BLUE WATERS

SUSTAINED PETASCALE COMPUTING

Atmospheric Sciences Plus Blue Waters Create Exceptional Insights

William Kramer
National Center for Supercomputing Applications, University of Illinois
http://bluewaters.ncsa.illinois.edu

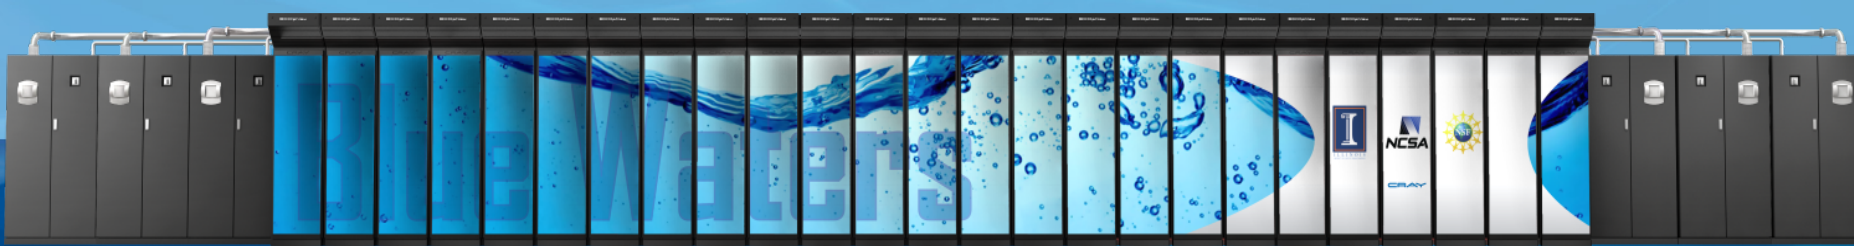

NCSA NSF GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION CRAY

What I like about this meeting

I get to think about the forest rather than lots of trees

My Summer Experience vis-à-vis CAS

Blue Waters Computing System

bluewaters.ncsa.illinois.edu

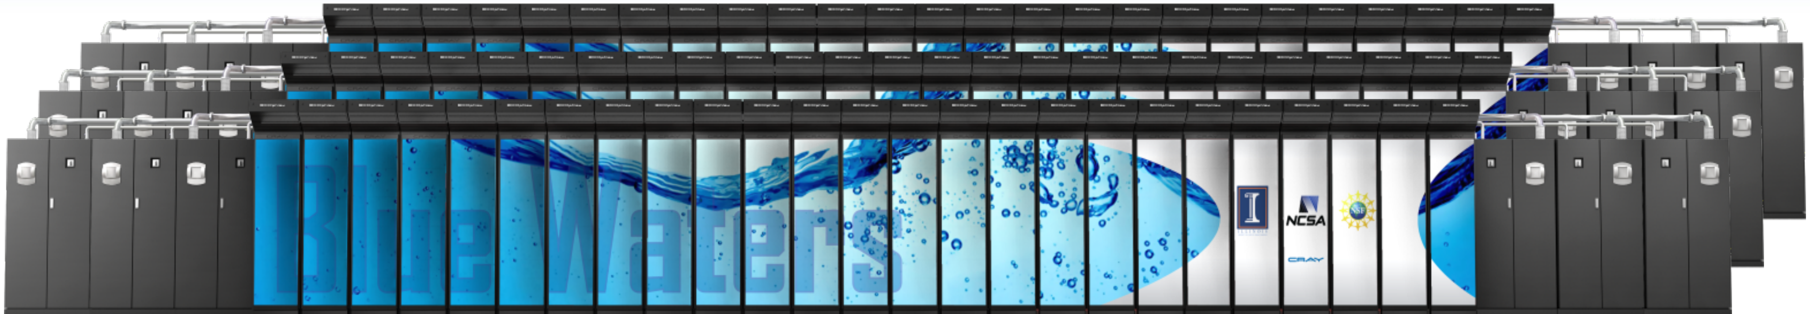

1.66PB Globally Addressable Memory
13.1 Peak PF

10/40/100 Gb Ethernet Switch

External Servers

IB Switch

>1 TB/sec

100 GB/sec

220+ Gb/sec
Going to 400+ Gb/sec by end 2015

- Full Scale Scale use across all ranges of research
- Measured, Sustained 1.3 PF/s over 14 benchmarks
- The largest System Cray has ever built – 45% larger than Titan
- Largest Memory of any system in open science – 1.66 PB
- Most networked facility in open science
- Not listed on the Top500 on purpose

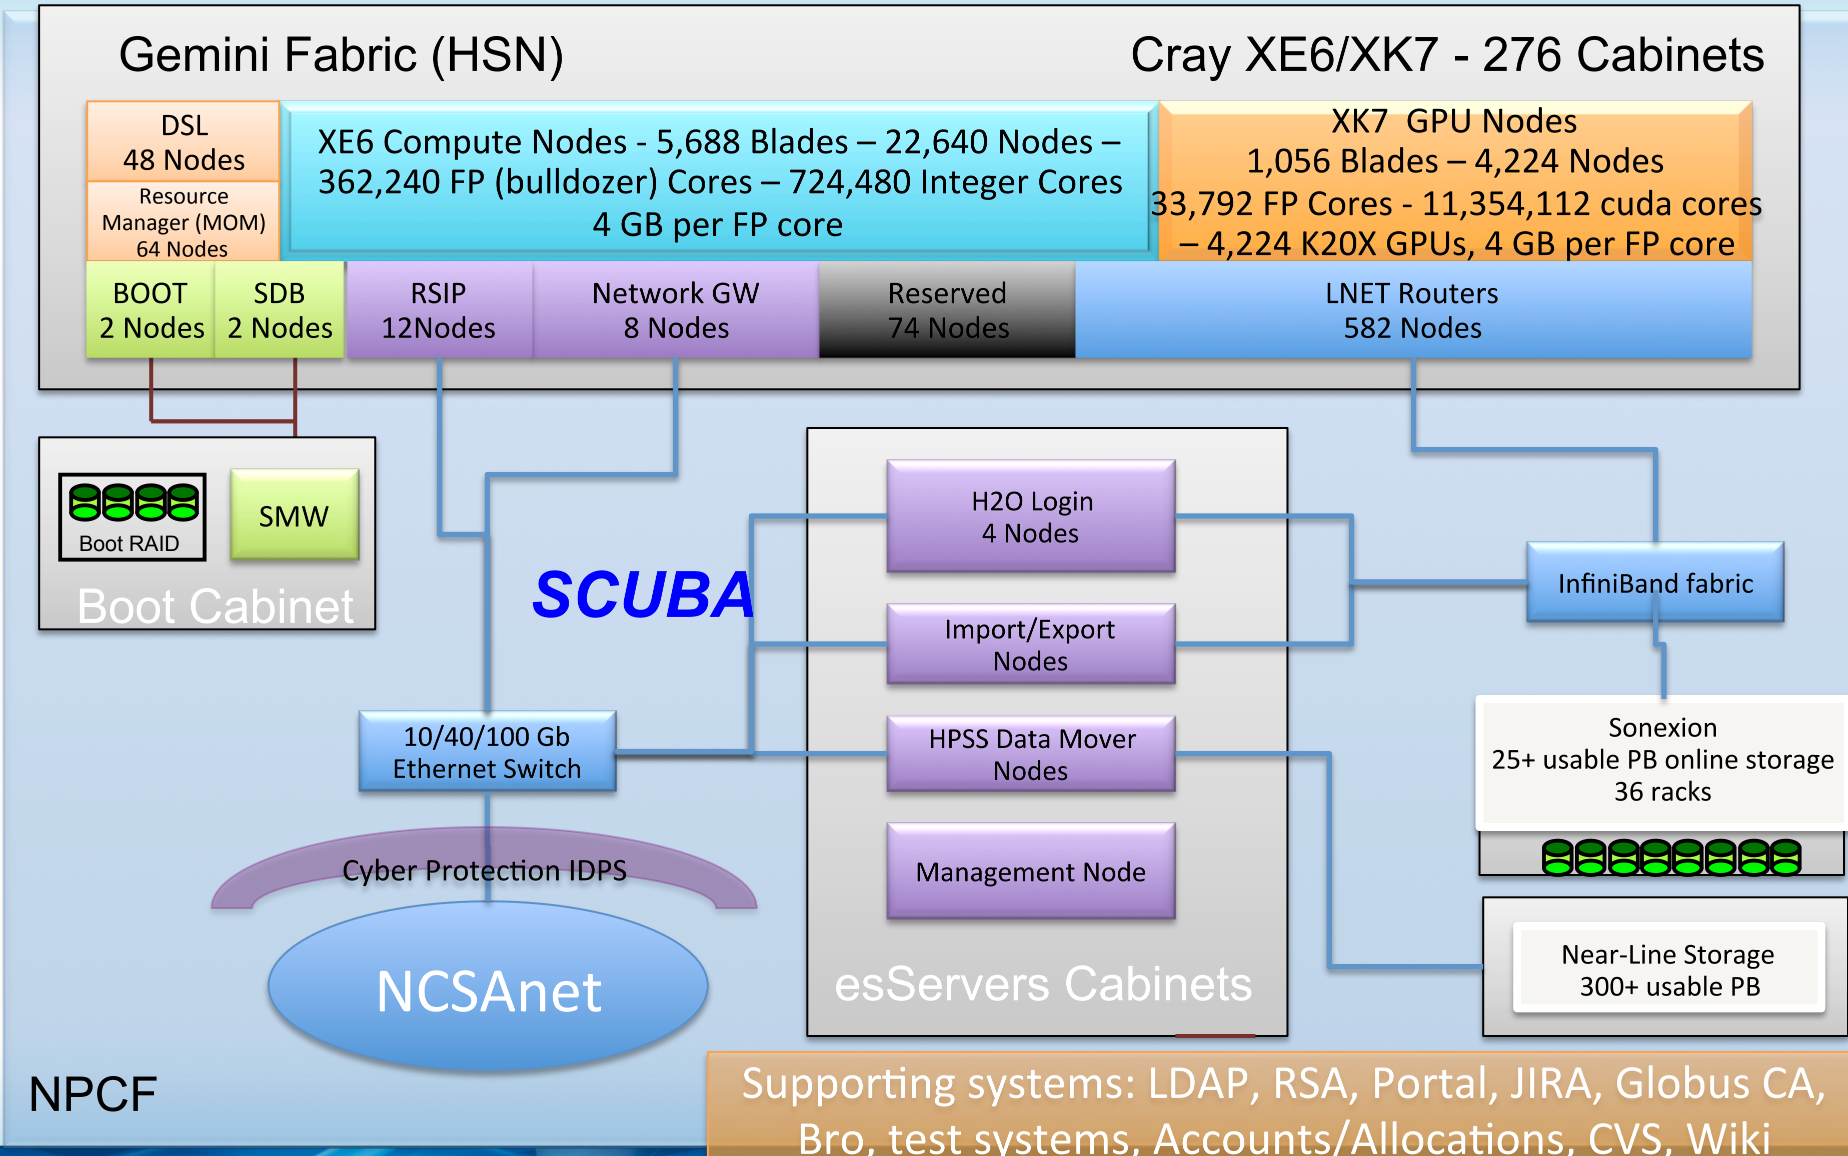
Gemini Fabric (HSN)
Cray XE6/XK7 - 276 Cabinets
DSL
48 Nodes
Resource Manager (MOM)
64 Nodes
XE6 Compute Nodes - 5,688 Blades – 22,640 Nodes – 362,240 FP (bulldozer) Cores – 724,480 Integer Cores
4 GB per FP core
XK7 GPU Nodes
1,056 Blades – 4,224 Nodes
33,792 FP Cores - 11,354,112 cuda cores – 4,224 K20X GPUs, 4 GB per FP core
BOOT
2 Nodes
SDB
2 Nodes
RSIP
12Nodes
Network GW
8 Nodes
Reserved
74 Nodes
LNET Routers
582 Nodes
Boot RAID
SMW
Boot Cabinet
SCUBA
H2O Login
4 Nodes
Import/Export
Nodes
HPSS Data Mover
Nodes
Management Node
esServers Cabinets
InfiniBand fabric
Sonexion
25+ usable PB online storage
36 racks
Near-Line Storage
300+ usable PB
10/40/100 Gb
Ethernet Switch
Cyber Protection IDPS
NCSAnet
NPCF
Supporting systems: LDAP, RSA, Portal, JIRA, Globus CA, Bro, test systems, Accounts/Allocations, CVS, Wiki

