U.S. DEPARTMENT OF ENERGY | Office of Science

A View From Washington

The Challenges and Opportunities for the Future of Scientific Computing

Richard Carlson
Richard.Carlson@science.doe.gov

ModSim

August 9-11 , 2017

'Spoilers'

- **The Advanced Scientific Computing Research (ASCR) program office is still developing it's strategic goals for emerging programs and future activities**
  - No decisions have been made on funding levels or specific program directions
  - Future programs will cross-cut several ASCR programs and SC program offices
  - Budget challenges will continue to exist
- **Developing the next generation (Exascale) supercomputer (NSCI) is a strategic goal for ASCR**
  - Future programs must align with, and build upon, this goal
  - Computationally intensive and Data intensive workflows are both part of the NSCI computing initiative

DOE/SC - ASCR

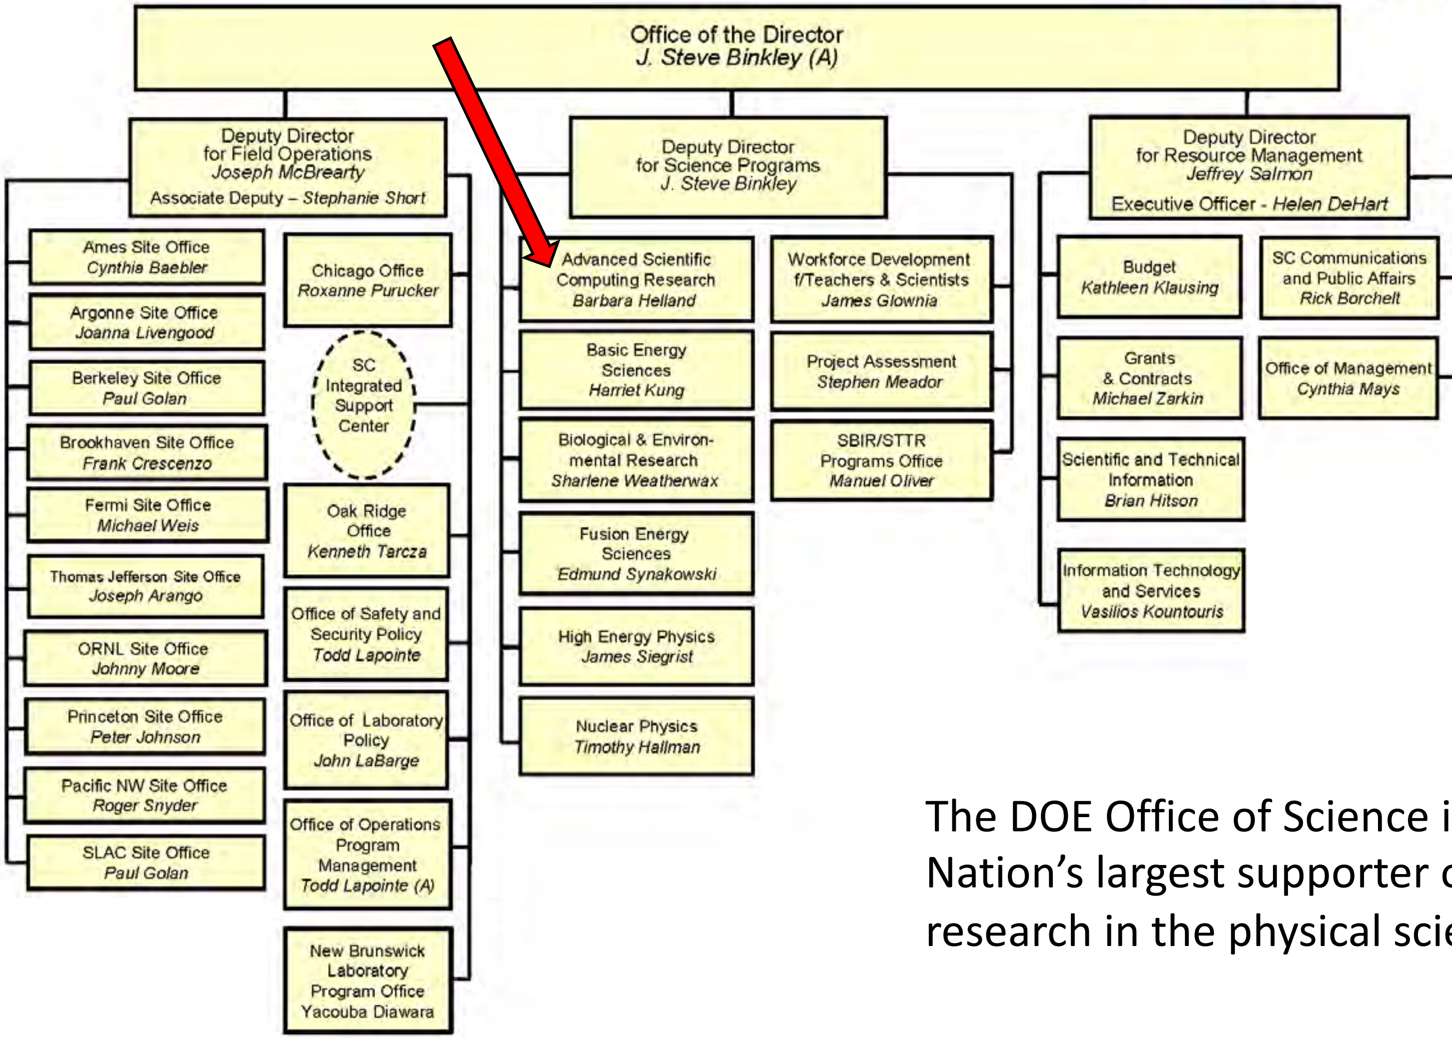

The DOE Office of Science is the Nation’s largest supporter of basic research in the physical sciences

ASCR at a Glance

Office of Advanced Scientific Computing Research

Associate Director – Barbara Helland
Phone: 301-903-7486
E-mail: Barbara.Helland@science.doe.gov

