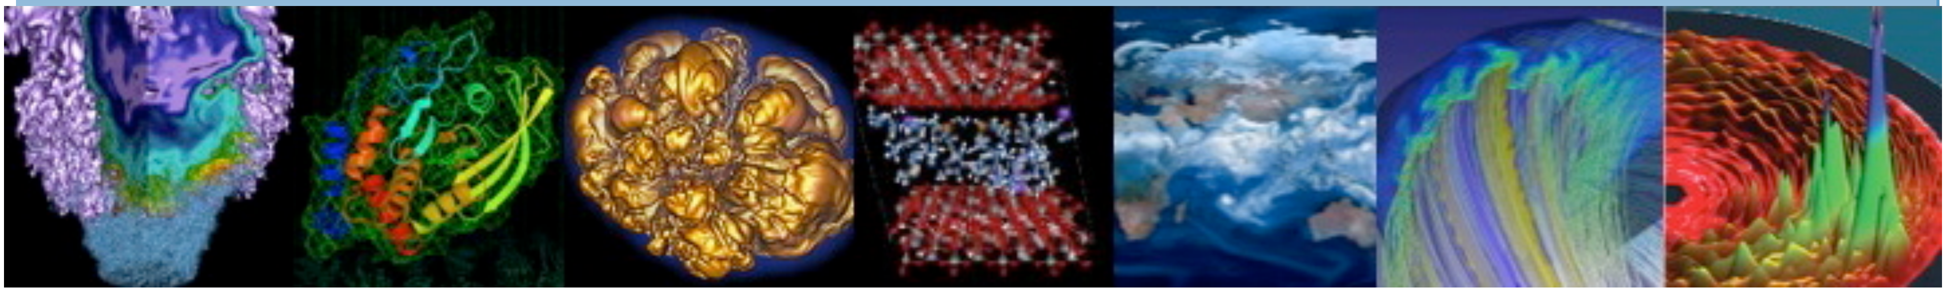

NERSC Accomplishments and Plans

Kathy Yelick

Associate Laboratory Director for Computing Sciences

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC Facility Leads DOE in Scientific Computing Productivity

Computing for science

- 4000 users, 500 projects
- 48 states; 65% from universities
- Hundreds of users each day
- 1500 publications per year

Systems for science

- 1.3PF Cray system, Hopper
- .5 PF in smaller systems, including those for JGI, HEP/NP

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Data Analysis Grows more Automated with the Explosion of Scientific Data Sets

NERSC used in 2011 Nobel Prize

- Type Ia supernovae used as “standard candles” to measure distance.
- Simulations at NERSC in late 90s modeled the appearance from Earth.

More recently: astrophysics discover early nearby supernova.

- Discovered within hours of its explosion, a rare glimpse at the supernova’s outer layers reveal what kind of star exploded.
- The last such supernova was in 1972. Before that: 1937, 1898 and 1572
- NERSC accepts ~300GB/night and runs machine learning algorithms to process images and detect new transients;

The research shows that the universe is expanding at an accelerating rate. The nature of the dark energy force behind this may be the most important problem in 21st century physics.

Leading Scientific Discoveries: MaterialsGenome @ NERSC

18 Years from new materials discovery to commercialization

- Reduce time from discovery to use
- Inverts materials design problem:
  - What properties do I want?
  - Which materials have them?
- Ab initio calculations of materials
- Share data through a searchable library
- Massive ensembles now running on our most energy-efficient, cost-effective machine: Hopper

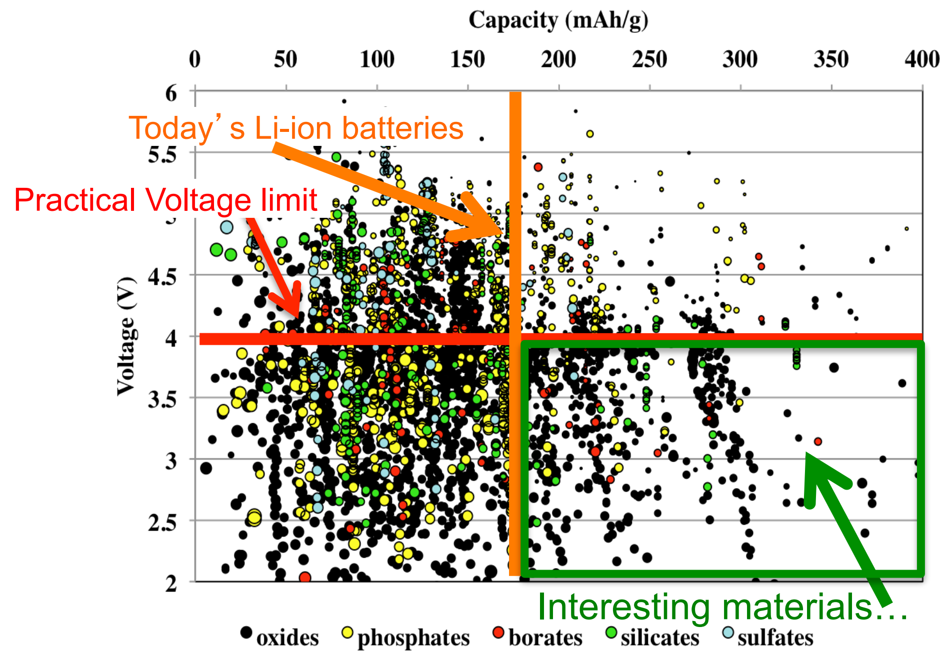

Ab-initio simulation of over 20,000 potential Li-ion cathode compounds. Above 4 Volts solvents becme unstable. This research is pushing the boundary of capacity upward beyond today's limit of 170 mAh/g.

PIs: Gerd Ceder and Kristen Persson; with NERSC and CRD/LBNL help

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Leading Scientific Discoveries: Molecular Dynamics and Protein Folds

- **Produced public catalog of the unfolding dynamics of 11,000 proteins, covering all 807 self-contained autonomous folds**
- **Simulations used 12M hours of NERSC on custom code and help from NERSC on load balancing, optimizations, and workflow**
- **Mined amyloid producing proteins and found common structural feature between normal and toxic forms.**
  - Custom-designed complementary compounds, which bind with toxic forms of proteins that cause multiple diseases, including Alzheimer's and mad cow.
  - Results suggest drug designs, screening for blood/food supply, and diagnostic tools for up to 25 amyloid diseases.

Valerie Daggett, PI, U. Washington

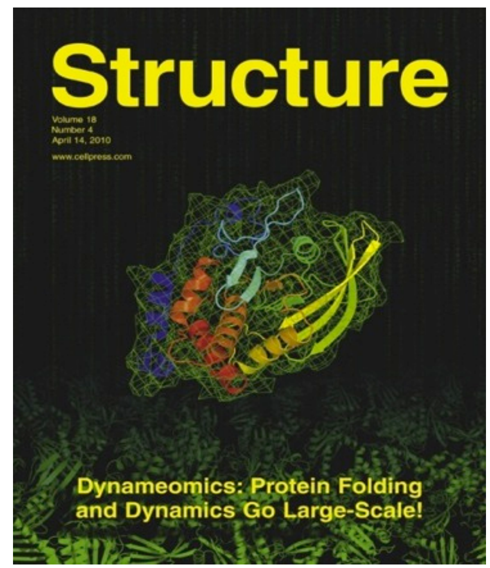

U.S. DEPARTMENT OF ENERGY | Office of Science

Science is in the Data!

