Leadership Computers as Instruments of Discovery

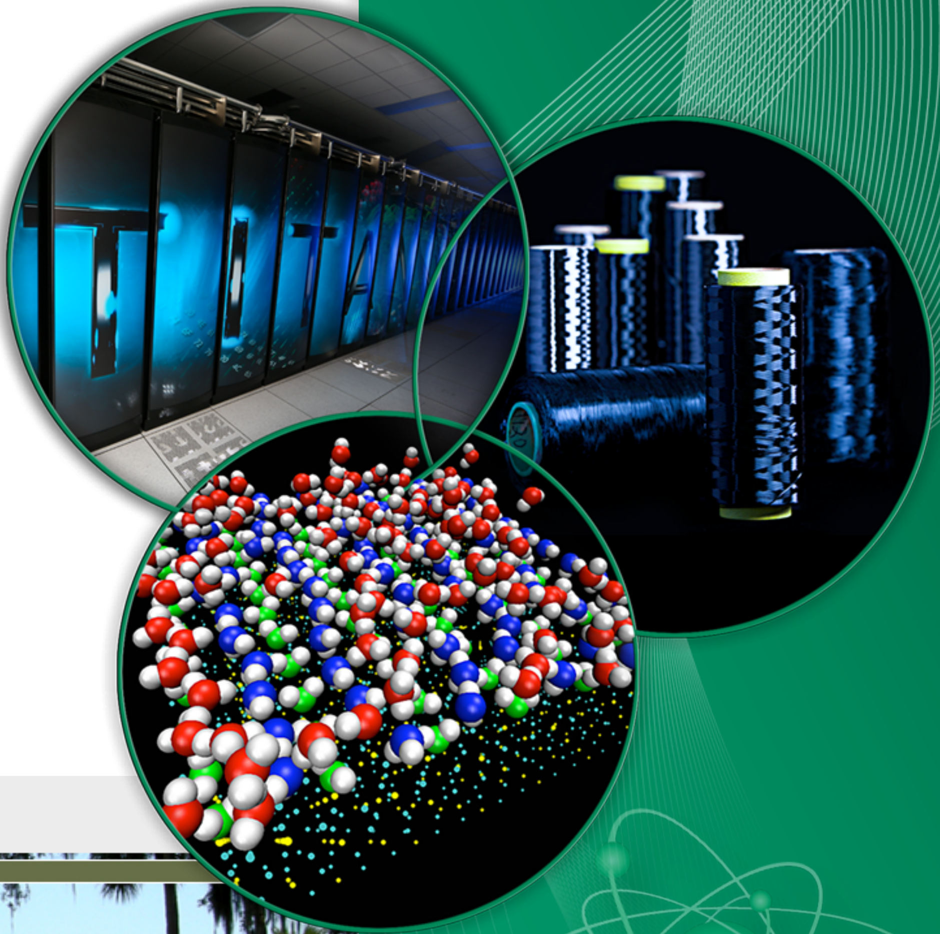

Jeff Nichols
Associate Laboratory Director
Oak Ridge National Laboratory

March 25 - 28, 2013

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE U.S. DEPARTMENT OF ENERGY

DOE is a world leader in HPC
- leaders in architecting, acquiring and deploying computers used as instruments of discovery

- A record of leadership in computing performance
- Thousands of users from government labs, universities, and industry
- Forefront computing facilities
  - Used to solve mission problems in nuclear defense, science, and engineering
  - Enabling prize-winning science (e.g., Nobel), engineering solutions, and thousands of publications

TOP500
November 2012

Titan at ORNL
(#1, 17+ PF)

Sequoia at LLNL
(#2, 16+ PF)

Mira at ANL
(#4, 8+ PF)

Cielo at LANL/SNL
(#18, 1+PF)

Hopper at LBNL
(#19, 1+PF)

Delivering advances through mission-focused partnerships

Advanced research in computing

- Computer science for systems and software
- Applied mathematics for modeling and analysis

Scientific partnerships for application development

- SciDAC: Exemplar of interdisciplinary work in computational science
- Co-Design: Partnerships for applications and major technology transition
- ASC: Integrating element of stockpile stewardship

Partnerships with vendors

- NRE investments to design future systems
- Research for technology innovation (FastForward, Design Forward)

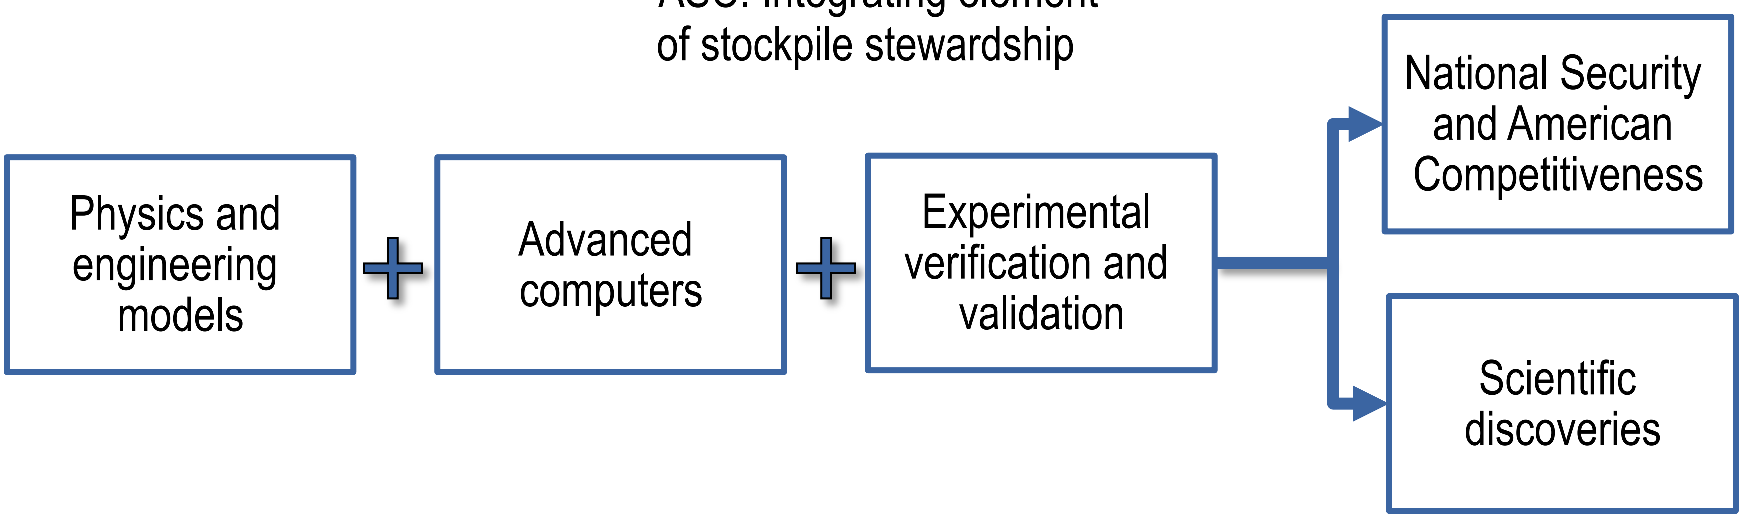

The nature of research has changed

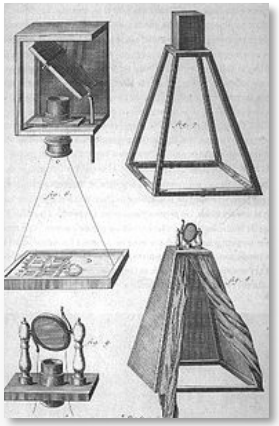

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi G\rho}{3} - K\frac{c^2}{a^2}$$

1,000 years ago

- Experimental
  - Description of natural phenomena
  - Experimental methods and quantification

Last 500 years

- Theoretical
  - Formulation of Newton's laws, Maxwell's equations, …

Last 50 years

- Computational
  - Simulation of complex phenomena

Today

- Data
  - Distributed communities unifying theory, experiment, and simulation with massive data sets from multiple sources and disciplines

Data joins experiment, theory, and simulation in the discovery process

- Sources:
  - Experiments (e.g., science facilities)
  - Output from simulations
  - Human activity (e.g., social media)
  - Sensors (e.g., satellite data)
- Potential applications:
  - Medicine,
  - Economics
  - Science
  - Disaster recovery
  - National security

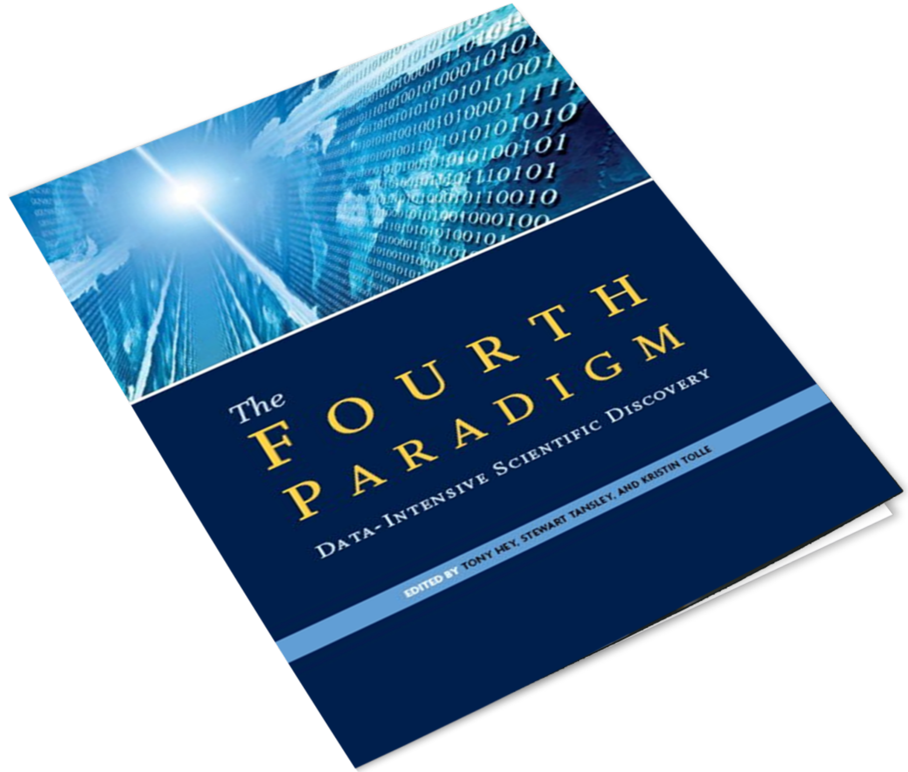

Bringing many types of disparate data to bear on a particular discovery process or mission outcome is one of the new frontiers of science and technology

Big data = volume, variety, velocity

Slide courtesy of Jeff Nichols

The data explosion

Every second one hour of video is uploaded to YouTube

Experiments

Simulations

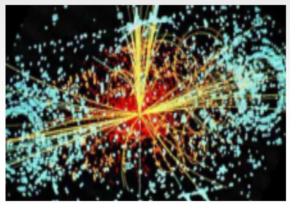

Archives

Social Media

Sensors

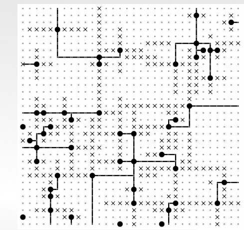

Information Technology

The Challenge
Enable Discovery

Petabytes
Exabytes
Zettabytes

The Response

Deliver capability to mine, search, and analyze this data in near real time

Science itself is evolving

Slide courtesy of Jeff Nichols

The “Instrument” integrates both data analysis and simulation capabilities

- Simulation and data are critical to DOE
- Both need more computing capability
- Both have similar HW technology requirements
  - High bandwidth to memory
  - Efficient processing
  - Very fast I/O
- Different machine balance may be required

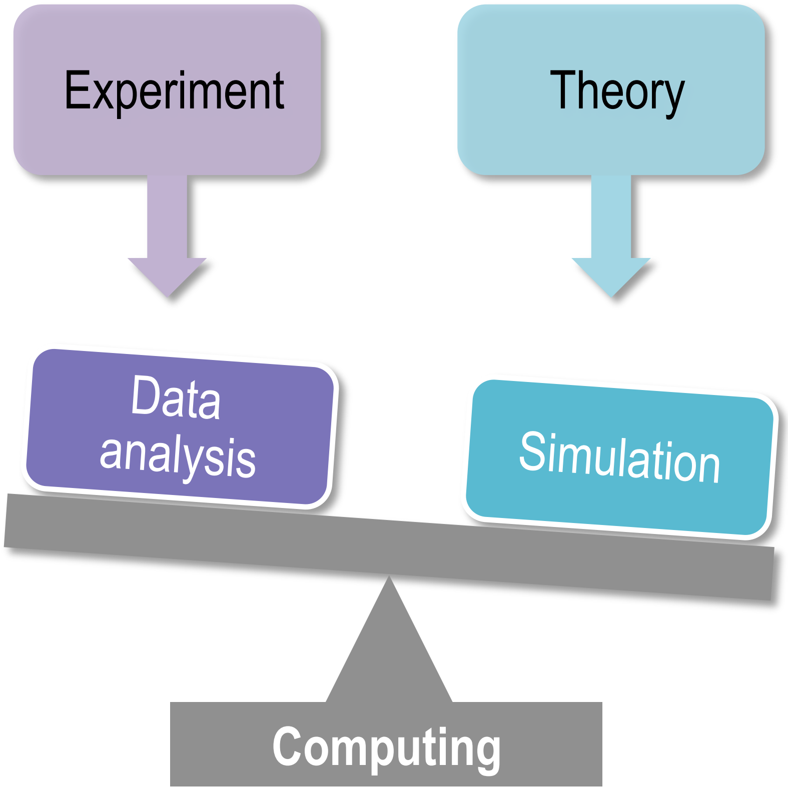

Big data
Analyzing and managing large complex data sets from experiments, observation, or simulation and sharing them with a community

