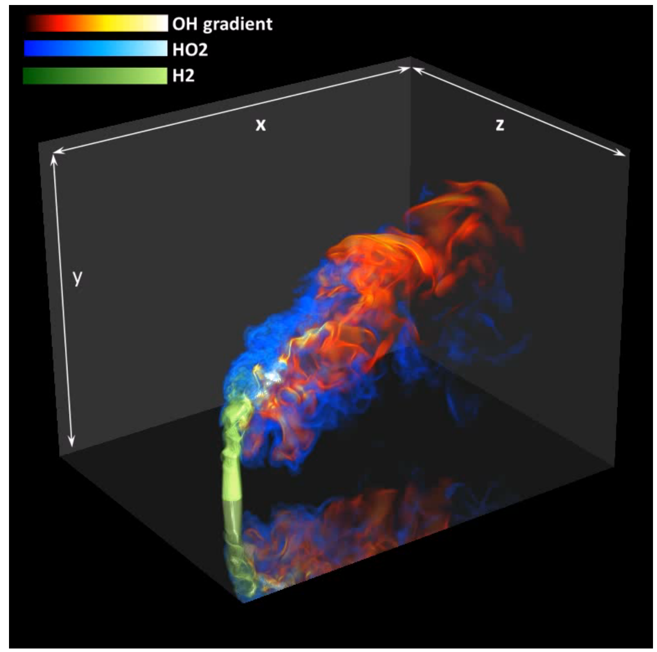

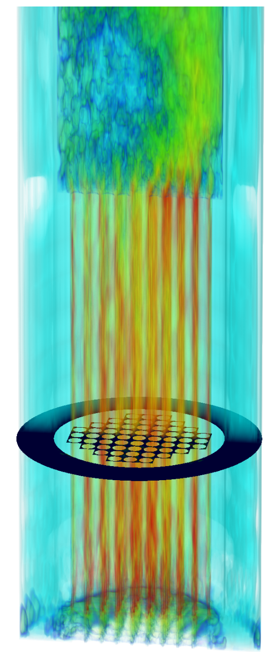

Exascale Computing Initiative

Proposed Timeline

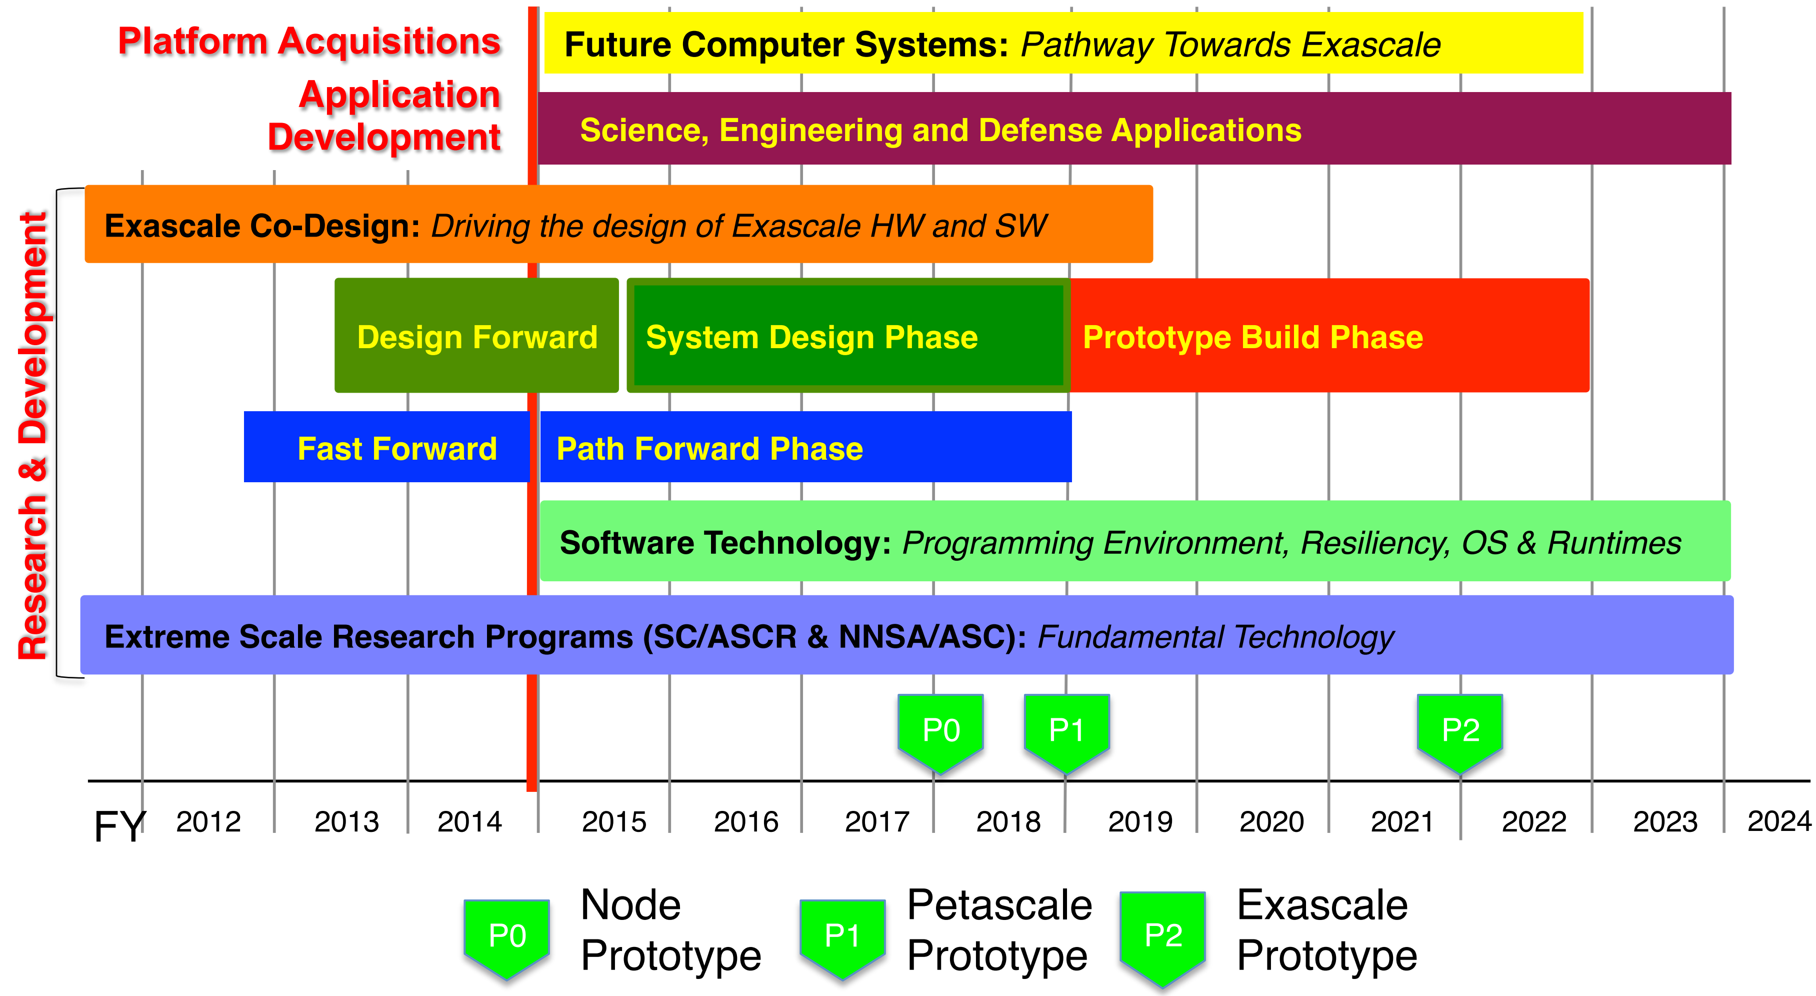

Summary
Four Primary Research Drivers

- **Ensuring a continued, sustainable, and strong technology base that is energy efficient and high performance**
  - Computer technology is a critical element in our daily life, economic vitality, scientific advances and national security
- **Increasing the productivity of computational scientists**
  - Making high-end computing "easier" will lower the barriers to broad adoption of HPC technologies by industrial and business user
- **Achieving the full potential of the scientific discovery environment**
  - Today's capabilities are the result of a piecemeal research approach
  - Need infrastructure for international collaboration
- **Enable scientific workflow ecosystems that supports sharing, collaboration, and reuse of workflows**