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC and ALS Combine for Scientific Discoveries

- **Collaboration between theoreticians and experimentalists.**
- **Milestone: Calculation of the X-ray spectrum of atomic nitrogen; confirmed by data collected in a unique ALS experiment.**
  - *Both the experiment and the calculations were a "tour de force," according to Physical Review Letters referees.*
  - *"Without access to such a petaflop enhanced XE6 Cray architecture like Hopper, the calculations could simply not have been done." – Brendan McLaughlin (Harvard).*
  - *Currently pushing Hopper to it limits on larger atoms*

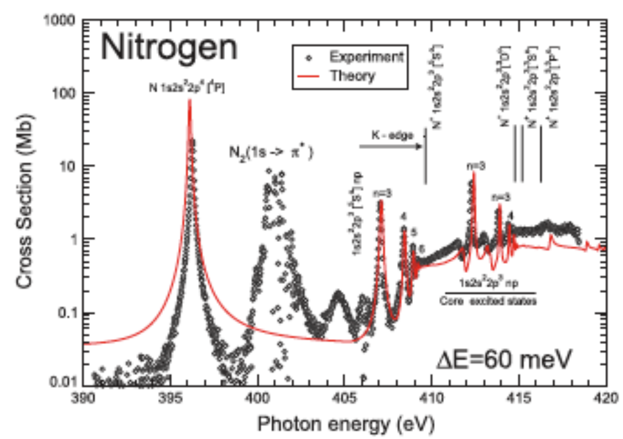

M. Pindzola (Auburn U)

Nitrogen photoionization cross section from Hopper and ALS observations

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC Systems

Large-Scale Computing Systems

Franklin (NERSC-5): Cray XT4
- 9,532 compute nodes; 38,128 cores
- ~25 Tflop/s on applications; 356 Tflop/s peak

Hopper (NERSC-6): Cray XE6
- 6,384 compute nodes, 153,216 cores
- 120 Tflop/s on applications; 1.3 Pflop/s peak

Clusters

140 Tflops total

Carver
- IBM iDataplex cluster

PDSF (HEP/NP)
- ~1K core cluster

Magellan Cloud testbed
- IBM iDataplex cluster

GenePool (JGI)
- ~5K core cluster

NERSC Global Filesystem (NGF)

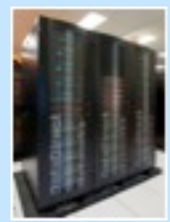

Uses IBM's GPFS
- 1.5 PB capacity
- 5.5 GB/s of bandwidth

HPSS Archival Storage

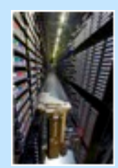

- 40 PB capacity
- 4 Tape libraries
- 150 TB disk cache

Analytics

Euclid

(512 GB shared memory)

Dirac GPU testbed (48 nodes)

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NX Provides Faster Remote Visualization

- **NX Servers plus client software**
- **Used worldwide for**
  - Scientific data visualization
  - Remote debugging with GUIs

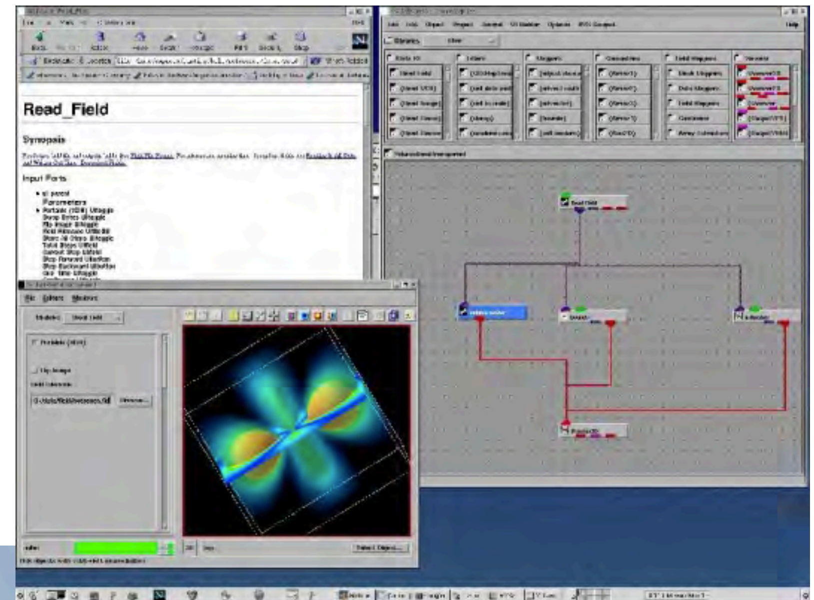

The Return of Deep Consulting

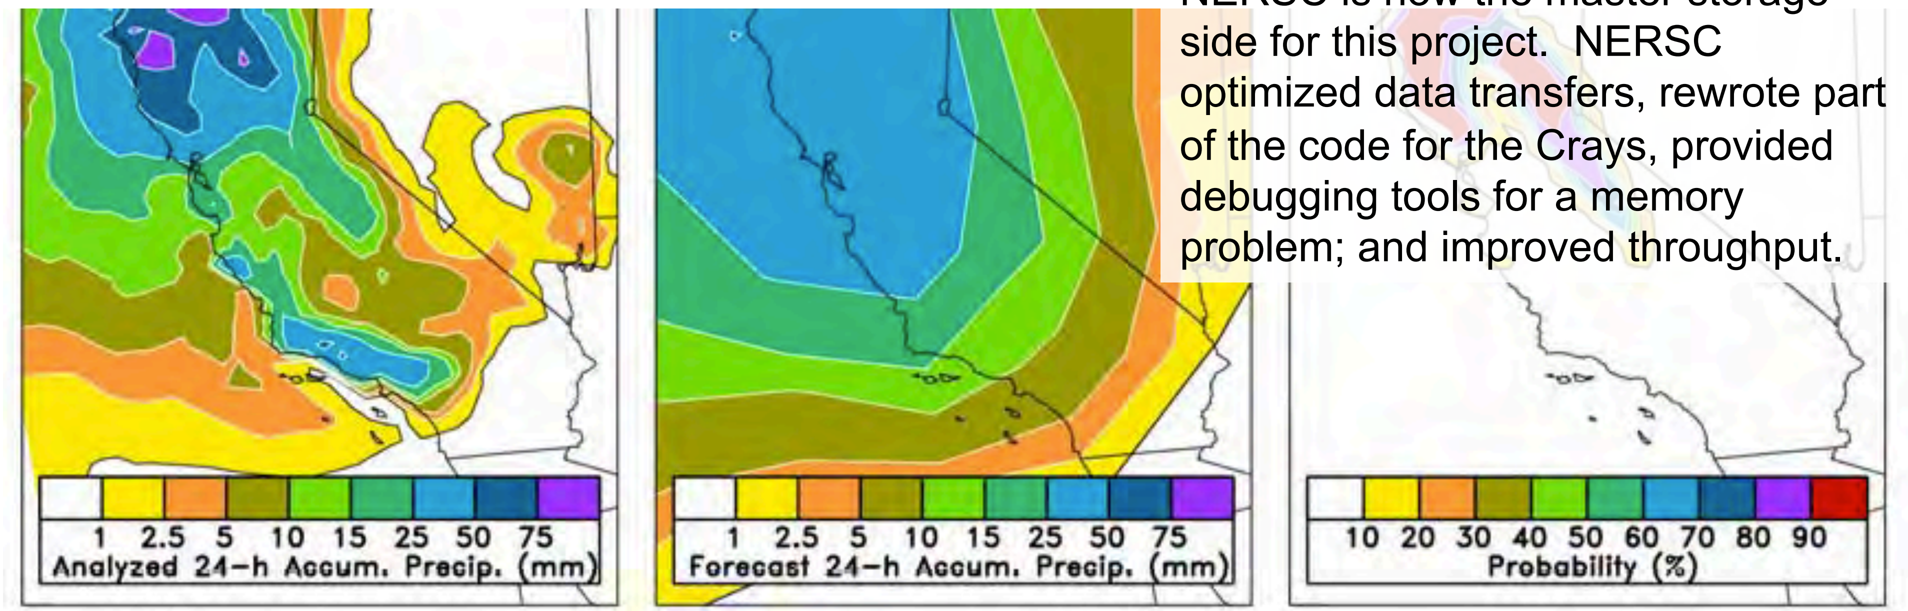

NERSC is now the master storage side for this project. NERSC optimized data transfers, rewrote part of the code for the Crays, provided debugging tools for a memory problem; and improved throughput.