Simulation
Used to implement theory; helps with understanding and prediction

Materials genome

JCESR | JOINT CENTER FOR ENERGY STORAGE RESEARCH

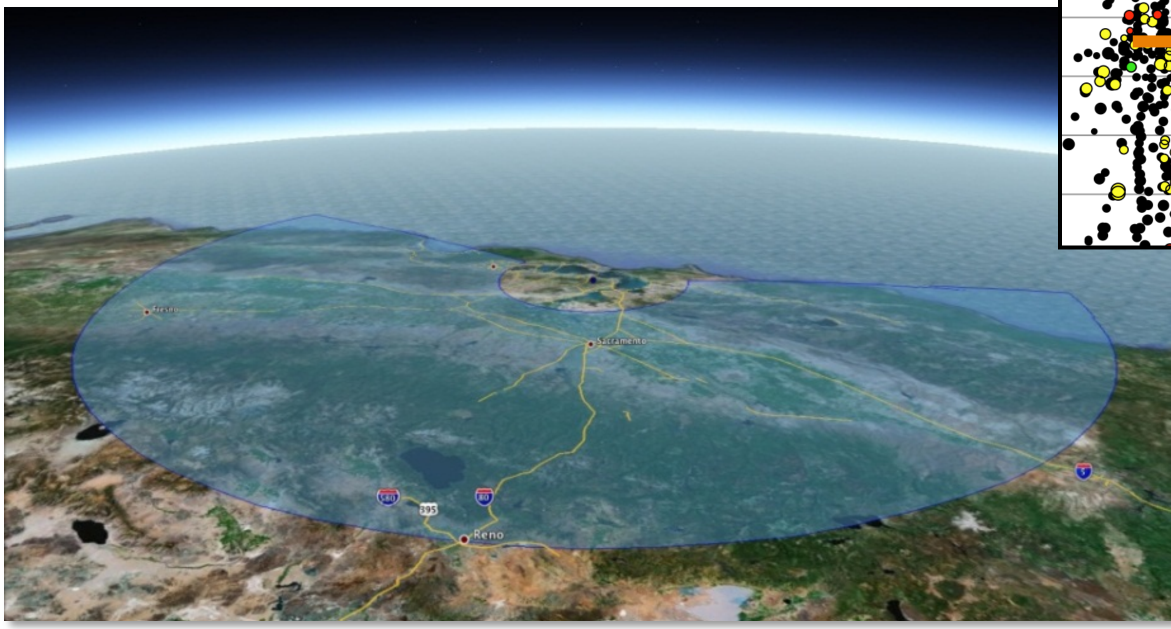

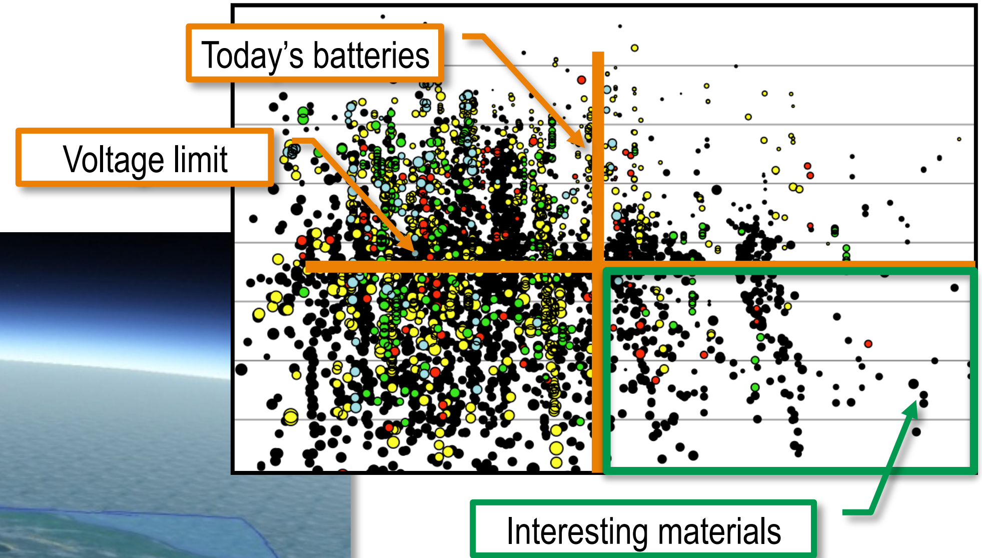

JCESR goals by 2018:

- Increase energy density (70 miles → 350 miles)
- Reduce battery cost per mile ($150 → $30)

Materials Project computing demands: 1000× today

- **Goals:** Increase U.S. competitiveness; cut in half 18 year time from discovery to market
- Key to energy storage hub

- **Simulations:** Thousands of simulations to screen potential materials; millions needed for new materials, extreme environments, etc.

- **Data:** Web-searchable data from simulations accessed by >3,000 users in first year; one-fourth from industry

Combustion

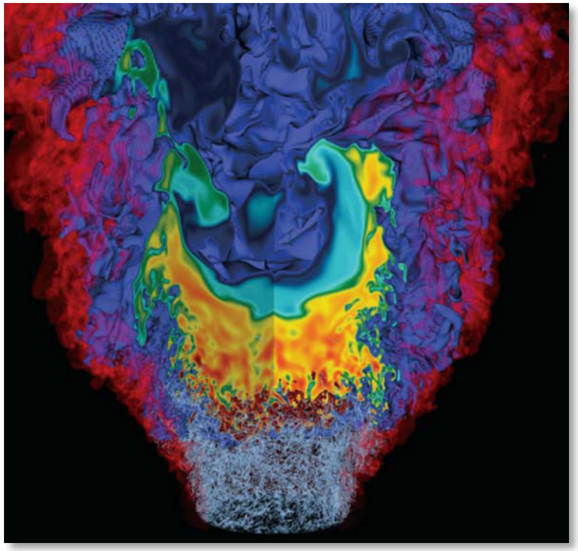

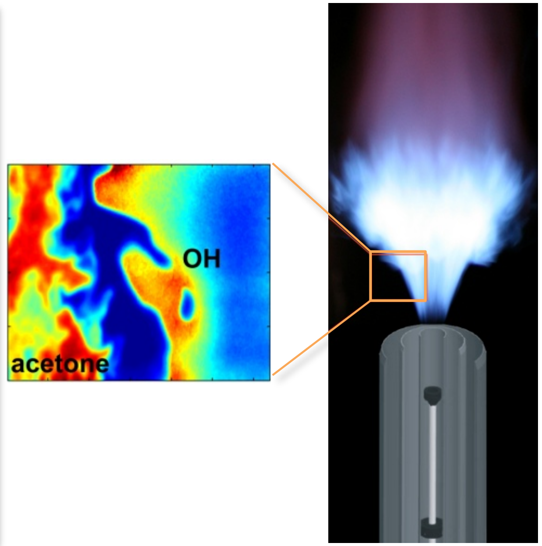

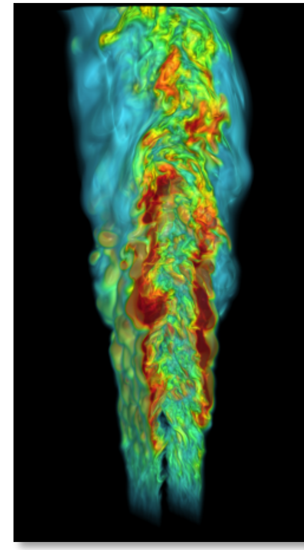

HPC facilities are needed to improve understanding of combustion

- **Goal:** 50% improvement in engine efficiency
- **Simulations:** Exascale needed for higher resolution, alternative fuels, design
- **Data:** Simulations help explain experimental data

Fundamental Science: Astronomy

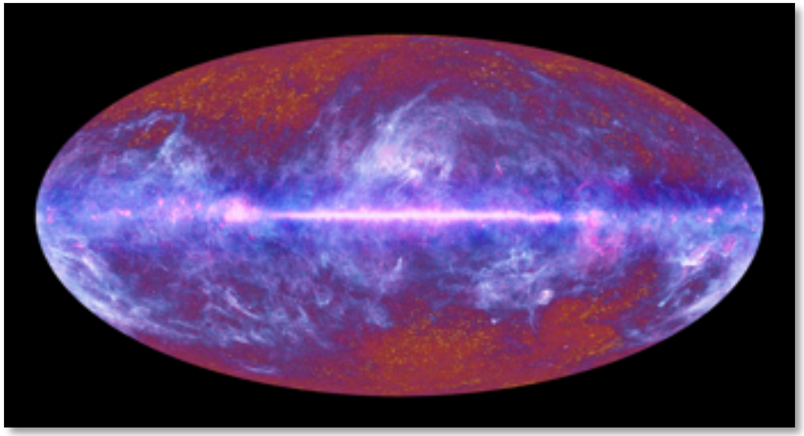

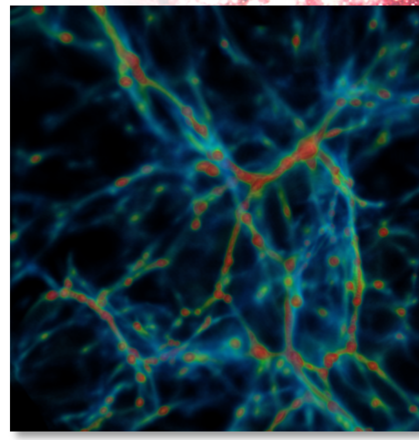

HPC facilities are needed to simulate missing data and analyze observational data

- **Goals:** Fundamental science of dark energy, the formation and ongoing expansion of the universe
- **Simulations:** Explain what is missing (behind Milky Way) and seen (how a supernova looks from earth)
- **Data:** Machine learning to identify supernova; massive networking/ analysis to measure cosmic background