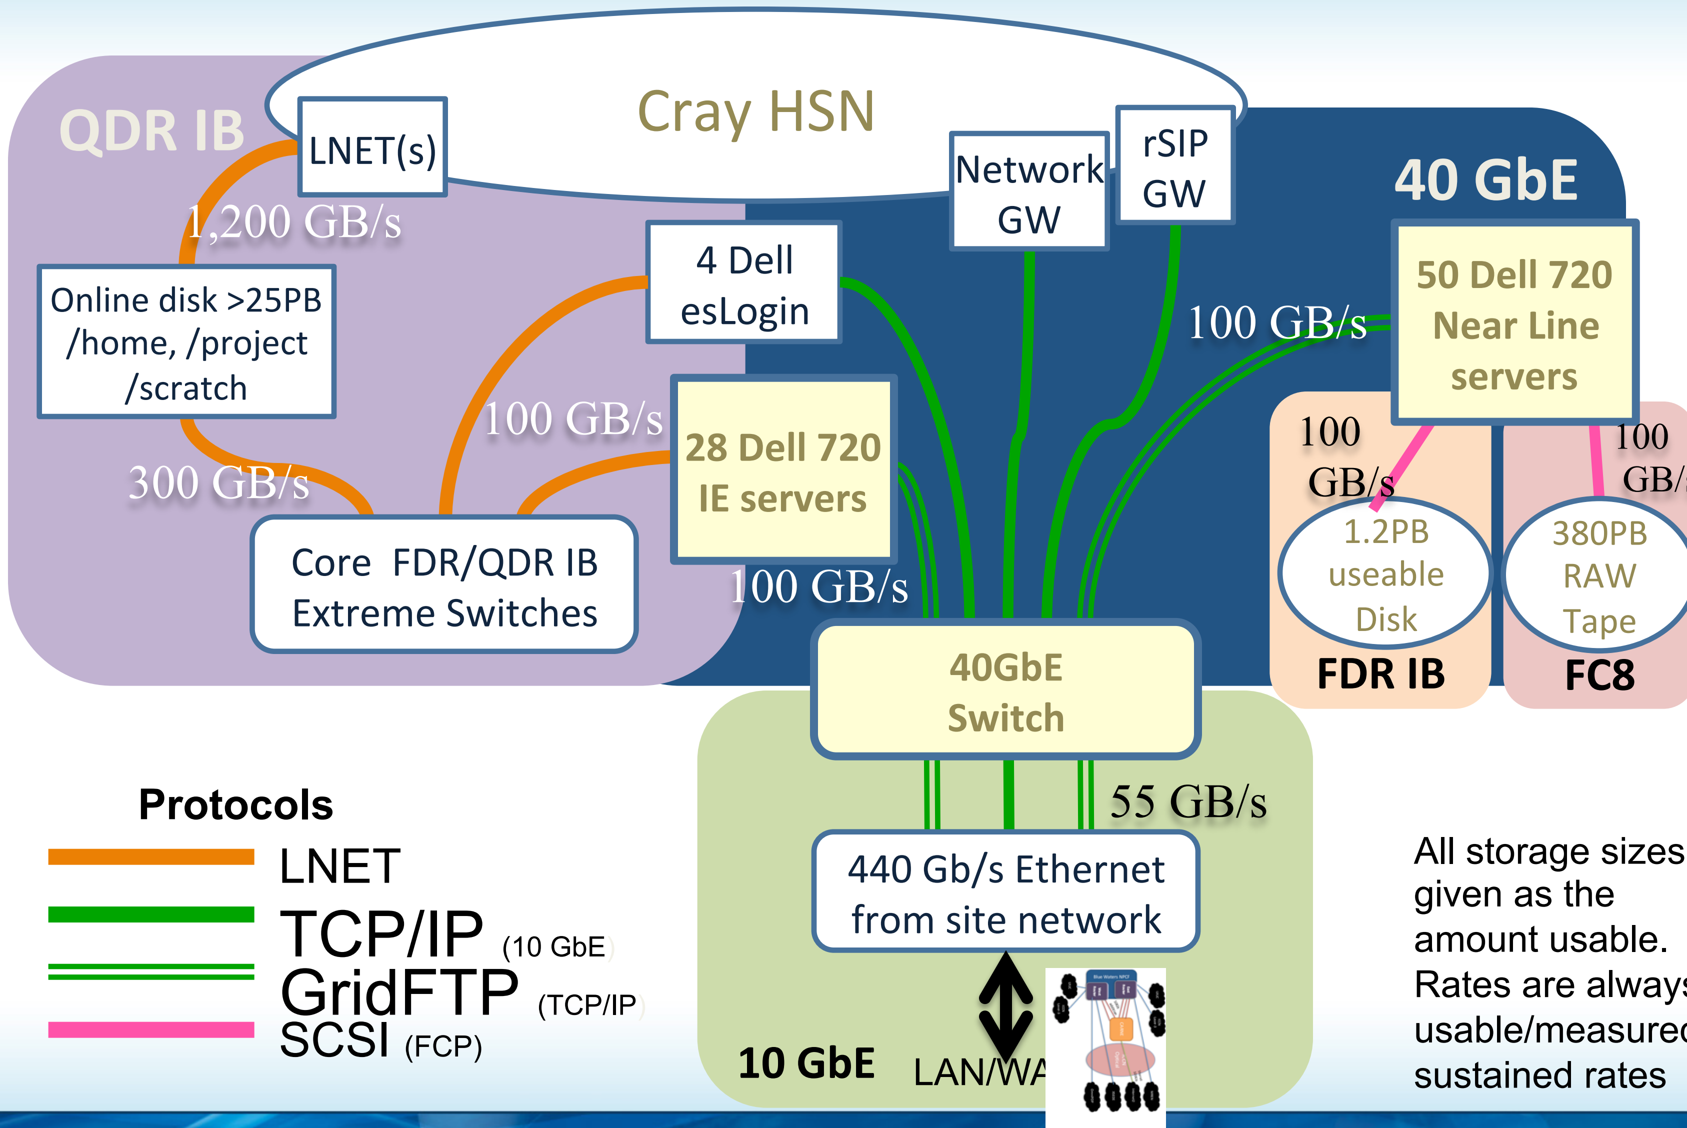

All storage sizes given as the amount usable.
Rates are always usable/measured sustained rates

July 2014-June 2015 Usage by Discipline

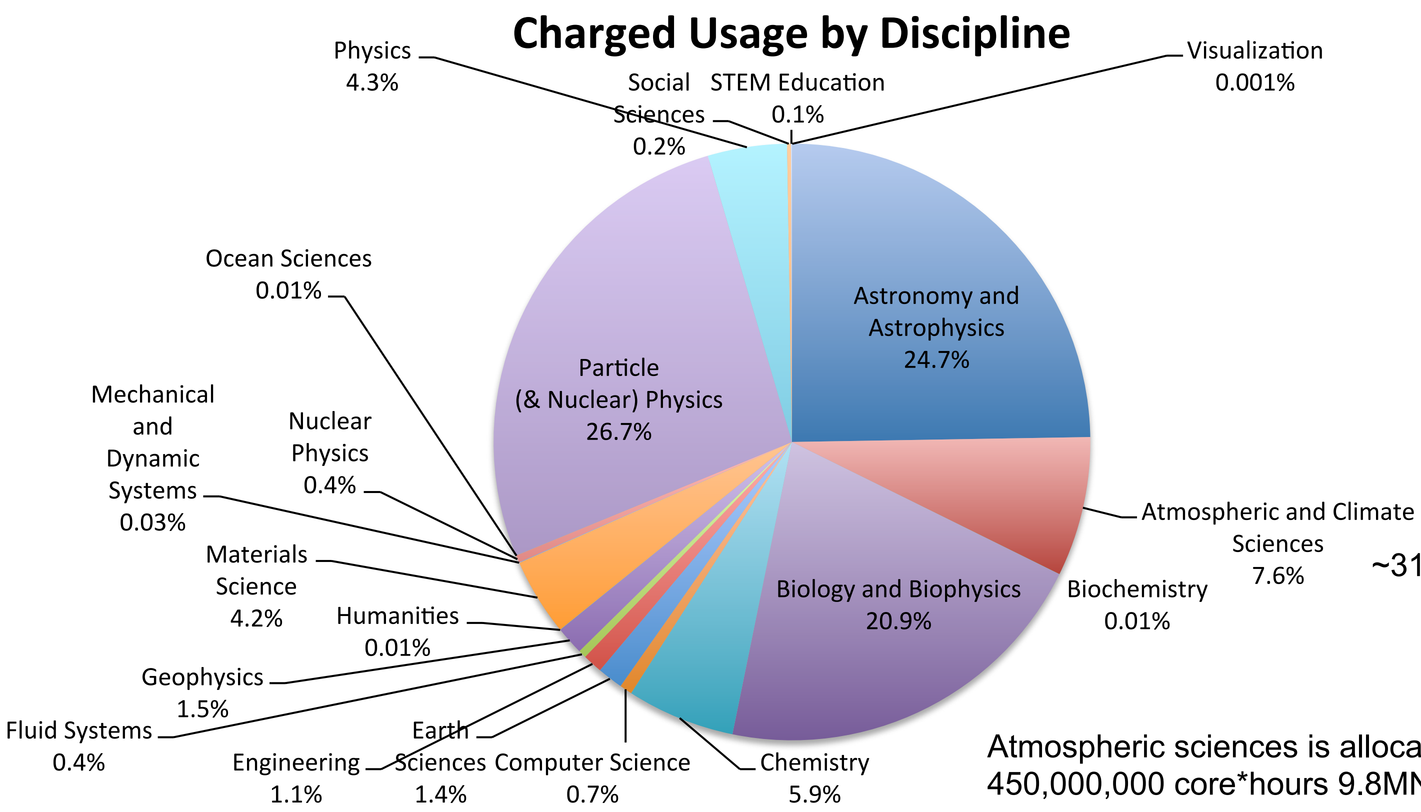

~313,000,000 core*hours

Atmospheric sciences is allocated 450,000,000 core*hours 9.8MNH

Wuebbles, Washington, et. al. – Climate Change Uncertainties

Blue Waters allows multiple, high resolution runs, 150+ years past and 100 years future, to characterize uncertainty.

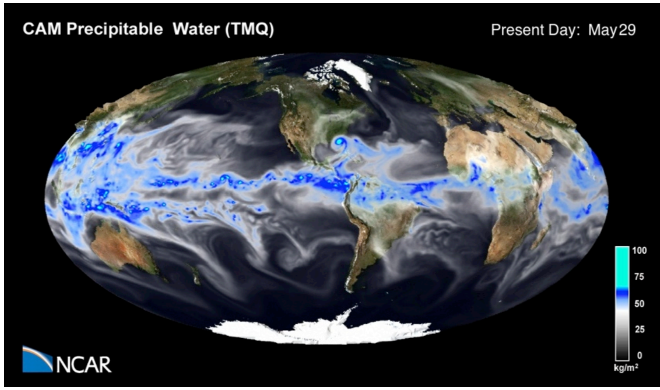

- Challenge Goal
  - Validated the effects of very high resolution (10-30 km horizontal resolution) in coupled climate models.
- Usage/Accomplishments
  - 3 present day AMIP (1979-2010) experiments were conducted using CAM5 at 0.25^{o} resolution with different atmosphere/ocean coupling. "After examining the simulations in detail we believe the modified coupling approach (flux calculations on the higher-resolution atmosphere grid) is correct, while the current default coupling is demonstrably unphysical in situations with strong wind curvature.
  - WRF with a resolution of ~1 degree, and dynamically downscale the data using weather research forecasting model (WRF) so we can view predicted atmospheric variables at 12 km resolution
  - Climate-Weather Research Forecasting model (CWRF, Liang et al. 2012) to examine uncertainties in the treatment of cloud, aerosol and radiative transfer processes
- PRAC 338 Million core hours

Rauber, Washington, et. al. – Climate Change Uncertainties

- Challenge Goals
  - Better quantifying future regional climate change by running a high-resolution (0.25 degree atmosphere, 1 degree ocean) global coupled climate model CESM)
  - In the framework of CMIP6 to meet the needs for the next generation assessments of climate change specific studies that understanding the effects of small-scale regional features and interactions across spatial scales in climate through even higher-resolution
- PRAC Allocation - 128M core hours

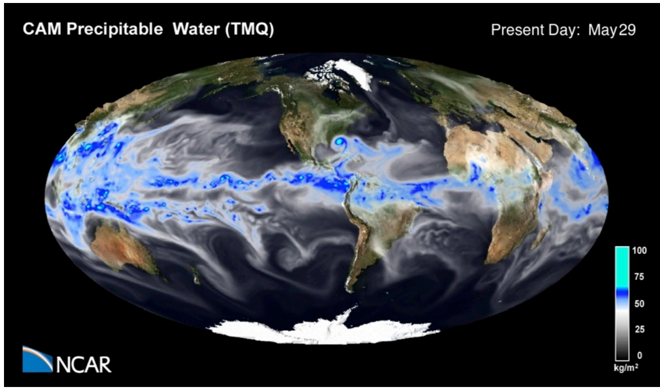

Wilhemson and Orf - Numerical Simulations of Supercells and Embedded Tornadoes

- Challenge Goal
  - first time to simulate the multiscale interaction between the supercell and a tornado using high resolution that is needed to adequately represent boundary layer impacts including the shallow inflow into the tornado, to represent the remarkably thin precipitation curtains that form the “hook echo” adjacent to the tornado, and to study surface friction effects on tornado development.
  - The CM1 numerical storm model will be used to study the long-lived EF-5 El Reno supercell/ tornado on May 24, 2011.
- Usage/Results
  - See Leigh Orf’s presentation at this conference
- Illinois Allocation – 52.2M core hours (1,632,299 MNH)

Reed - Optimizing EOS Satellite Design

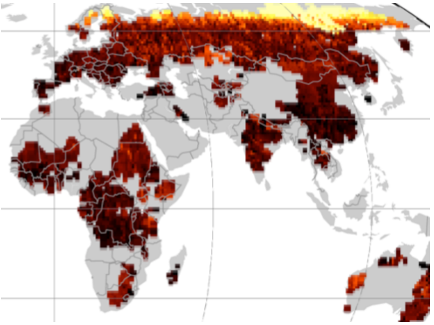

Blue Waters made possible "… the largest Monte Carlo simulation experiment for evaluating the required satellite frequencies and coverage to maintain acceptable global forecasts of terrestrial hydrology (especially in poorer countries)."

- Challenge Goals
  - Understand and optimize space based sensor systems for hydrology
- Usage/Accomplishments to Date
  - At 524,288 cores, our search approaches theoretically ideal performance. These results are the best benchmark results ever attained
  - Draim results reveal that carefully optimizing an initial orbital geometry to exploit natural perturbations (e.g., effects of sun, moon, etc.) to maintain continuous global coverage performance
  - Show how limits in satellite-based precipitation observations propagate to uncertainties in surface runoff, evaporation, and soil moisture at distinctly different locations globally.
- Blue Waters Help
  - Debugged application after unintentional changes to code affected file system metadata responsiveness when run at 3,080 nodes (49,280 FP cores). Required coordinated effort from storage and application groups.
  - IO library (netcdf) issues.
- Impacts
  - World Hydrological Research and Response
  - Cost Effectiveness of EOS
- NSF Allocation – 152M core*hours used

MSMP

BIG DATA

Morin - ArcticDEM

Morin - Never envisioned being able to image an entire country in a weekend

- Challenge Goals
  - Produce a 2-8m digital stereo surface model of the Arctic during the US Chairmanship of the Arctic Council using DigitalGlobe stereo imagery licensed by the National Geospatial-Intelligence
  - Presidential Announcement
  - Multi Agency
    - National Geospatial-Intelligence Agency – Manages the Digital Globe project and providing data and tasking to NSF/PGC; National Science Foundation - Providing Blue Waters HPC resources and funding to PGM to generating DEMs and managing the project; NASA – Provided funding for the development of the DEM extraction software packages ASP and SETSM; USGS – Overall coordination; Google and ESRI – Data distribution
- Allocations – Illinois - 32M core hours initial – expect 400M cores by end

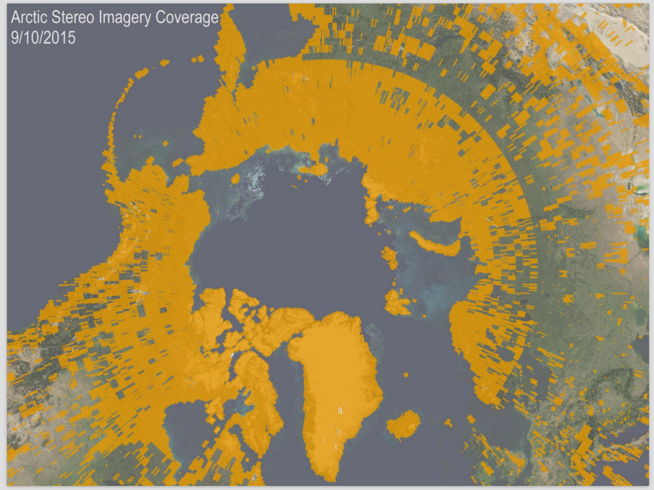

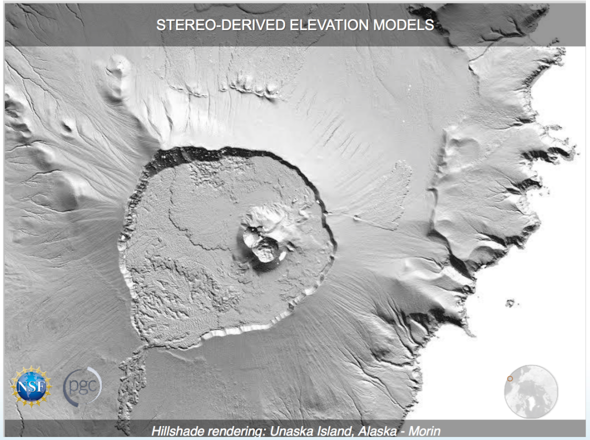

BIG DATA

Larry Di Girolamo - 3D Radiative Transfer Model Coupled to the Weather Research and Forecasting Model

- Challenge Goals
  - Understood aspect of the weather and climate system is the impact of Earth's clouds on solar and terrestrial radiative forcing of the Earth–Atmosphere system, and the subsequent feedback onto the dynamics of the system.
  - Couple a full 3D RT model with a high resolution cloud dynamics model to study and understand how errors from the crude RT approximations used in the past feed back on cloud properties and their evolution
  - Migrate and reprocess MODIS satellite data to consistent HDF5 formats – 6 DAACs