- **INCITE program created at NERSC, adopted by LCFs, by DOE in 2010**
  - NISE 2012 will be large projects moving toward exascale
  - Considering data-intensive "red carpet" program

Support for Data Intensive Computing

- **NERSC has several projects, some joint with Cray, to broaden systems support**
  - Task farmer for running massive numbers of jobs with unpredictable times
  - MySGE (Sun Grid Engine) for complex workflows
  - Filecacher caches large random-access files in memory

NERSC History of Production High Performance Computing

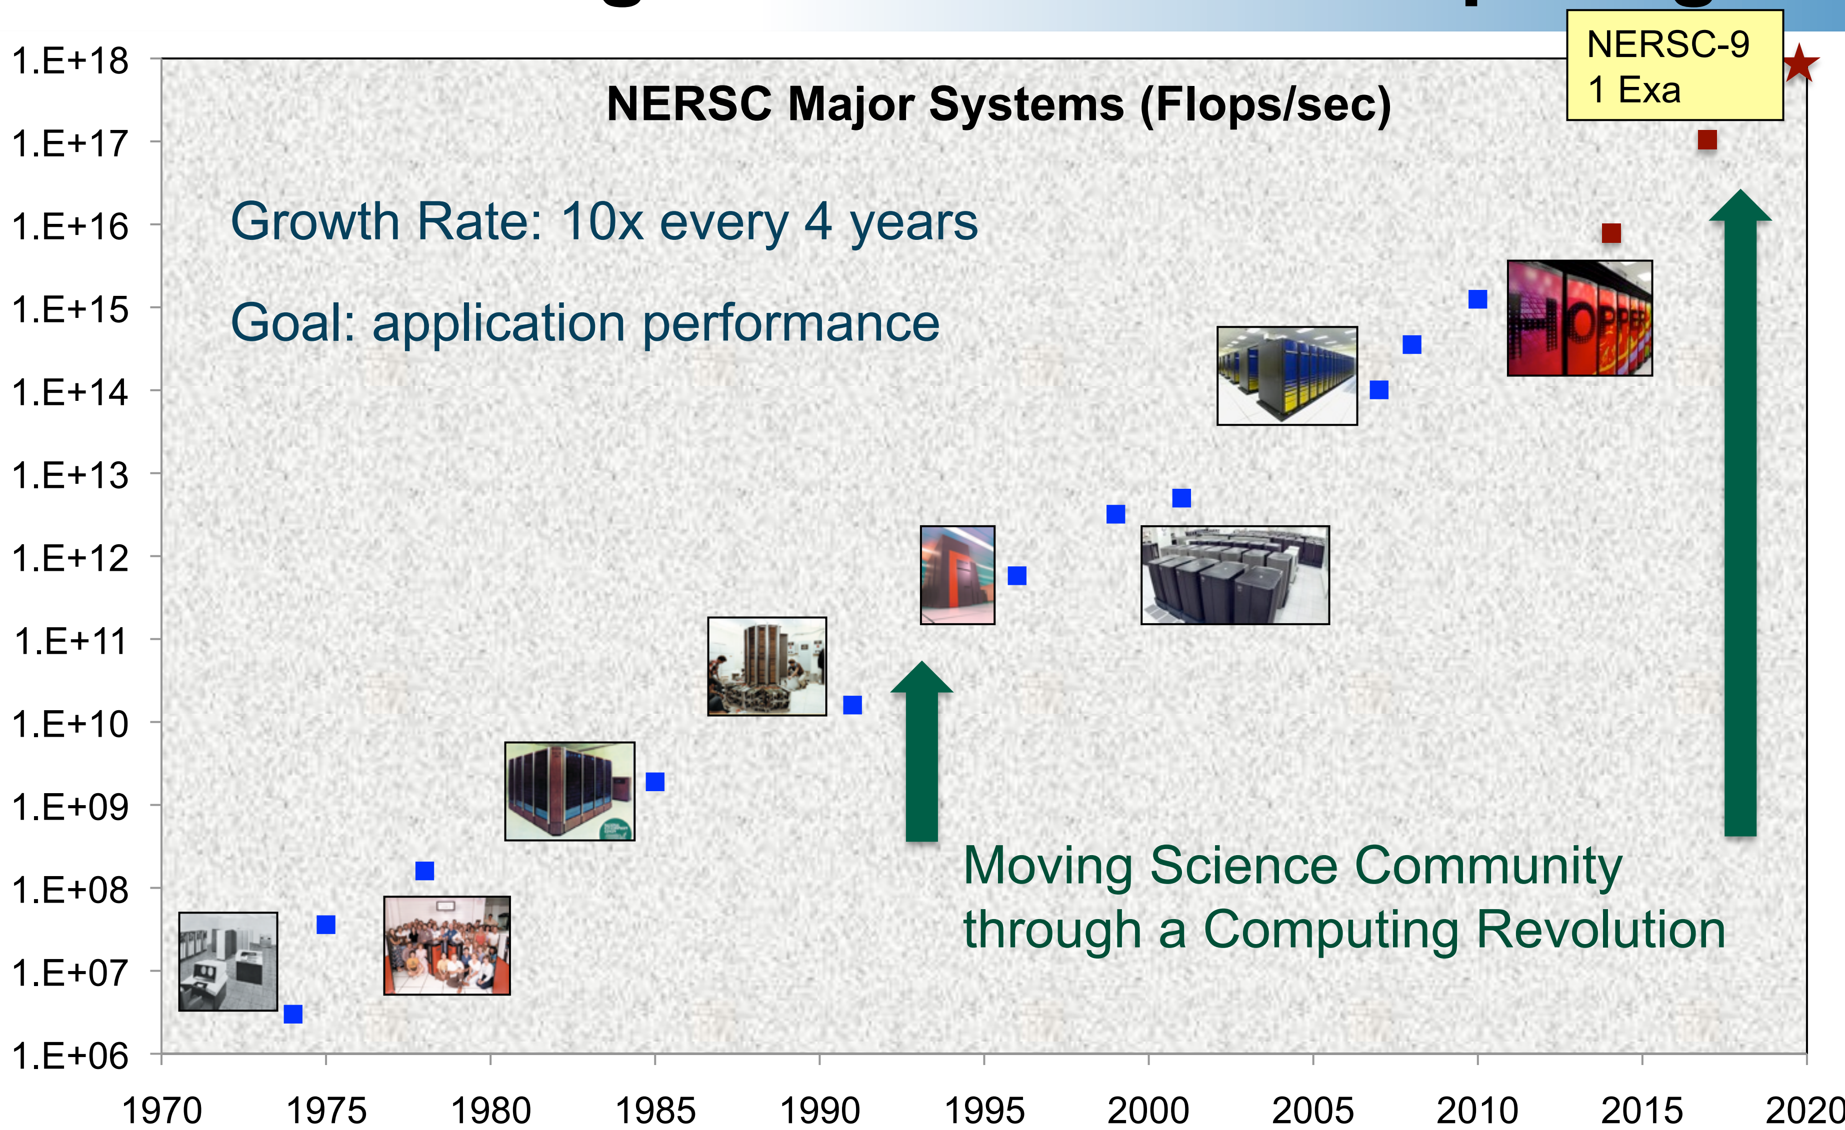

Computing Crisis is Not Just about Exascale

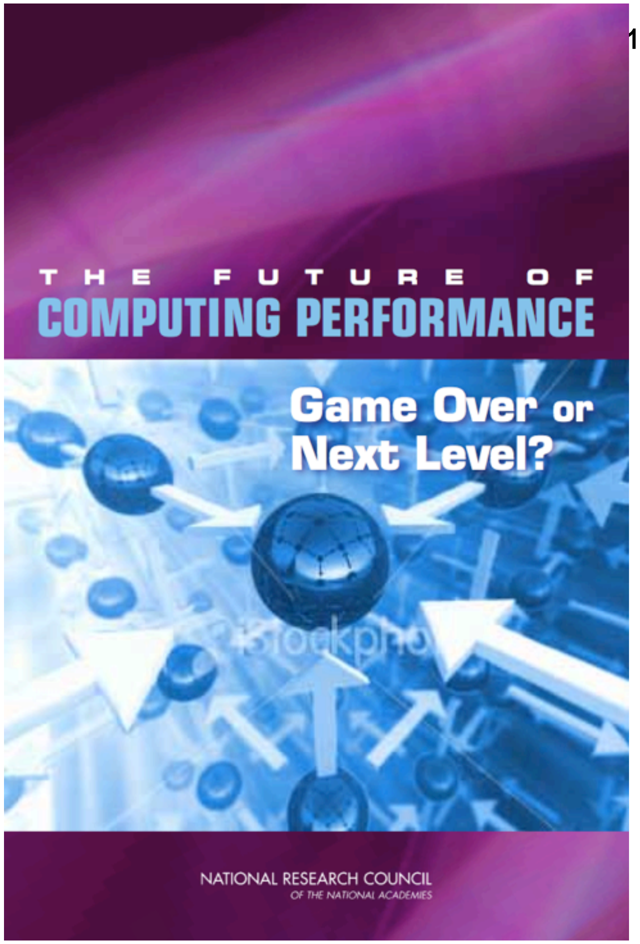

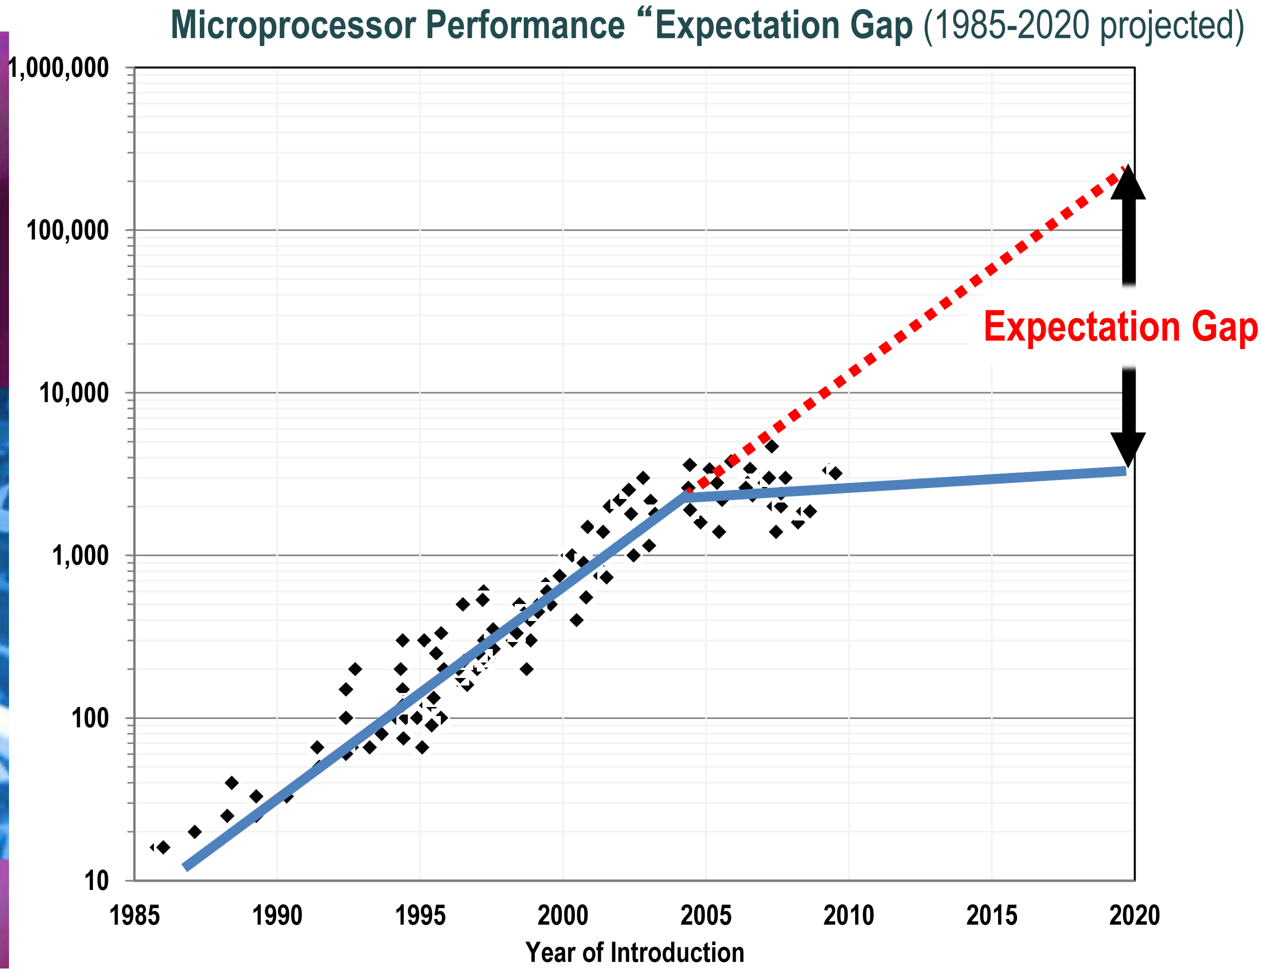

Industry motivated, path forward is unclear

National Research Council (NRC) – Computer Science and Telecommunications Board (CSTB.org)

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC Workshops Confirm this Gap for Science