BACKUP SLIDES

Exascale Computing
We Need to Reinvent Computing

Traditional path of 2x performance improvement every 18 months has ended

- For decades, Moore's Law plus Dennard scaling provided more, faster transistors in each new process technology
- This is no longer true – we have hit a power wall!
- The result is unacceptable power requirements for increased performance

We cannot procure an exascale system based on today's or projected future commodity technology

- Existing HPC solutions cannot be usefully scaled up to exascale
- Energy consumption would be prohibitive (~300MW)

Exascale will require partnering with U.S. computing industry to chart the future

- Industry at a crossroads and is open to new paths
- Time is right to push energy efficiency into the marketplace

Exascale Challenges

- **Four primary challenges must be overcome**
  - Parallelism / concurrency
  - Reliability / resiliency
  - Energy efficiency
  - Memory / Storage
- **Productivity issues**
  - Managing system complexity
  - Portability
  - Generality
- **System design issues**
  - Scalability
  - Efficiency
  - Time to solution
  - Dependability (security and reliability) must be integrated at all levels of the design
  - Most provide optimal performance for work flow

Exascale Computing Initiative (ECI)
Target System Characteristics

20 pJ per average operation

~x 40 improvement over today's efficiency

Billion-way concurrency

Ecosystem to support new application development and collaborative work, enable transparent portability, accommodate legacy applications

High reliability and resilience through self-diagnostics and self-healing

Programming environments (high-level languages, tools, …) to increase scientific productivity