Research

Division Director –

~~**Stephen Lee (Acting); 301-903-5800; Stephen.Lee@science.doe.gov**~~

Ceren Susut (Acting); 301-903-5800
Ceren.Susut-Bennett@science.doe.gov

Facilities

Division Director –

Christine Chalk (Acting); 301-903-9958
Christine.Chalk@science.doe.gov

Research and Evaluation Prototypes Division

Explore technologies beyond Moore's Law (Quantum, Neuromorphic, …)

Fundamental Scientific Research

- **Applied Mathematics:** **Algorithms and software to solve complex science problems;**
- **Computer Science:** **Advanced Operating Systems, runtime architectures, and analysis methods to achieve exascale based science;**
- **Computational Partnerships:** **CoDesign to pioneer the future of scientific applications;**
- **Next Generation Networks for Science:** **Enabling the future of collaborative and distributed science**

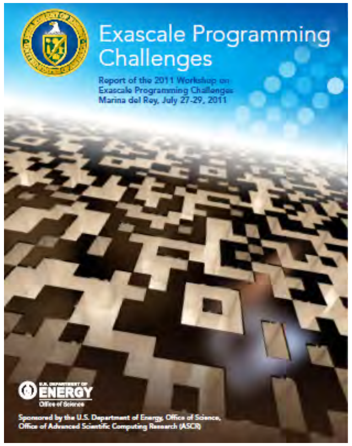

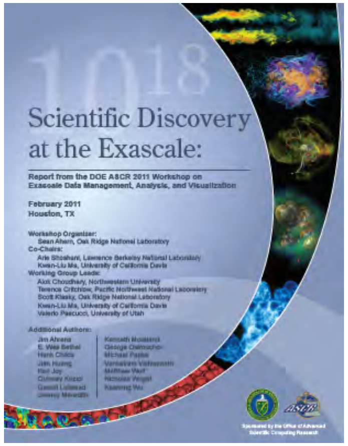

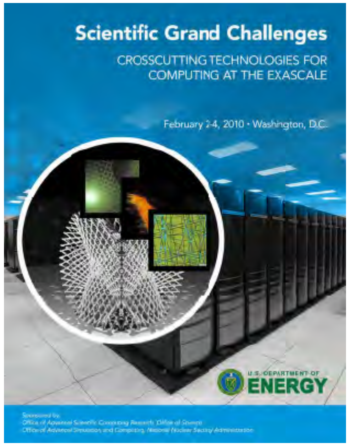

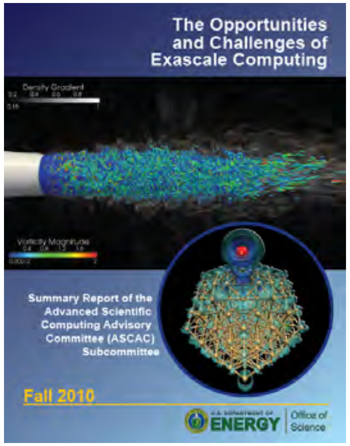

World Class Facilities

- **High Performance Production Computing for the Office of Science**
  - Characterized by a large number of projects (over 400) and users (over 4800)
- **Leadership Computing for Open Science**
  - Characterized by a small number of projects (about 50) and users (about 800) with computationally intensive projects
- **Linking it together – ESnet**
- **Investing in the future – R&E Prototypes**