Climate change analysis

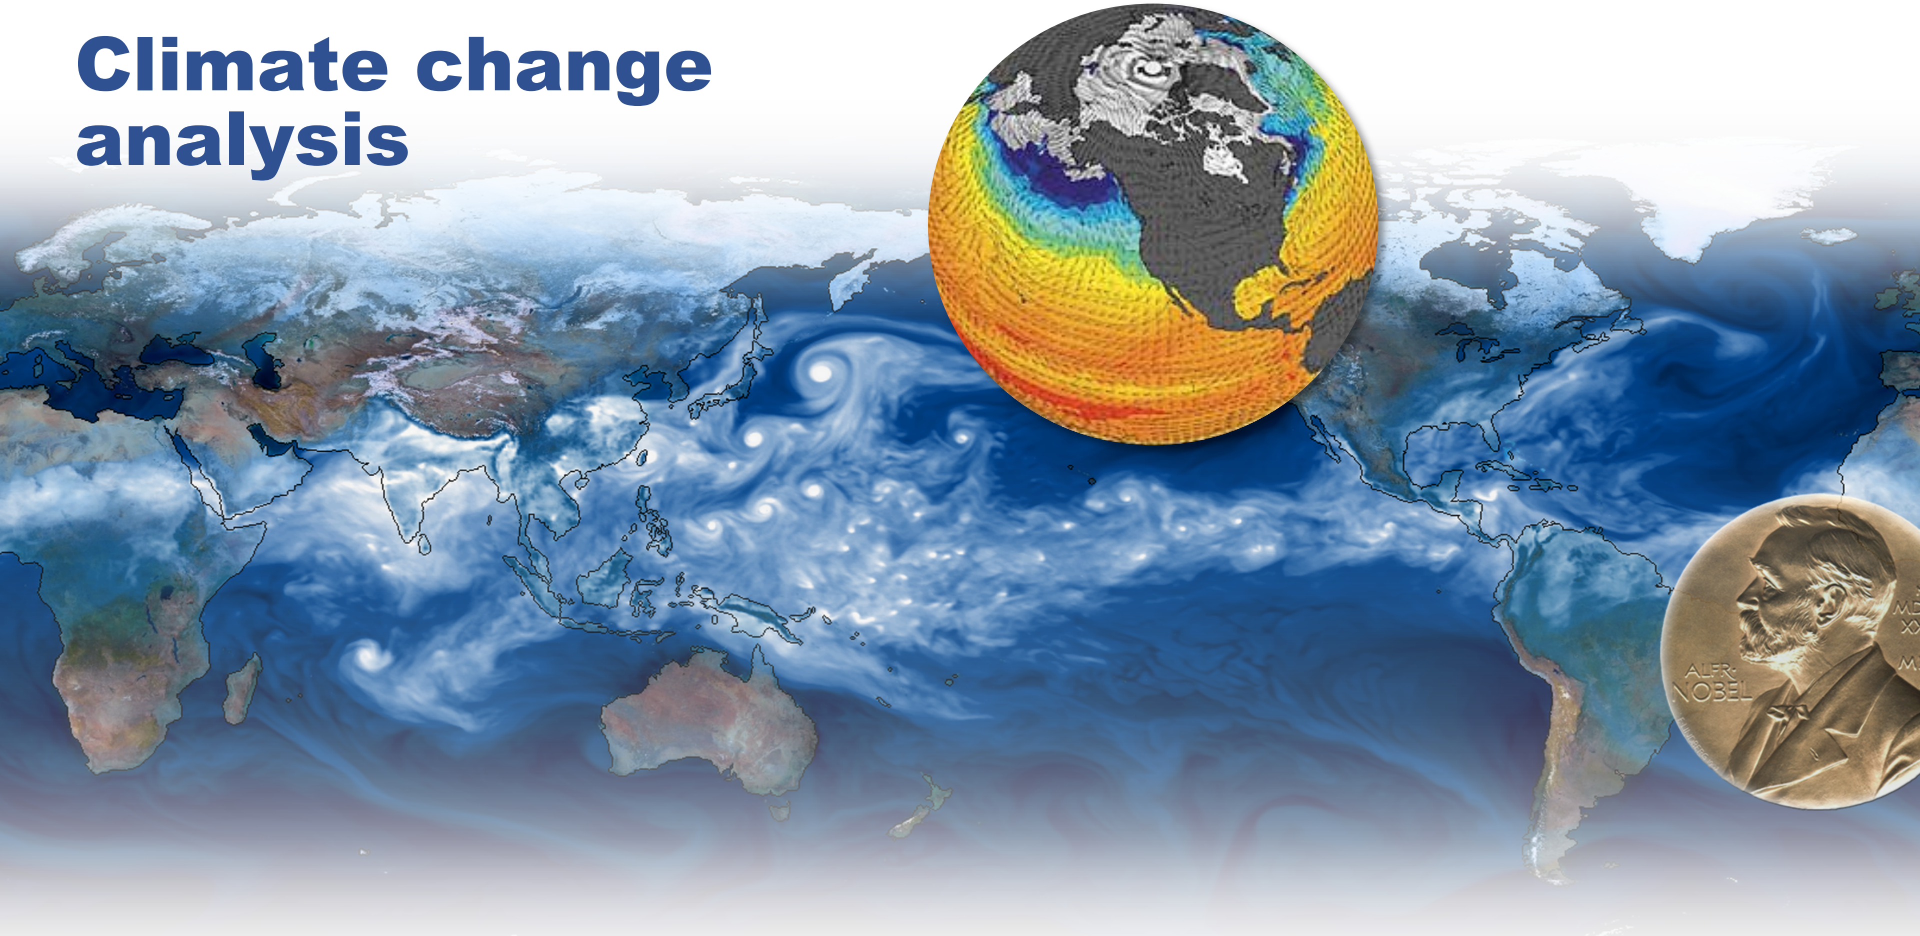

HPC facilities are needed to analyze climate data and understand impacts of change

- **Goals:** Resolve clouds, identify tipping points, quantify confidence, quantify agriculture and economics
- **Simulations:** 100× computing needed for more resolution and more scenarios
- **Data:** Reanalysis of 100+ years of weather data; 100× more computing for this "in situ" analysis

Industry benefits from advanced computing instruments innovated at DOE labs

Auto industry success demonstrates how HPC is used to compete and win in 21st century world economy

- Bring new products to market faster combining use of more advanced modeling for testing
- Reduce modeling time by allowing teams to try multiple scenarios and determine optimal solutions early in a program
- Permit use of larger, more detailed models, reducing design time
- Simulations that took months can now be run in days
- HPC: Key enabler for product development to deliver quality products faster, meeting the time-to-market customers expect

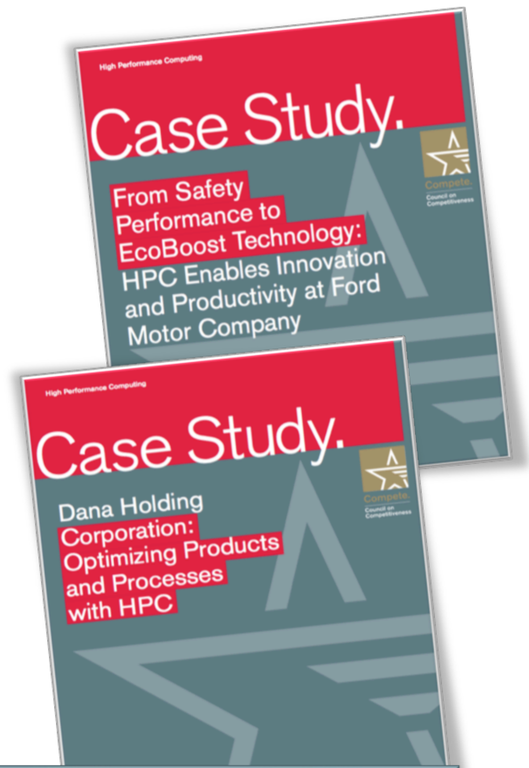

But continued American dominance in developing advanced HPC environments and in using these effectively is essential to maintaining competitive advantage

When industry engages with DOE lab expertise, the HPC advantage grows

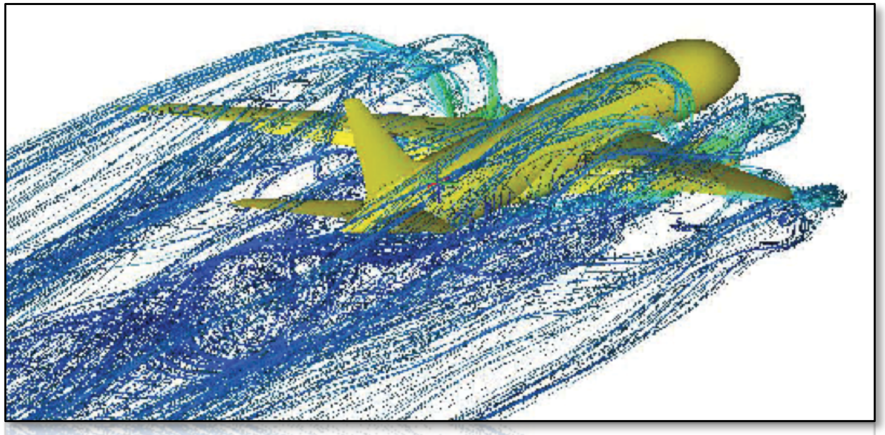

"The world has become a really smart place … we no longer enjoy the advantage we had in the 1950s and 1960s . . . Today our competitors are just as capable of making innovative discoveries as we are."

". . . when we were designing the 767 . . ., we built and tested about 77 wings. By using supercomputers to simulate the properties of the wings on recent models such as the 787 and the 747-8, we only had to design seven wings, a tremendous savings in time and cost, especially since the price tag for wind tunnel testing has skyrocketed over the past 25 years."

"Our work with supercomputers allows us to get a better product out the door faster, which makes us more competitive—it's as simple as that."

HPC drives low-cost aerodynamic designs that impact environment and economy

Powerful computer platforms + Advanced CFD codes and DOE expertise → Identification of critical drag producing regions

NASA NAVISTAR

Wind tunnel tests at NASA Ames validate simulation predictions

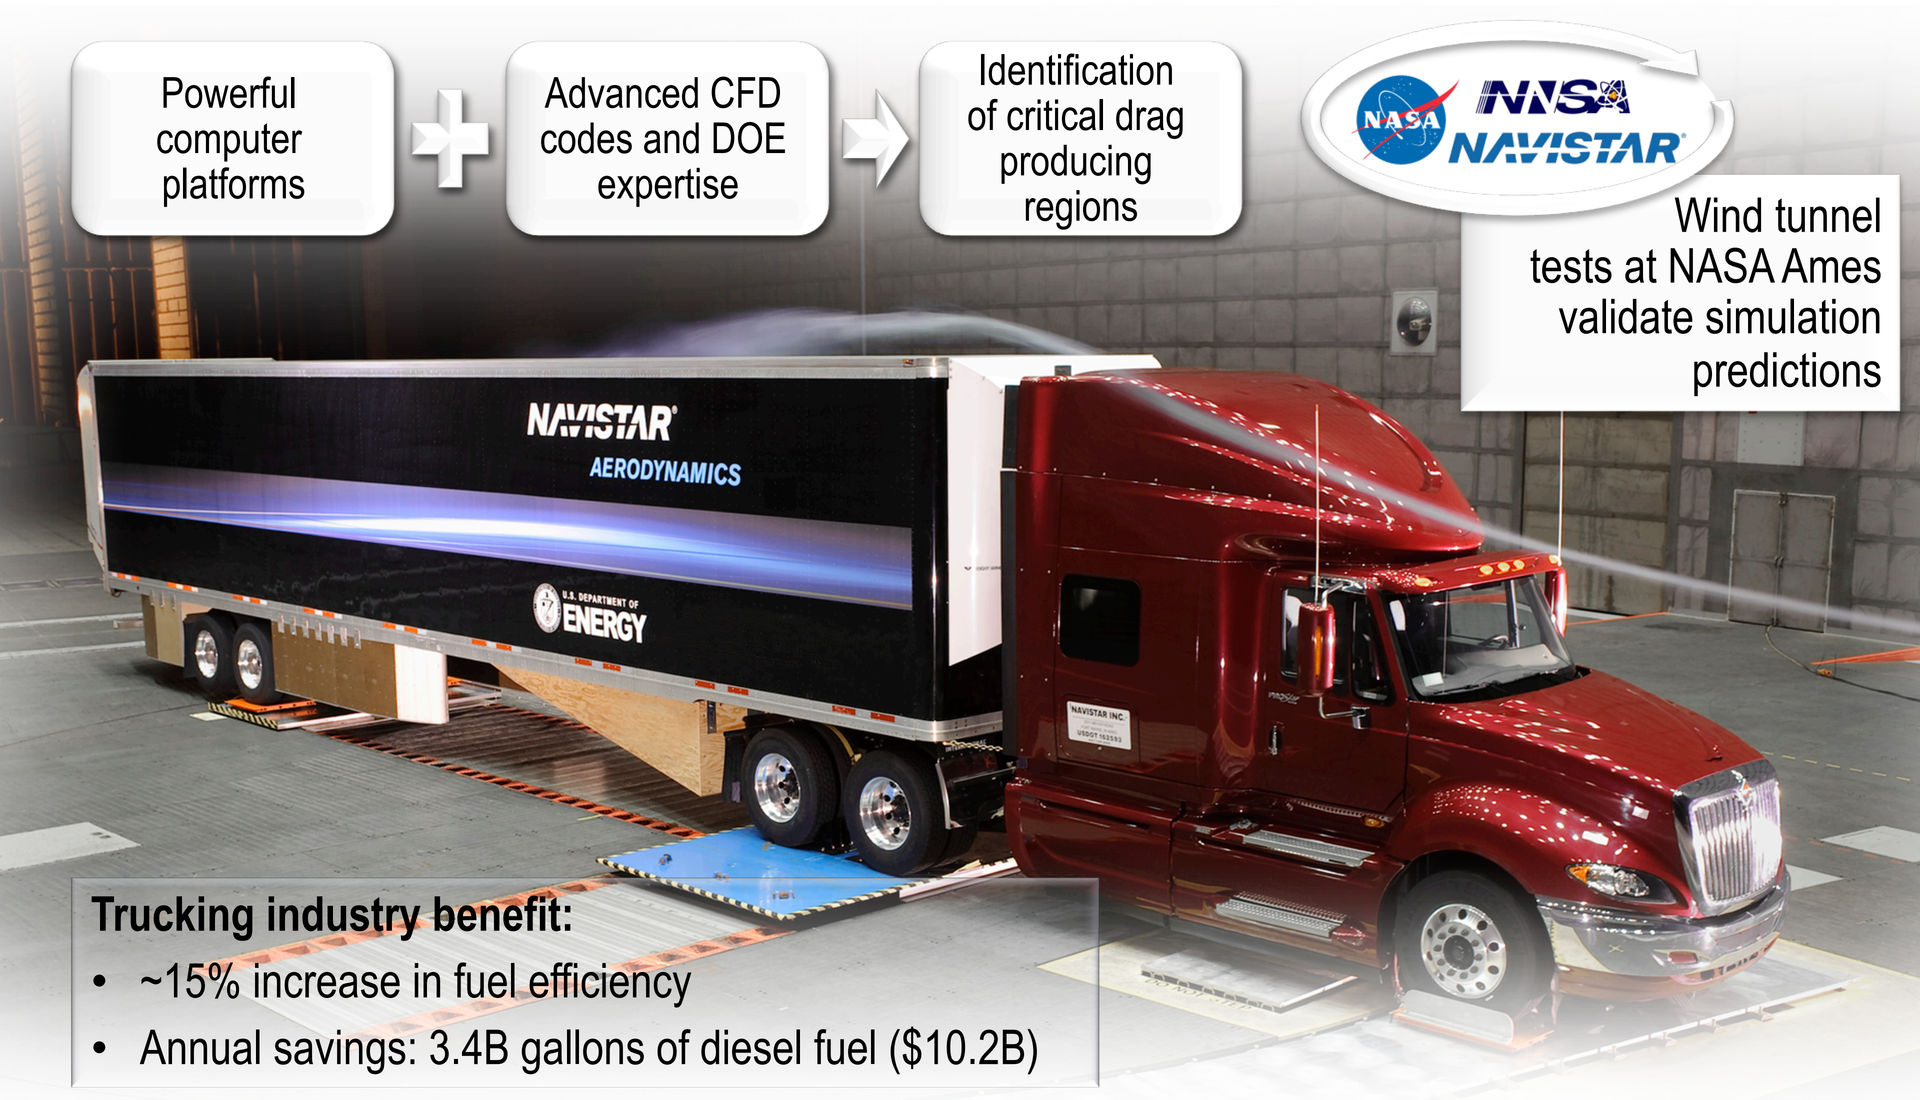

Trucking industry benefit:

- ~15% increase in fuel efficiency
- Annual savings: 3.4B gallons of diesel fuel ($10.2B)

https://publicaffairs.llnl.gov/news/news_releases/2010/NR-10-02-08.html

HPC innovation through DOE–industry partnerships

United Technologies Research Center	RAMGEN POWER SYSTEMS	GM	FM Global	GE	Ford
Catalysis	**Design innovation**	**Gasoline engine injector**	**Industrial fire suppression**	**Turbo machinery efficiency**	**Underhood cooling**
Demonstrated biomass as a viable, sustainable feedstock for hydrogen production for fuel cells; showed that nickel is a feasible catalytic alternative to platinum	Accelerates design of shock wave turbo compressors for carbon capture and sequestration	Optimization of injector hole pattern design for desired in-cylinder fuel-air mixture distributions (4–40× potential improvement in workflow throughput via 100s of ensemble simulations)	Developing high-fidelity modeling capability for fire growth and suppression; fire losses account for 30% of U.S. property loss costs	Simulated unsteady flow in turbo machinery, opening new opportunities for design innovation and efficiency improvements	Developed a new, efficient and automatic analytical cooling package optimization process leading to one of a kind design optimization of cooling systems

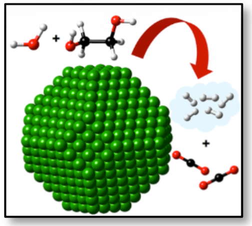

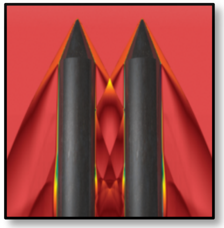

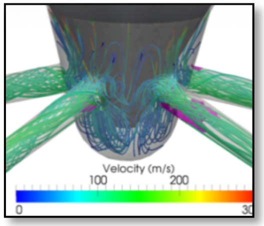

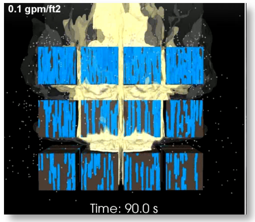

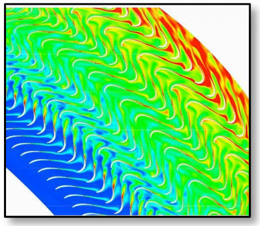

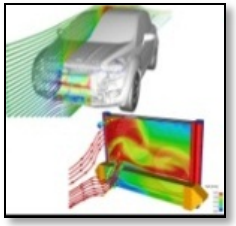

Summary:
Future success depends on U.S. control over HPC platform solutions and development of the most aggressive applications

- **Exascale** will provide the next quantum leap in modeling capability
- American industry will follow in the path defined by DOE/NNSA, once we have invented and hardened it
- This remaining American bastion of critical leadership can't be left to the Chinese (or even more friendly players)

"Results from our reliability technology partnership with Los Alamos will reduce P&G costs by $1.5B annually."

Mark Peterson, Procter & Gamble

"P&G would not have terascale informing our innovations and preserving jobs today if the National Laboratories had not been researching terascale hardware and software ten years ago. We have kept pace with their leading edge, and if the Laboratories' leading edge ever fails to keep advancing, P&G will fall behind every year as well."

Tom Lange, Procter & Gamble

DOE has significant expertise in delivering software solutions

- Applications are often built using scientific libraries
- Applications represent the investment of many millions of dollars and must all be moved to exascale technologies

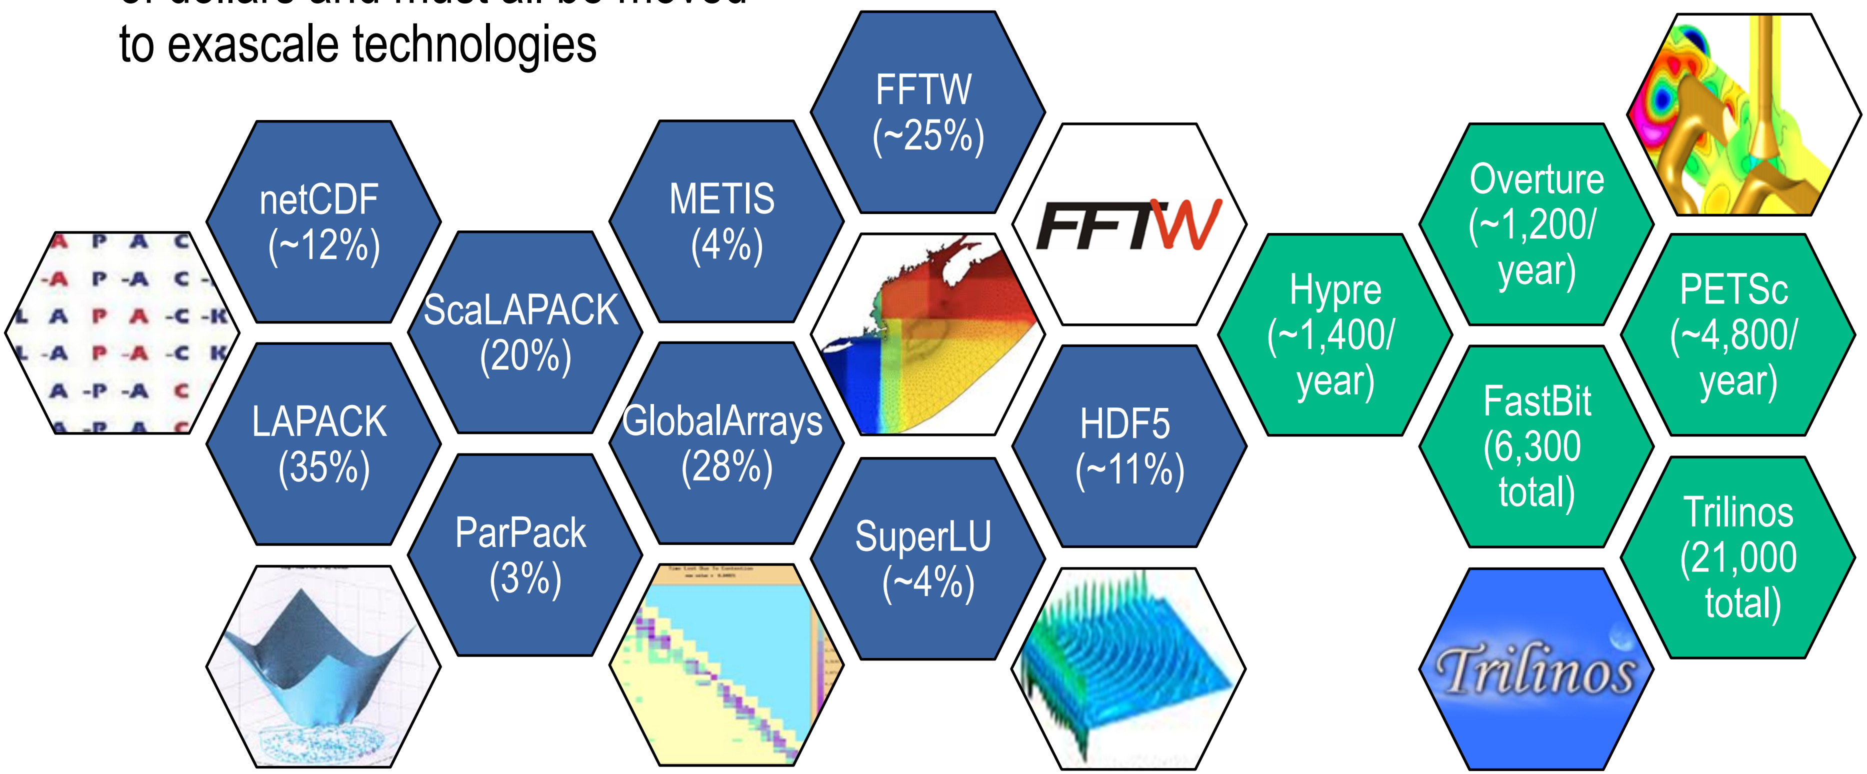

Libraries (percentage of NERSC projects using library)

Downloads

Moving HPC forward will require overcoming additional software challenges

New programming models	Dynamic runtime systems	Locality-aware, energy-efficient strategies	Interoperability
Manage parallelism and data movement through innovations in interfaces	Adapt to changing application goals and system conditions	Manage locality and minimize energy consumption	Facilitate interoperability across different HPC languages and interfaces, as well as cross-cutting execution models

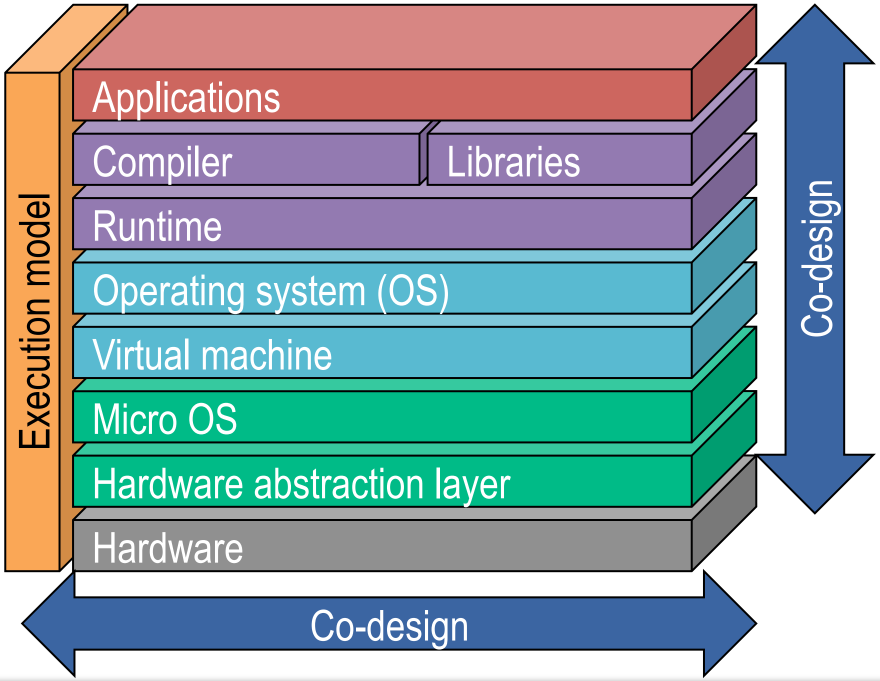

Moving HPC forward requires application developers and hardware architects to "co-design" solutions

- Scientific problem requirements guide computer architecture and system software design
- Technology capabilities and constraints inform formulation and design of algorithms and software
- Shared global perspective across the design-space establishes shared conceptual framework for co-design and interoperability
  - Parallelism
  - Latency
  - Dependability

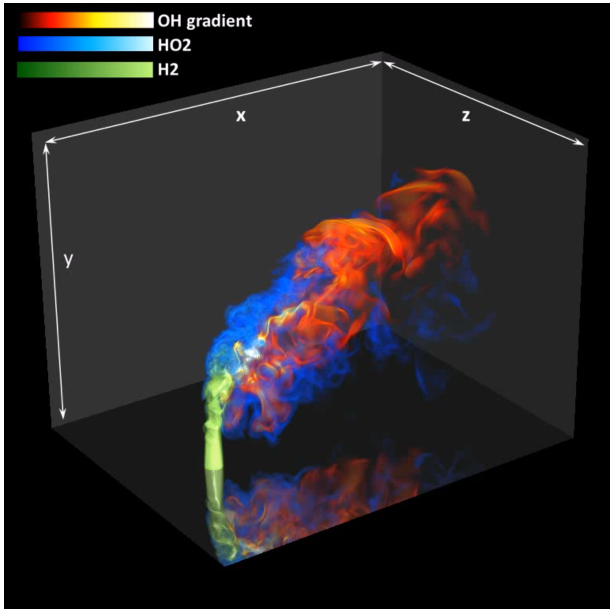

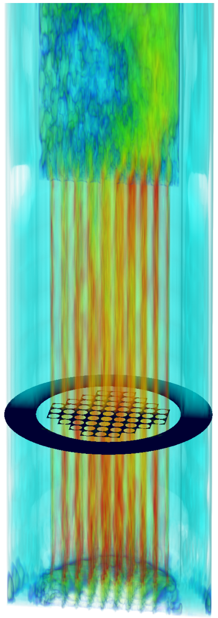

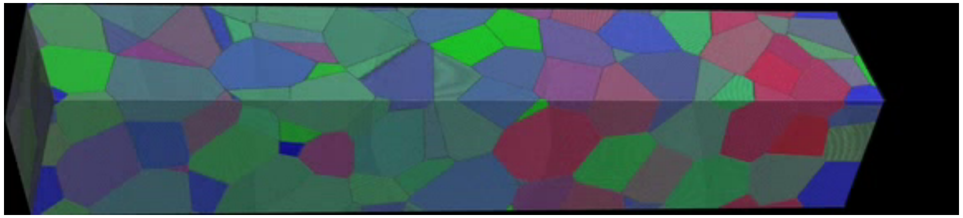

ORNL's "Titan" Hybrid System: Installation complete and undergoing acceptance testing