- Illinois Allocation – 15.3 M core hours

Hurricane Igor in the Atlantic Ocean 09-17-2010

Karimbadi - Understanding of Collisionless Plasmas Enabled through Petascale Kinetic Simulations

- Challenge Goal
  - Better understand Earth's space environment and its interaction with the Sun
  - Global fully kinetic and hybrid simulations to understand the response of the magnetosphere to external solar perturbations.
    - Such simulations are critical for development of accurate space weather forecasting capabilities.
- Major results to date:
  - Global fully kinetic simulations of magnetic reconnection
  - First large-scale 3D simulations of decaying collisionless plasma turbulence
  - 3D global hybrid simulations addressing coupling between shock physics & magnetosheath turbulence

NSF Allocation – 441.6 M core*hours used

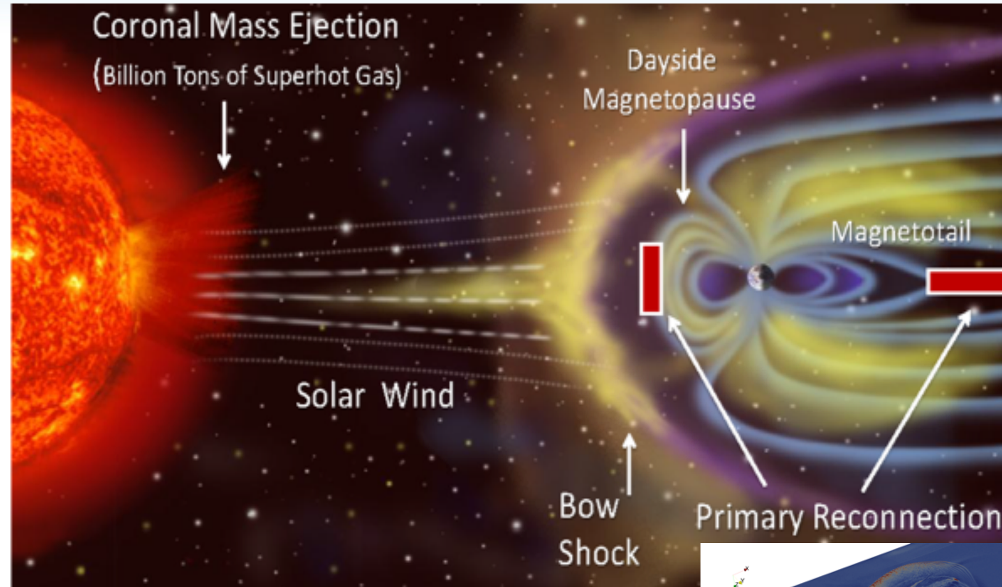

90 million miles or ~ 100 Suns

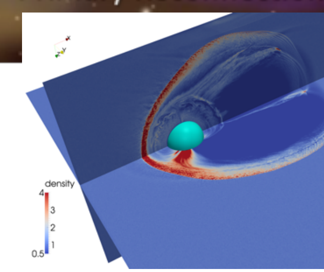

Fully kinetic simulation (all species kinetic; code: VPIC)
~up to 10^{10} cells
~up to $4x10^{12}$ particles
~120 TB of memory
~10^7 CPU-HRS
~up to 500,000 cores

Large scale hybrid kinetic simulation: (kinetic ions + fluid electrons; codes: H3D, HYPERES)
~up to $1.7x10^{10}$ cells
~up to $2x10^{12}$ particles
~130 TB of memory

MSMP

CADENS

BIG DATA

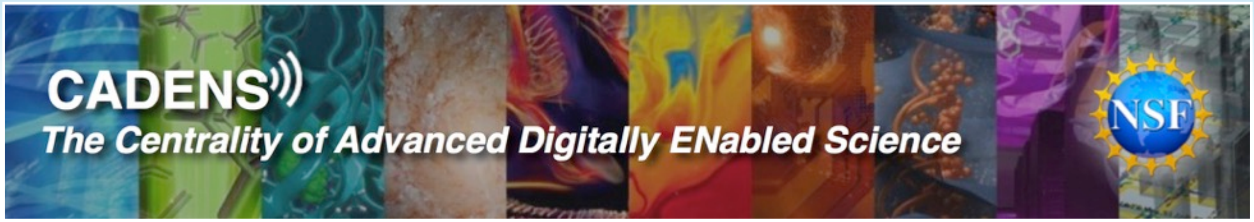

Primary Goals:

(1) To increase digital literacy and raise public awareness about the Centrality of Advanced Digitally ENabled Science (CADENS)

(2) To expand visualization opportunities for science teams and expand outreach opportunities for visualization experts

Develop and widely distribute
3 Digital Fulldome Museum Shows
Develop online supplemental materials for teachers, K-16, and make them available for public; translate into languages for global distribution

Develop and widely distribute
9 High-Definition TV Documentaries

Promotion and co-marketing campaign

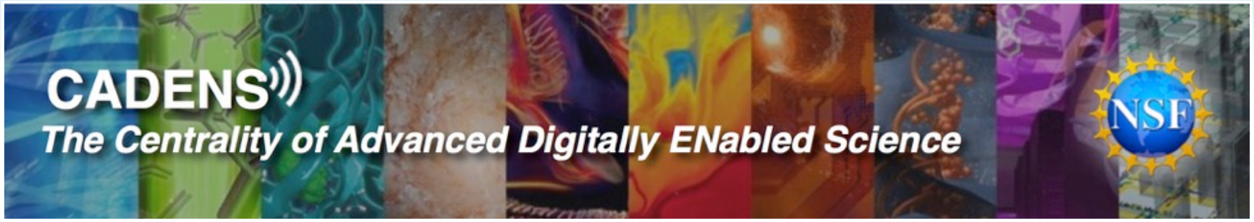

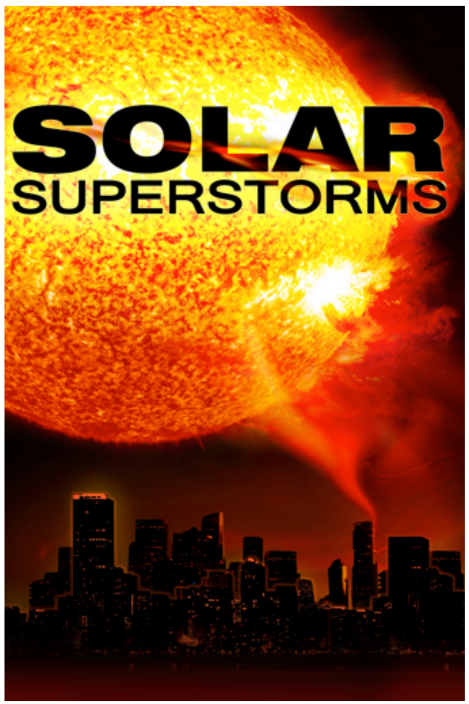

"Solar Superstorms"
First CADENS Digital Museum Show
Release Summer 2015
Narrated by Benedict Cumberbatch

HOMA KARIMABADI

SOLAR WIND AND THE MAGNETOSPHERE

INSTITUTION	UNIVERSITY OF CALIFORNIA AT SAN DIEGO
SPATIAL SCALE	100 KILOMETER GRID CELLS
TIME SCALE	~30 MINUTES
RESOLUTION	2D: 2048 X 8192 PIXELS 3D: 2048 X 4096 X 2048
DATA SIZE	2D: 220 TIMESTEPS, 30 GB 3D:47 TIMESTEPS,12 TB
COMPUTED AT	BLUE WATERS, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

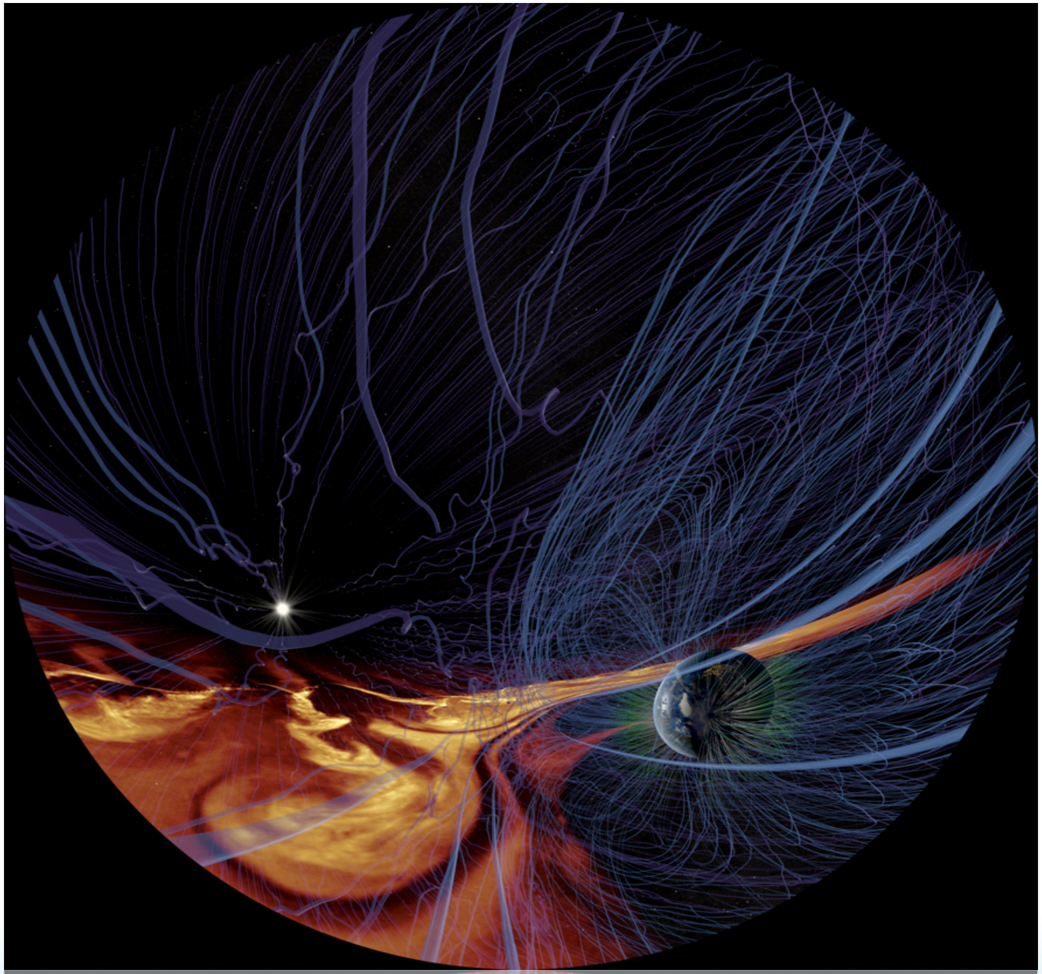

- Simulation is massive so it cannot be rerun for better temporal resolution
- Interpolation of 2D and derived 3D data
- 3D rendered with Yt on Blue Waters supercomputer
- Best fit for magnetic field lines
- Blue Marble Earth with lit nighttime cities