New Processors Means New Software

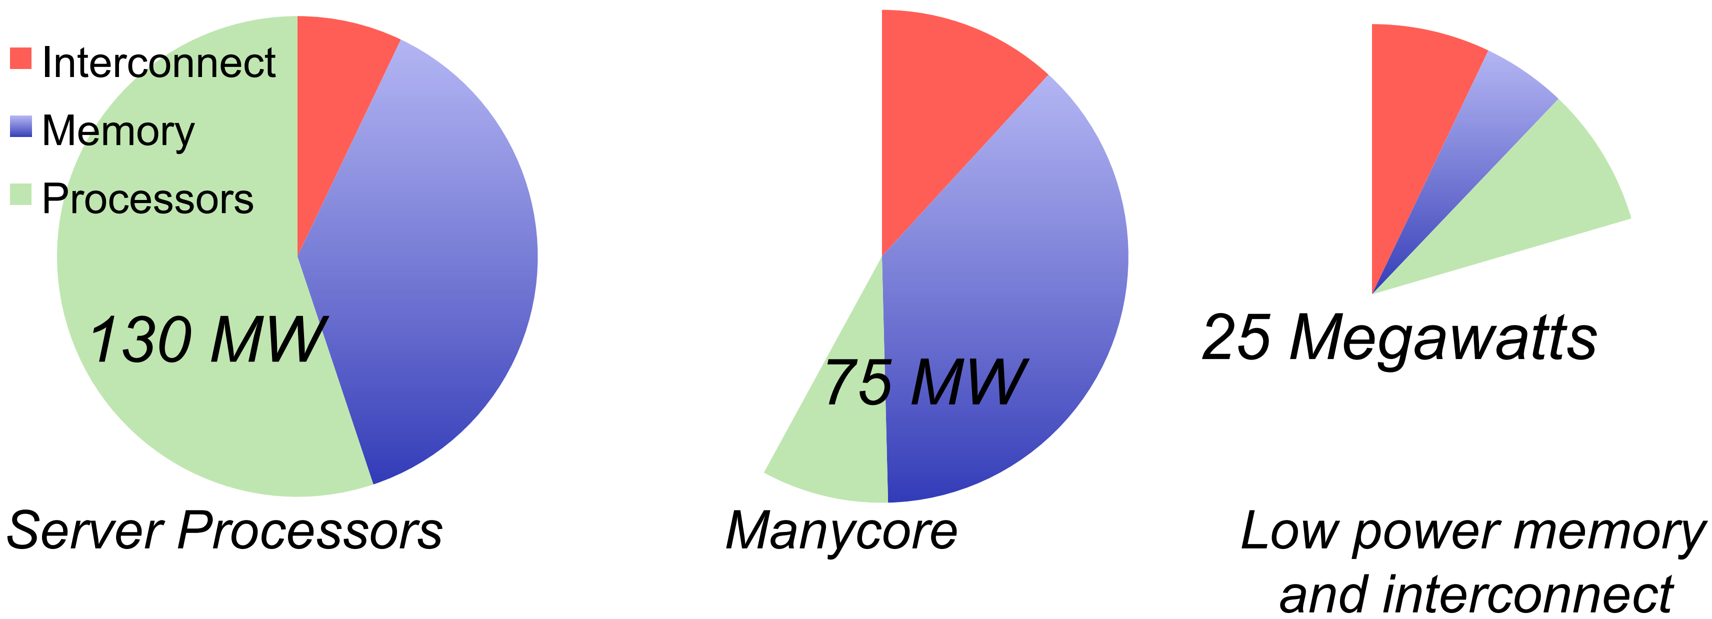

- **Exascale will have chips with thousands of tiny processor cores, and a few large ones**
- **Architecture is an open question:**
  - sea of embedded cores with heavyweight “service” nodes
  - Lightweight cores are accelerators to CPUs
- **Low power memory and storage technology are key**

Exascale for Thousands of Users

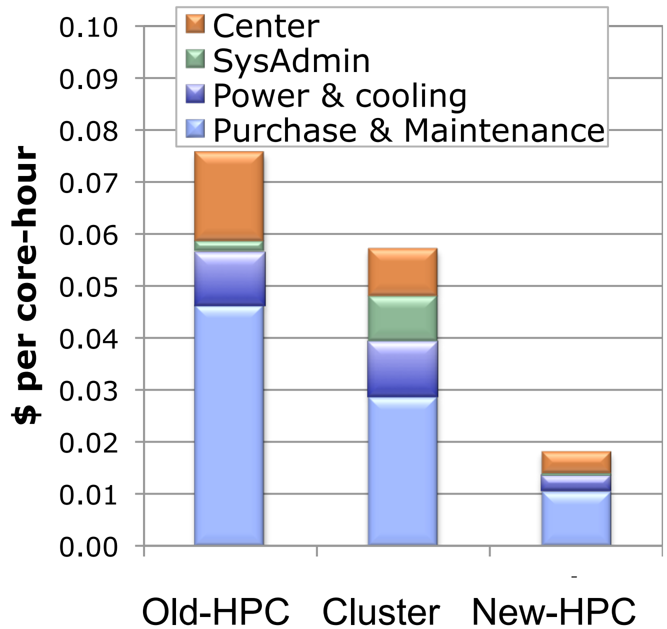

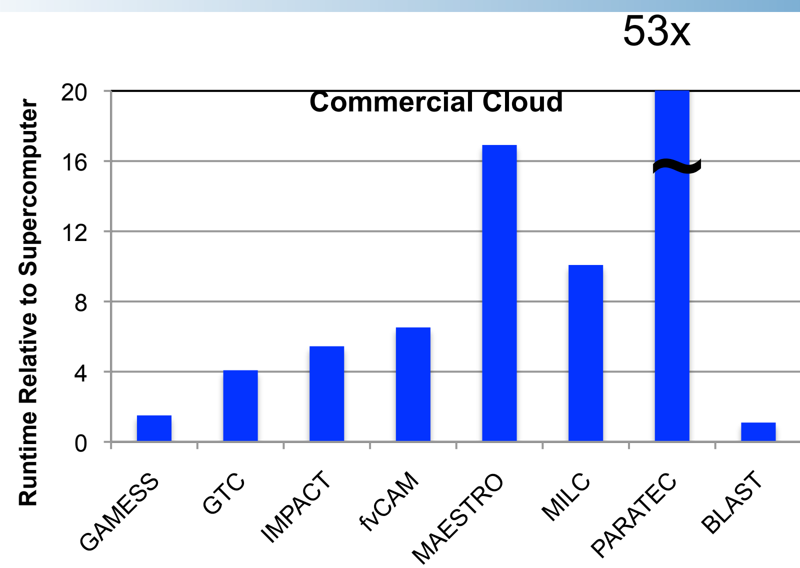

- **Minimum cost per core (or app flop) are:**
  - **Newest machines with largest core count per node (power)**
  - **Largest machine: amortize personnel costs**
- **But commercial clouds are slower & more expensive**
  - **Price not dropping with Moore's Law (18% in 5 years)**
  - **6-7x cost to buy NERSC compute + storage in 2011 cloud**