#1 TOP500® SUPERCOMPUTER SITES
Nov. 2012

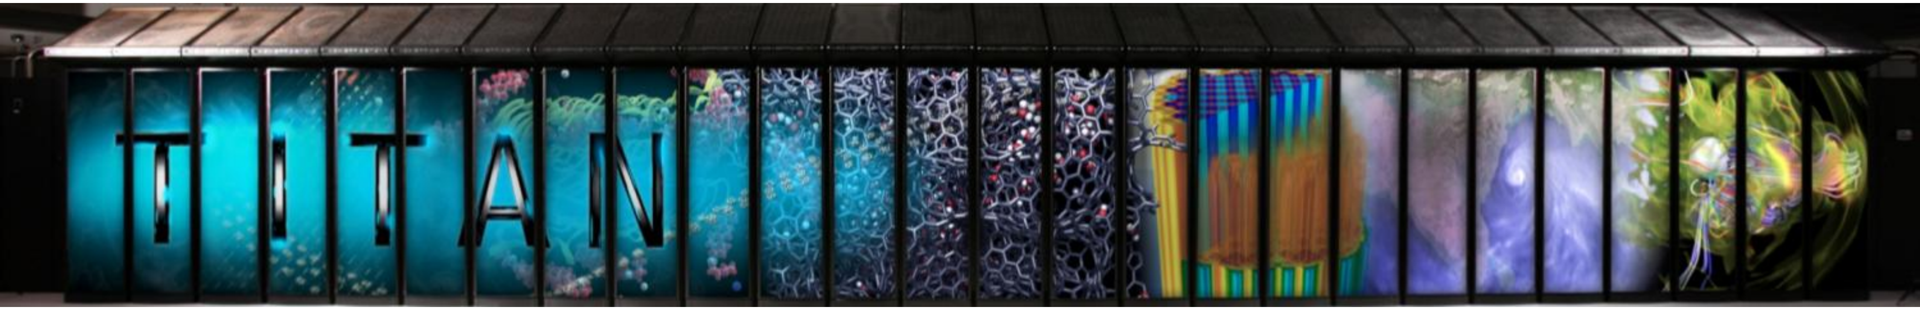

Current Status:

- Passed functionality and performance tests
- Fixing a stability problem that is requiring board repair at Cray
- Users are on the system running on the CPUs
- GPUs will be available to users in mid-March
- Expect to complete acceptance testing in April/May, after all boards are repaired

SYSTEM SPECIFICATIONS:

- Peak performance of 27.1+ PF
    - 24.5 GPU + 2.6 CPU
- 18,688 Compute Nodes each with:
    - 16-Core AMD Opteron CPU
    - NVIDIA Tesla "K20x" GPU
    - 32 + 6 GB memory
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect
- 8.9 MW peak power

The National Center for Computational Sciences is one of the world's most powerful computing facilities

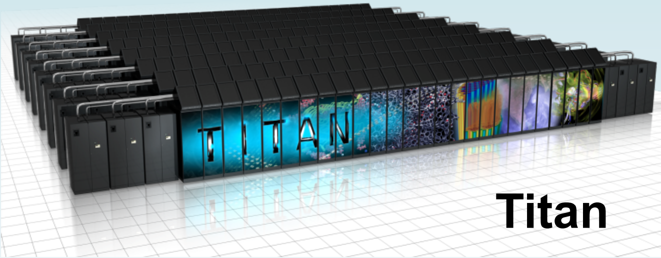

Titan

Peak performance	27 PF/s
Memory	710 TB
Disk bandwidth	240 GB/s
Square feet	5,000
Power	8.8 MW

Data Storage

- Spider File System
  - 10 PB capacity
  - 240 GB/s bandwidth
- HPSS Archive
  - 240 PB capacity
  - 5 Tape libraries

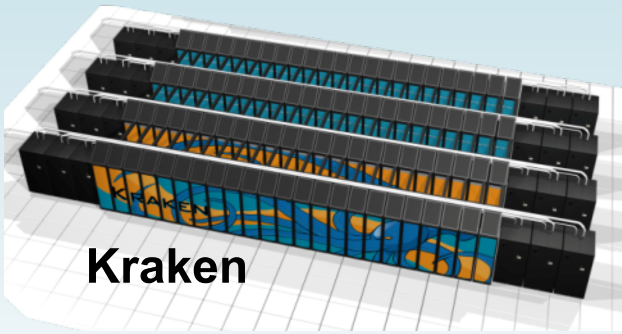

Kraken

Peak performance	1.17 PF/s
Memory	147 TB
Disk bandwidth	> 50 GB/s
Square feet	2,300
Power	3.5 MW

Data Analytics & Visualization

- LENS cluster
- Ewok cluster
- EVEREST facility
- uRiKA data appliance

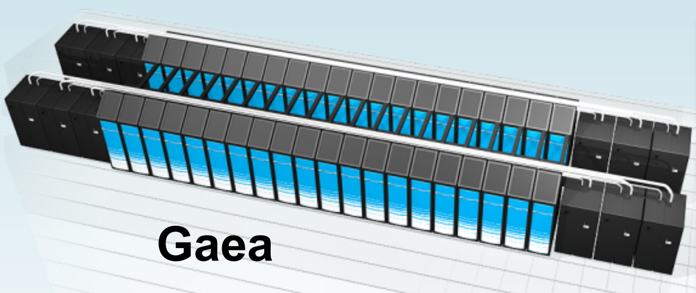

Gaea

Peak Performance	1.1 PF/s
Memory	240 TB
Disk Bandwidth	104 GB/s
Square feet	1,600
Power	2.2 MW

Networks

- ESnet – 100 Gbps
- Internet2 – 10 Gbps
- XSEDEnet – 10 Gbps
- Private dark fibre

The Second Generation Spider File System

- An External Lustre 2.4 File System for all NCCS Compute Platforms
  - More than 300 Lustre OSS and MDS Servers
  - 32,000 meters of fiber optic cable

- Block Storage Based on 36 DDN SFA12000 Storage System
  - 20,160 Nearline SAS Spindles provide more than 32PB of RAID 6 (8+2) storage
  - Peak read/write performance of 1TB/s
- Mellanox-based FDR Infiniband fabric with more than 1,600 56Gbps ports
- Connectivity to the Cray XK7 Titan over 420 separate Connect-IB 56GBps HCAs
- Production-ready in Summer 2013

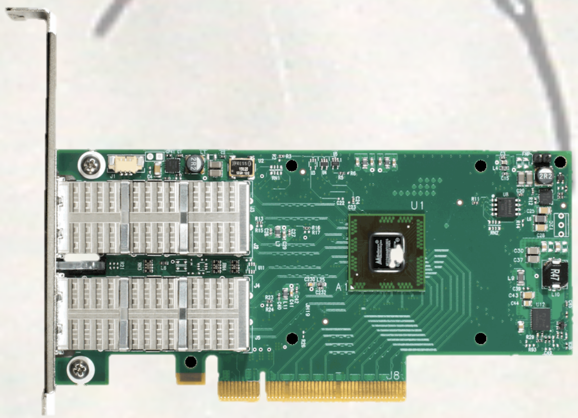

Everest Visualization Laboratory

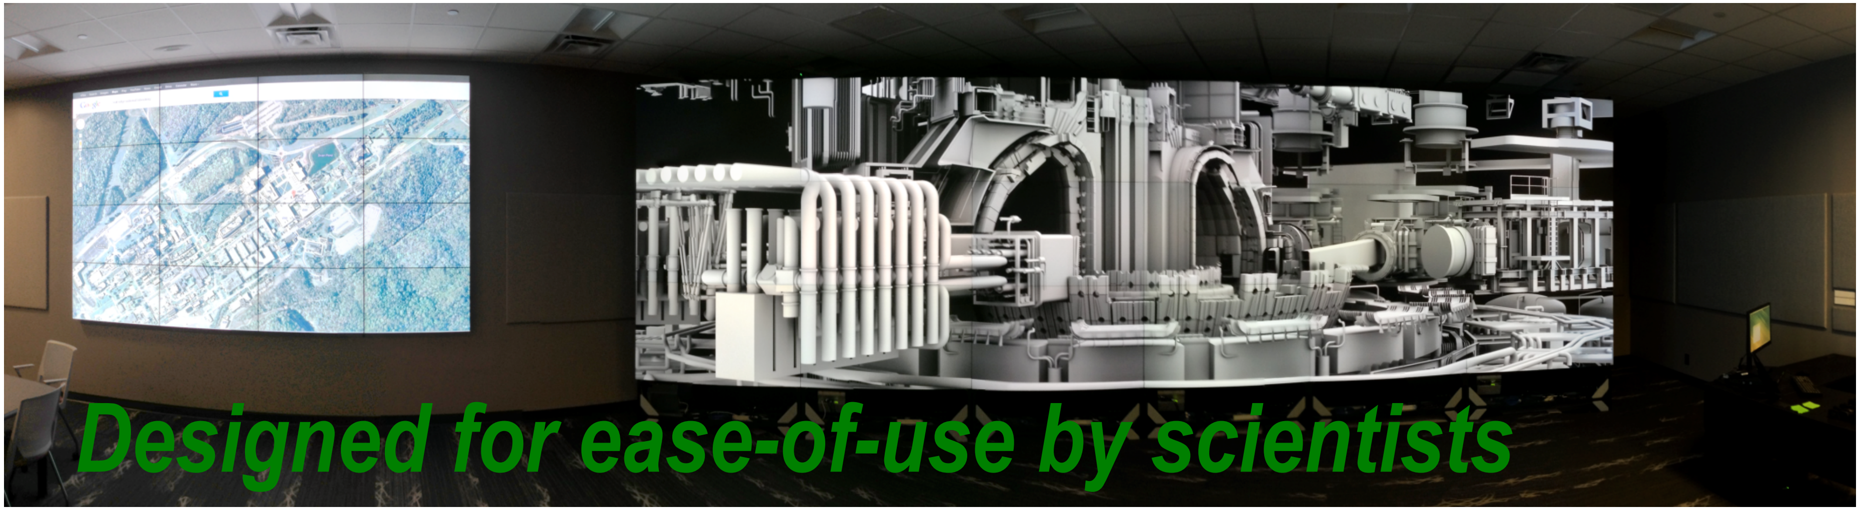

- *33 megapixel wall*
- *4x4 Planar LCD panels*
- *Complete room remodel begun in fall 2012*
- *Returned to production January 2013*

We are achieving breakthroughs in Knowledge Discovery Science

- Actionable insights from massive, dynamic, disparate data
- Ability to detect, understand, and predict processes underlying the data

Systems

- Sharing and trust
- Social media
- Streaming
- Architecture
- Sensors
- Mobile
- Workflow

Data analytics

- Text Analysis
- Multi-modality fusion
- Clock-constrained
- Large-scale
- Geo-temporal
- Social networks

Modeling and simulation

- Discrete-event
- Agent-based
- Predictive
- Real-time
- Physics-based

Cyber security

- HPC-based
- KD-based
- Cyber-physical
- Quantum-based

Mobile platforms (cell phones, iPad, PDA, UAV)
Desktops, clusters, and cloud computing (homogeneous and heterogeneous)
Distributed (sensor networks, computational platform mixes)
Extreme-scale computing (HPC)