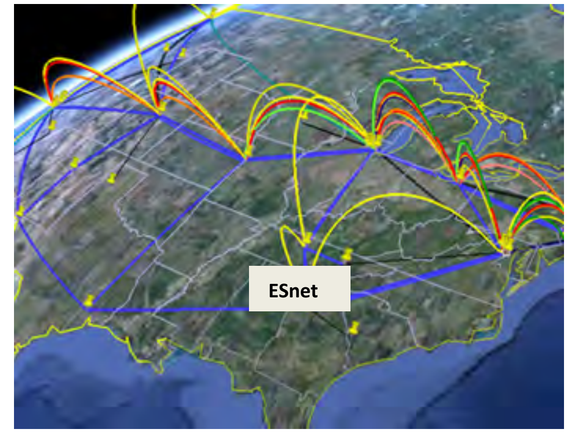

Titan at ORNL

Mira at ANL

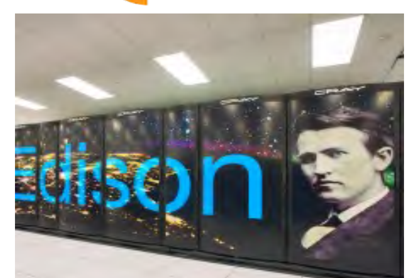

Edison at LBNL

DOE/SC FY17 Budget

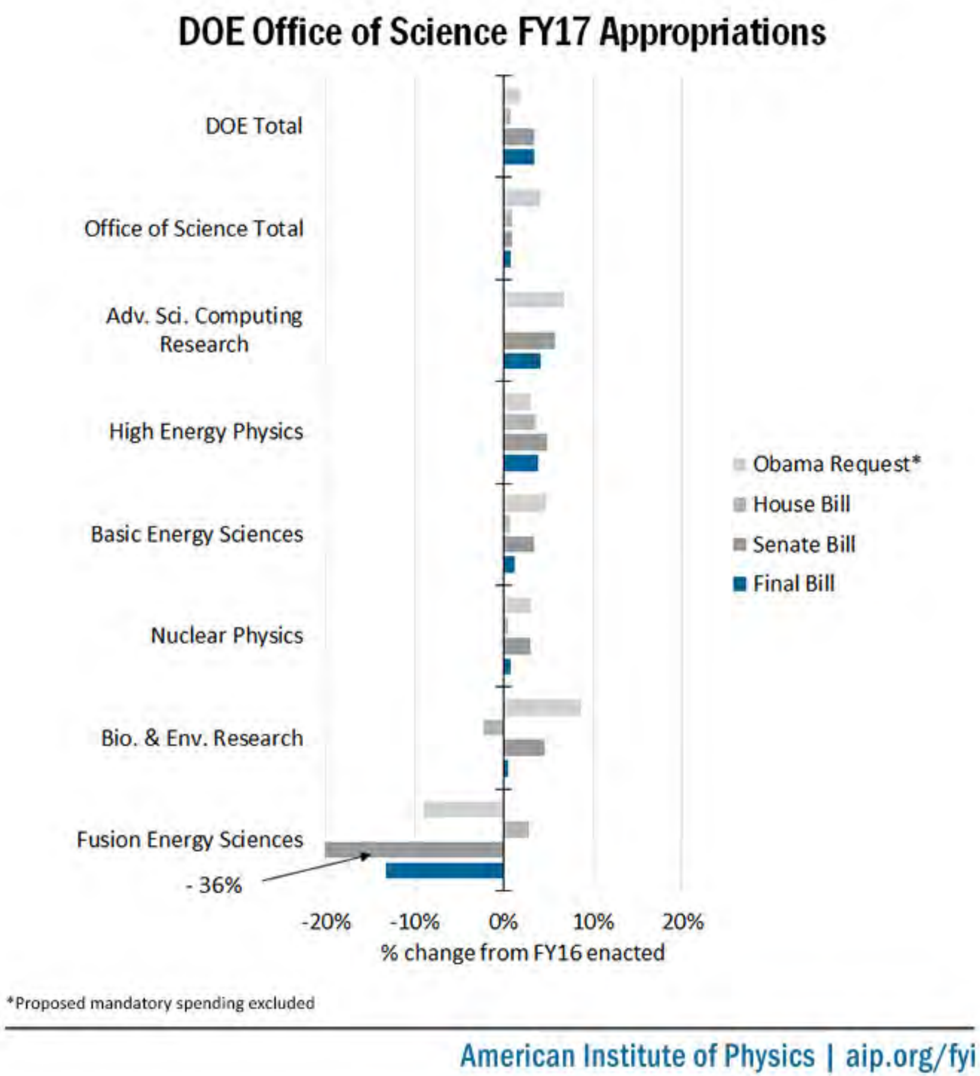

ASCR FY17 Budget

- **ASCR FY17 Budget details**
  - Administration Request $663M (+$42M)
  - House Appropriations Bill $621M (+0)
  - Senate Appropriations Bill $656M (+$35M)
  - Final Enacted Bill: $647 (+$26M)

- **Exascale Computing Budget**
  - Final Enacted Bill: $164M

- **Facilities and Research**
  - Final Enacted Bill: Facilities: $327M; Research: $156M

Most Recent Lab Announcements

2017 Mathematical Multifaceted Integrated Capability Centers (MMICCs)
Announcement Number: LAB 17-1766
Post Date: May 5, 2017
Close Date: July 11, 2017

Exploratory Research for Extreme-Scale Science: Quantum Algorithm Teams (QATs) 17-EXPRESS
Announcement Number: LAB 17 1758
Post Date: May 9, 2017
Close Date: July 21, 2017

Quantum Testbed Pathfinder
Announcement Number: LAB 17-1774
Post Date: May 22, 2017
Close Date: July 21, 2017

Scientific Discovery Through Advanced Computing (SciDAC) Institutes
Announcement Number: LAB 17-1787
Post Date: June 15, 2017
Close Date: July 26, 2017