Extreme Scale Science Work Flow

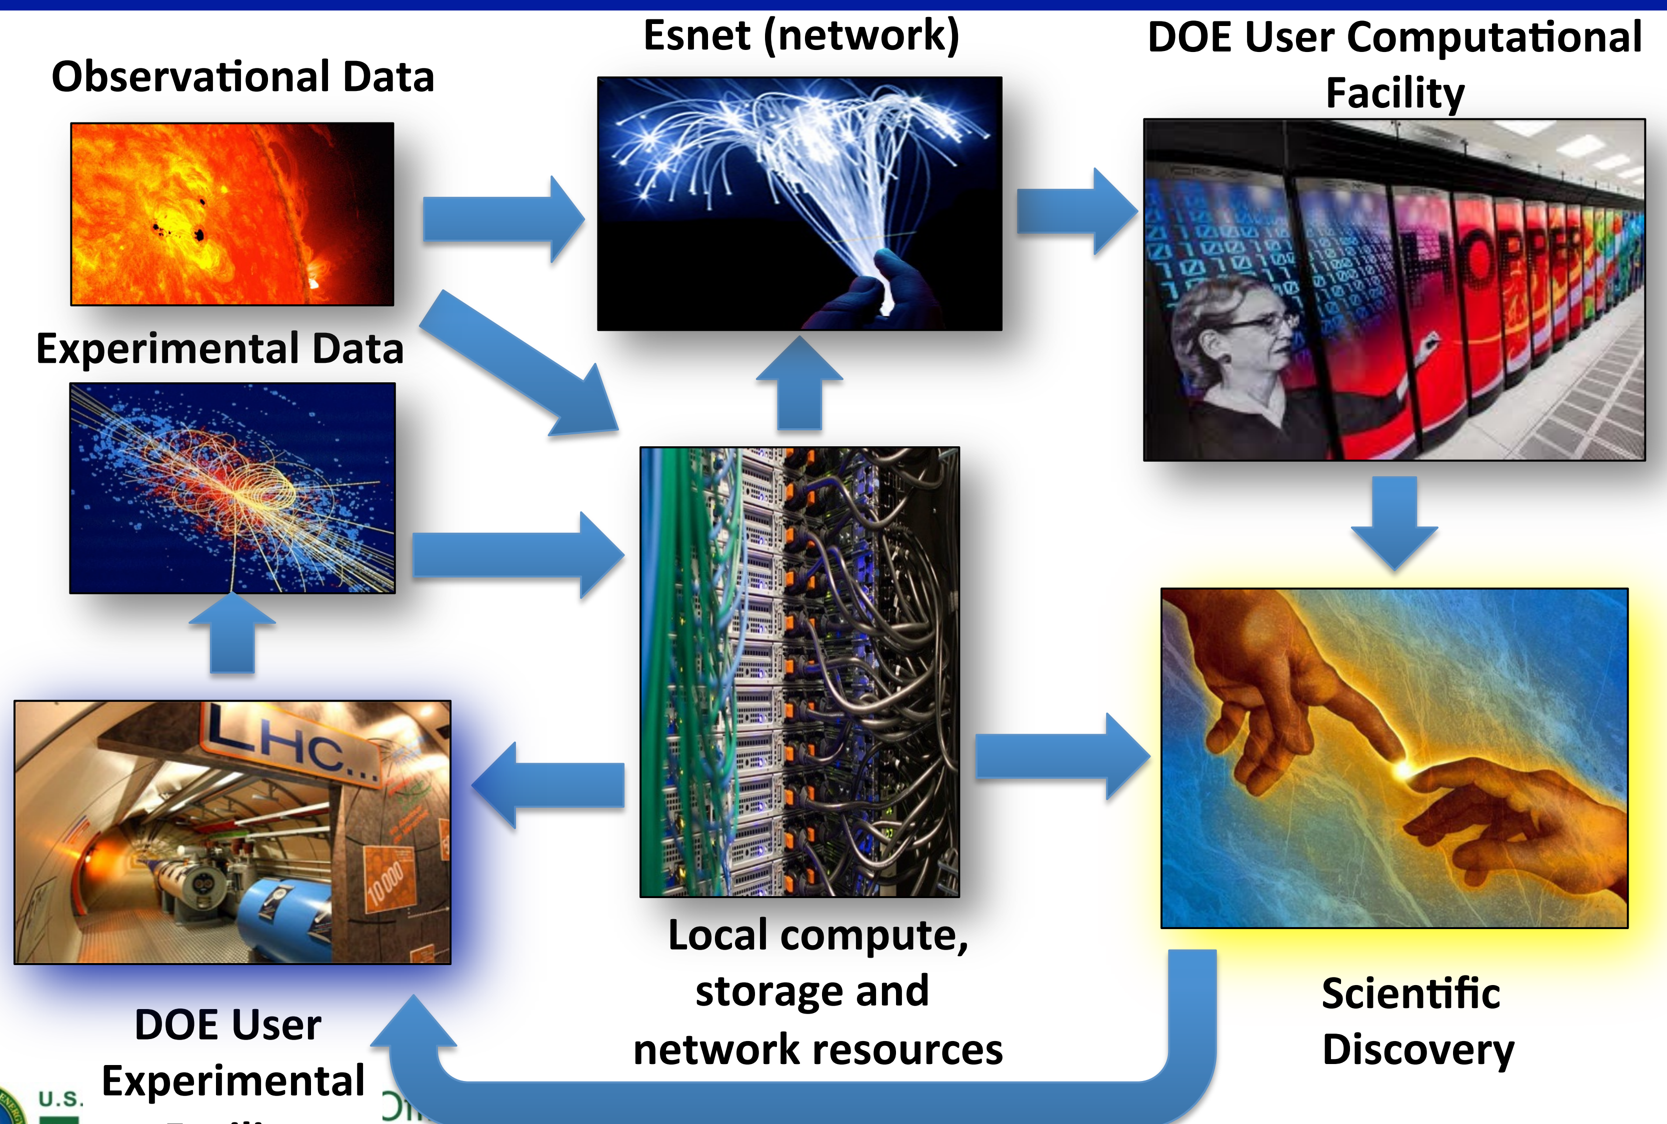

Scientific Data Issues

For purely *in situ* analysis, need to know a priori what analysis is needed

- Consider the value of automating exploratory analysis or hypothesis generation.

How to capture provenance and make the information usable and useful for validation of results.

Need resilient data management, analysis and visualization software

How to help scientists understand tremendous volumes of data.