Addressing the need: Simplified Knowledge-Discovery Stack

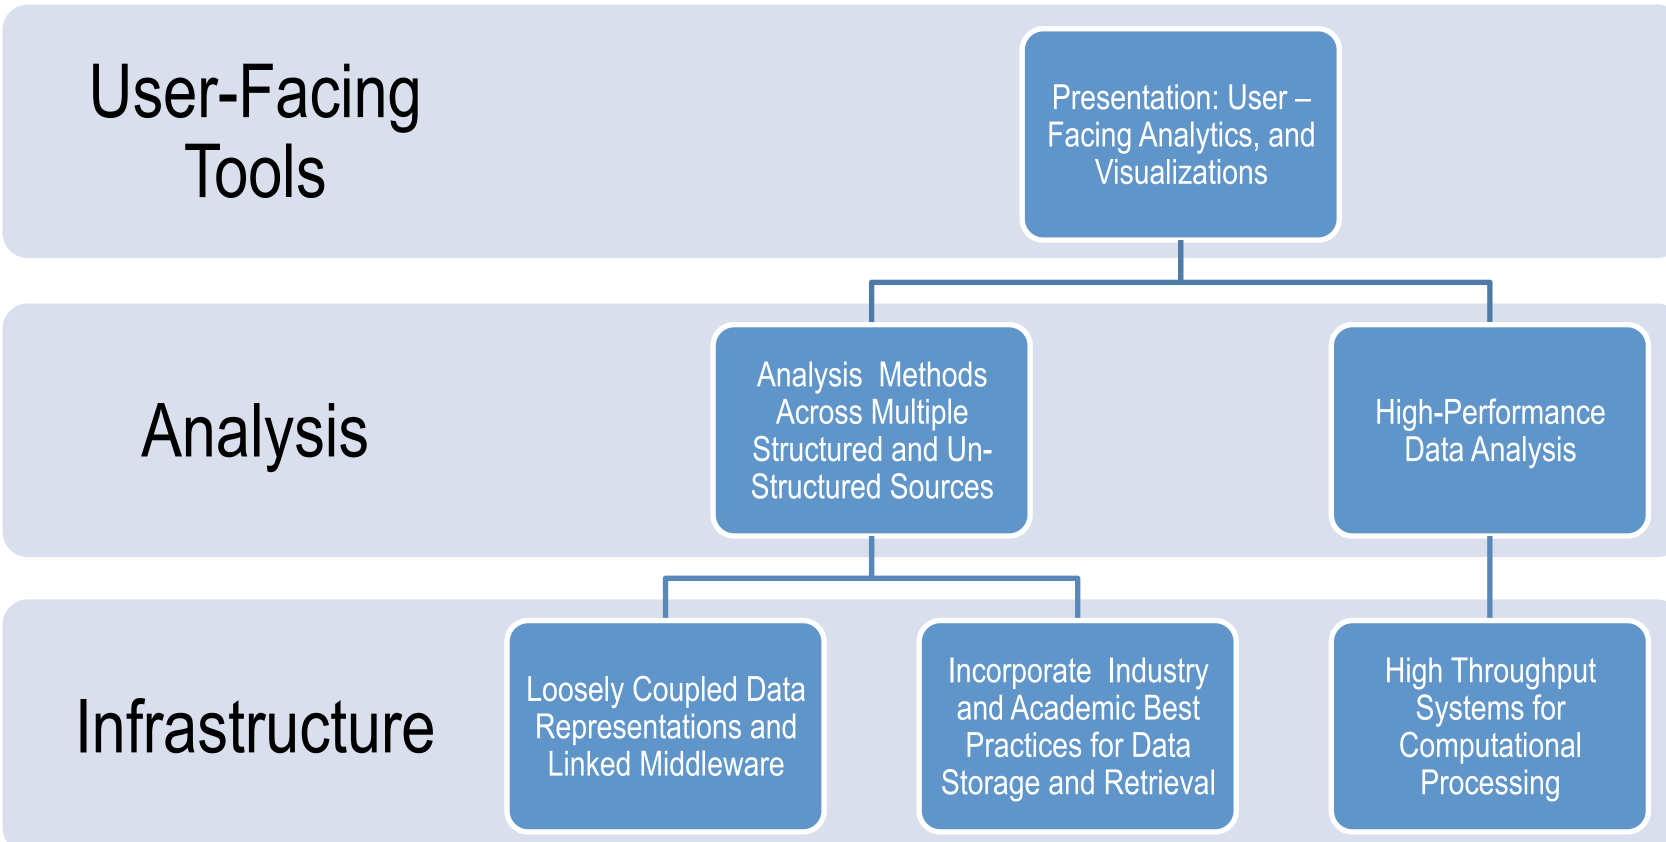

Slide courtesy of Jeff Nichols

Big Data Application Examples

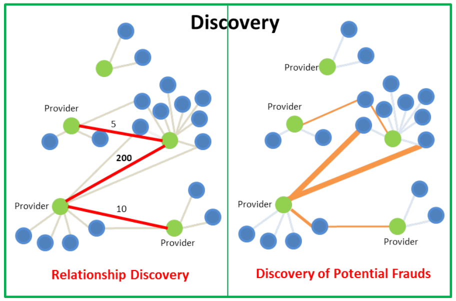

Fraud and Threat Analysis
Analytics of large connected relationships.

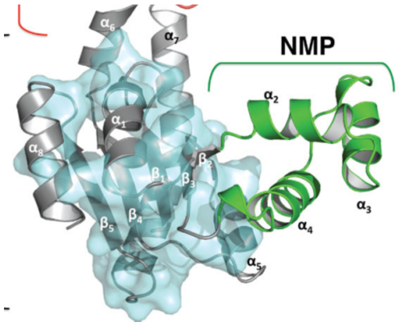

Biological System Dynamics
Whole-system modeling, simulation and analytics

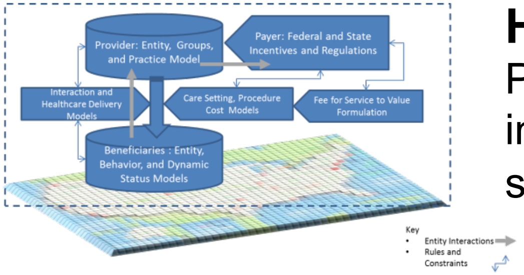

Healthcare
Policy interventions scenario analytics

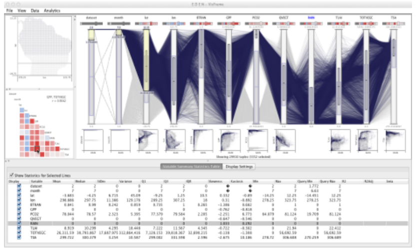

Sensor Data Modeling
Visual-analytics for environment modeling

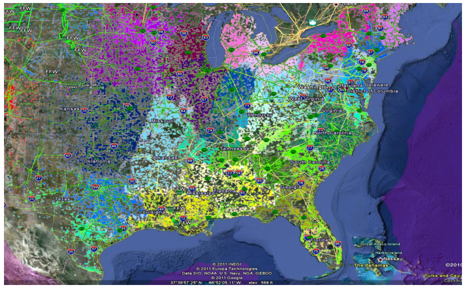

Electric Grid
Sensed grid data analytics for tighter real-time inference and control

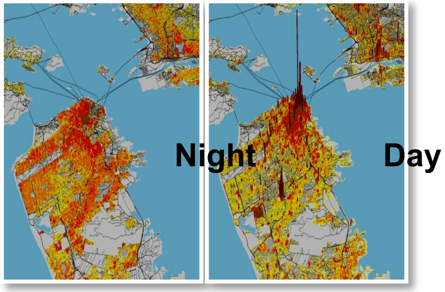

Large-Scale Data Fusion
Multi-modal datasets analysis

The Titan Ecosystem is a Unique Scientific Instrument

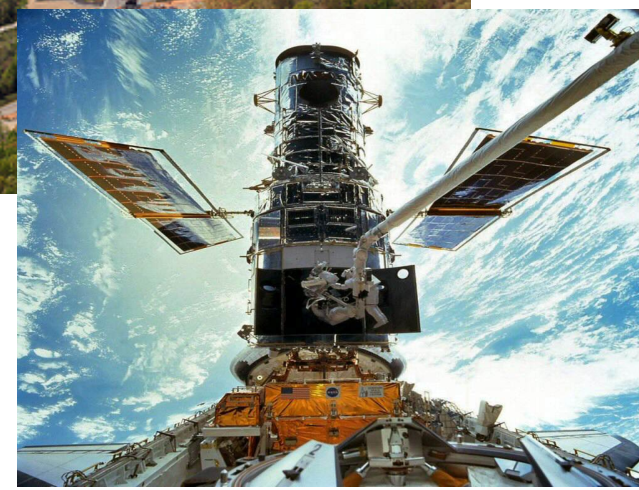

- Sheer computational power is its most conspicuous trait, but the ecosystem surrounding the machine – filesystems, visualization resources, expertise – is where, “science gets done.”
- Data analysis and visualization are the pieces of the scientific workflow that lead directly to insight.

Large Data Volumes are Being Produced by Simulations Today

Example: Climate

- ***Ultra High Resolution Project (J. Hack, PI)***
  - 70M hours using the *Community Earth System Model*
  - 35,000 files, ranging in sizes from 21 MB to 110 GB
  - The total volume of model output will be 300+ TB
- ***Early Science using CAM-SE (J-F. Lamarque, PI)***
  - *CAM-SE* with *MOZART* Chemistry (106 tracers)
  - 450M hours on Titan using GPU accelerators.
  - 250+ TB of simulation results
- ***Climate End Station (W. Washington, PI)***
  - ~ 200M hours in 3 years and 200+ TB data

ADIOS

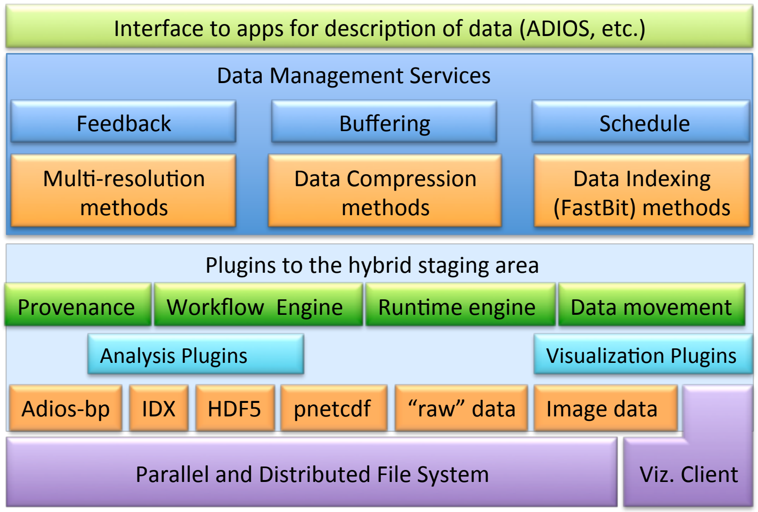

The Adaptable IO System (ADIOS) provides a simple, flexible way for scientists to describe the data in their code that may need to be written, read, or processed outside of the running simulation. By providing an external to the code XML file describing the various elements, their types, and how you wish to process them this run, the routines in the host code (either Fortran or C) can transparently change how they process the data.

The in code IO routines were modeled after standard Fortran POSIX IO routines for simplicity and clarity. The additional complexity including organization into hierarchies, data type specifications, process grouping, and how to process the data is stored in an XML file that is read once on code startup. Based on the settings in this XML file, the data will be processed differently. For example, you could select MPI individual IO, MPI colletive IO, POSIX IO, an asynchronous IO technique, visualization engine, or even NULL for no output and cause the code to process the data differently without having to either change the source code or even recompile.

The real goal of this system is to give a level of adaptability such that the scientist can change how the IO in their code works simply by changing a single entry in the XML file and restarting the code. The ability to control at a per element basis and not just a data grouping such as a restart, diagnostic output, or analysis output makes this approach very flexible. Along with this detail level, a user can also just change which transport method is used for a data type such as a restart, analysis, or diagnostic write.

For the transport method implementer, the system provides a series of standard function calls to encode/decode data in the standardized .bp file format as well as "interactive" processing of the data by providing direct downcalls into the implementation for each data item written and also callbacks when processing a data stream once a data item has been identified along with its dimensions and a second callback once the data has been read giving the implementation the option to allocate memory and process the data as close to the data source as is reasonable.

- An I/O abstraction framework
- Provides portable, fast, scalable, easy-to-use, metadata rich output
- Change I/O method on-the-fly
- Abstracts the API from the method
- http://www.nccs.gov/user-support/center-projects/adios/
- Typical speeds up Writes/Reads for simulations on OLCF by >10X over “typical” parallel libraries
- Allows for the creation of in-transit analytics

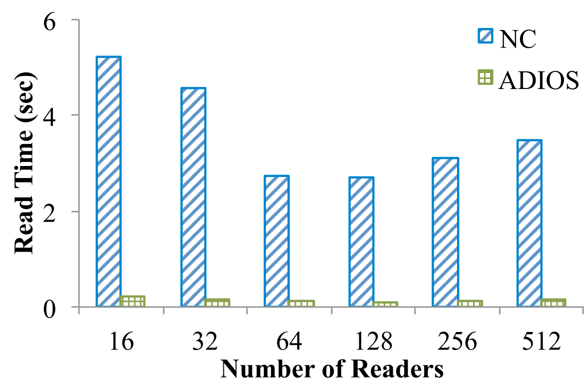

Read times for GEOS-5 data on OLCF reading space-time slice

I/O performance of the Combustion S3D code (96K cores), the SCEC PCML3D (30K cores), and the Fine/Turbo (4K cores) codes.