DOE/SC FY18 Proposed Budget

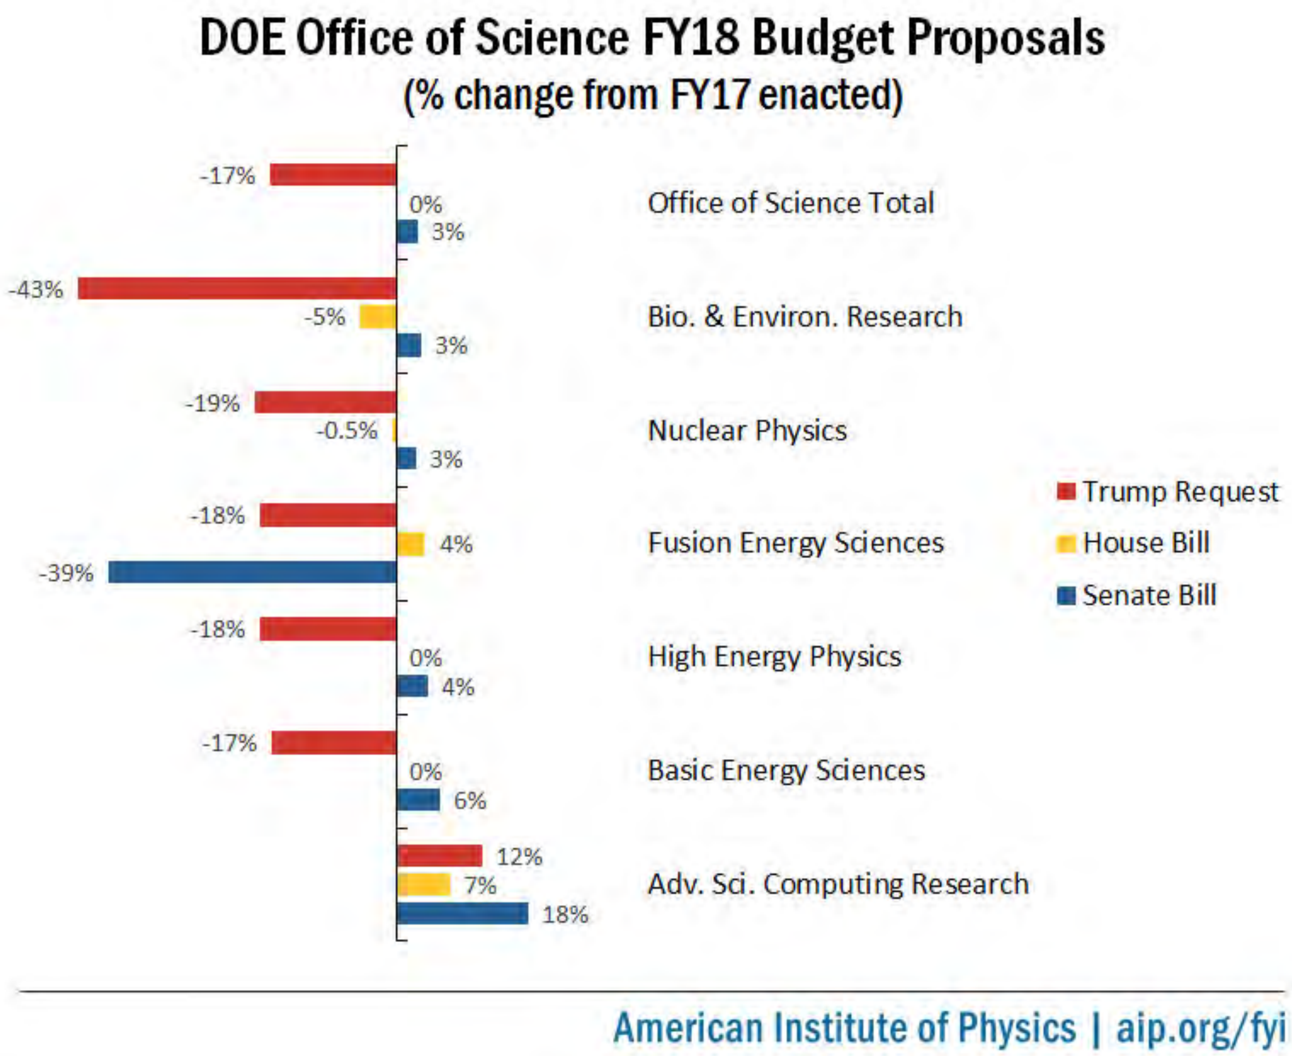

ASCR FY18 Proposed Budget

- **ASCR FY18 Budget details**
  - Administration Request $722M (+$75M)
  - House Appropriations Bill $694M (+$47M)
  - Senate Appropriations Bill $763M (+$116M)

- **Exascale Computing Budget**
  - Administration Request $197M (+$33M)
  - House Appropriations Bill $170M (+$6M)
  - Senate Appropriations Bill $184M (+$20M)

- **Facilities and Research**
  - House Appropriations Bill: Facilities: $369M; Research: $155M
  - Senate Appropriations Bill: Facilities: $423M; Research: $156M
    - Per Senate: research includes CSGF - $10M; R&E Prototypes - $24.6M

DOE Science Computing Communities

- **SciDAC demonstrated that large computationally complex problems are best solved by in-depth collaborative work between ASCR researchers and Domain scientists**
  - Requires years of close collaboration to extract good performance out of an application
  - Requires Leadership class computers and staff to operate these facilities
  - Numerous examples of outstanding science discoveries have resulted

- **Large coherent Science communities (HEP, FES) can offer compute services to members, doing most things by themselves**

- **An emerging class of Experimental/Observations scientists are experiencing a dramatic growth in data and they need help to convert this data into knowledge**
  - Must allow scientists to use experimental and compute facilities without requiring years of training
  - Must deal with geographically distributed facilities, resources, and scientists
  - Potential for enormous impact on a wide array of science communities

Computationally Intensive - Combustion Simulations

- **Goal: 50% improvement in engine efficiency**
- **Center for Exascale Simulation of Combustion in Turbulence (ExaCT)**
  - Combines Modeling, Simulation, and Experimentation
  - Uses new algorithms, programming models, and computer science

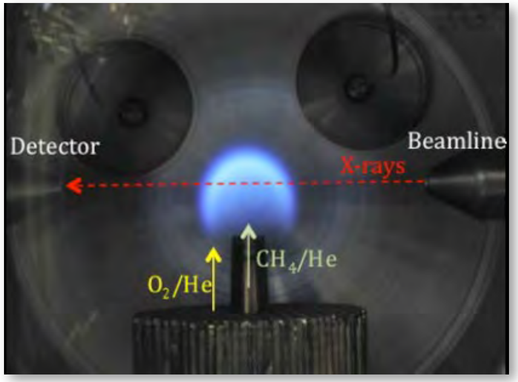

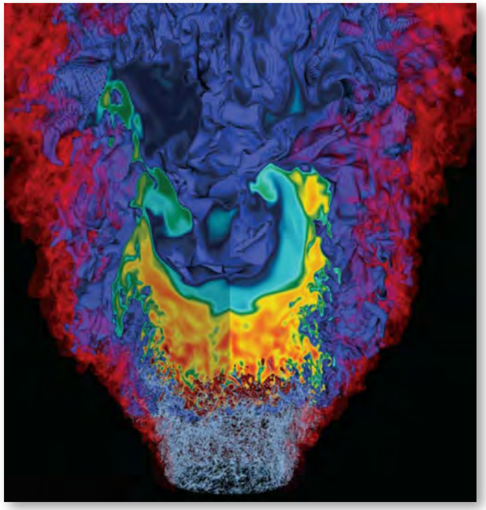

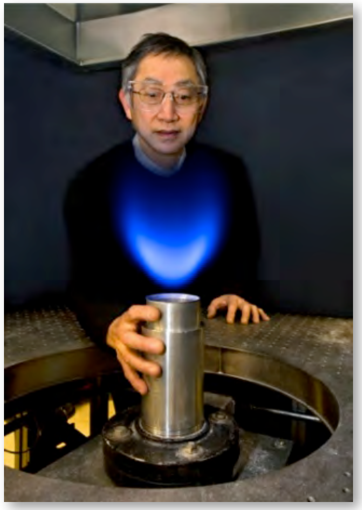

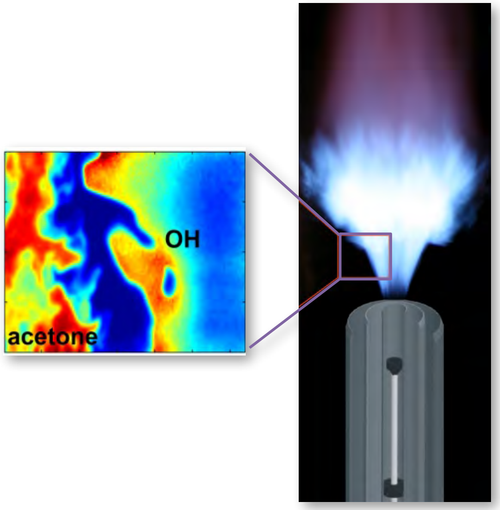

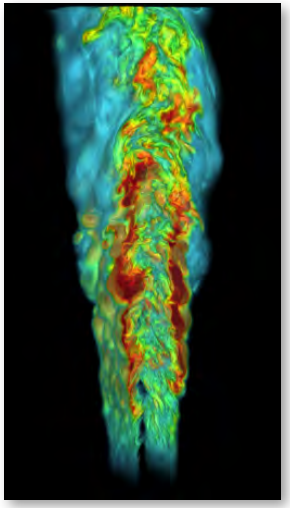

Experimentally Intensive – Belle II

- **International effort to improve existing measurements and to search for new physics**
- **At designed collision rates it will generate ~25PB of raw data per year (90x10^{9} events per year)**
- **Total data stored is expected to reach ~350PB**