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC Roadmap

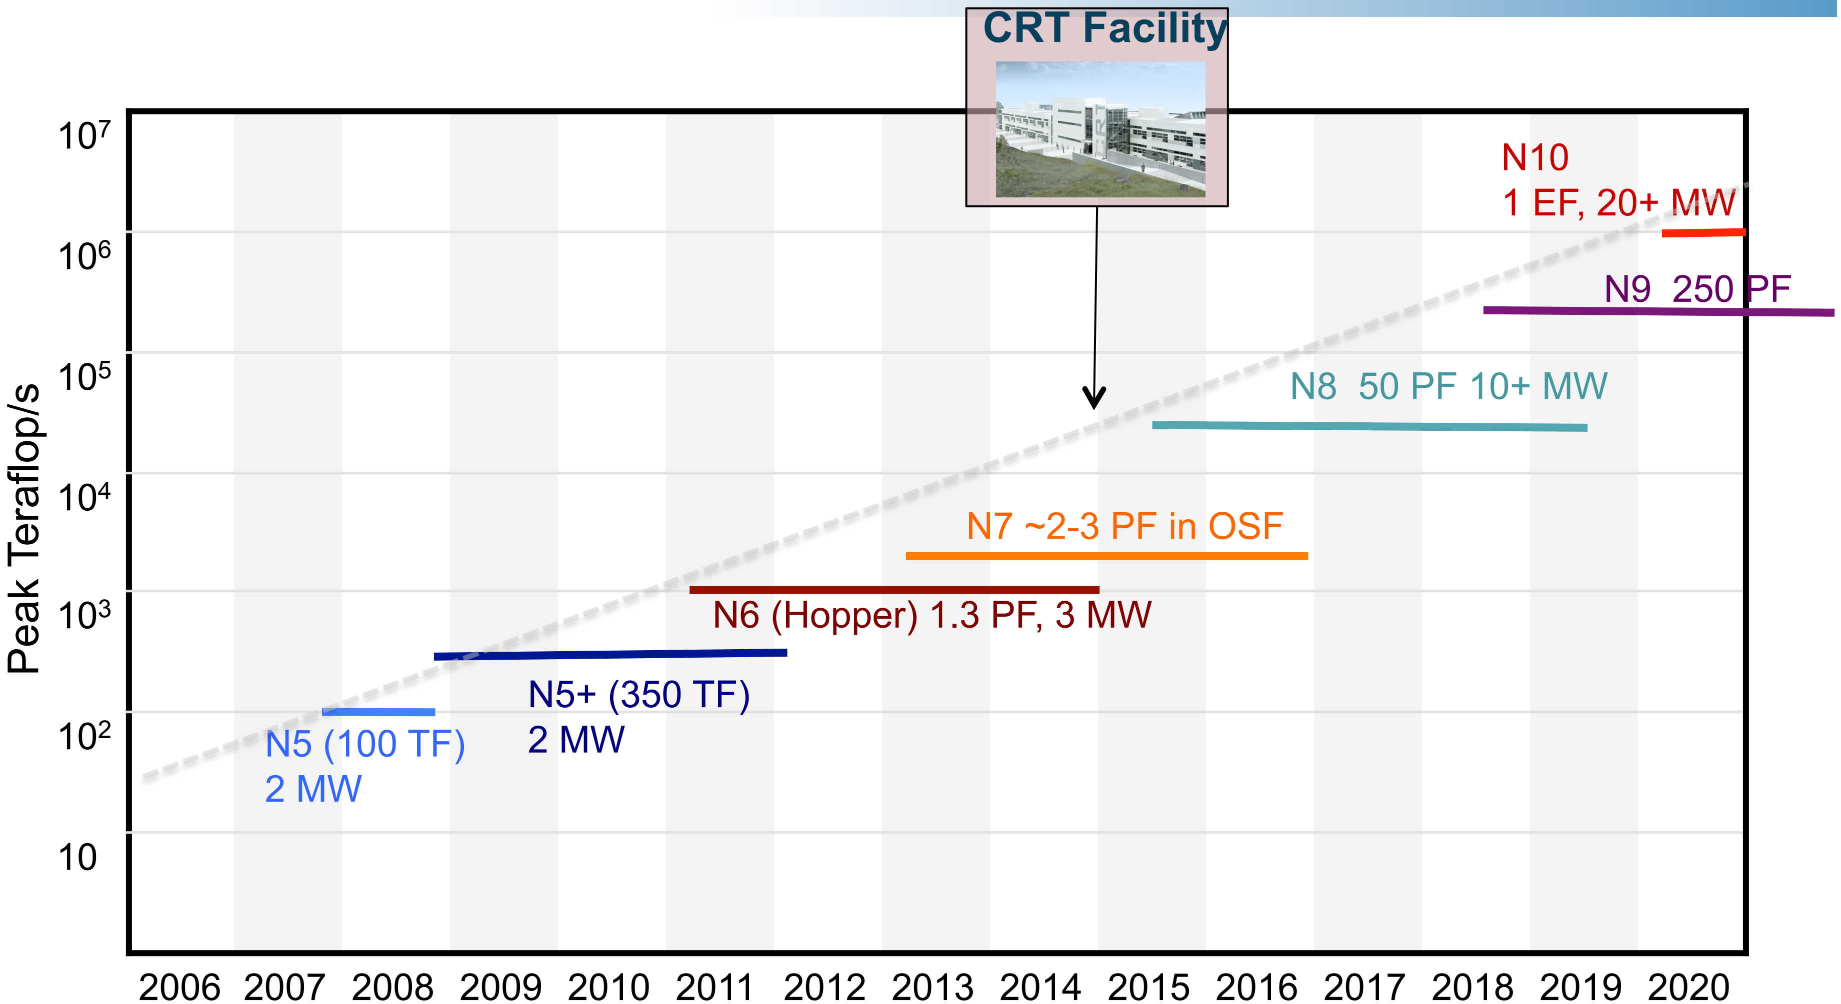

NERSC performance has traditionally grown at 10x every 3-4 years

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

DOE Facilities Require Exascale Computing

Astronomy

Particle Physics

Chemistry and Materials

Genomics

Fusion

Petascale to Exascale

- Petabyte data sets today, many growing exponentially
- Processing grows super-linearly
- Need to move entire DOE workload to Exascale

Data-Intensive Computing

- **Goal: Grow storage, transfer & analysis capability for DOE facilities**
  - Data generation exceeds storage and process rate
  - Needs energy efficient computing, memory & I/O
- **Berkeley Lab provides essential expertise**
  - ESnet *bandwidth reservations* for large data sets
  - NERSC has dedicated computing for HEP/NP/BER: JGI, STAR, Alice, Atlas, Daya Bay, IceCube, Planck
  - VisIt: Petascale visualization on 300 TB
  - FastBit algorithm: produced 1000x speedup in data analysis on Wakefields and Fusion
- Need for increased investments in
  - Data provenance & management
  - Data analytics algorithms, software, and systems

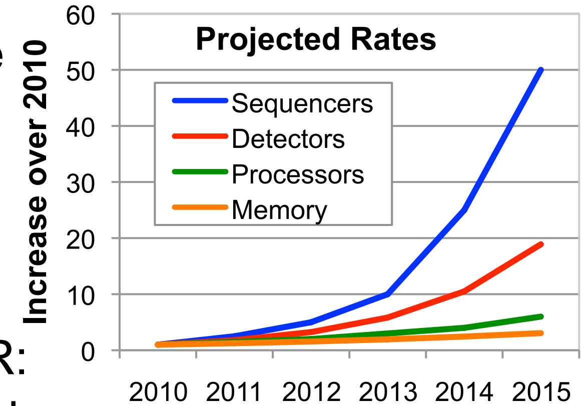

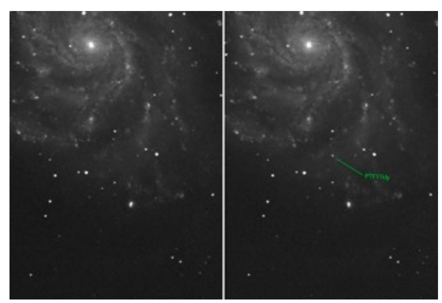

"Youngest nearby Supernovae discovered

Antihelium-4 discovery in STAR used ESnet and NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

ESnet: DOE's Leadership and Production Network

DOE Science Network:

- 72% annual traffic growth
- International collaborations
- Bandwidth reservations and monitoring

Advanced Networking Initiative (ANI):

- 100 Gbps network
- Contract to Internet2 signed in July
- Demo planned in late 2011
- Separate network research testbed serving 17 projects