HOMA KARIMABADI

SOLAR WIND AND THE MAGNETOSPHERE

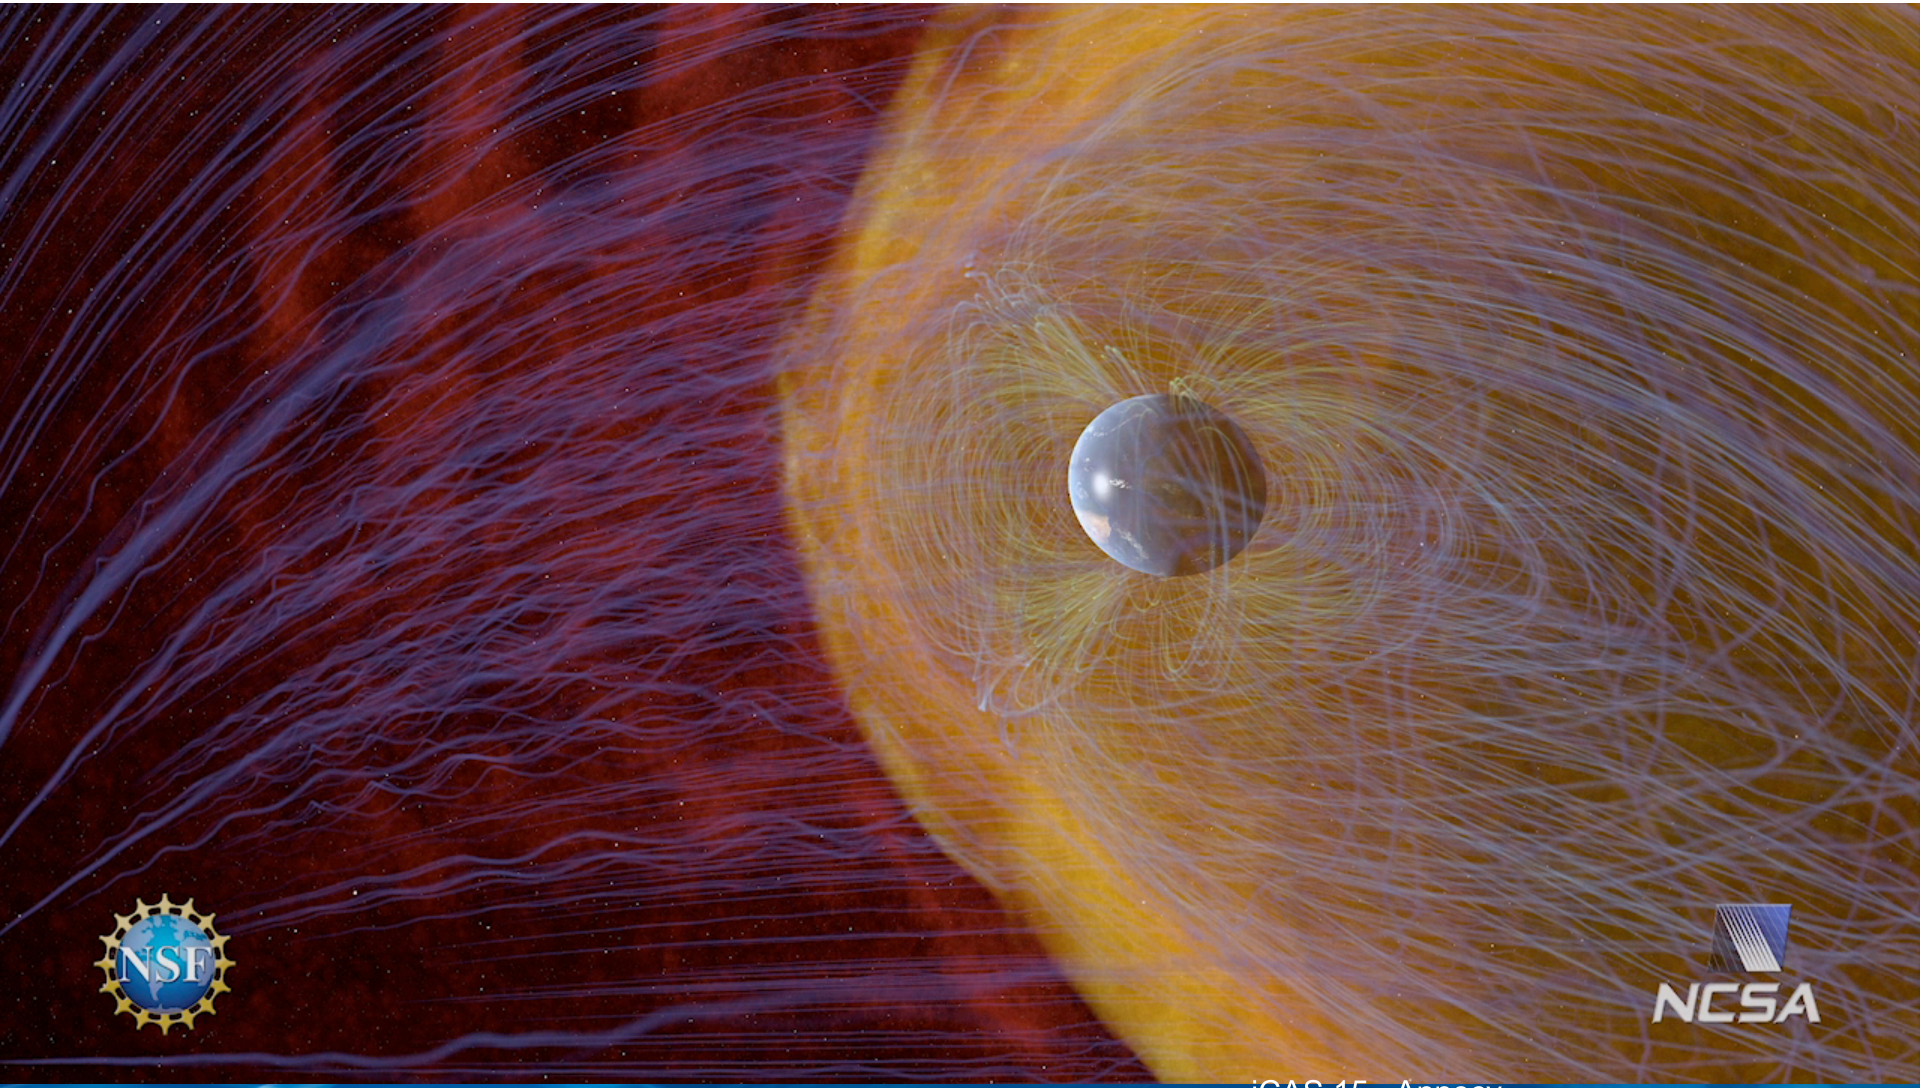

JOHN H. WISE

SOLAR WIND AND THE MAGNETOSPHERE

INSTITUTION	GEORGIA TECH
SPATIAL SCALE	UP TO 50,000 PARSECS ACROSS
TIME SCALE	200 MILLION YEARS
RESOLUTION	AMR SPARSE GRID 256^3 12 LEVELS OF REFINEMENT UP TO 330,000 PARTICLES
DATA SIZE	330 TIMESTEPS, 2 TB
COMPUTED AT	BLUE WATERS, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

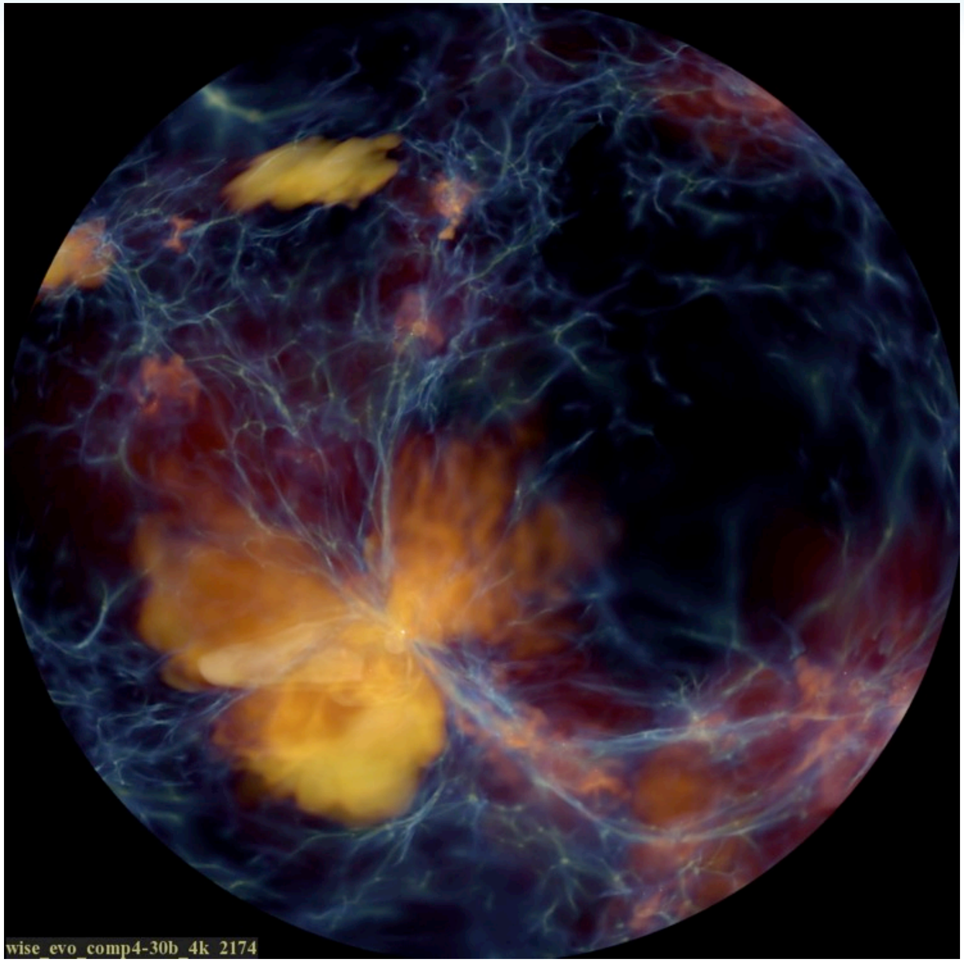

- Star-like particles
- Nested resolution volume
- Many data attributes
- Rendering on Blue Waters with Yt

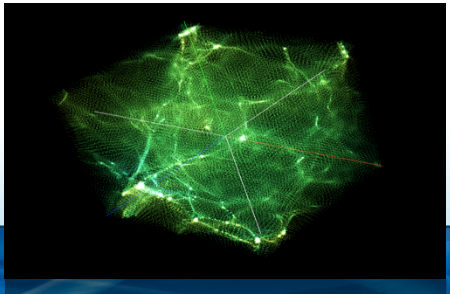

JOHN H. WISE

SOLAR WIND AND THE MAGNETOSPHERE

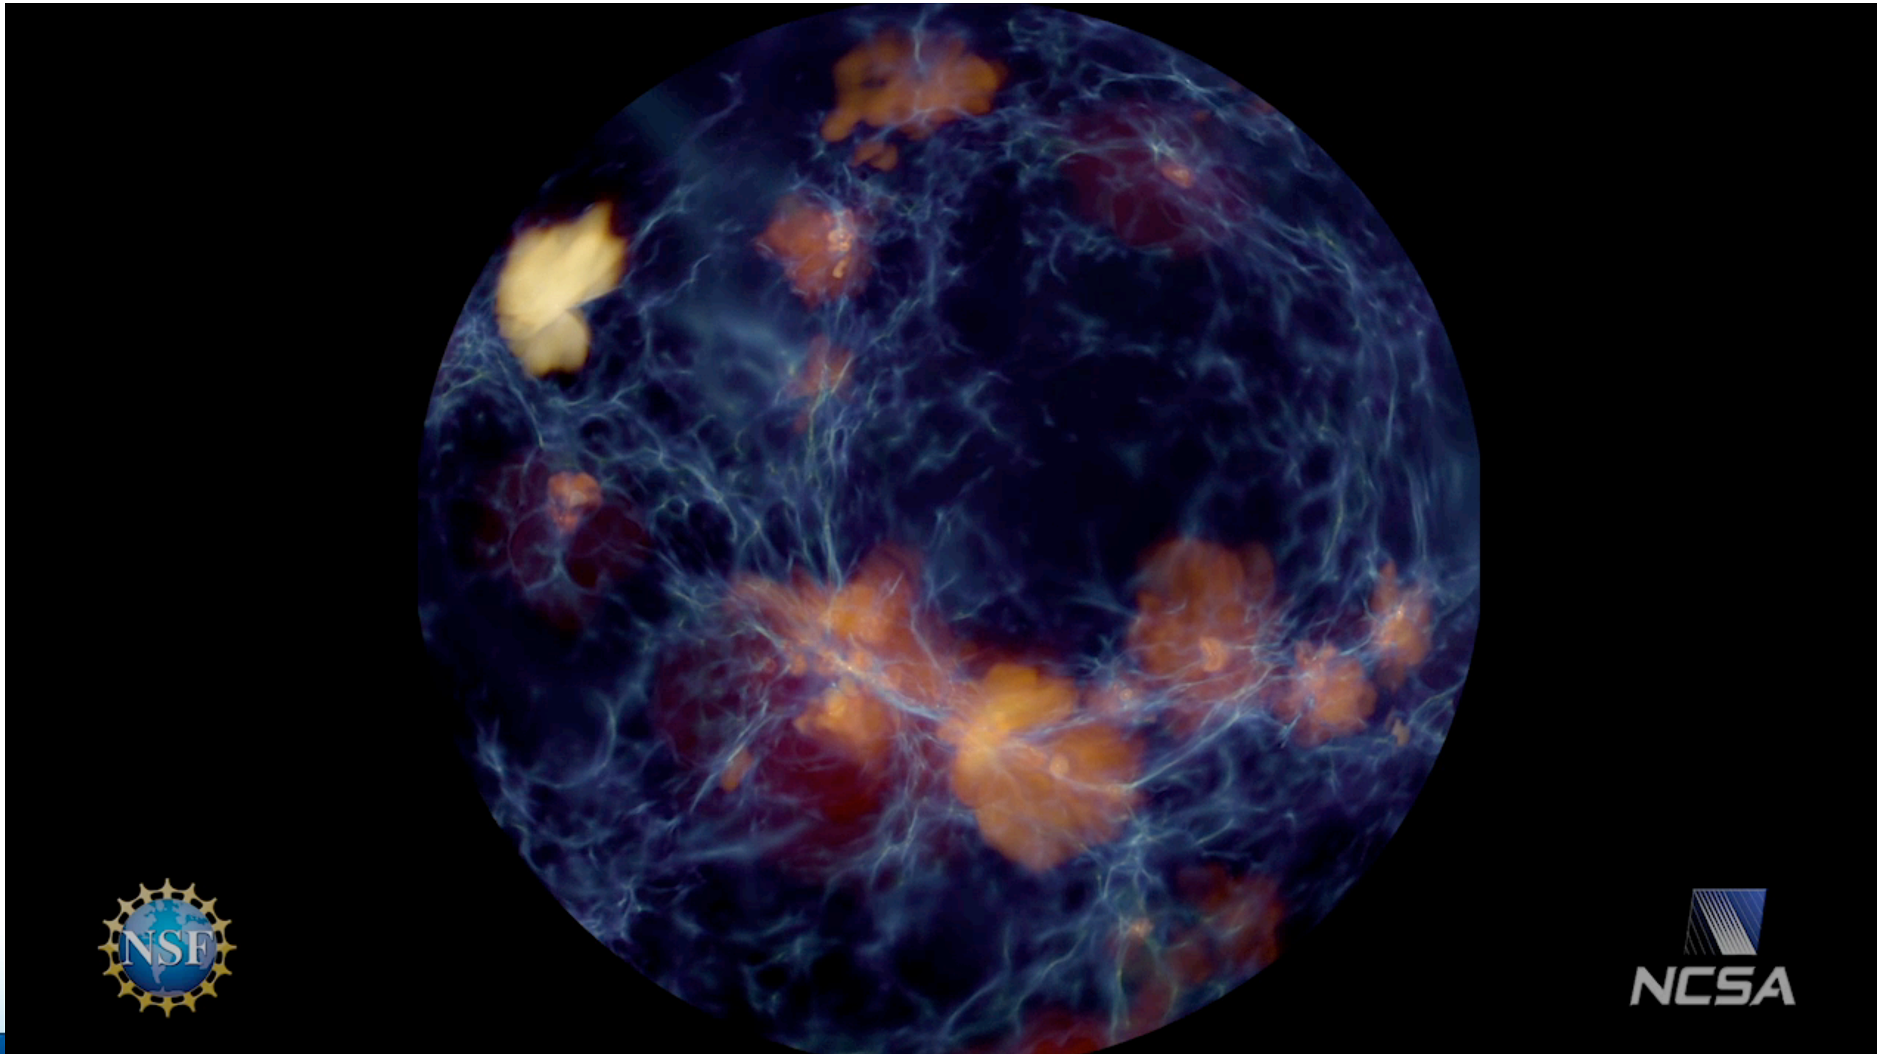

Jordan - Earthquake System Science

"Using the well-balanced system capabilities of Blue Waters to complete CyberShake calculations within weeks rather than months."

Challenge Goals

Develop 3D physics-based earthquake simulations for the urban regions of California that are more accurate than the empirical NSHMP standard

Impact

- Building Codes
- Seismic Retrofit of 400+ dams and other infrastructure
- Disaster Response

Usage/Approach

- Accurate 3D model with the 1D model illustrates the importance of complex geological structures in governing the amplitudes of strong ground motions.
- Reducing the total CyberShake makespan from ~61 days to ~14 days.
- Hazard maps for 2% probability of exceedance in 50 years from the CyberShake 14.2 study are shown in Figure 1.

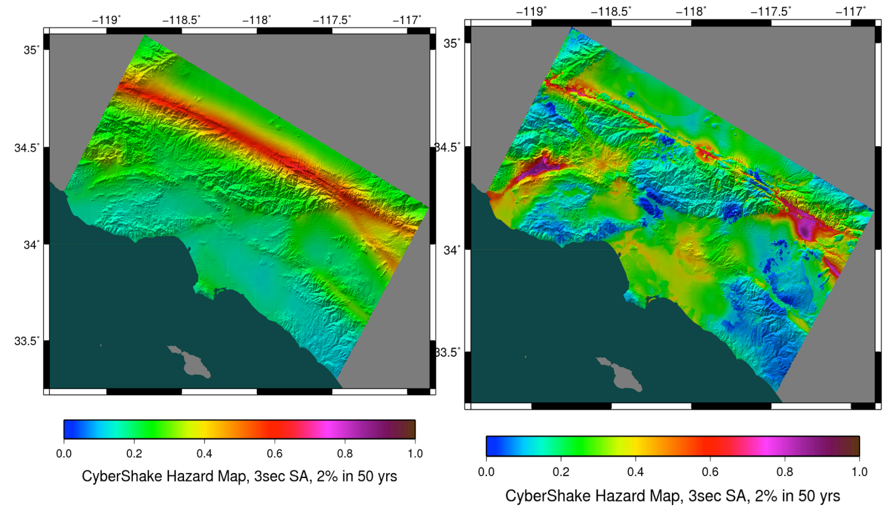

Next Generation Challenges

- Higher Frequency to model single level structures
- Expanded areas
- Higher Fidelity

NCSA BLUE WATERS' REACH

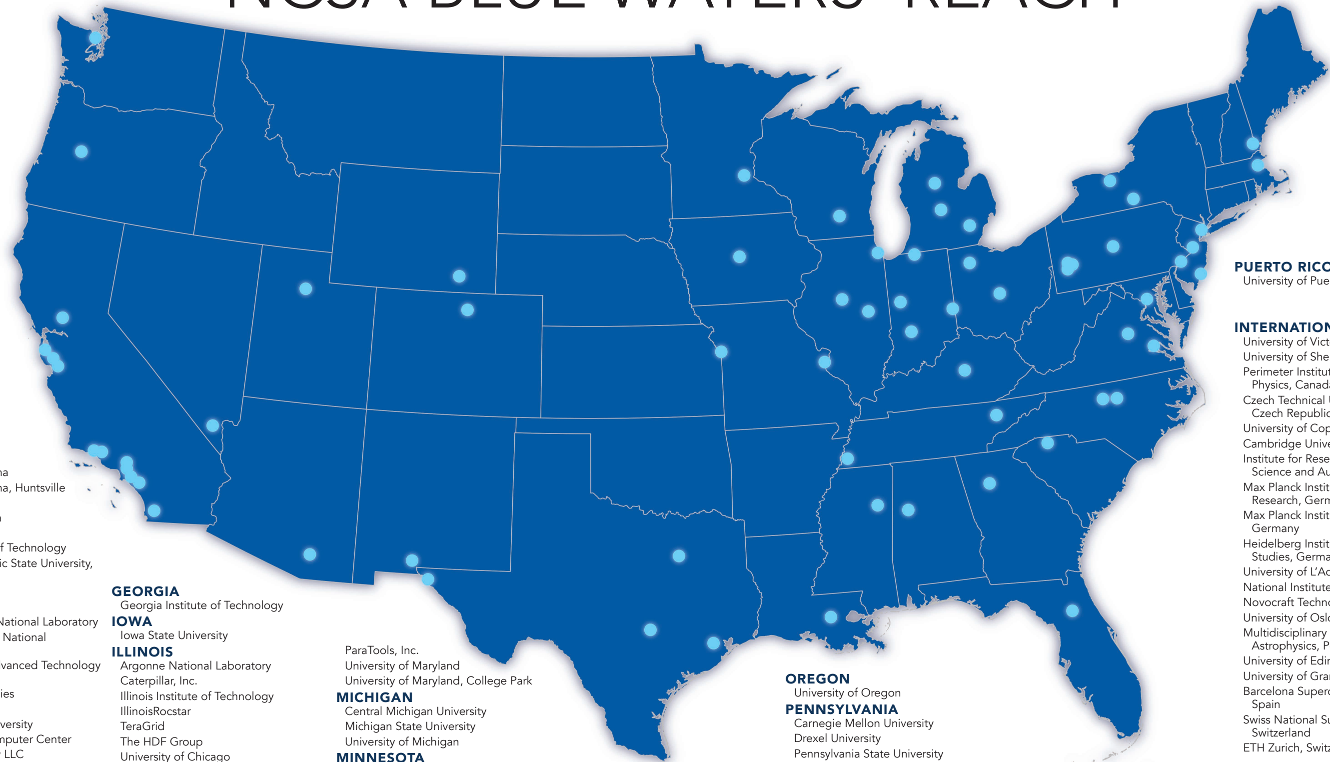

ALABAMA
University of Alabama
University of Alabama, Huntsville

ARIZONA
University of Arizona

CALIFORNIA
California Institute of Technology
California Polytechnic State University, San Luis Obispo
City of Hope
Hoover Institution
Lawrence Berkeley National Laboratory
Lawrence Livermore National Laboratory
Lockheed Martin Advanced Technology Center
Mellanox Technologies
Nvidia Corporation
San Diego State University
San Diego Supercomputer Center
Seagate Technology LLC
Stanford University
University of California, Berkeley
University of California, Davis
University of California, Irvine
University of California, Los Angeles
University of California, San Diego
University of California, San Francisco
University of California, Santa Barbara
University of California, Santa Cruz
University of Southern California
Xyratex

COLORADO
Allinea
Colorado School of Mines
Colorado State University
National Center for Atmospheric Research
University Corporation for Atmospheric Research

CONNECTICUT
Yale University

FLORIDA
University of Florida

GEORGIA
Georgia Institute of Technology

IOWA
Iowa State University

ILLINOIS
Argonne National Laboratory
Caterpillar, Inc.
Illinois Institute of Technology
IllinoisRocstar
TeraGrid
The HDF Group
University of Chicago
University of Illinois at Chicago
University of Illinois at Urbana-Champaign
University of Southern Denmark
Western Illinois University

INDIANA
Earlham College
Indiana University
Purdue University
University of Notre Dame

KENTUCKY
Centre College
University of Kentucky

LOUISIANA
Louisiana State University

MASSACHUSETTS
Harvard University

MARYLAND
Center for Ocean-Land-Atmosphere Studies
Institute of Global Environment and Society
Johns Hopkins University
ParaTools, Inc.
University of Maryland
University of Maryland, College Park

MICHIGAN
Central Michigan University
Michigan State University
University of Michigan

MINNESOTA
Fond du Lac Tribal and Community College
Mayo Clinic
University of Minnesota
University of Minnesota, Twin Cities

MISSOURI
University of Missouri, Kansas City

MISSISSIPPI
Mississippi State University

NORTH CAROLINA
North Carolina State University at Raleigh
Shodor Education Foundation, Inc.
University of North Carolina, Chapel Hill
University of North Carolina, Charlotte

NEW HAMPSHIRE
University of New Hampshire

NEW JERSEY
Princeton Plasma Physics Laboratory
Princeton University
Richard Stockton College of New Jersey
Rutgers, the State University of New Jersey

NEW MEXICO
Los Alamos National Laboratory
New Mexico State University
Sandia National Laboratories

NEVADA
University of Nevada-Las Vegas

NEW YORK
Brookhaven National Laboratory
Cornell University
IBM
Kitware, Inc.
Memorial Sloan Kettering Cancer Center
Rochester Institute of Technology
SUNY at Stony Brook

OHIO
Bluffton University
Ohio State University
Ohio Supercomputer Center
Procter & Gamble Co

OKLAHOMA
University of Tulsa

OREGON
University of Oregon

PENNSYLVANIA
Carnegie Mellon University
Drexel University
Pennsylvania State University
Pittsburgh Supercomputing Center
Slippery Rock University
Spiralgen, Inc.
Temple University
University of Pittsburgh

RHODE ISLAND
Dassault Systems Simulia Corp

SOUTH CAROLINA
Wofford College

SOUTH DAKOTA
South Dakota State University

TENNESSEE
National Institute for Computational Sciences
Oak Ridge National Laboratory
University of Memphis
University of Tennessee, Knoxville
University of Tennessee, Oak Ridge

TEXAS
Southern Methodist University
University of Houston-Clear Lake
University of Texas at Austin
University of Texas at El Paso
University of Texas Medical Branch at Galveston

UTAH
Adaptive Computing, Inc.
Brigham Young University
University of Utah

VIRGINIA
Aerospace Corporation
College of William and Mary
Jefferson Laboratory
Old Dominion University
University of Mary Washington
University of Virginia
Virginia Polytechnic Institute and State University

WASHINGTON
Cray, Inc.
University of Washington

WEST VIRGINIA
West Virginia University

WYOMING
University of Wyoming

PUERTO RICO
University of Puerto Rico, Mayaguez

INTERNATIONAL
University of Victoria, Canada
University of Sherbrooke, Canada
Perimeter Institute for Theoretical Physics, Canada
Czech Technical University in Prague, Czech Republic
University of Copenhagen, Denmark
Cambridge University, England
Institute for Research in Computer Science and Automation, France
Max Planck Institute for Solar System Research, Germany
Max Planck Institute for Astrophysics, Germany
Heidelberg Institute for Theoretical Studies, Germany
University of L'Aquila, Italy
National Institute for Astrophysics, Italy
Novocraft Technologies, Malaysia
University of Oslo, Norway
Multidisciplinary Center for Astrophysics, Portugal
University of Edinburgh, Scotland
University of Granada, Spain
Barcelona Supercomputing Center, Spain
Swiss National Supercomputing Centre, Switzerland
ETH Zurich, Switzerland

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

2014-2015 XE Scale

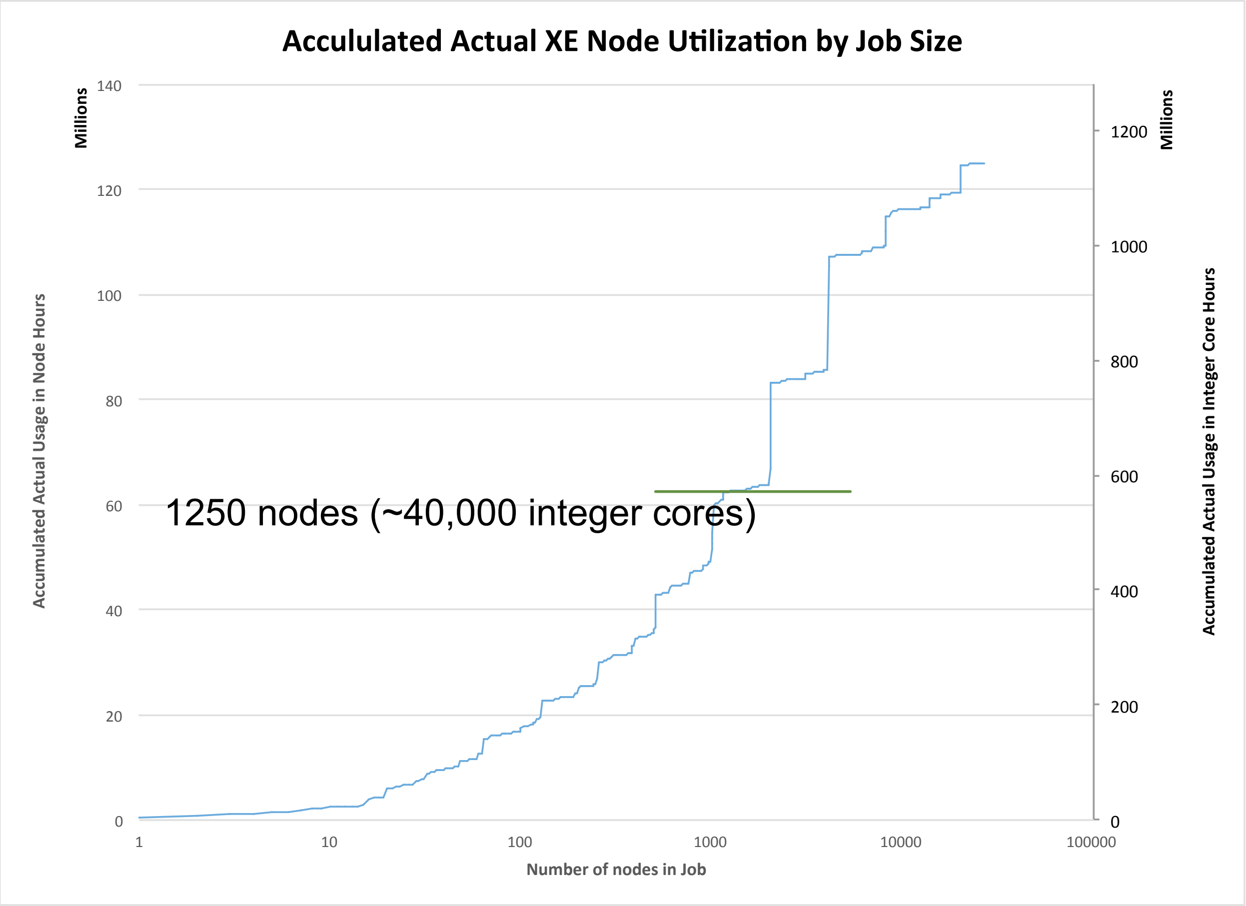

2014 XK Scale

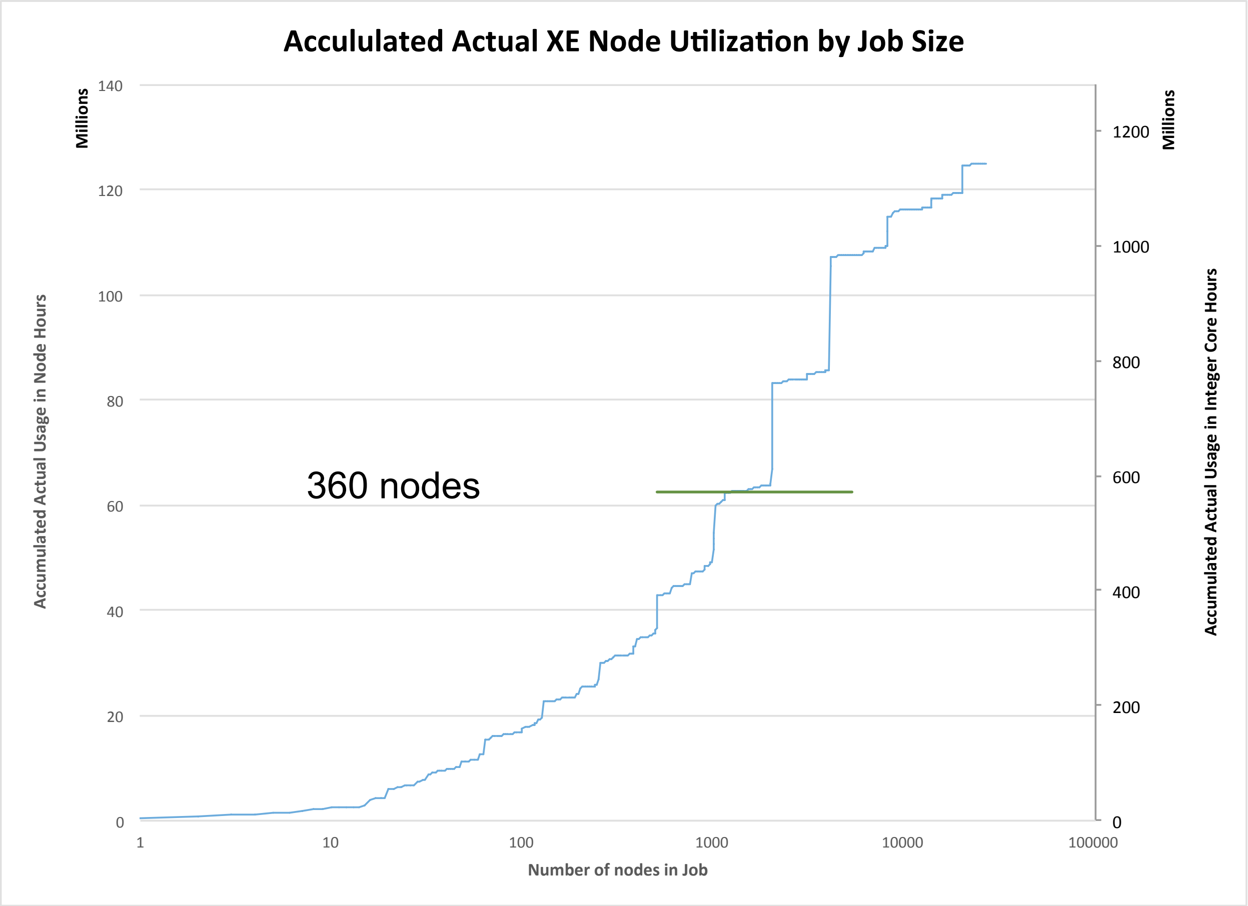

Long Running Very Large Applications

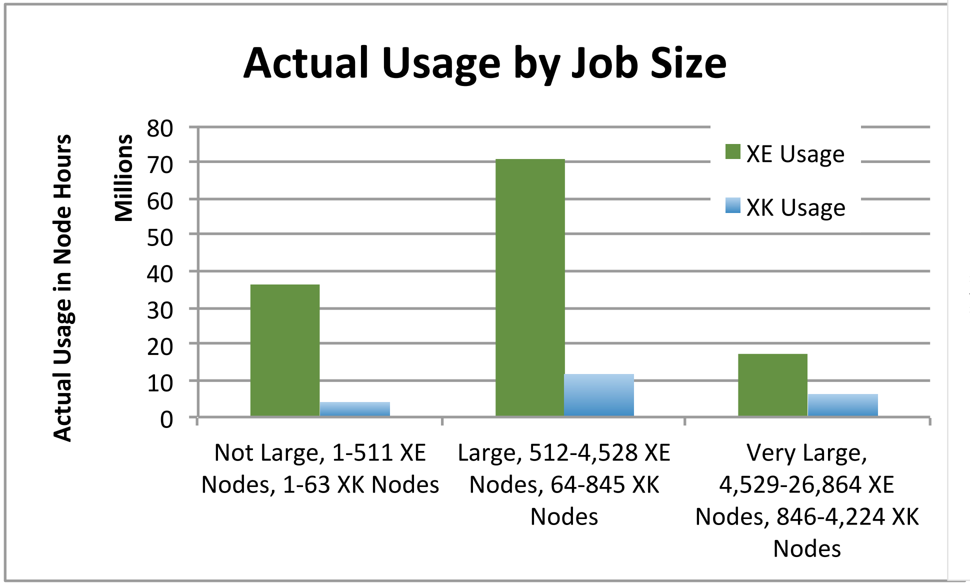

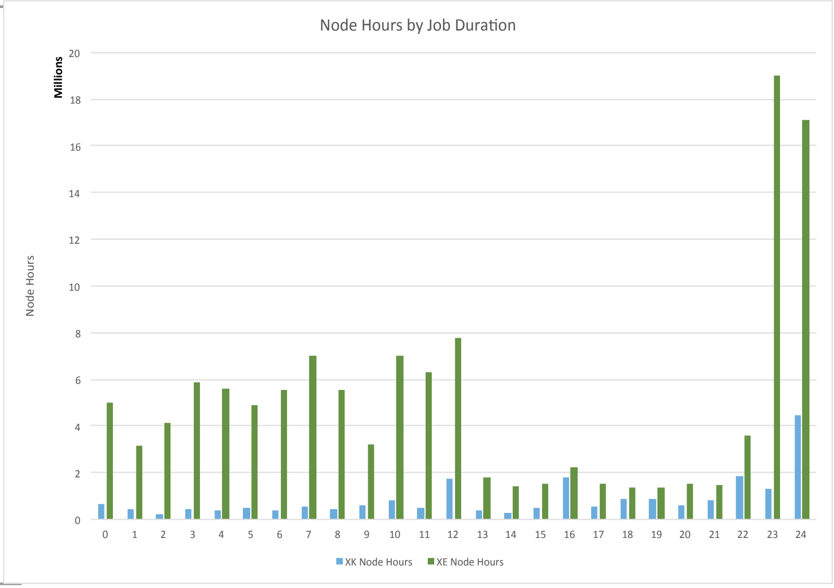

Total Time – April 2, 2013 to April 1, 2015

Blue Waters

CPU Time (h)		Wait Time (h)	Wall Time (h)		Processors	
Total:	Avg (Per Job):	Avg (Per Job):	Total:	Avg (Per Job):	Max:	Avg (Per Job):
10,151,114,503.3	9,127.85	6.30	3,046,024.5	2.74	859,648	3,525

XSEDE

Service (XD SU)		CPU Time (h)		Wait Time (h)	Wall Time (h)		Processors	
Total:	Avg (Per Job):	Total:	Avg (Per Job):	Avg (Per Job):	Total:	Avg (Per Job):	Max:	Avg (Per Job):
9,663,323,328.0	1,119.41	2,748,979,021.8	318.45	4.02	27,621,089.1	3.20	98,304	110

Usage

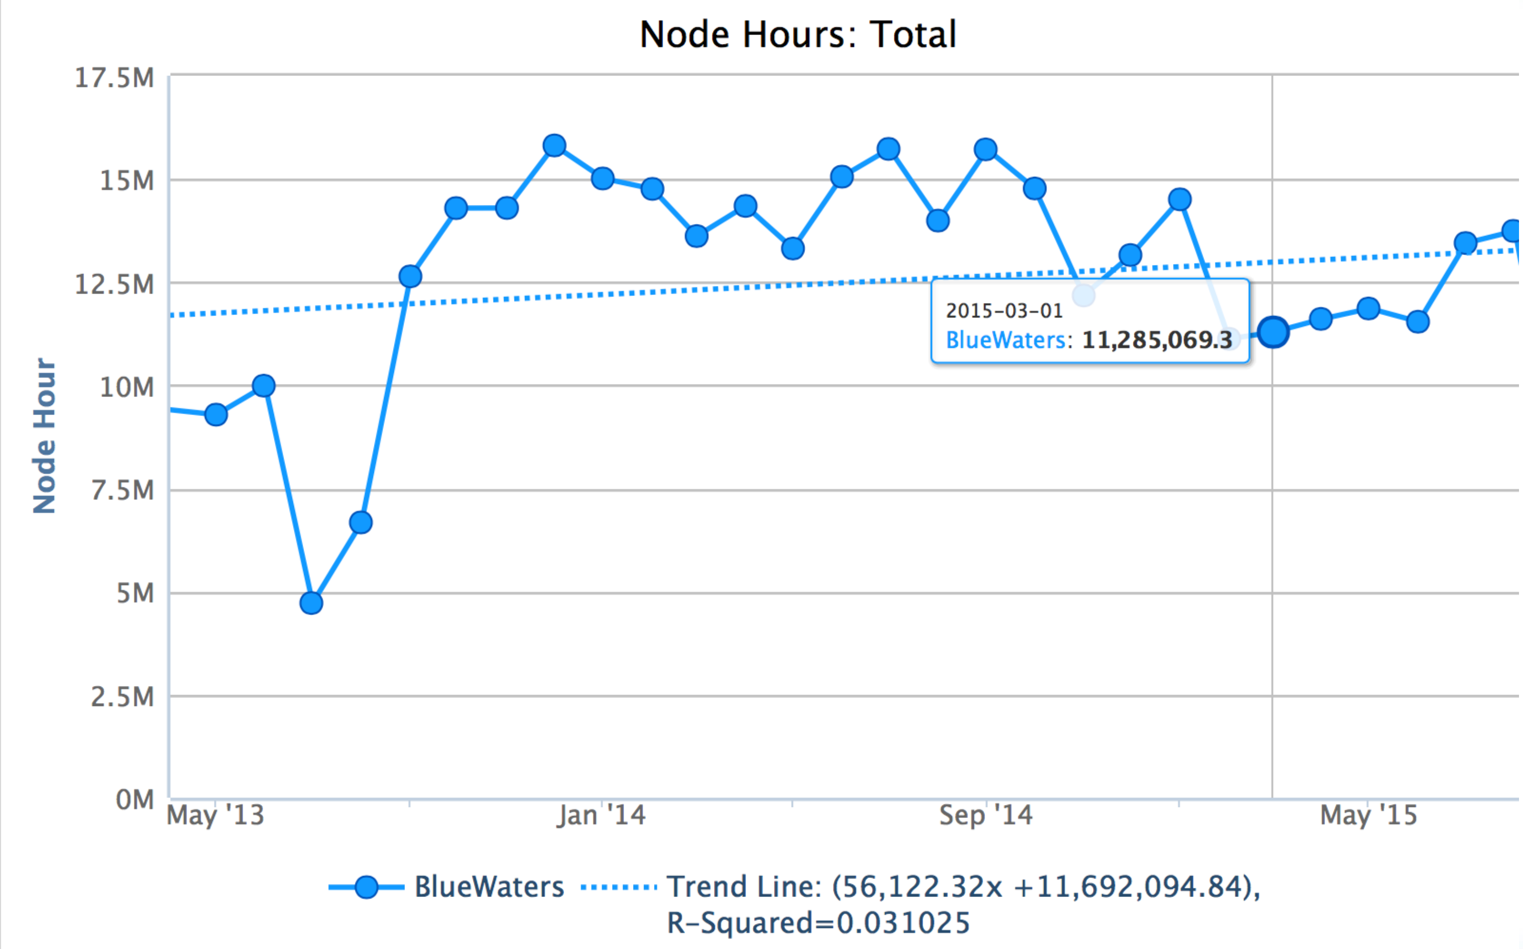
Node Hours: Total
17.5M
15M
12.5M
10M
7.5M
5M
2.5M
0M
Node Hour
May '13
Jan '14
Sep '14
May '15
2015-03-01
BlueWaters: 11,285,069.3
BlueWaters
Trend Line: (56,122.32x +11,692,094.84),
R-Squared=0.031025

Largest Jobs

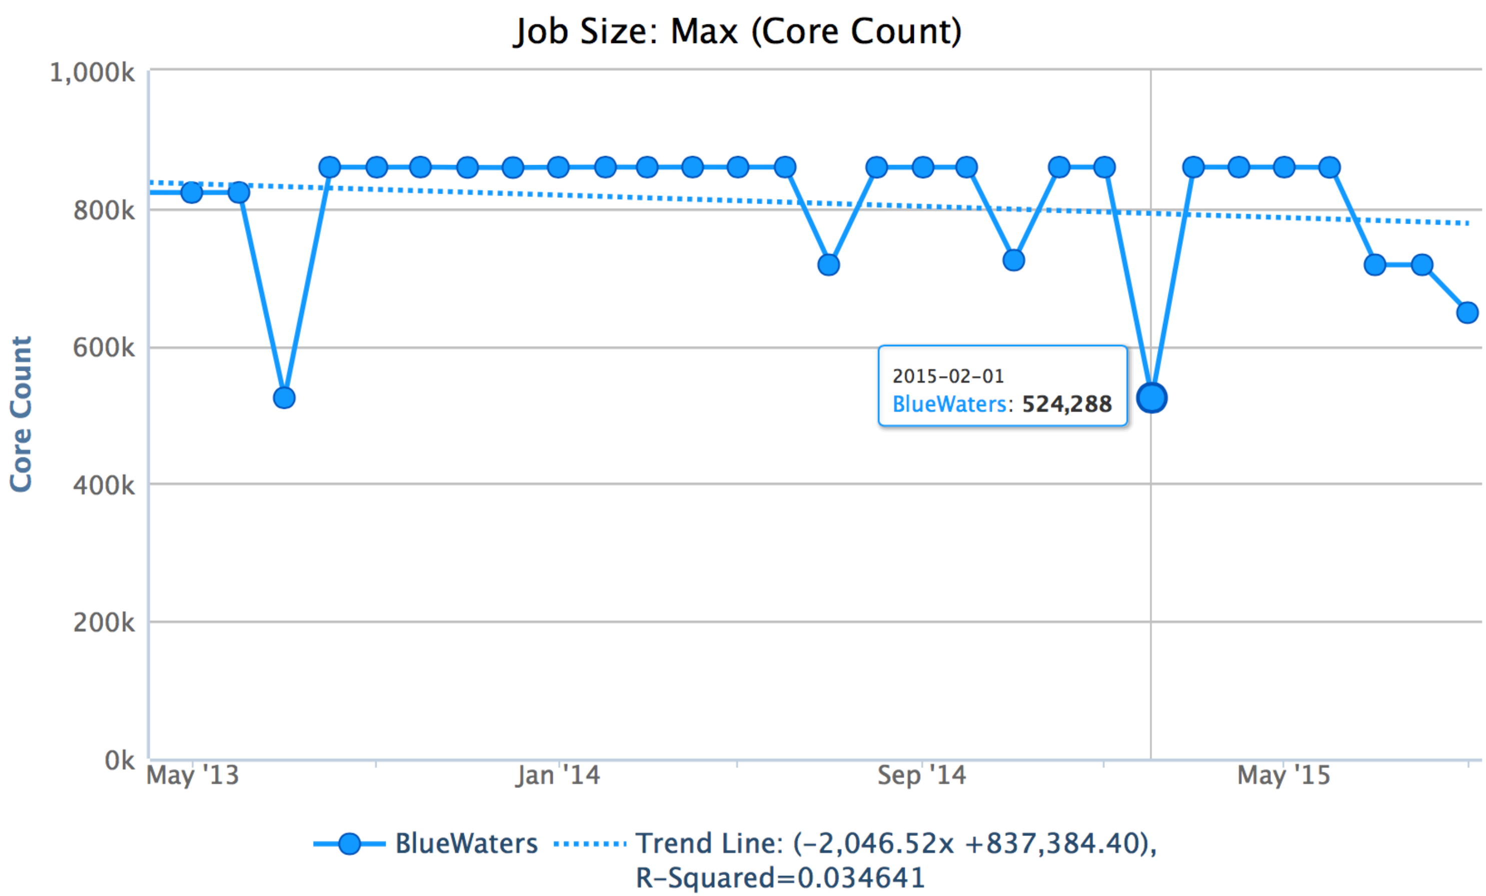

Example Views of Job Behavior

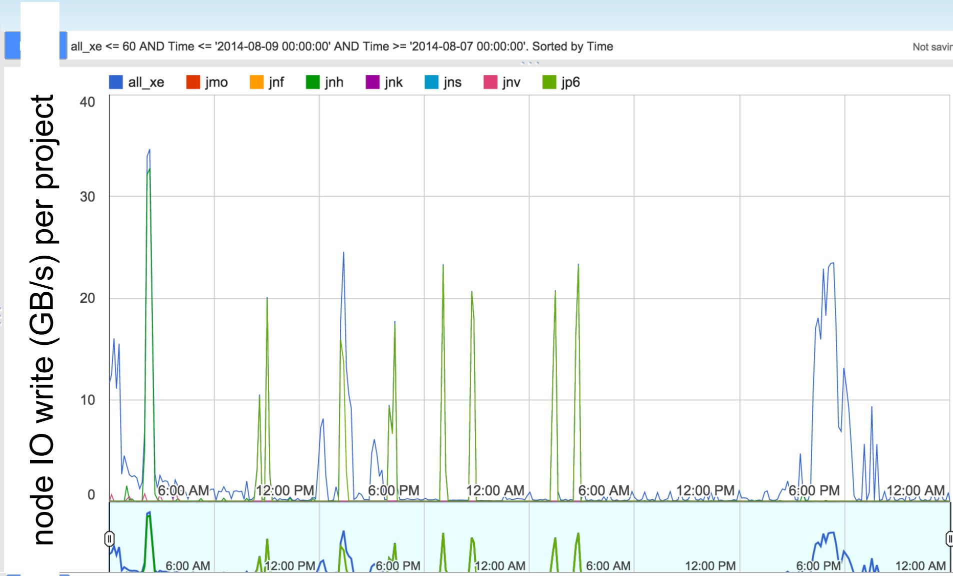

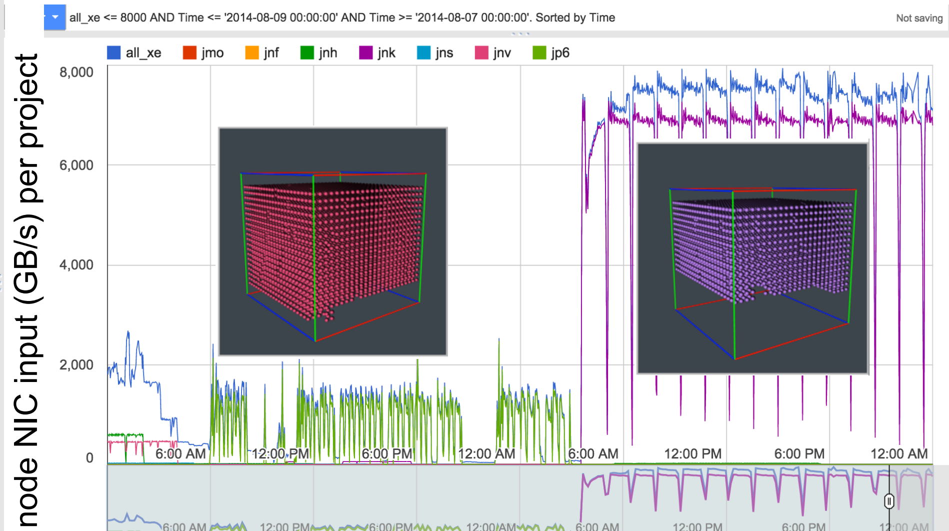

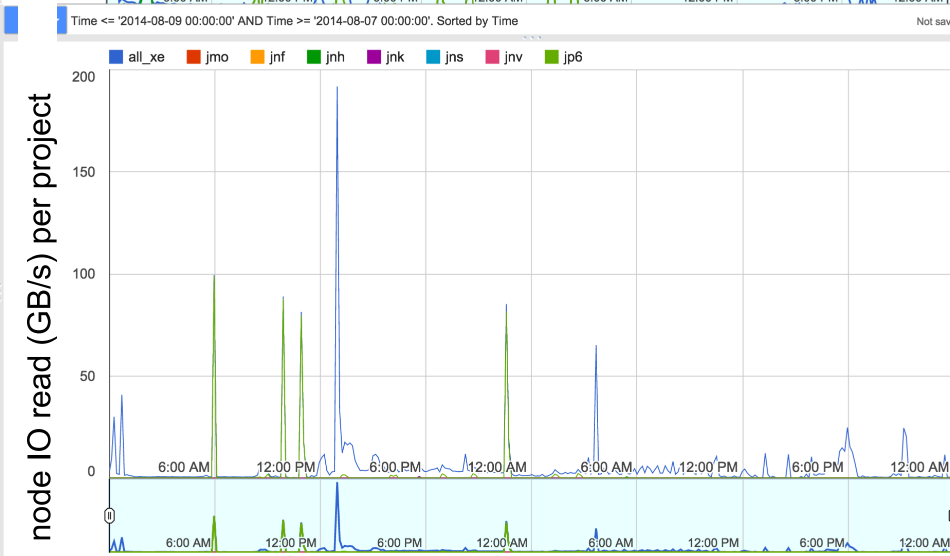

20,250 XE node P-gadget job 15,872 XE node namd job

Usage by NSF PRAC team – A Behavior Experiment

An observed experiment – teams self select what is most useful

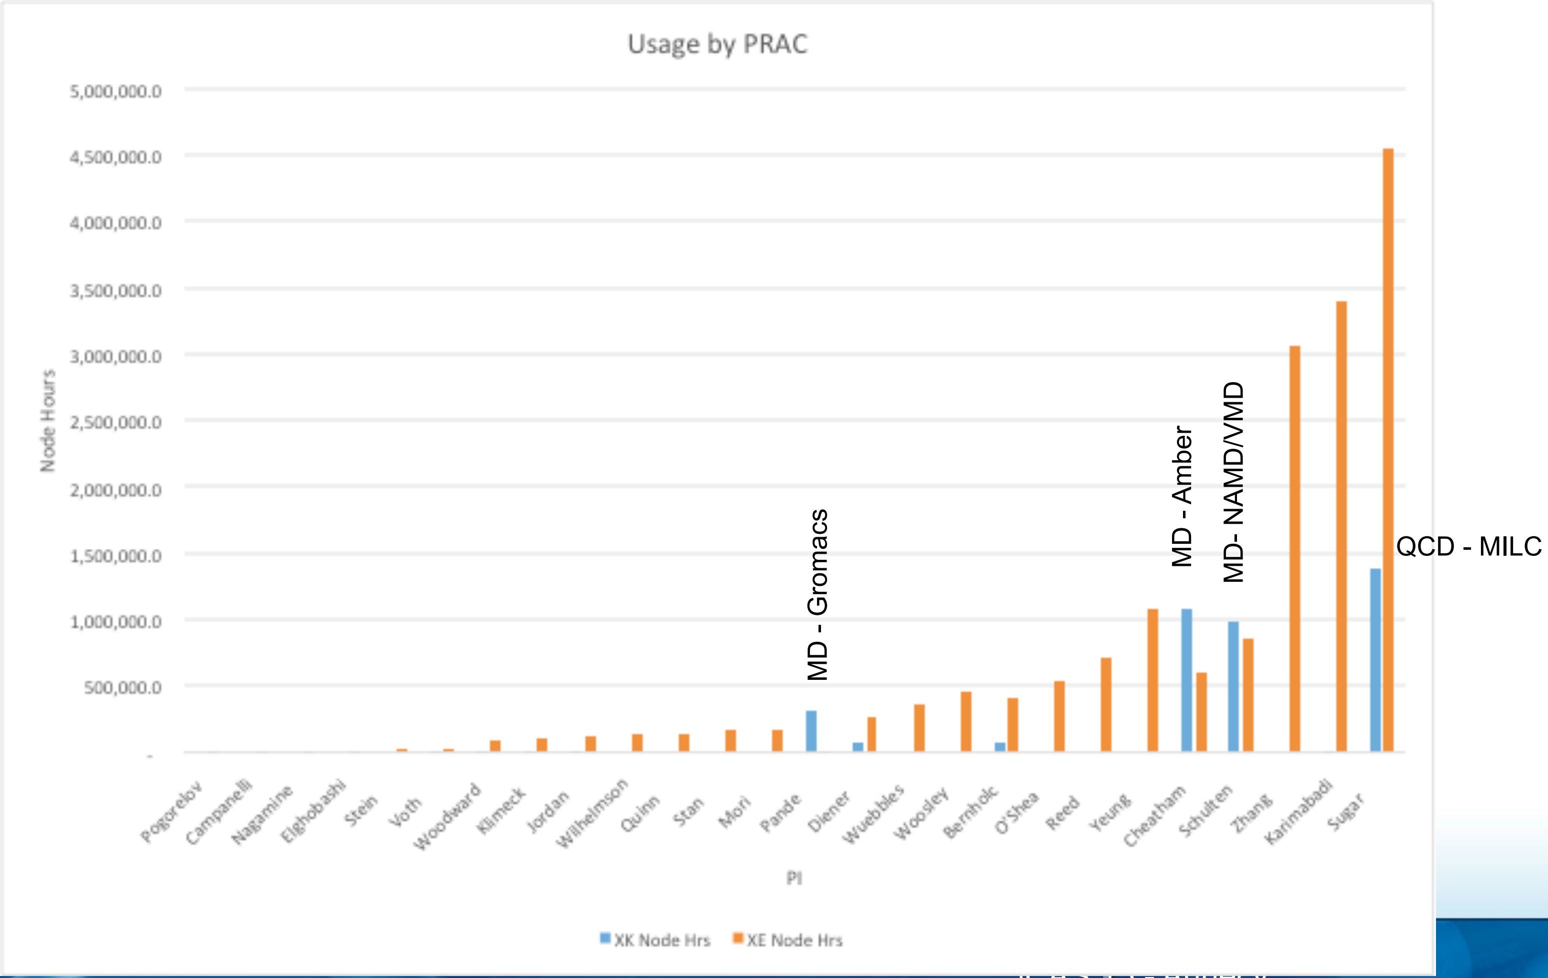

Usage by NSF PRAC team – A Behavior Experiment

An observed experiment – teams self select what is most useful

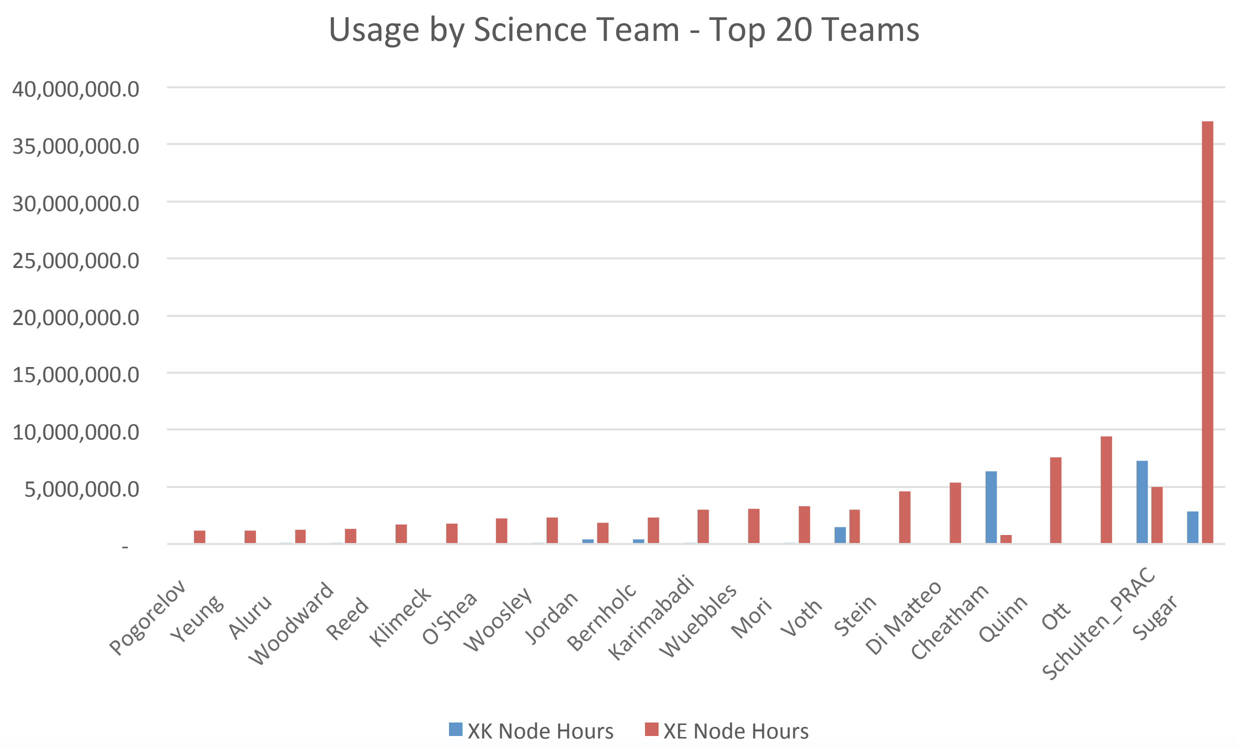

Recent and Near Future Work

- Topology Aware Scheduling – a major improvement for applications and systems
  - Faster Applications
  - Better Consistency
  - E.g >25% increase in injected network
  - Currently assessing the tradeoffs
- Collecting over 8 billions datums a day for every aspects of the system
- Data usage and preservation
- Lustre HSM to HPSS
- Network Upgrade
- Data Sharing Service

Characteristics for Future @Scale Use

- Increasing range of use and needs
- Dramatically increase fidelity in models and simulations to improve insights and address new problems.
- Longer simulated time periods
- Increased number of problems to address
- Changing workflow methods
- Increased integration with data sources and increased use of simulation data products.
- Changing algorithmic methods
- Industrial use models

Inference Spiral of System Science and Research

A wide range of coupling of data, computing and analysis

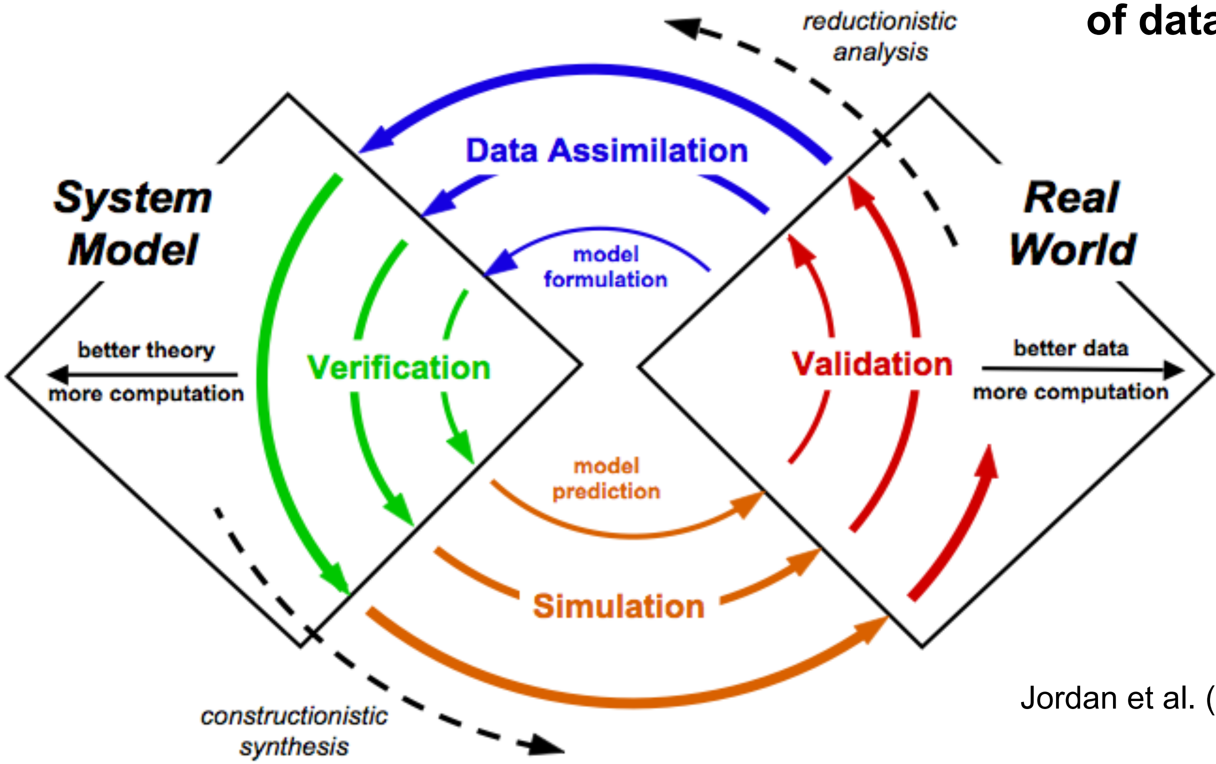

Jordan et al. (2010)

- **As models become more complex and new data brings in more information, increasing computational and analysis resources are necessary**

Slide courtesy of T Jordan - SCEC

Flow of Influence

- It often takes tremendous computing and analysis power to develop new ways to solve the most challenging problems
- Very specialized approaches are needed
- Improving the algorithms (methods of solving problems) decreases the time it takes to solve a problem at least as much as new hardware.
- What is done on a high end systems typically becomes common practices a decade later on other systems, and is used for many standard things within another decade

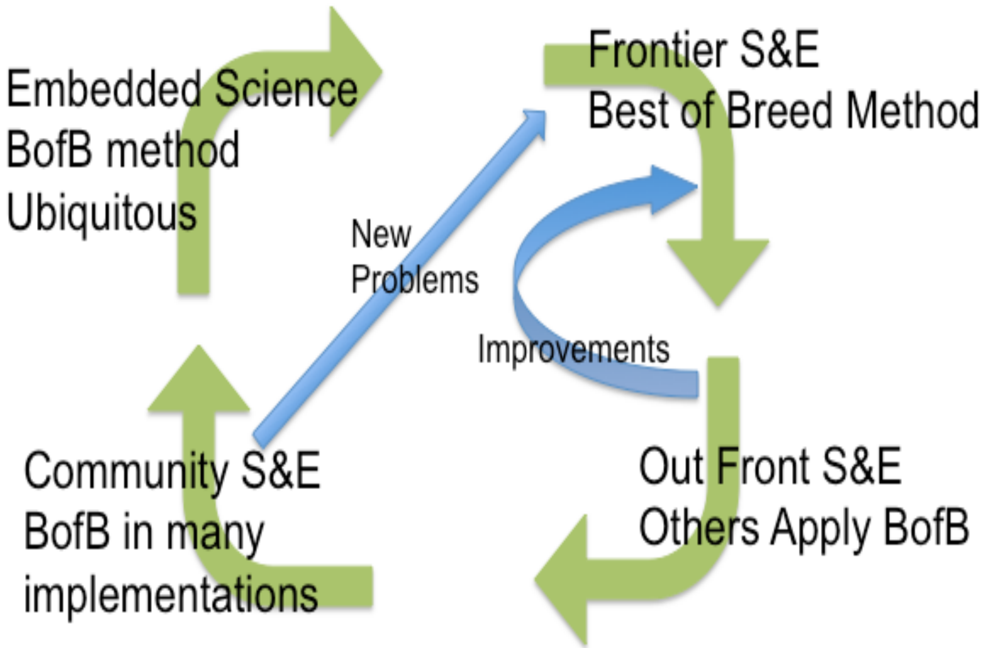

National Strategic Computing Iinitative

– Slides from Irene Qualters NSF ACI Director

National Strategic Computing Initiative

Executive Order Signed July 29, 2015

- National
  - "Whole of government" approach
  - Public/private partnership with industry and academia
- Strategic
  - Leverage beyond individual programs
  - Long time horizon (decade or more)
- Computing
  - HPC as advanced, capable computing technology
  - Multiple styles of computing and all necessary infrastructure
  - Scope includes everything necessary for a fully integrated capability
- Initiative
  - Above baseline effort
  - Link and lift efforts

Enhance U.S. strategic advantage in HPC for security, economic competitiveness, and scientific discovery

NSCI Executive Order calls on NSF to play a leadership role

- Scientific discovery advances
- The broader HPC ecosystem for scientific discovery
- Workforce development

Co-lead with DOD and DOE

The Government's Role in NSCI

- DOD + DOE
  - Capable exascale program
  - Analytic computing to support missions: science and national security
- NSF
  - Scientific discovery
  - Broader HPC ecosystem
  - Workforce Development
- IARPA + NIST
  - Future computing technologies
- NASA, FBI, NIH, DHS, NOAA
  - Deployment within their mission contexts

NSCI objectives

1. Accelerate delivery of a capable exascale computing system (hardware, software) to deliver approximately 100X the performance of current 10PF systems across a range of applications reflecting government needs
2. ***Increase coherence between technology base used for modeling and simulation and that used for data analytic computing.***
3. ***Establish, over the next 15 years, a viable path forward for future HPC systems in the post Moore's Law …***
4. ***Increase the capacity and capability of an enduring national HPC ecosystem, employing a holistic approach … networking, workflow, downward scaling, foundational algorithms and software, and workforce development.***
5. Develop an enduring public-private partnership to assure that the benefits .. are transferred to the U.S. commercial, government, and academic sectors

2. Increase coherence between technology base used for modeling and simulation and that used for data analytic computing

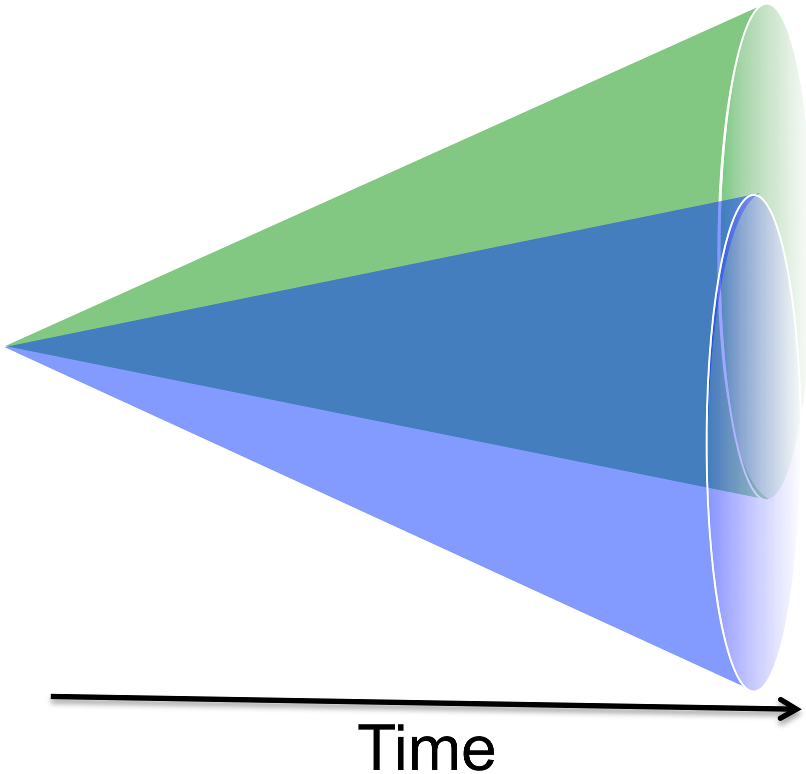

Modeling and Simulation

- Multi-scale
- Multi-physics
- Multi-resolution
- Multidisciplinary
- Coupled models

Data Science

- Data Assimilation
- Visualization
- Image Analysis
- Data Compression
- Data Analytics

NSF Role: Support foundational research and research infrastructure within and across all disciplines (across all NSF directorates)

Aspirations for Convergence

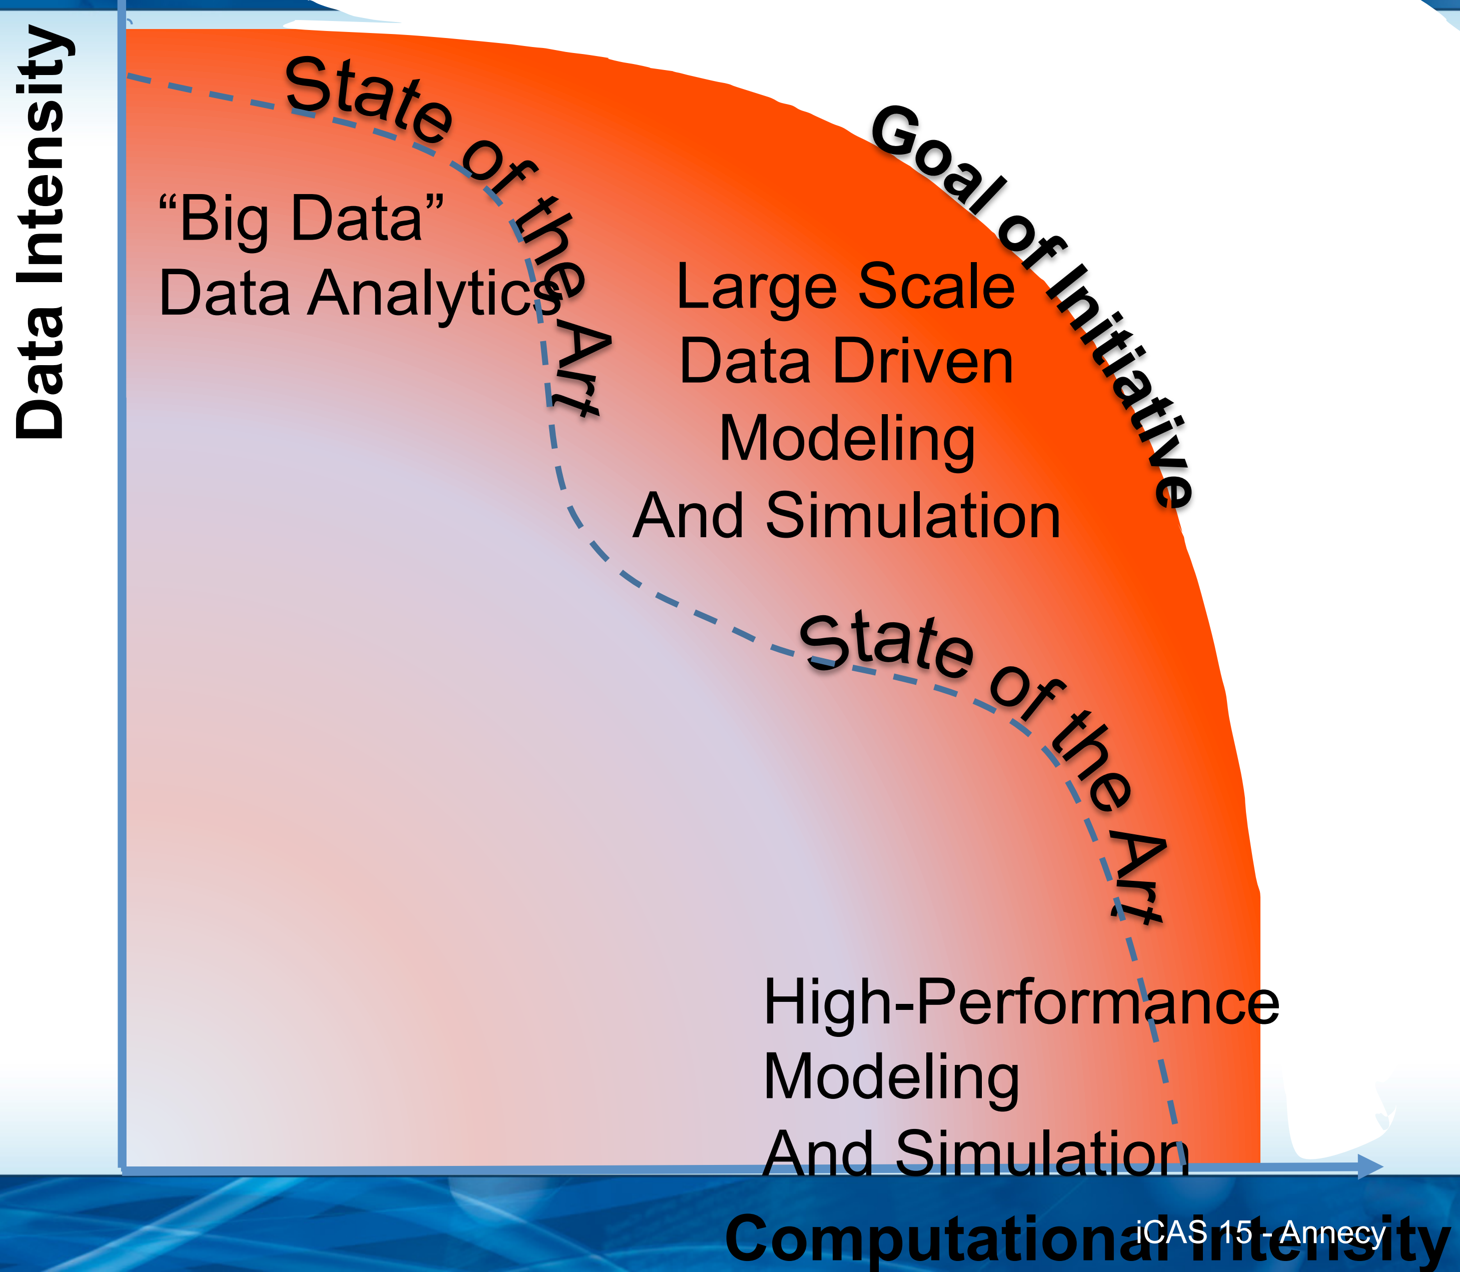

NSF role in NSCI: Enduring Computational Ecosystem for Advancing Science and Engineering

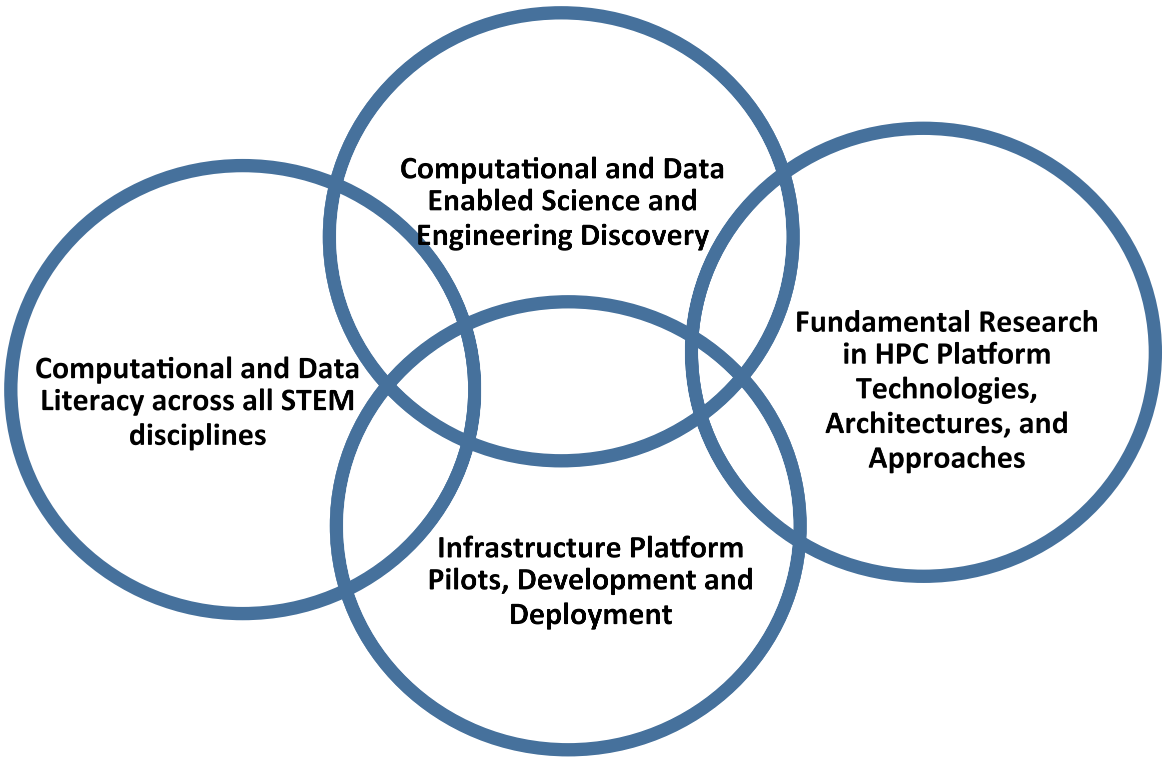

Blue Waters continues to make major impacts

- *"Using the well-balanced system capabilities of Blue Waters to complete CyberShake calculations within weeks rather than months." – T. Jordan*
- *Blue Waters allows" multiple, high resolution runs, 150+ years past and 100 years future, to characterize uncertainty. " – D. Weubbles*
- *Blue Waters made possible "… the largest Monte Carlo simulation experiment for evaluating the required satellite frequencies and coverage to maintain acceptable global forecasts of terrestrial hydrology (especially in poorer countries)." – P. Reed*
- Blue Waters "*results show that 3D magnetorotational core-collapse supernovae are fundamentally different from what has been anticipated on the basis of axisymmetric simulations.*" – C. Ott
- Blue Waters "*… allows us to make direct comparisons with the [Hubble Ultra Deep Field] survey.*" – T. Quinn
- *"To obtain results in a reasonable time requires using as many processors as possible. Currently, only Blue Waters provides a significant number of useable processors." – R Stein*
- *"Blue Waters is unique in the NSF computing infrastructure, enabling large and detailed simulations allowing for the discovery of a global oscillation of shell hydrogen ingestion."* – P Woodward
- *"Work completed in weeks that normally would take 6 months." – T. Cheatham*
- "*Blue Waters … critical for the implementation and deployment of the UCG-MD software*" , extremely large number of parallel compute nodes and a close collaborative environment. G Voth
- *"To the best of our knowledge, this is the first time a supercell producing a long-track EF5 tornado has ever been simulated." – R. Wilhelmson*

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.