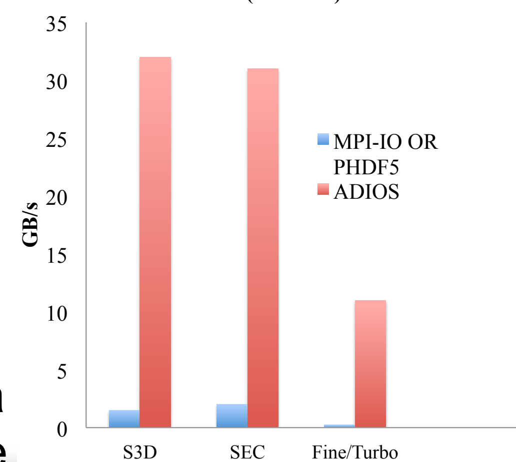

Simulations with and without ADIOS

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE U.S. DEPARTMENT OF ENERGY

New method to stream data over the WAN

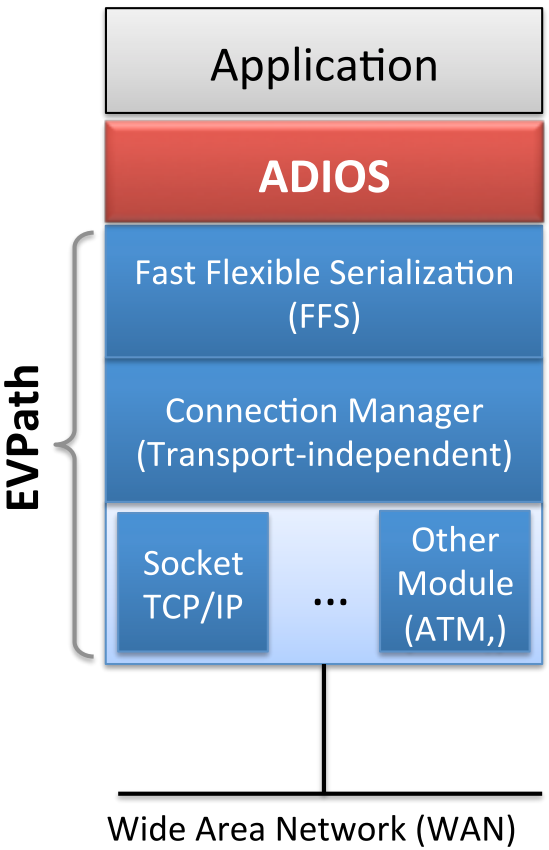

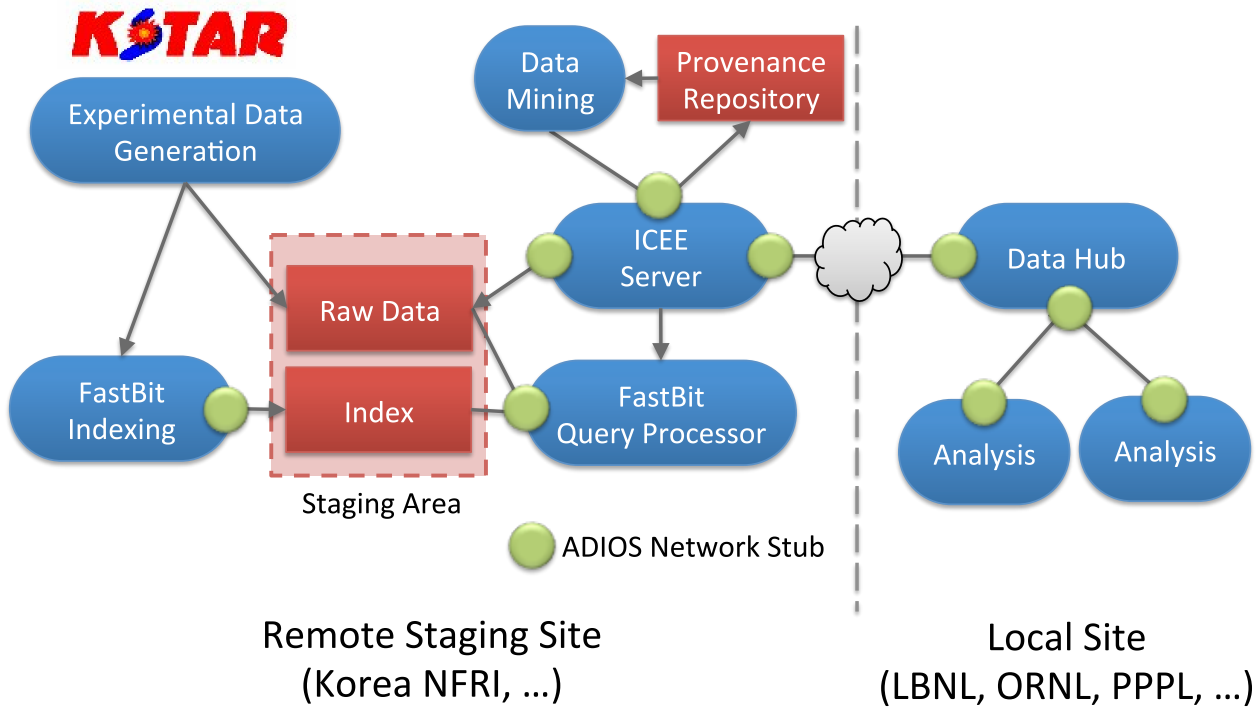

- Abstraction allows workflows to be executed transparently from experiments to LCFs
- Allows data streams to be processed in-transit
- Service Oriented Architecture allows “researchers” to contribute new services which can be executed during the experiments

Collaboration w/ Experimental Facilities

OLCF doing streaming data analysis of SNS experiments (ADARA)

Production software developed by OLCF and SNS teams to immediately analyze beamline data, display it to scientist and archive results in HPSS

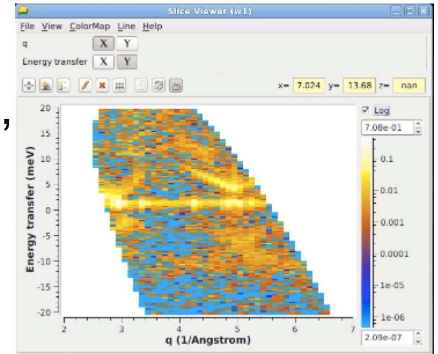

Coupling simulation and experiment

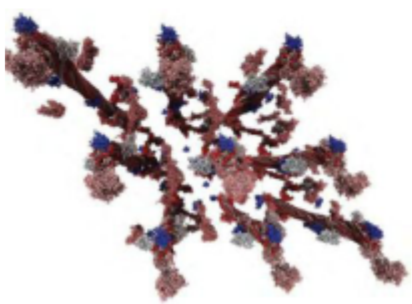

Overlapping length scales of small angle neutron scattering (HFIR) results and OLCF MD calculations enable validating results describing lignin folding. ***Phys Rev E (2011)***

Extreme-scale Data Science

Using unique capabilities of OLCF parallel file system, analyzed 300 TB of global simulation data along with paleoclimate proxy data to explain lagged temperature response to changes in CO_2. ***Nature (2012)***

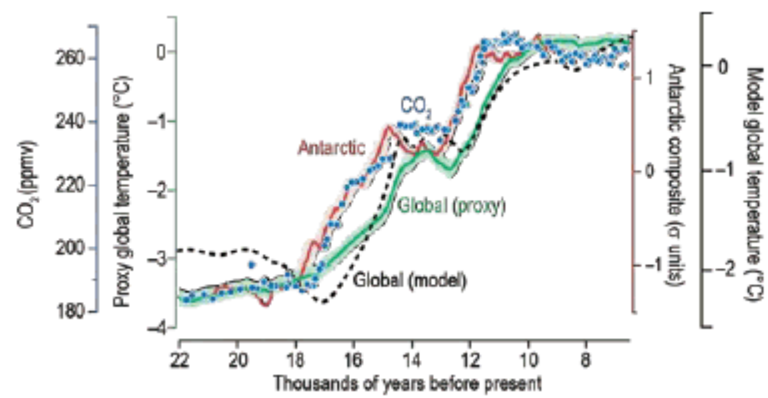

Analysis – Fitting/Modeling/Simulation

$$I_{calculated}(\vec{Q}_0, E_0) = \int_{\Delta\vec{Q}=-\infty}^{\infty} \int_{\Delta E=-\infty}^{\infty} S(\vec{Q}_0 + \Delta\vec{Q}, E_0 + \Delta E) \times R(\Delta\vec{Q}, \Delta E)\, d\Delta\vec{Q}\, d\Delta E$$

$$S(Q,E) \propto \sum_s \sum_\tau \frac{1}{E_s} \left| \sum_d \frac{\overline{b_d}}{\sqrt{M_d}} \exp(i\mathbf{Q}\cdot\mathbf{r}_d) \exp(-W_d) (\mathbf{Q}\cdot\epsilon_{d,s}) \right|^2 \times \langle n_s + 1 \rangle \delta(E - E_s) \delta(\mathbf{Q} - \mathbf{q} - \tau)$$

$$W_d = \frac{\hbar^2}{4M_d} \sum_s \frac{\left|\mathbf{Q}\cdot\epsilon_{d,s}\right|^2}{E_s} \langle 2n_s + 1 \rangle$$

Some areas to help:

- Fitting/visualization of models/data
- Uncertainty quantification
- Optimum design of experiments
  data → model & model → data
- ← e.g. Efficient math calculations
- + others

Not just models:

- What about "first principles" simulation/calculation?
- Compute neutron scattering laws from simulation?
- How do we parameterize the simulations so that we can "fit" against experiment?
- Resources/scheduling/user interface

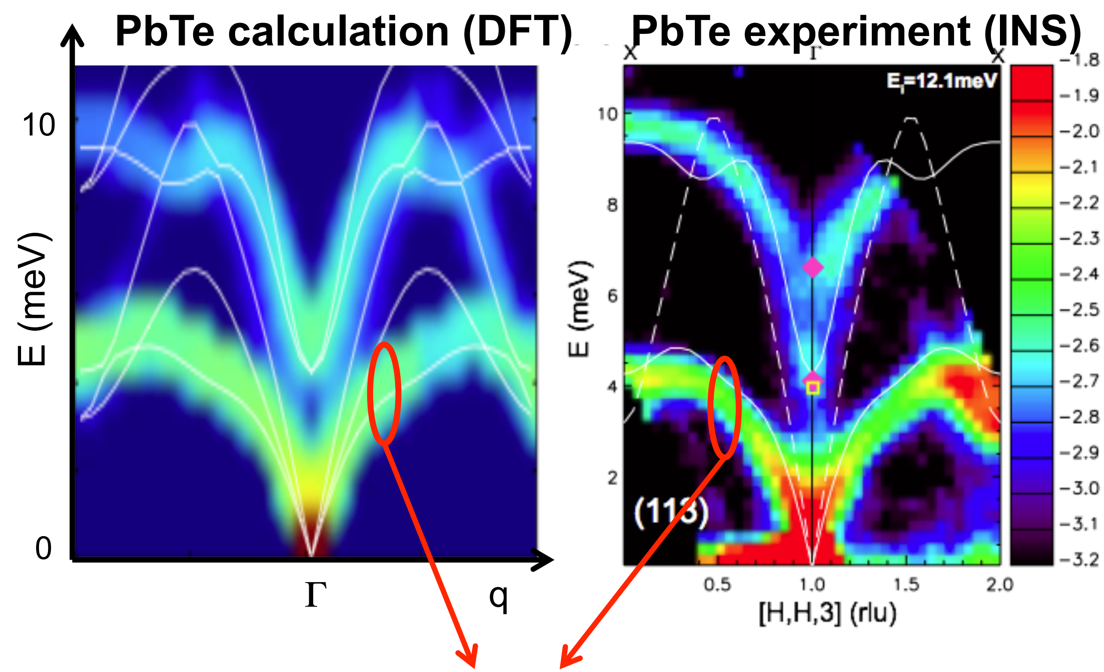

Focus on *width* of dispersions:

- measured linewidths need corrections for instrument resolution effects.
- will support development of reliable linewidth predictions.

Delaire *et al.*, Nature Materials (2011).