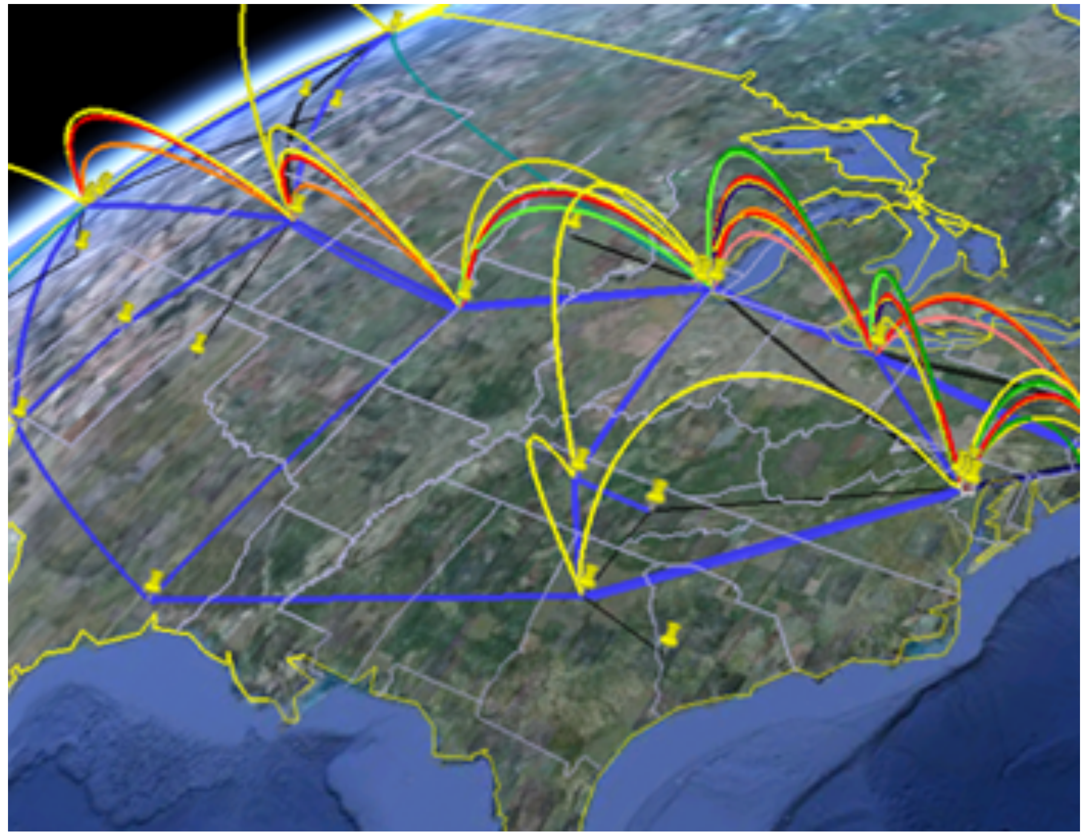

ESnet+ANI, DOE will be the world leader in networking for science.

Data Challenges

- Streaming data sets
- Complex workflows (not batch jobs)
- Real-time constraints
- Large science communities with shared access to data
- Different machine requirements:
  - Frontend processing (massively serial)
  - Information retrieval (graphs)
  - Storage systems (map reduce)