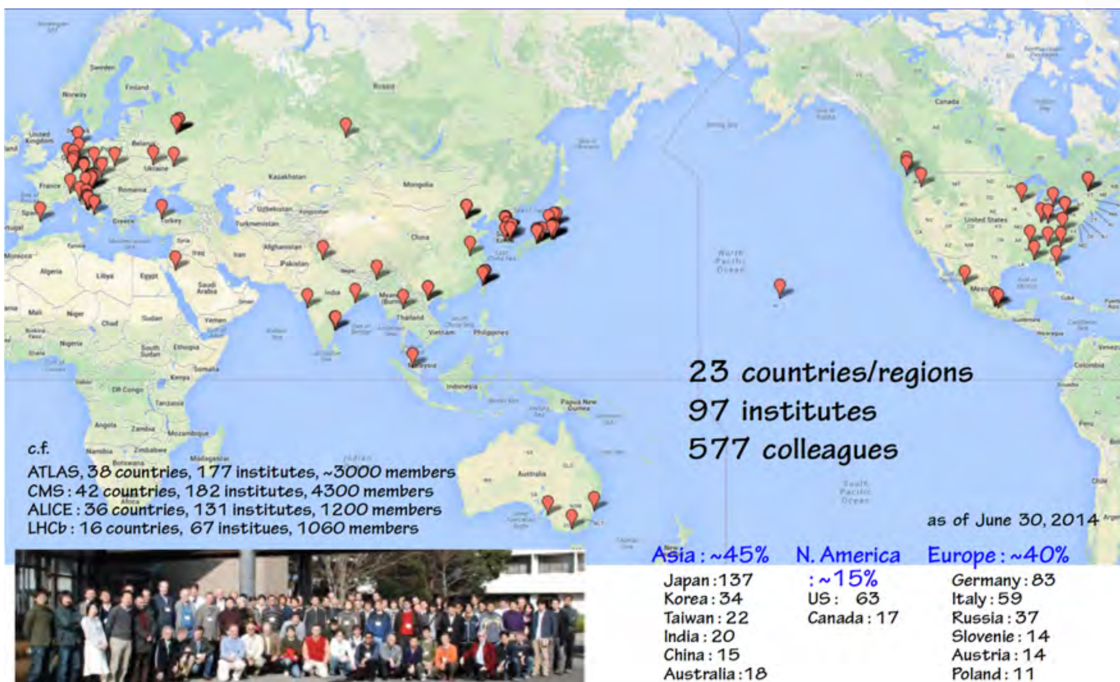

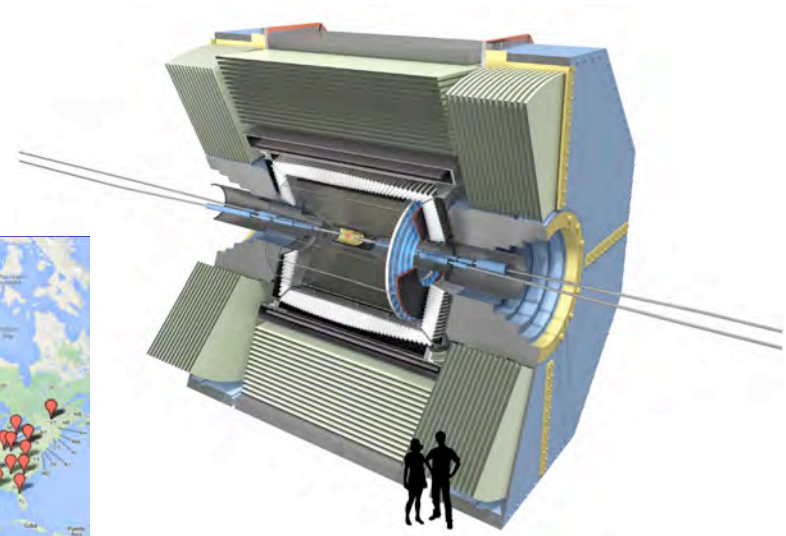

Lead: Malachi Schram

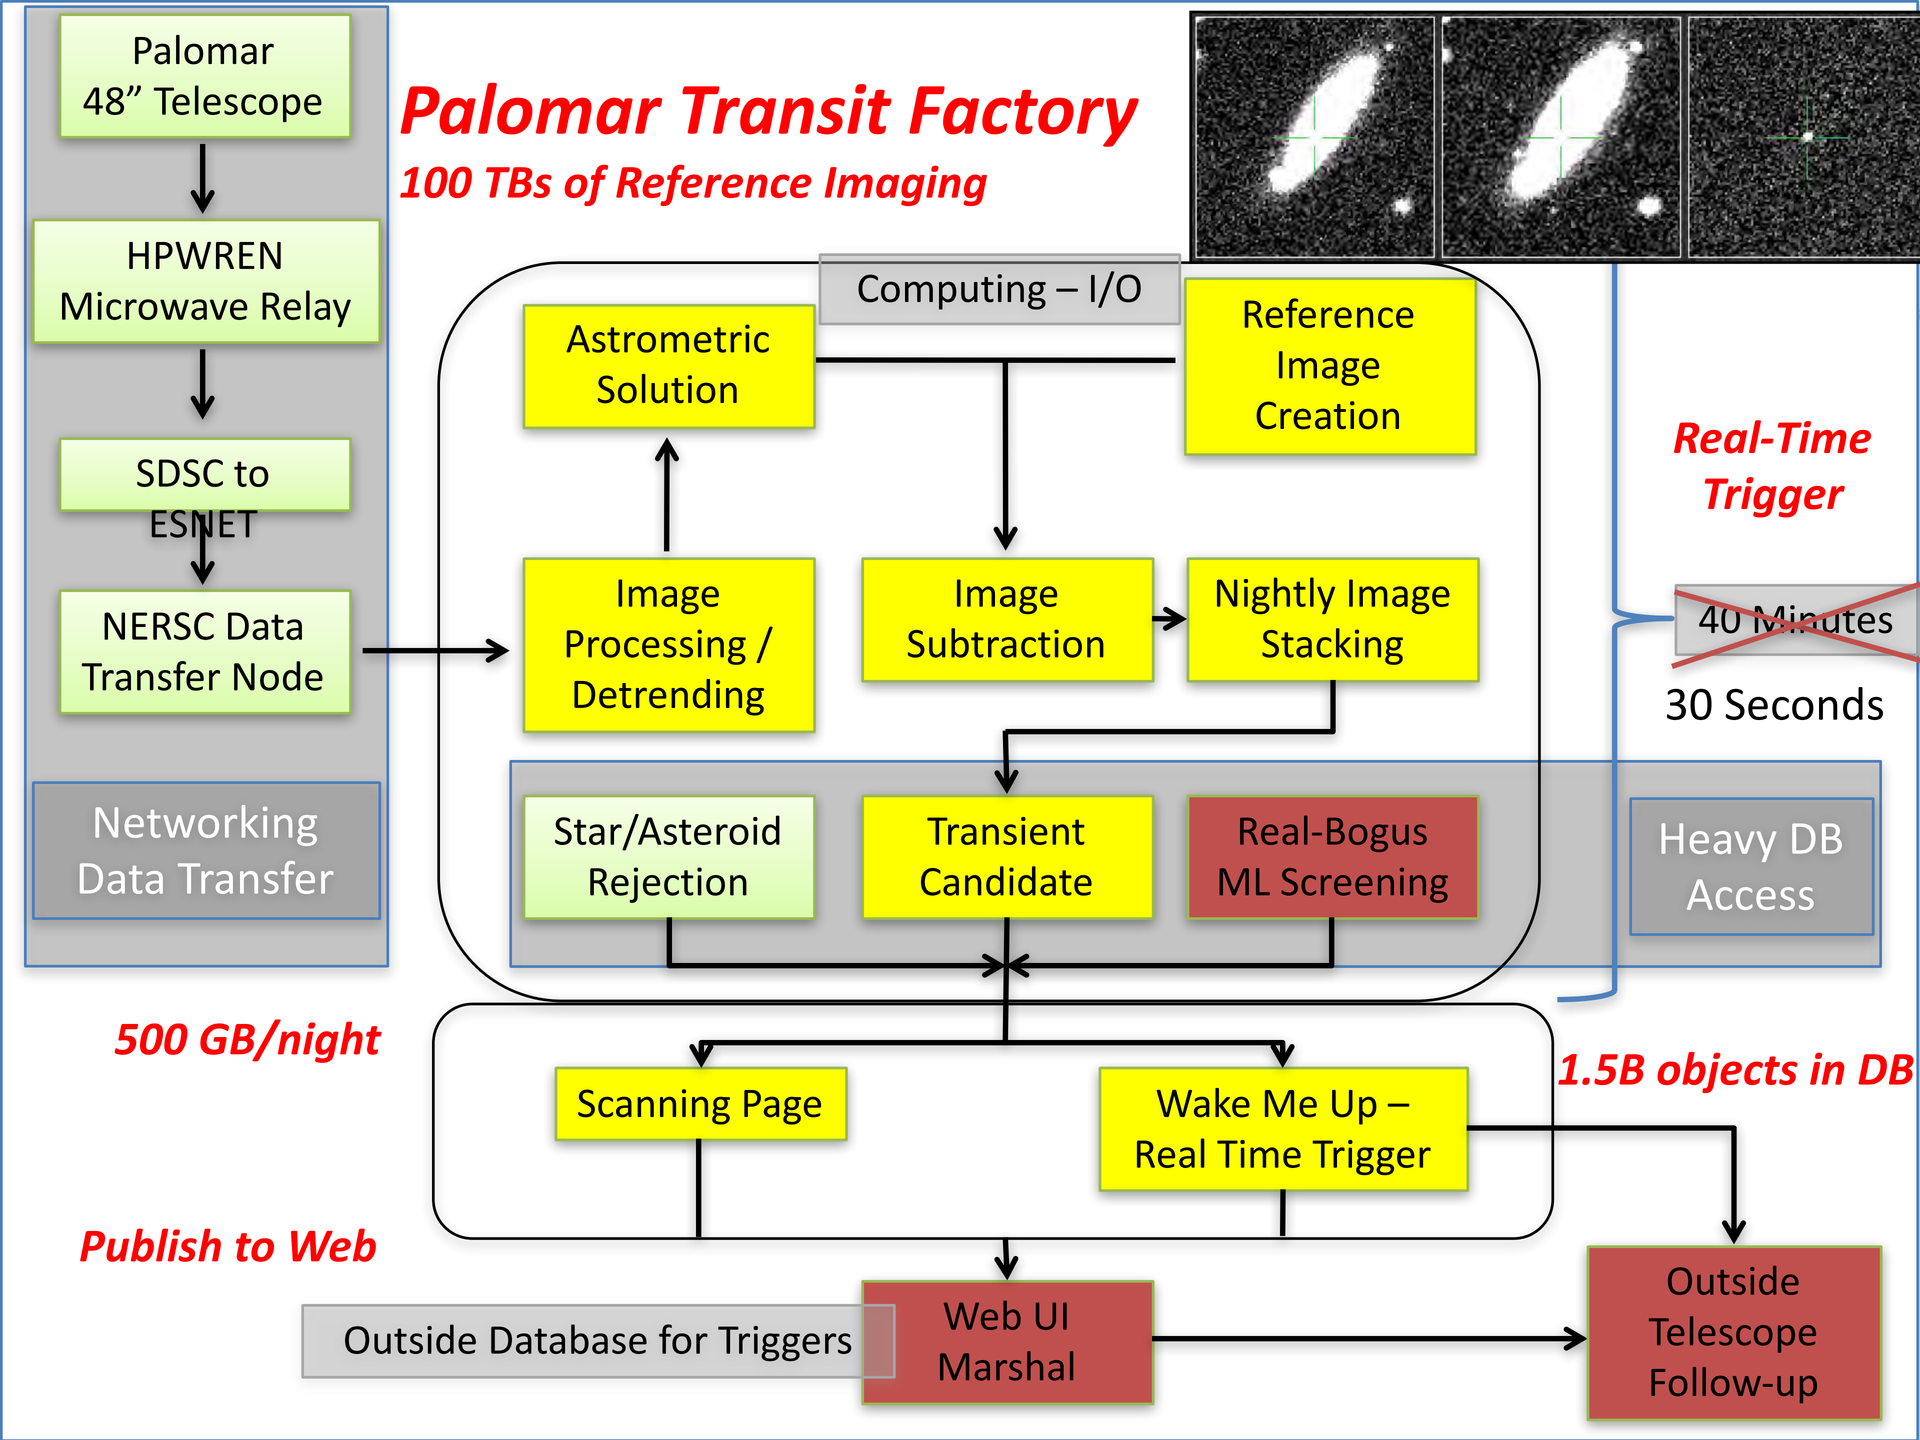
Palomar Transit Factory
100 TBs of Reference Imaging
Palomar 48” Telescope
HPWREN Microwave Relay
SDSC to ESNET
NERSC Data Transfer Node
Networking Data Transfer
Computing – I/O
Astrometric Solution
Reference Image Creation
Image Processing / Detrending
Image Subtraction
Nightly Image Stacking
Star/Asteroid Rejection
Transient Candidate
Real-Bogus ML Screening
Heavy DB Access
Real-Time Trigger
40 Minutes
30 Seconds
500 GB/night
1.5B objects in DB
Scanning Page
Wake Me Up – Real Time Trigger
Publish to Web
Outside Database for Triggers
Web UI Marshal
Outside Telescope Follow-up

Scientific Computing Tomorrow

- **Computationally intensive codes generate more data and adopt in-situ analysis methodologies**
- **Experimental instruments generate more data requiring near-real-time processing**
- **Single-user scientists collaborating with facility staff requiring fast access to data analysis and visualization services**

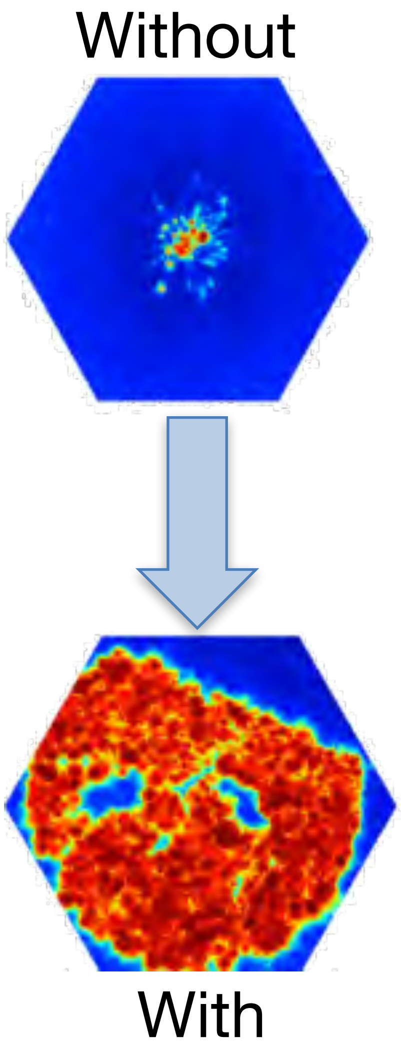

Impact of using immediate assessment of sample alignment in a near-field, high-energy diffraction microscopy experiment

Collaboratively Intensive – Material Structures

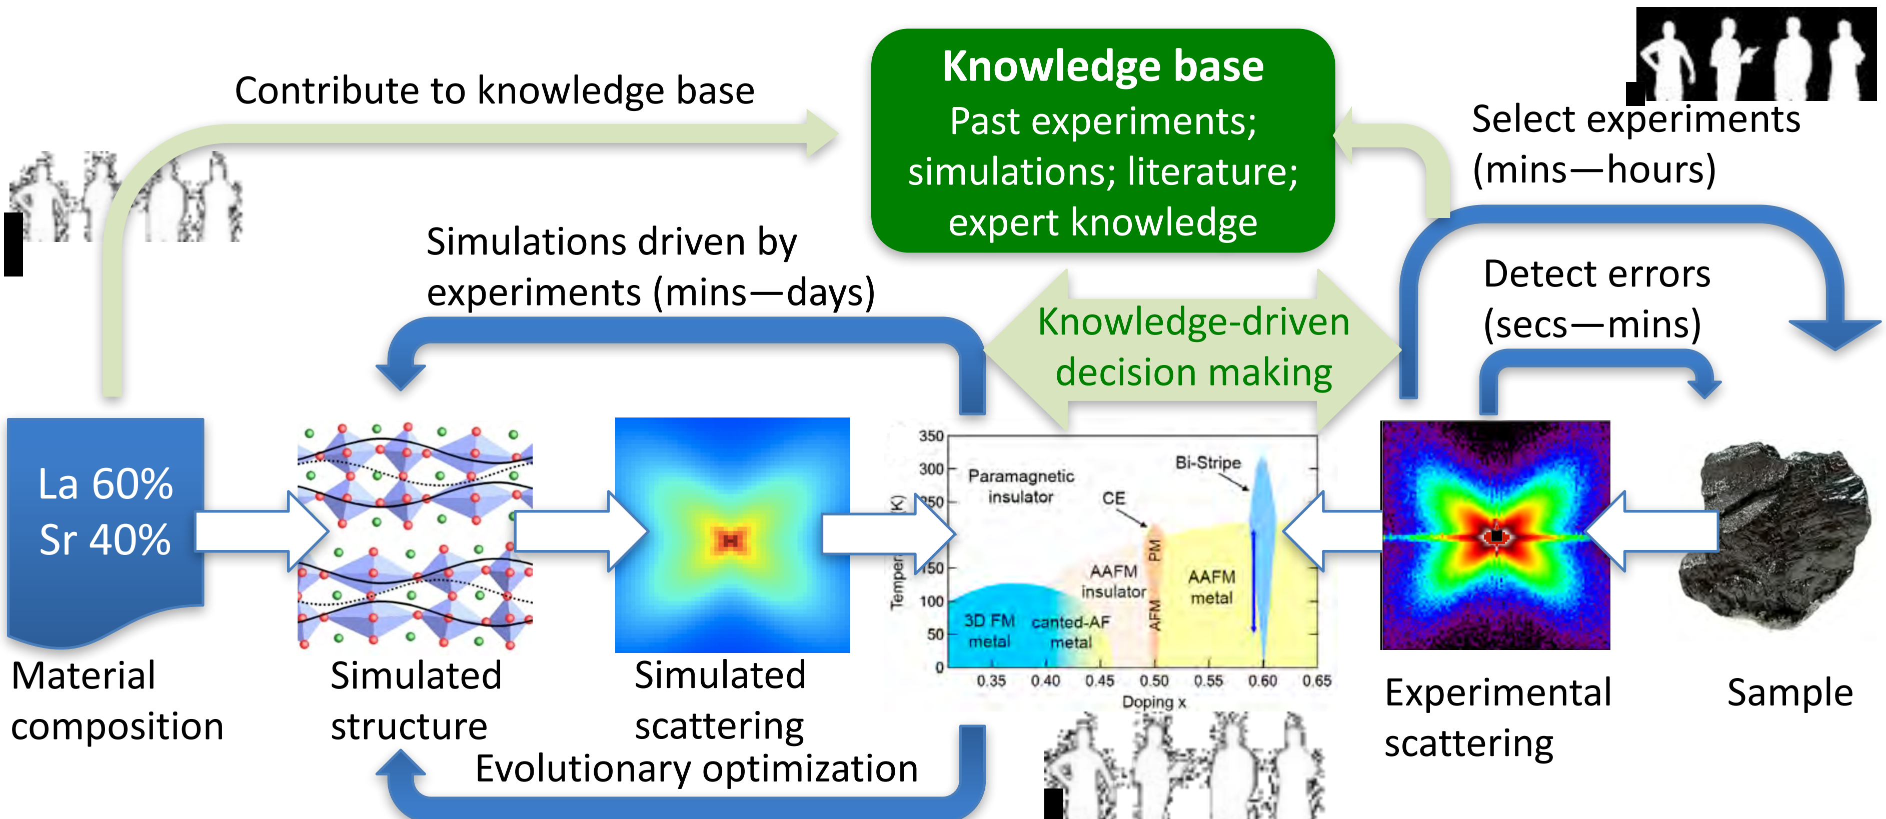

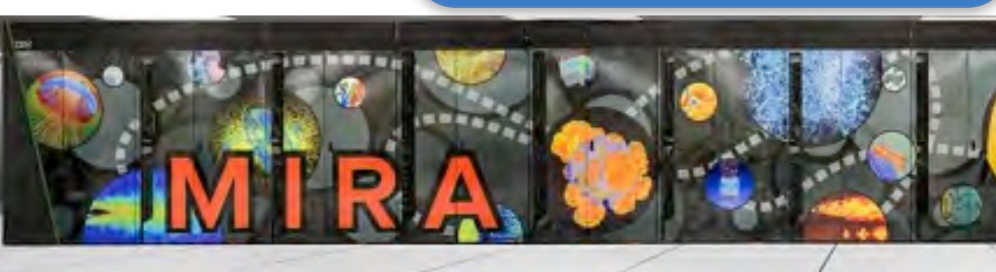

Comparing Software Environments

Data Intensive	Compute Intensive
Programmers are focused on delivering functionality	***Programmers are focused on delivering performance***
Software responds to elastic resource demands	After allocation, ***resources static until termination***
Data access often ***fine-grained***	Data access is ***large bulk*** (aggregated) requests
Services are resilient to fault	***Applications restart after fault***
Often ***customized*** programming models	Widely ***standardized*** programming models
Libraries help ***move computation to storage***	Libraries help ***move data to CPUs***
Users routinely deploy their own environments	***Users almost never deploy customized environments***

Comparing Data Environments

Data Intensive	Compute Intensive
Data coordination occurs within a ***globally distributed environment***	Data coordination occurs within a ***single computer***
Data generated at **multiple locations** and possibly in several **different formats** and structures	Data generated **inside computer** with **format** and structure **defined** by application developer
Multiple independent data sets share ***common network path***	Multiple coordinated data sets share ***allocated compute resources***
Variations in delay and jitter require data ***prestaging minutes to days in advance***	Variations in delay and jitter require data ***prefetching msec in advance***
Authentication performed by ***multiple facilities or administrative domains***	Authentication performed at ***single facility or administrative domain***

ASCR Investment in ModSim

- **Long history of Modeling and Simulation activities involving domain science problems**
  - Partnerships between ASCR and other DOE/SC program offices
- **Long history of Modeling and Simulation of advanced computers**
  - Partnership between ASCR and NNSA to understand how applications will perform on new supercomputers
- **Expand efforts to explore Modeling and Simulation of complex workflows and distributed science discovery ecosystems**
  - Multi-agency activity to identify research challenges that drive future Modeling and Simulation (ModSim) programs

Emerging Areas for Future Study

- **Scientific Workflow Systems designed specifically to compose and execute a series of computational or data manipulation steps**
  - Identify Scientific Research aspects of Workflows
    - Move beyond Engineering tasks to find/fix problems that occur
  - Develop testable theories, experiments, and simulations that demonstrate a deep understanding of both the workflow and the distributed infrastructure being used
- **Streaming Data from detectors and sensors**
  - Near Real-Time processing to determine state of experiment
  - Comparison of Experimental Analysis and Simulation Data
  - Simultaneous Real-Time exploration (zoom, pan, subset) of Data by multiple scientists

Emerging Areas for Future Study

- **Extreme Heterogeneity**
  - New compute paradigms
    - Quantum, Neuromorphic, Probabilistic
  - Specialized accelerators
    - GPU, FPGA, TPU, ASCIs
  - Merged with traditional CPU based computing

- **Integrating SMART components and services into Scientific Computing environments**
  - Power efficient SMART components
  - Composable services created dynamically by a scientist using SMART components

Conclusions

- **Scientific Computing is expanding in multiple directions, including Hardware/Software Heterogeneity and user experience levels**
- **Scientific Workflows are an enabling technology to support science discovery in the 21st century**
- **Real-Time and Near Real-Time supercomputing will expand, not replace, traditional batch oriented services**