Designing porous 3D carbon nanotube sponge-like materials

- Boron doping in nanotube lattice during growth
- Influences elbow joints and branching
- Leads to 3D supra-sponge structure
- Density comparable to aerogel
- Mechanically and thermally stable
- Electrically conducting
- Magnetic
- Strongly oleophillic
- Able to absorb large quantities organic solvents and oil

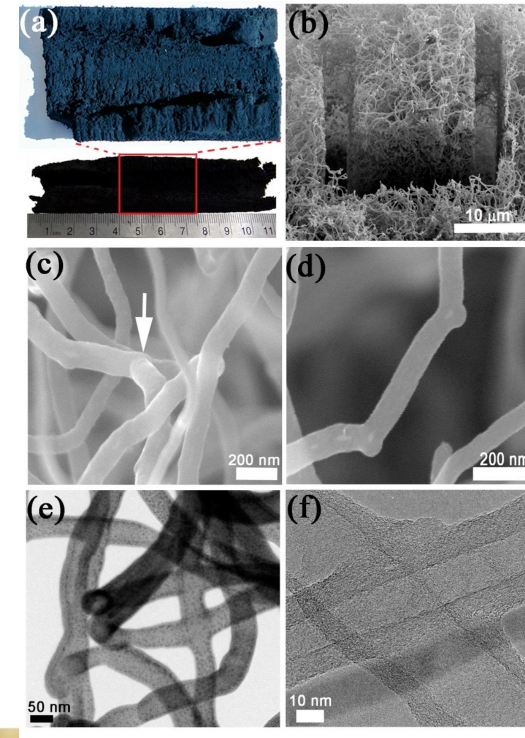

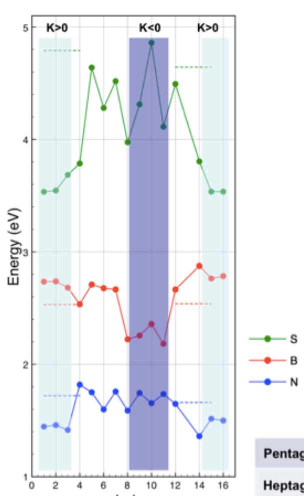

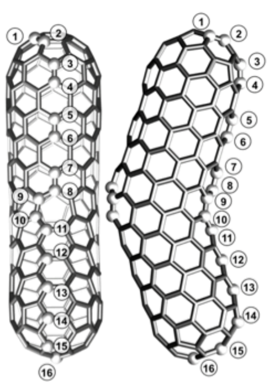

	Boron	Nitrogen	Sulfur
Pentagon	0.48 eV	-0.55 eV	-0.44 eV
Heptagon	-0.67 eV	0.03 eV	-0.28 eV

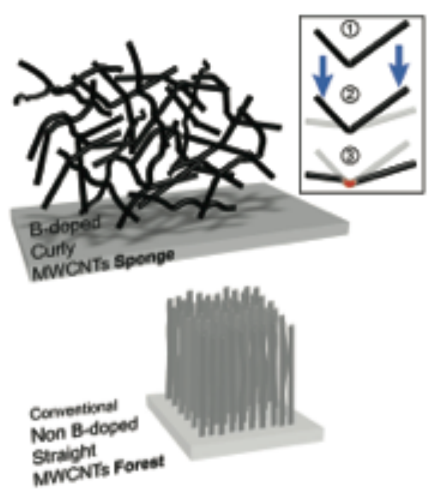

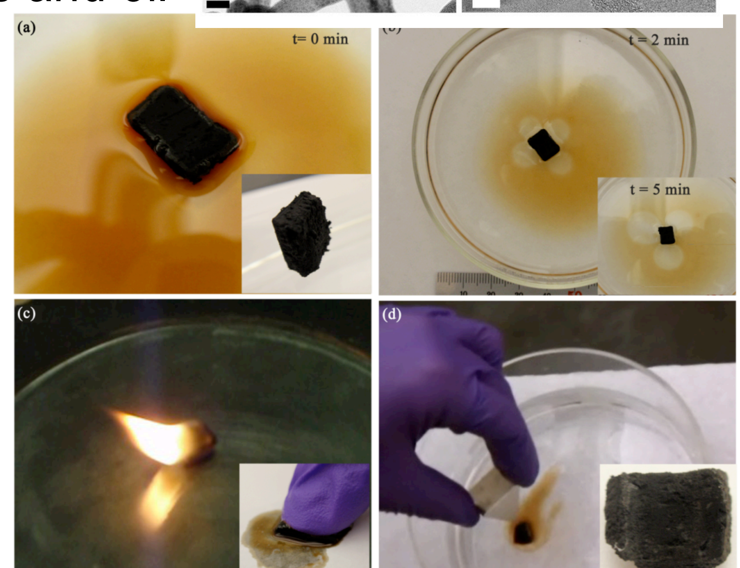

Nature Sci. Rep **2**, 363 (2012)

We are building a Health Data Science Institute

A collaborative R&D environment to explore large-scale data-driven transformation of healthcare

Claims

Clinical

Omics

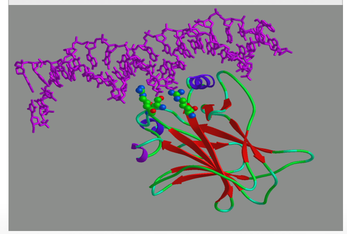

Social

Devices

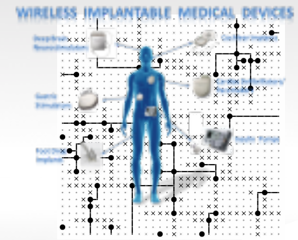

Data Sources

The Challenge
Enable Discovery

ORNL
Infrastructure

Outcomes

OAK RIDGE National Laboratory

Science challenges in the next decade …

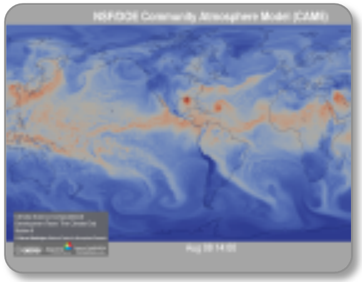

Climate Change Science
Understand the dynamic ecological and chemical evolution of the climate system with uncertainty quantification of impacts on regional and decadal scales.

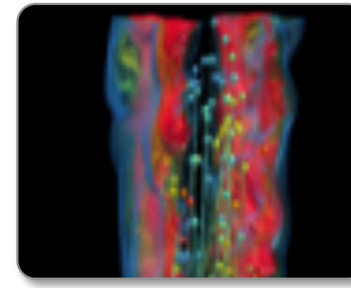

Combustion Science
Increase efficiency by 25%-50% and lower emissions from internal combustion engines using advanced fuels and new, low-temperature combustion concepts.

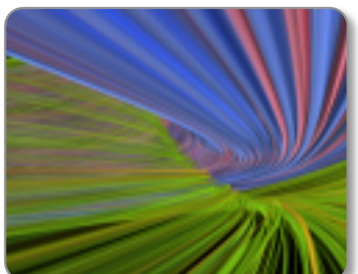

Fusion Energy/ITER
Develop predictive understanding of plasma properties, dynamics, and interactions with surrounding materials.

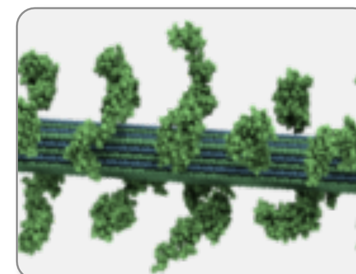

Biomass to Biofuels
Enhance the understanding and production of biofuels for transportation and other bio-products from biomass.

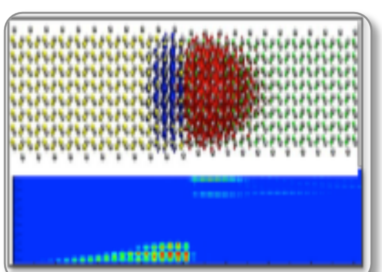

Solar Energy
Improve photovoltaic efficiency and lower cost for organic and inorganic materials.

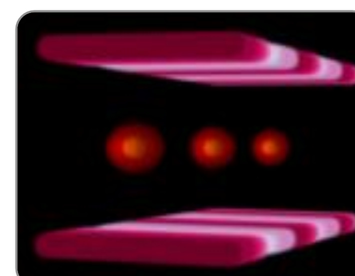

Globally Optimized Accelerator Designs
Optimize designs as the next generations of accelerators are planned, increasingly detailed models will be needed to provide a proof of principle and a cost-effective method to design new light sources.

OAK RIDGE National Laboratory

Energy Storage advances will be enabled by HPC over the next decade

Key science challenges: Gain the fundamental understanding of reaction processes at the atomic and molecular level required for predictive design of new materials for energy storage and predictive engineering of safe, large-format rechargeable batteries.

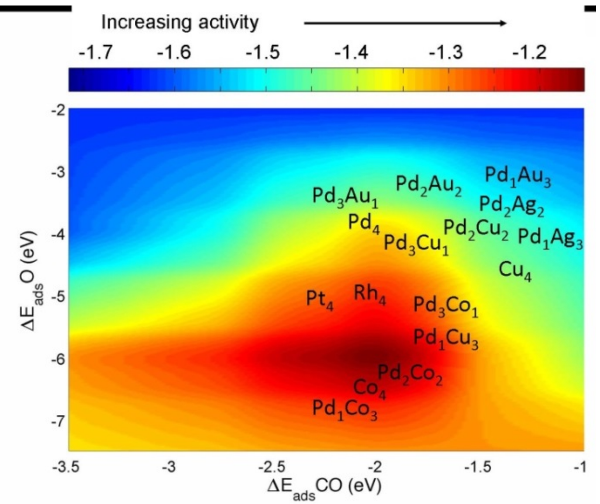

screening of catalysts for Li-air batteries

Science enabled by HPC Capabilities

2013-2016

- Computational screening of thousands of candidate materials simulating ionic transport and thermodynamic properties along with limited simulations of electrochemical processes at interfaces.
- Multi-scale modeling to improve efficiency and cyclability of electrochemical processes in Li-air batteries.

2016-2020

- Inverse methods for materials design, enabling the discovery of specific materials with higher energy and power densities, better stability and safety, and longer lifetimes.
- Achieve broad impact on current Li-ion technologies and enable significant exploration of solutions beyond Li-ion.

OAK RIDGE
National Laboratory

Nuclear Energy advances will be enabled by HPC over the next decade

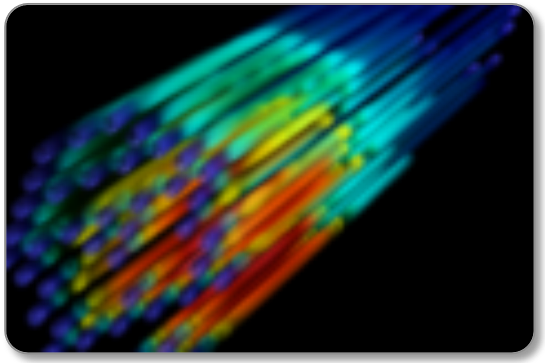

Key science challenges: For existing reactors provide safe, increased fuel utilization, power upgrades, and reactor lifetime extensions. Design new, safe, cost-effective reactors.

Science enabled by HPC Capabilities

2013-2016	2016-2020
• Coupled-physics simulation of fluid flow, heat transfer, neutron transport, and fuel behavior for fuel assemblies. • Single-physics and low-fidelity coupled-physics analysis of full cores. • Predict behavior of nominal reactor operation and existing nuclear fuels with limited quantification of uncertainties.	• Coupled-physics simulation of fluid flow, heat transfer, neutron transport, and fuel behavior for multiple fuel assemblies and limited full-core analysis of nominal reactor operation. • Initial capability for analysis of reactor transients, including some accident scenarios. Improved quantification of uncertainties.

Leadership Facility Strategy for the future
Roadmap and Timeline

Through a three-phase plan executed over the next decade, the LCF will deploy a series of ever-more-powerful, balanced, scalable, HPC and data resources to support the most challenging computational problems of the nation.

Phase 1: Procure and operate pre-Exascale systems (2016) Three-way RFP w/ ORNL, ANL, LLNL
Phase 2: Procure and operate Exascale systems (2020)
Phase 3: Procure and operate Second generation Exascale systems (2024)

Computer System requirements for each Leadership Computing Center

	2012	2016	2020	2024
Peak FLOP/s	10-20 PF	100-200 PF	500-2000 PF	2000-4000 PF
Memory	0.5-1 PB	5-10 PB	32-64 PB	50-100 PB
Burst Buffer	N/A	500 TB	3 PB	5 PB
Storage Disk +tape	20+100 PB	100+1000 PB	1+10 EB	5+50 EB
Power & Space	6-12 MW 5,000-10,000 ft^2	15-20 MW 8,000-15,000 ft^2	20-30 MW 20,000 ft^2	25-35 MW 25,000 ft^2

Leadership Facility Strategy (cont.)
Software and Data

System Software and Libraries

- Just as critical as the hardware will be software that can exploit the full scale and capability of the hardware to enable knowledge discovery.
- LCF will bear an increasing portion of the responsibility for ensuring that users are provided with a stable and effective software environment.

Major External Initiatives Impacting Strategy

- Exascale Initiative (including FastForward and DesignForward).
- Big Data Initiative.

Integrated Data Science Infrastructure across DOE Facilities

- Growing demand by DOE/SC programs for data science resources that match the scale of LCF computational sciences resources.
- Tightly coordinated multisite data facility built as an extension of the ASCR Facilities at Argonne, Berkeley and Oak Ridge.
- Required 2014 resources are projected to be 100 PB with a 3 Petaop/s data analysis capability growing to 5 Exabytes with a 300 Petaops/s data analysis capability by 2020.

Challenges

Enable fundamentally new methods of scientific discovery by building stronger collaborations with experimental facilities as well as programs that have large computation and data science challenges.

There is a growing class of **large data science problems** where the volume and velocity of the data **require the computational and data resources only available at todays Leadership Computing Facilities.**

- **Coupling of simulation and experiment:** The big data generated by large science experiments will be fed directly to the Leadership computer simulations and the output used to drive the experiment in a feedback loop that has the potential to revolutionize discovery.
- **Analysis and data exploration of huge volumes** of disparate data from sensors, satellites, and experimental data will require LCF computers with the largest amounts of internal memory of any computers in the world.
- **As generators of big data from simulations**, the LCF will be responsible for the protection and collaborate with the data creators to disseminate this data to scientists around the world through data portals or other means.

Questions?