LBNL Organization

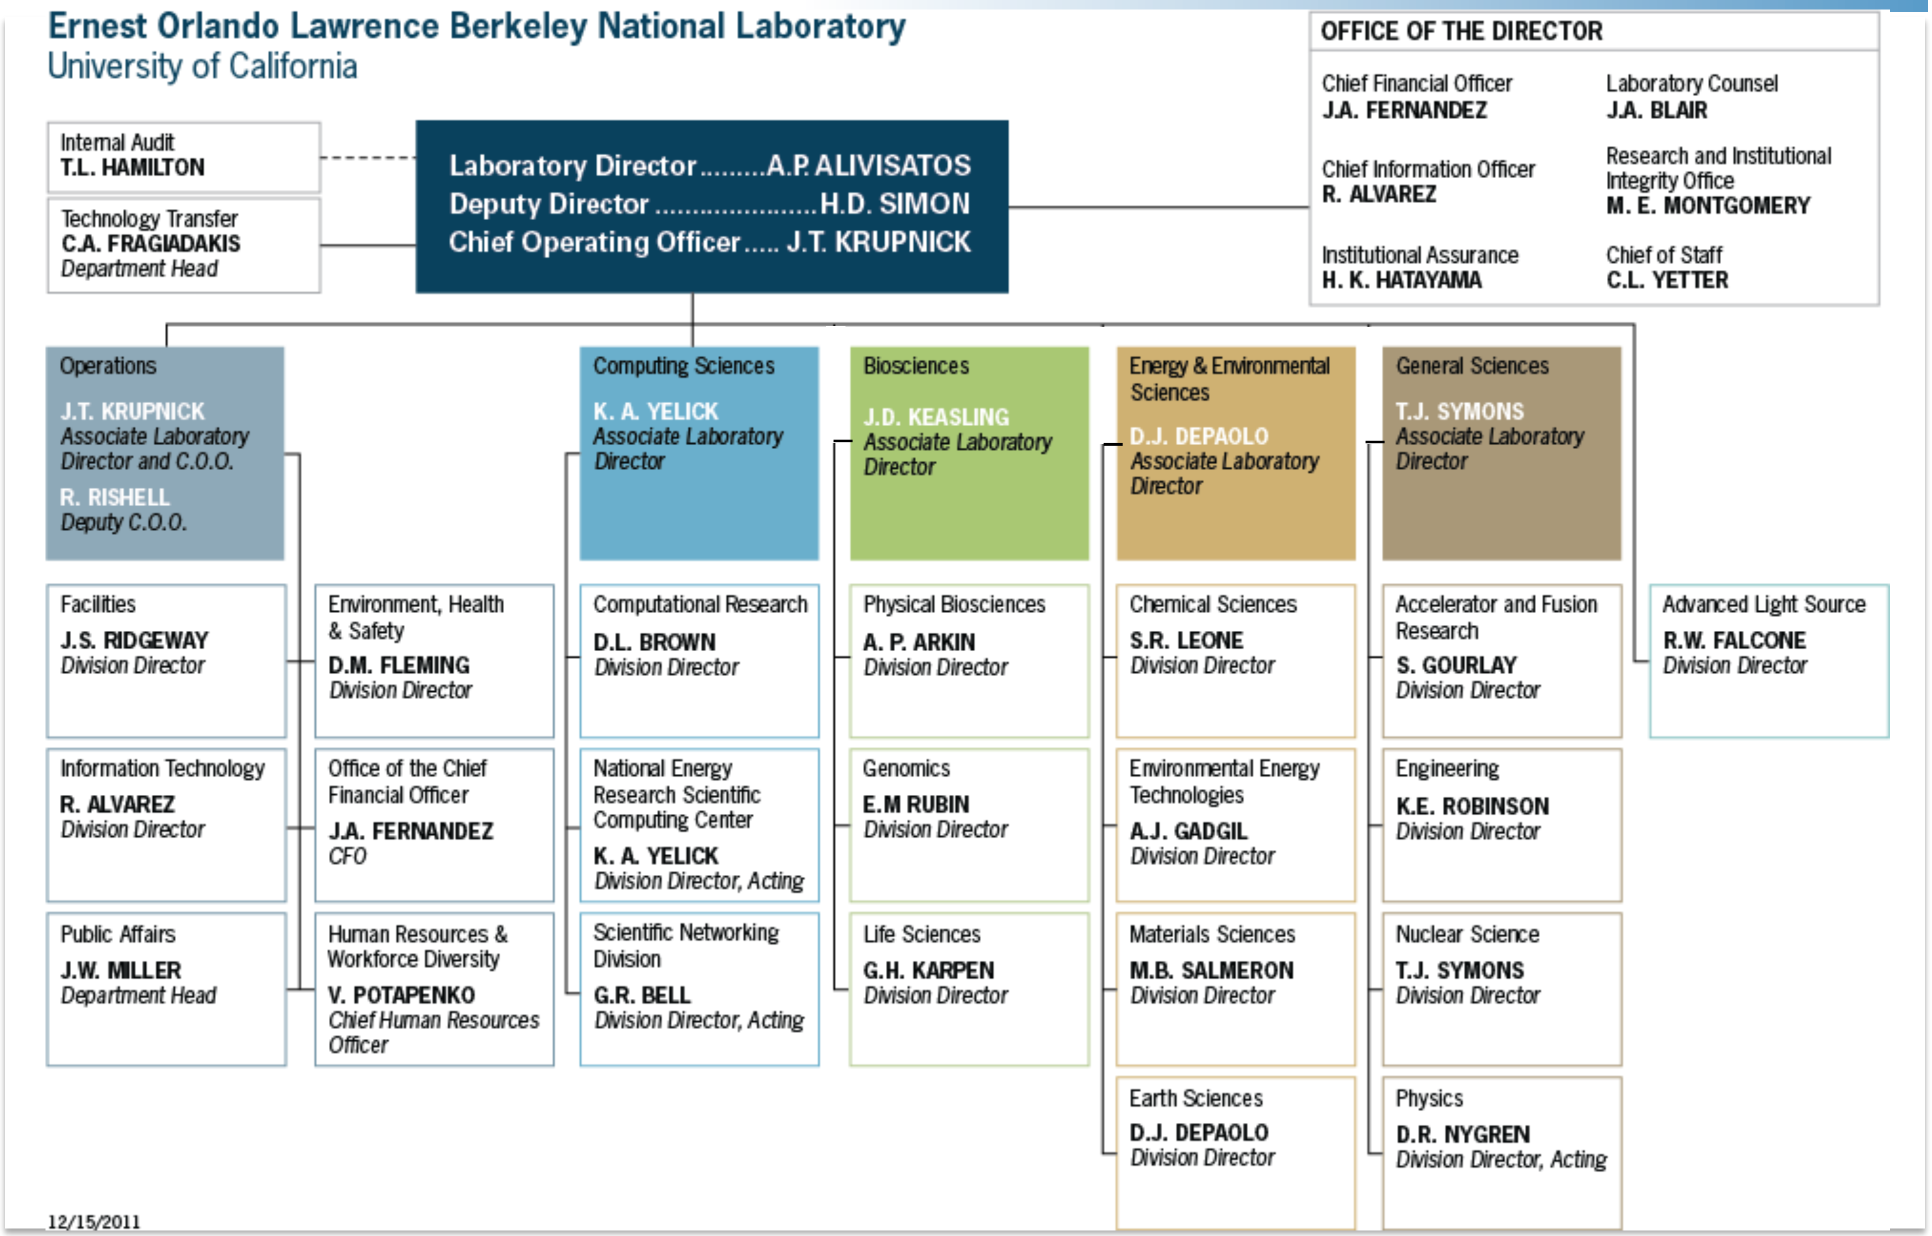

LBNL VIEW OF SCIENCE IN 2020

Societal needs for technical solutions to energy and environment problems will intensify

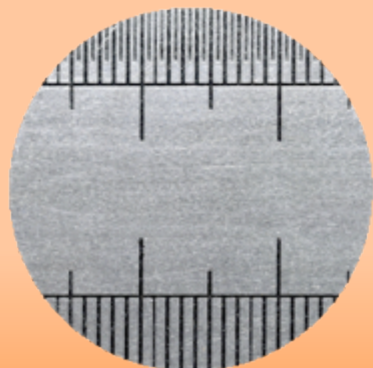

Measurement tools will open realms of inquiry

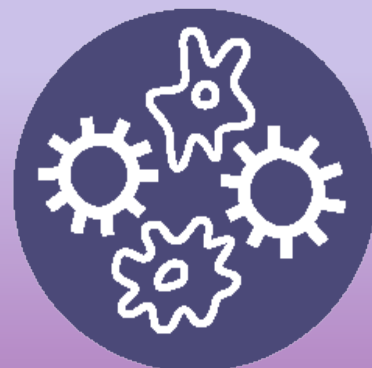

Biology revolution will impact other disciplines

Reliance on computation will grow with data

Fundamental discoveries in basic science will strengthen our foundation and understanding.

RKELEY LAB

NERSC Strategy: *Science First*

- **Support computational science:**
  - Provide effective **machines that support fast algorithms**
  - Deploy with **flexible systems software** to run a **broad range** of applications
  - Help users with **expert services**
  - Develop **tools** to make systems more accessible
- **NERSC future priorities are driven by science:**
  - Increase application capability: **“usable Exascale”**
  - **Simulation and data analysis** of simulated and experimental data

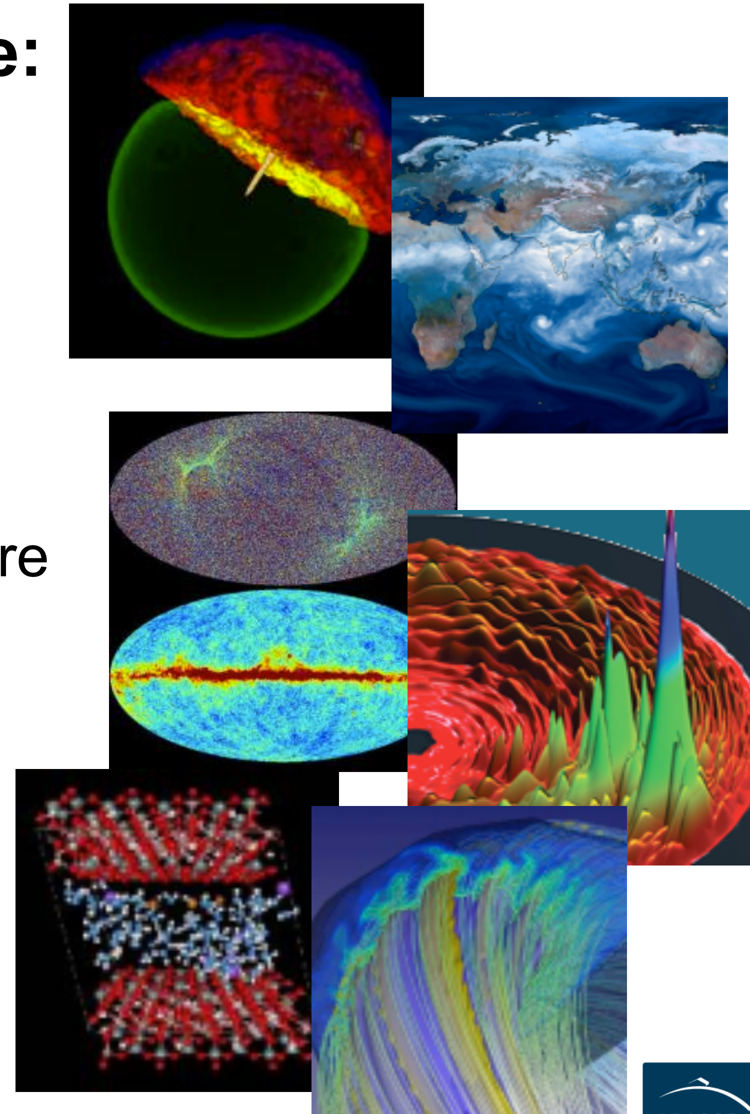

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory