Supercomputers as Scientific Instruments: Drivers for Integrated Compute and Data Infrastructure

20 Years of Excellence in Computational Science
OLCF
Oak Ridge Leadership Computing Facility
1992–2012

Jack C. Wells
Director of Science
National Center for Computational Sciences
Oak Ridge National Laboratory

SOS18, St. Moritz, 18 March 2014

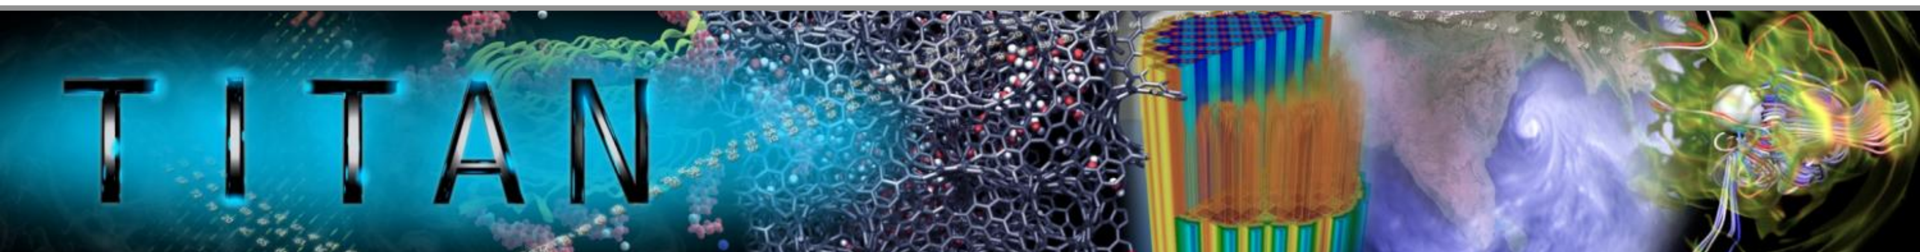

Outline

- Intro to U.S. DOE Leadership Computing Program
- OLCF Requirements Process
- Integrating Data and Compute Infrastructure with SNS
  - ADARA: Accelerating Data Acquisition, Reduction, and Analysis
  - CAMM: Center for Accelerated Materials Modeling

What is the Leadership Computing Facility (LCF)?

- Collaborative DOE Office of Science program at ORNL and ANL
- Mission: Provide the computational and data resources required to solve the most challenging problems.
- 2-centers/2-architectures to address diverse and growing computational needs of the scientific community
- Highly competitive user allocation programs (INCITE, ALCC).
- Projects receive 10x to 100x more resource than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).

OAK RIDGE National Laboratory

Argonne NATIONAL LABORATORY

The OLCF has delivered five systems and six upgrades to our users since 2004

- Increased our system capability by 10,000x
- Strong partnerships with computer designers and architects
- Worked with users to scale codes by 10,000x
- Science delivered through strong user partnerships to scale codes and algorithms

Phoenix X1
- Doubled size
- X1e

2004

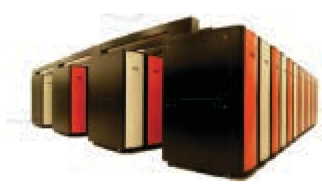

Jaguar XT3
- Dual core upgrade

2005

Jaguar XT4
- Quad core upgrade

2007

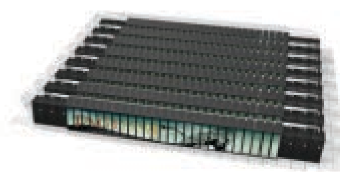

Jaguar XT5
- 6 core upgrade

2008

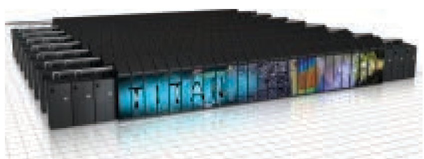

Titan XK7
- GPU upgrade

2012

Science breakthroughs at the OLCF:

SELECTED science and engineering advances over the period 2003 - 2013

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

Astrophysicists discover supernova shock-wave instability, ***Astrophys. J.*** *(2003)*
<u>254 citations, 3/2014</u>

Researchers solved the 2D Hubbard model and presented evidence that it predicts HTSC behavior, ***Phys. Rev. Lett*** *(2005)*
<u>105 citations, 3/2014</u>

First-Principles Flame Simulation Provides Crucial Information to Guide Design of Fuel-Efficient Clean Engines, ***Proc. Combust. Insti.*** (2007)
<u>78 citations, 3/2014</u>

Largest simulation of a galaxy's worth of dark matter, showed for the first time the fractal-like appearance of dark matter substructures, ***Nature*** (2008)
<u>326 citations, 3/2014</u>

World's first continuous simulation of 21,000 years of Earth's climate history, ***Science*** *(2009)*
<u>116 citations</u>

Biomass as a viable, sustainable feedstock for hydrogen production for fuel cells, ***Nano Letters*** *(2011)*
J. Phys. Chem. Lett. *(2010)*
<u>71 & 74 citations, respectively</u>

Demonstrated that three-body forces are necessary to describe the long lifetime of ^{14}C
Phys. Rev. Lett. (2011)
<u>28 citations, 3/2014</u>

Calculation of the number of bound nuclei in nature, ***Nature*** *(2012), 36 citations, 3/2014* , <u>36 citations, 3/2014</u>

Global Warming preceded by increasing carbon dioxide concentrations during the last deglaciation, ***Nature*** *(2012).*
<u>64 citations, 3/2014</u>

MD simulations show selectivity filter of a trans-membrane ion channel is sterically locked open by hidden water molecules, ***Nature*** *(2013)*

High-impact science across a broad range of disciplines

For example in 2013:

Paleoclimate Science
"Northern Hemisphere forcing of Southern Hemisphere climate during the last deglaciation,"
Feng He (UW Madison), *et al.*, *Nature*, February (2013)

Molecular Biology
MD simulations show selectivity filter of a trans-membrane ion channel is sterically locked open by hidden water
Jared Ostmeyer, et al. (U. Chicago) *Nature*, Sept. (2013)

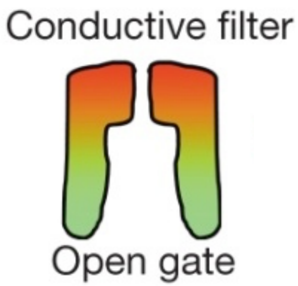

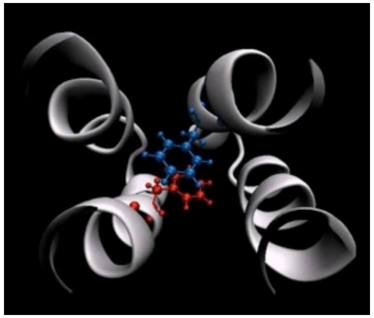

Molecular Biology
"A phenylalanine rotameric switch for signal-state control in bacterial chemoreceptors"
D. Ortega (UTK),
Nature Communications
December (2013)

Superconductivity
"Doping dependence of spin excitations and correlations with high-temperature super-conductivity in iron pnictides,"
Meng Wang(IOP CAS Beijing),
Nature Communications.
December (2013)

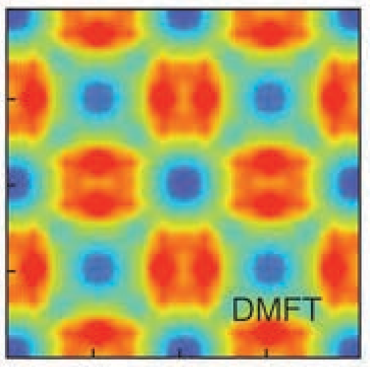

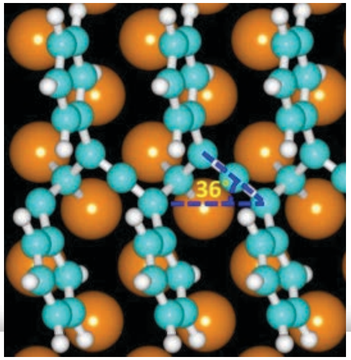

Polymer Science
"Self-Organized and Cu-Coordinated Surface Linear Polymerization"
Qing Li, B. Sumpter (ORNL),
Nature Scientific Reports.
July (2013)

Complex Oxide Materials
"Atomically resolved spectroscopic studyof Sr2IrO4: Experiment and theory," Qing Li (ORNL), E.G. Eguiluz (UTK)
Nature Scientific Reports.
October (2013)

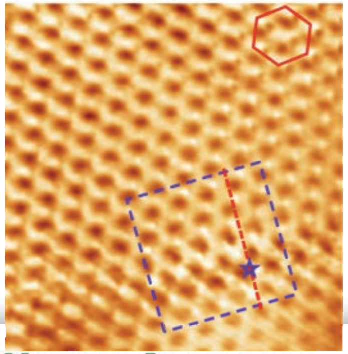

Science challenges for LCF in next decade

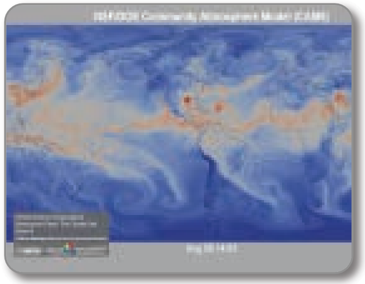

Climate Change Science
Understand the dynamic ecological and chemical evolution of the climate system with uncertainty quantification of impacts.

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

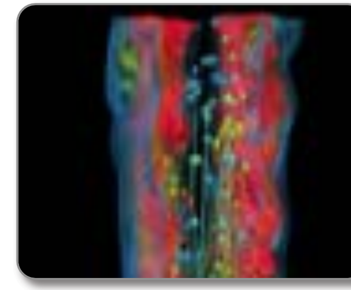

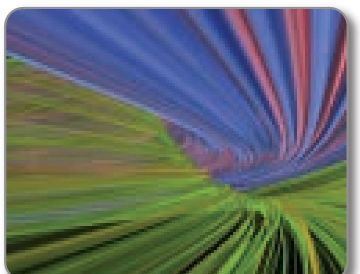

Fusion Energy
Develop predictive understanding of plasma properties, dynamics, and interactions with surrounding materials.

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

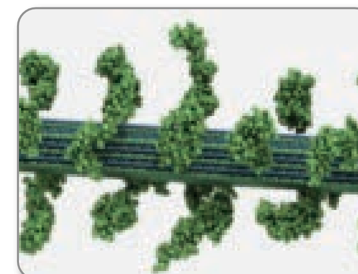

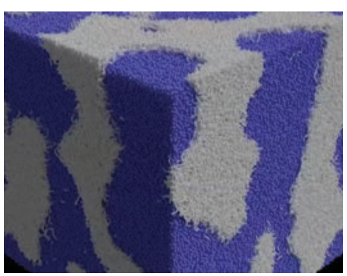

Solar Energy
Improve photovoltaic efficiency and lower cost for organic and inorganic materials.

Optimized Accelerator Designs
Optimize designs as the next generations of accelerators . Detailed models are needed to provide efficient designs of new light sources.

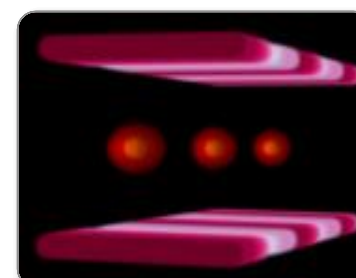

Requirements gathering

DOE/SC and LCFs support a diverse user community

- Science benefits and impact of future systems are examined on an ongoing basis
- Leadership Computing Facility staff have been actively engaged in community assessments of future computational needs and solutions
- Computational science roadmaps are developed in collaboration with leading domain scientists
- Detailed performance analyses are conducted for applications to understand future architectural bottlenecks
- Analysis of INCITE, ALCC, Early Science, and Center for Accelerated Application Readiness (CAAR) projects history and trends

Science Category	Represented Research Areas
	Bioinformatics Biophysics Life Sciences Medical Science Neuroscience Proteomics Systems Biology
	Chemistry

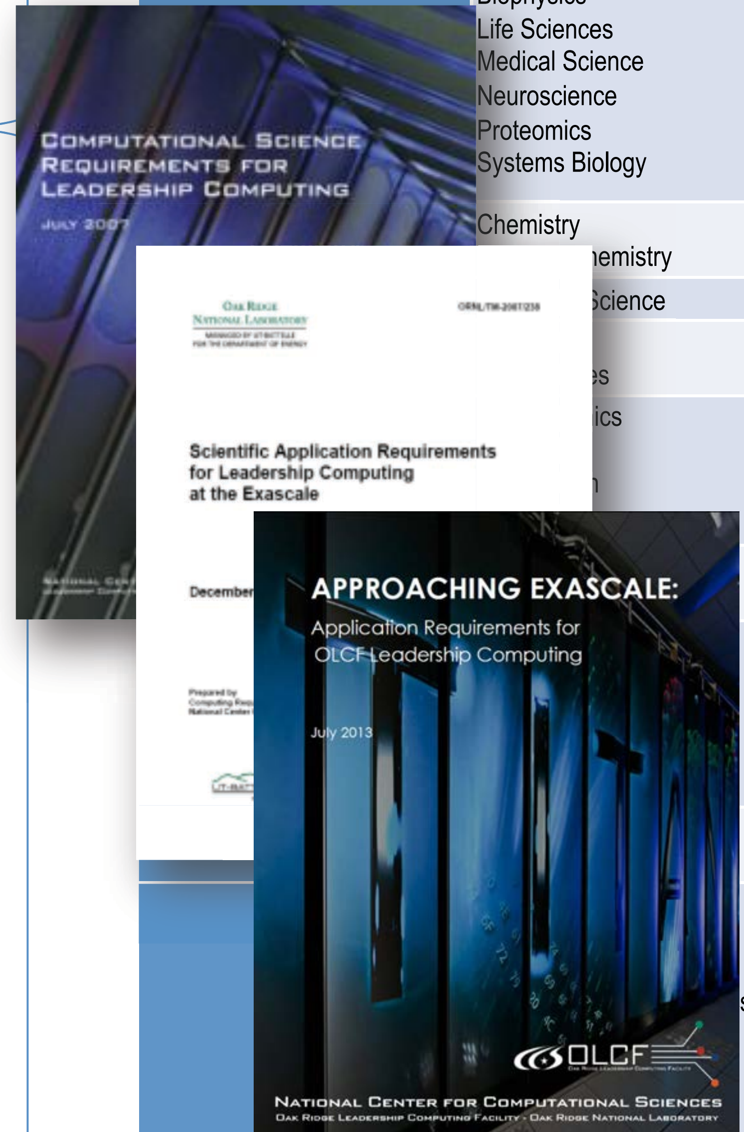

Oak Ridge National Laboratory
Managed by UT-Battelle for the U.S. Department of Energy

Embedded expertise at LCFs help gather requirements

- Early Science and CAAR programs worked with science teams to port applications to Mira and Titan
- These teams have analyzed the architectural features of these codes in detail for many-core and accelerator-based systems and use this to develop the technical requirements
- The OLCF Liaisons and ALCF Catalysts are scientists assigned to each INCITE project as part of the research teams
- They have a deep understanding of the codes that run on each system and the long-term requirements of these science teams

Peak FLOPS per Node was #1 Hardware Requirement in 2009 User Survey

System Attribute	Climate	Astrophysics	Fusion	Chemistry	Combustion	Accelerator physics	Biology	Materials science
Node peak flops	Red	Red	Red	Red	Red	Red	Red	Red
MTTI	Grey	Grey	Pink	Grey	Pink	Grey	Pink	Grey
WAN network bandwidth	Pink	Pink	Grey	Grey	Grey	Grey	Grey	Grey
Node memory capacity	Grey	Red	Red	Red	Red	Red	Red	Pink
Local storage capacity	Grey	Pink	Pink	Red	Red	Pink	Red	Pink
Archival storage capacity	Pink	Grey	Grey	Grey	Pink	Grey	Grey	Pink
Memory latency	Pink	Pink	Grey	Pink	Grey	Pink	Grey	Red
Interconnect latency	Red	Grey	Red	Red	Pink	Pink	Red	Red
Disk latency	Grey	Grey	Grey	Grey	Grey	Grey	Grey	Grey
Interconnect bandwidth	Red	Red	Red	Pink	Red	Red	Pink	Pink
Memory bandwidth	Red	Red	Pink	Pink	Pink	Red	Pink	Red
Disk bandwidth	Pink	Pink	Pink	Pink	Grey	Pink	Pink	Grey

(Priority: Red = high, pink = medium, grey = low.)

"Preparing for Exascale" OLCF Application Requirements and Strategy, December 2009

Roadmap: Fusion Energy/ITER

Key science challenges: Effectively model and control the flow of plasma and energy in a fusion reactor, scaling up to ITER-size. Develop predictive understanding of plasma properties, dynamics, and interactions with surrounding materials. Mitigate plasma disruptions.

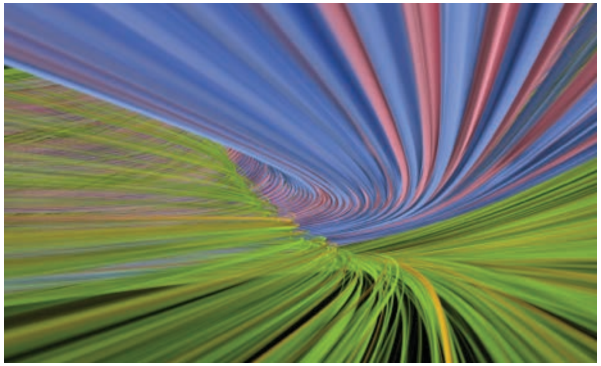

A global particle-in-cell simulation to show core turbulence in a tokamak. Image S. Ethier, PPPL

Science enabled by LCF Capabilities

2013-2016	2016-2020
• Perform high-fidelity simulation of edge plasma turbulent transport in tokamak from first-principles to address DIII-D and JET-scale plasmas with a goal of understanding high-confinement physics. • Increase simulation of tokamak edge plasma to ITER scale. Coupled simulations of plasma edge with core and chamber wall interactions. Control edge-localized modes and other destructive mechanisms.	• Perform integrated first-principles simulation including all the important multiscale physical processes to study fusion-reacting plasmas in realistic magnetic confinement geometries. • Produce an experimentally validated simulation capability for ITER that will be used to design and produce early results for DEMO, the anticipated follow-on facility, to solve the engineering issues necessary for practical electricity production with fusion plasmas.

The Bleeding 'Edge' of Fusion

INCITE Program
PI: CS Chang
Princeton Plasma Physics Laboratory
Hours used: 102 million

Science Objectives and Impact

- Fusion efficiency of ITER is expected to be determined by edge plasma transport
- First-principles simulation of edge physics on Titan can increase understanding of ITER
- Predictions validated and new insight gained through analysis of DIII-D tokamak results
- Success with DIII-D plasma simulations gives confidence for larger JET- and ITER-scale campaigns within the new program year

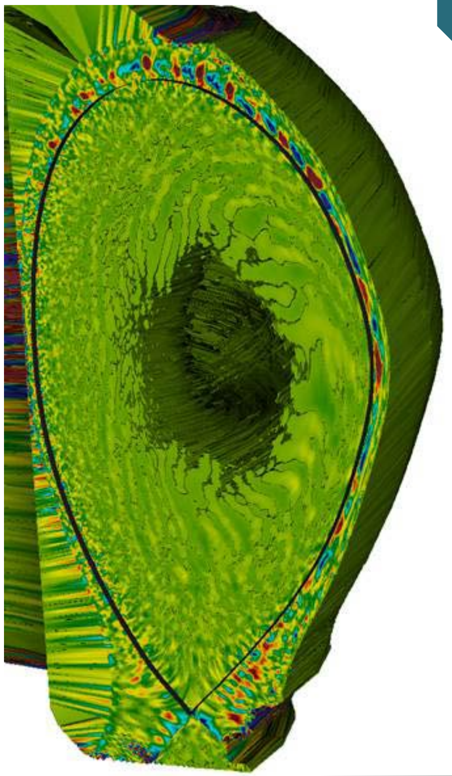

This figure shows the turbulence front from the plasma edge in the DIII-D reactor being spread inward in multiscale interaction with the evolving background profile under the central heat source. Eventually, the whole volume becomes turbulent, with the spatial turbulence amplitude distribution being just enough to produce the outward heat transport to expel the centrally deposited heat to the edge. This is how the plasma profile, the heat source and the turbulence self-organize.
- Visualization by David Pugmire ORNL)

Performance & OLCF Contribution

- OLCF provided the necessary HPC power; Up to 90% of Titan's max capability has been utilized for the study
- OLCF liaison contributed to algorithmic improvement, which made the XGC1 code use GPUs and CPUs efficiently, with linear scalability to maximal size of Titan
- OLCF personnel provided the Vis support

Science Results

- Simulated nonlinear coherent turbulence structures (called "blobs") in the plasma edge of DIII-D reactor
- Identified the momentum source and its inward transport process
- Predicted the divertor heat load distribution

All for the first time in first-principles calculations of a tokamak reactor

Integrating Compute & Data Capabilities: Creating a Pathway for Scientific Discovery

- Accelerating discovery in materials science
- Enhancing predictive capabilities

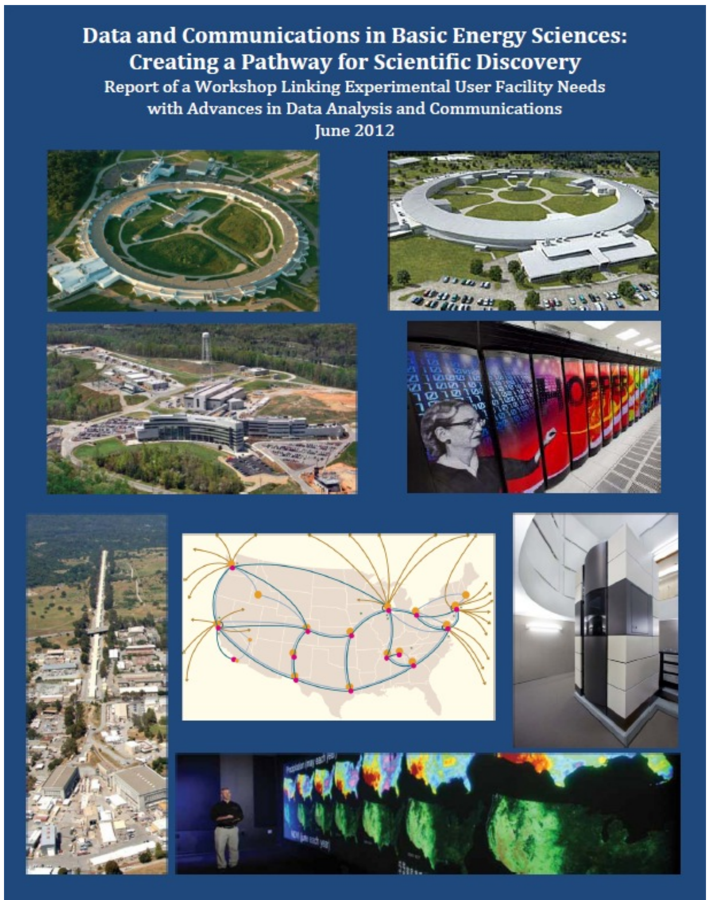

- Theory and analysis components should be integrated seamlessly within experimental workflow.
- Move analysis closer to experiment – future possibility of experiment steering.
- Match data management access and capabilities with advancements in detectors and sources.

ORNL has key strengths to address the materials challenge

High Flux
Isotope Reactor:
Intense steady-stae
neutron flux
and a high-brightness
cold neutron source

Spallation
Neutron Source:
World's most powerful
accelerator-based
neutron source

UT-ORNL
Joint Institute for
Neutron Sciences:
User gateway
for SNS and HFIR

Center for
Nanophase Materials
Science

Titan – Worlds fastest
supercomputer

Neutron Data Life Cycle

Feedback

DAS
- Neutron events
- Events from sample environment
- Other triggers

Reduction
- Corrected reduced data (histograms, S(Q,E), ..)
- Merging, reconstruction of data
- Instrument/technique dependent
- Need for ’real’ time reduction

Analysis
- Multi dimensional fitting
- Advanced visualization
- Comparison to simulation / feedback
- Field dependent, large variety of approaches

Simulation Modeling
- Multitude of techniques (DFT, MD, ..)
- Advanced simulation of experiments
- ‘Refinement’ using experimental data
- Multiple experiments / probes

User Facility

- Variety of experiments, topics, methods and ‘computer literacy’ of users are significant challenge.

Improving Productivity = Changing the Workflow

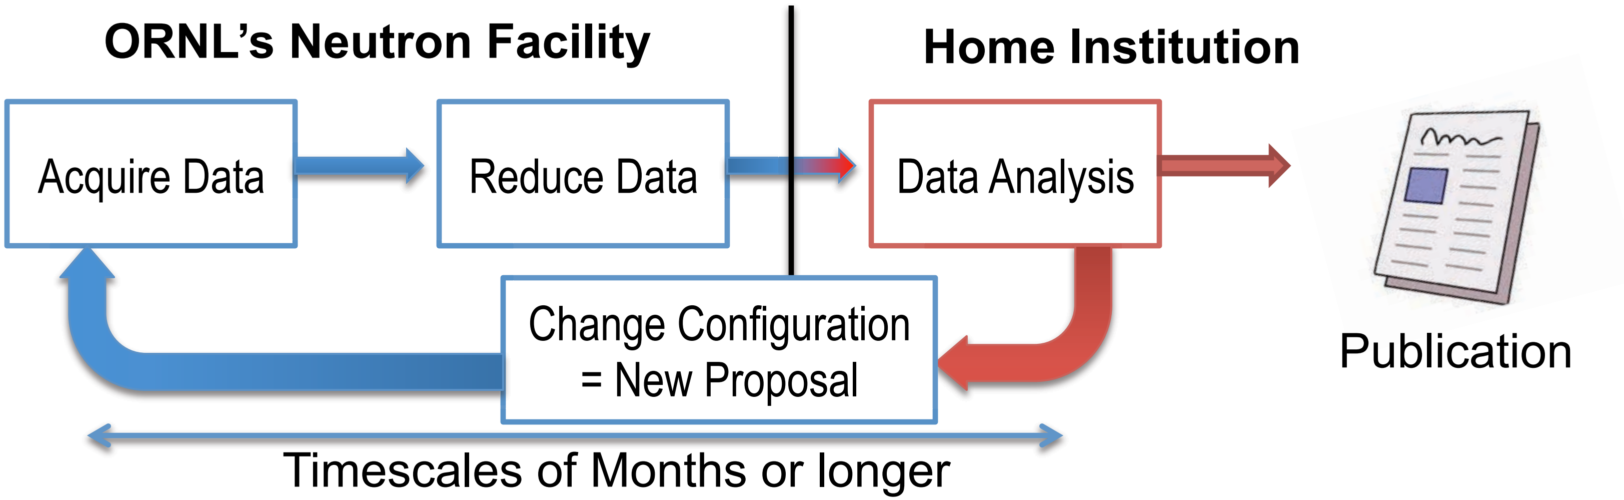

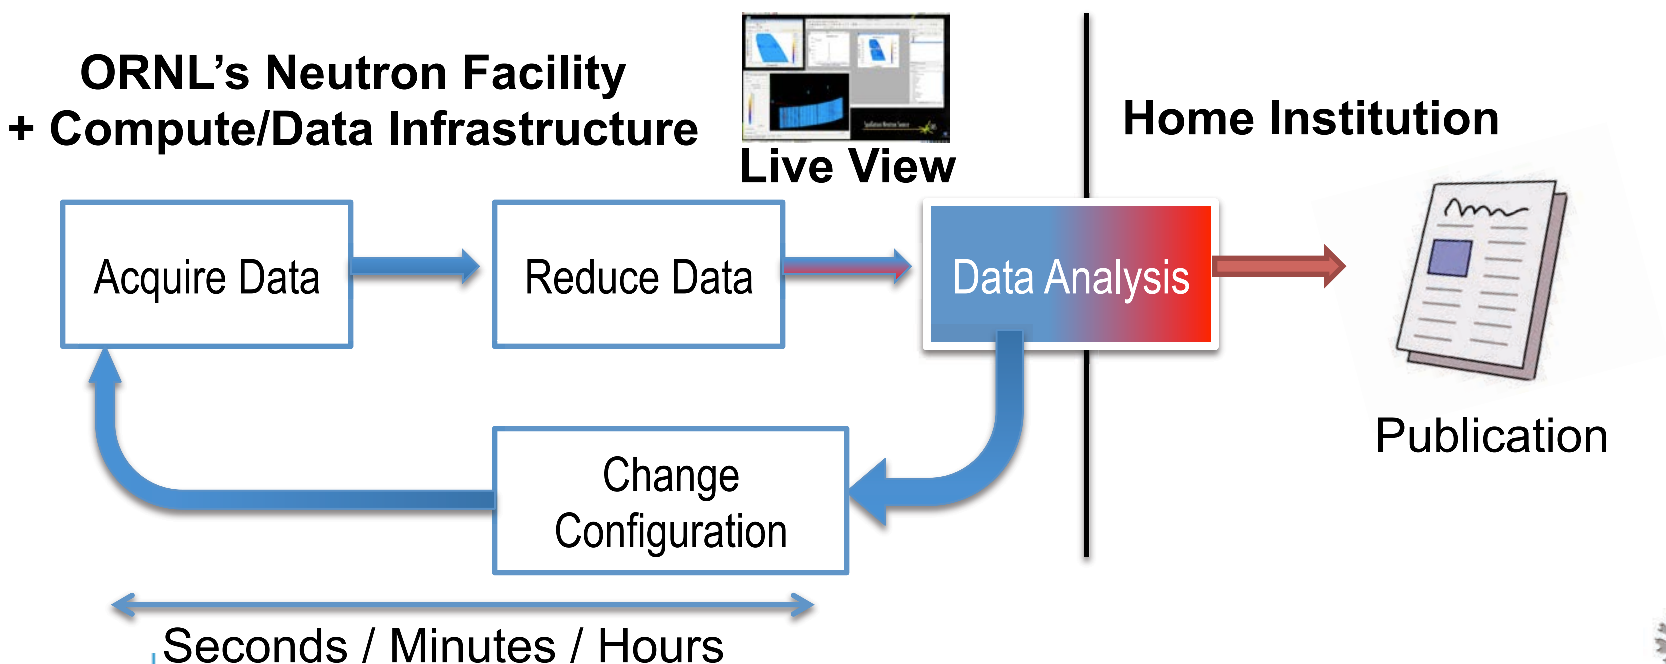

How can we do it?

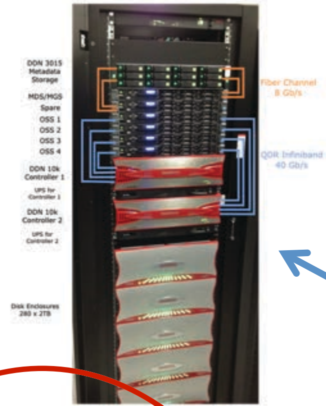

Chadwick (196 cores)
Fermi (512 cores)
Bragg (2048 cores)

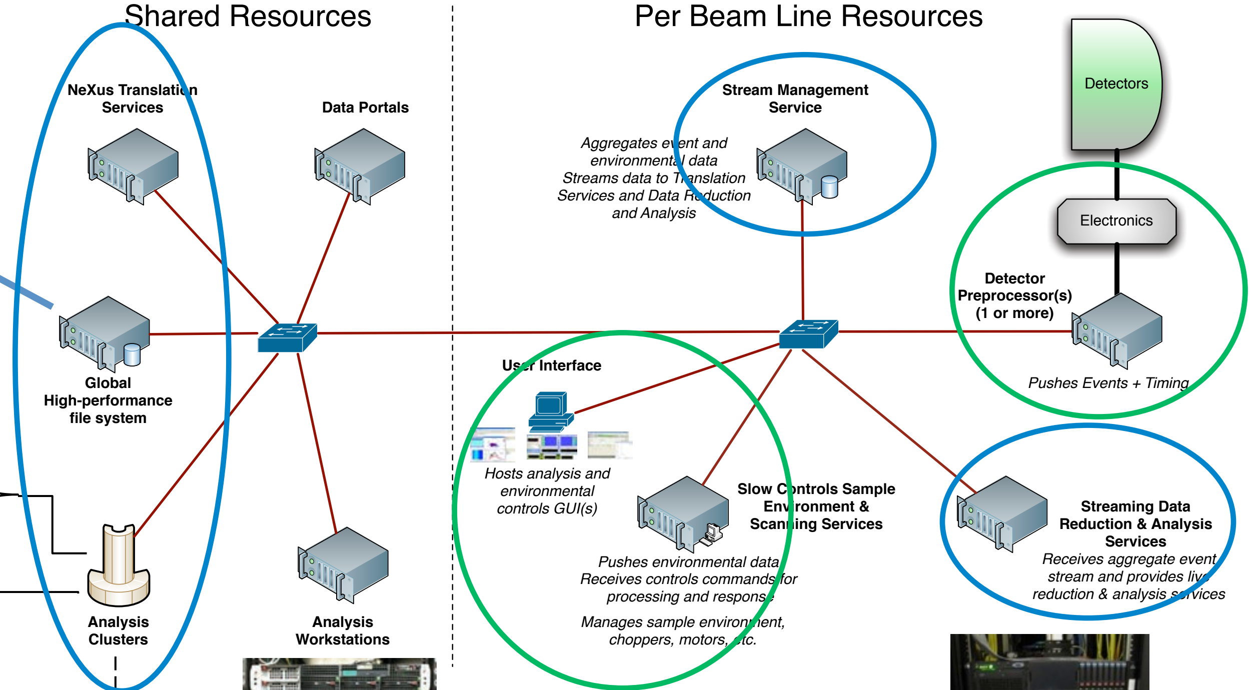

Login node workstations
32 cores, 256GB

Beam line workstations
32 cores, 256GB

Multiple (Integrated) Projects

- The ADARA project
- Center for Accelerating Materials Modeling
- Upgrades to the Data Acquisition System

ADARA is enabling real-time feedback from experiment and analysis through compute and data infrastructure

ADARA

- Leverages our multi-disciplinary capabilities at ORNL coupling Neutron Sciences Directorate with Computing and Computational Sciences Directorate.
- The ADARA Project lets us stream data to computational resources and provide live feedback from experiment in real-time S(Q,E).
- Provides a high performance data backplane for reduction, analysis, and coupling with simulation forming the basis for future work to integrate experiment and simulation.
- Prototype running on HYSPEC instrument. Deployment to other beamlines in 2013/2014.

ORNL has launched the Center for Accelerating Materials Modeling (CAMM)

- The CAMM will integrate materials modeling/simulation (MD/DFT) directly into the chain for neutron scattering data analysis, **offline** and **online** (in near real time)
- Developing workflows for refinement, integration of MD codes, **neutron scattering corrections** ..
- The CAMM is working with ORNL's Materials Science and Technology Division to study coarse grained MD simulations of polymers PEO-AA (CNMS), *ab-initio* MD simulations for ferroelectrics/thermoelectrics

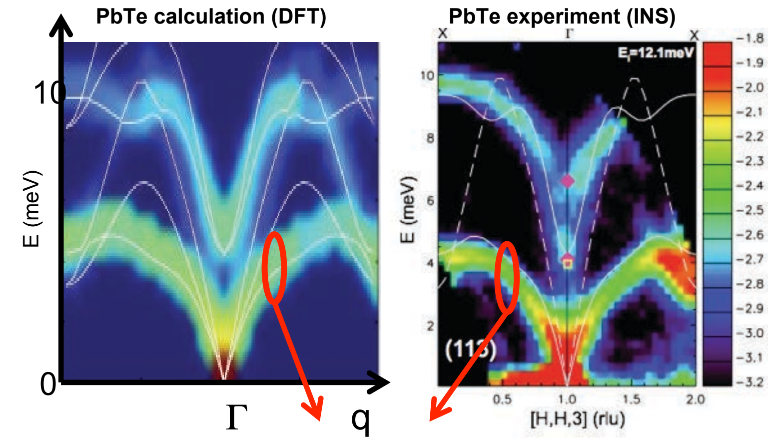

Example: *ab-initio* MD simulations for ferroelectrics/thermoelectrics. Focus on *width* of dispersions

The Center for Accelerating Materials Modeling (CAMM)

- *Partnership between ORNL's Neutron Sciences, Physical Sciences and Computing and Computational Sciences Directorates*
- *ORNL SEED money and DOE funds provided to study force field refinement from quasi-elastic and inelastic neutron scattering data*
- *CAMM formed in response to BES proposal call for* Predictive Theory and Modeling

What Mechanisms are Needed?

- Ability to stream neutron data to any number of downstream consumers
- Live *in-situ* reduction of neutron data as it is being generated – Instant feedback to scientists
- Translation of neutron data to NeXus files as it is being generated - NeXus files the instant the experiment is complete
- NeXus files created and stored in a high-performance parallel file system co-located with a high-performance compute system for timely post-analysis
- High performance parallel reduction and analysis operators

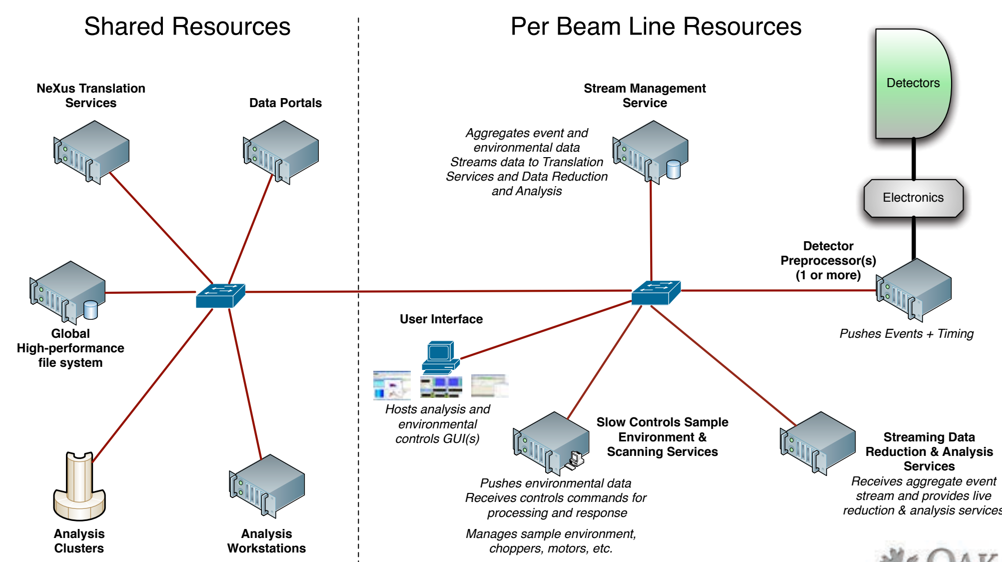

ADARA

ADARA – Accelerating Data Acquisition, Reduction and Analysis
Collaboration between SNS and Computing at ORNL

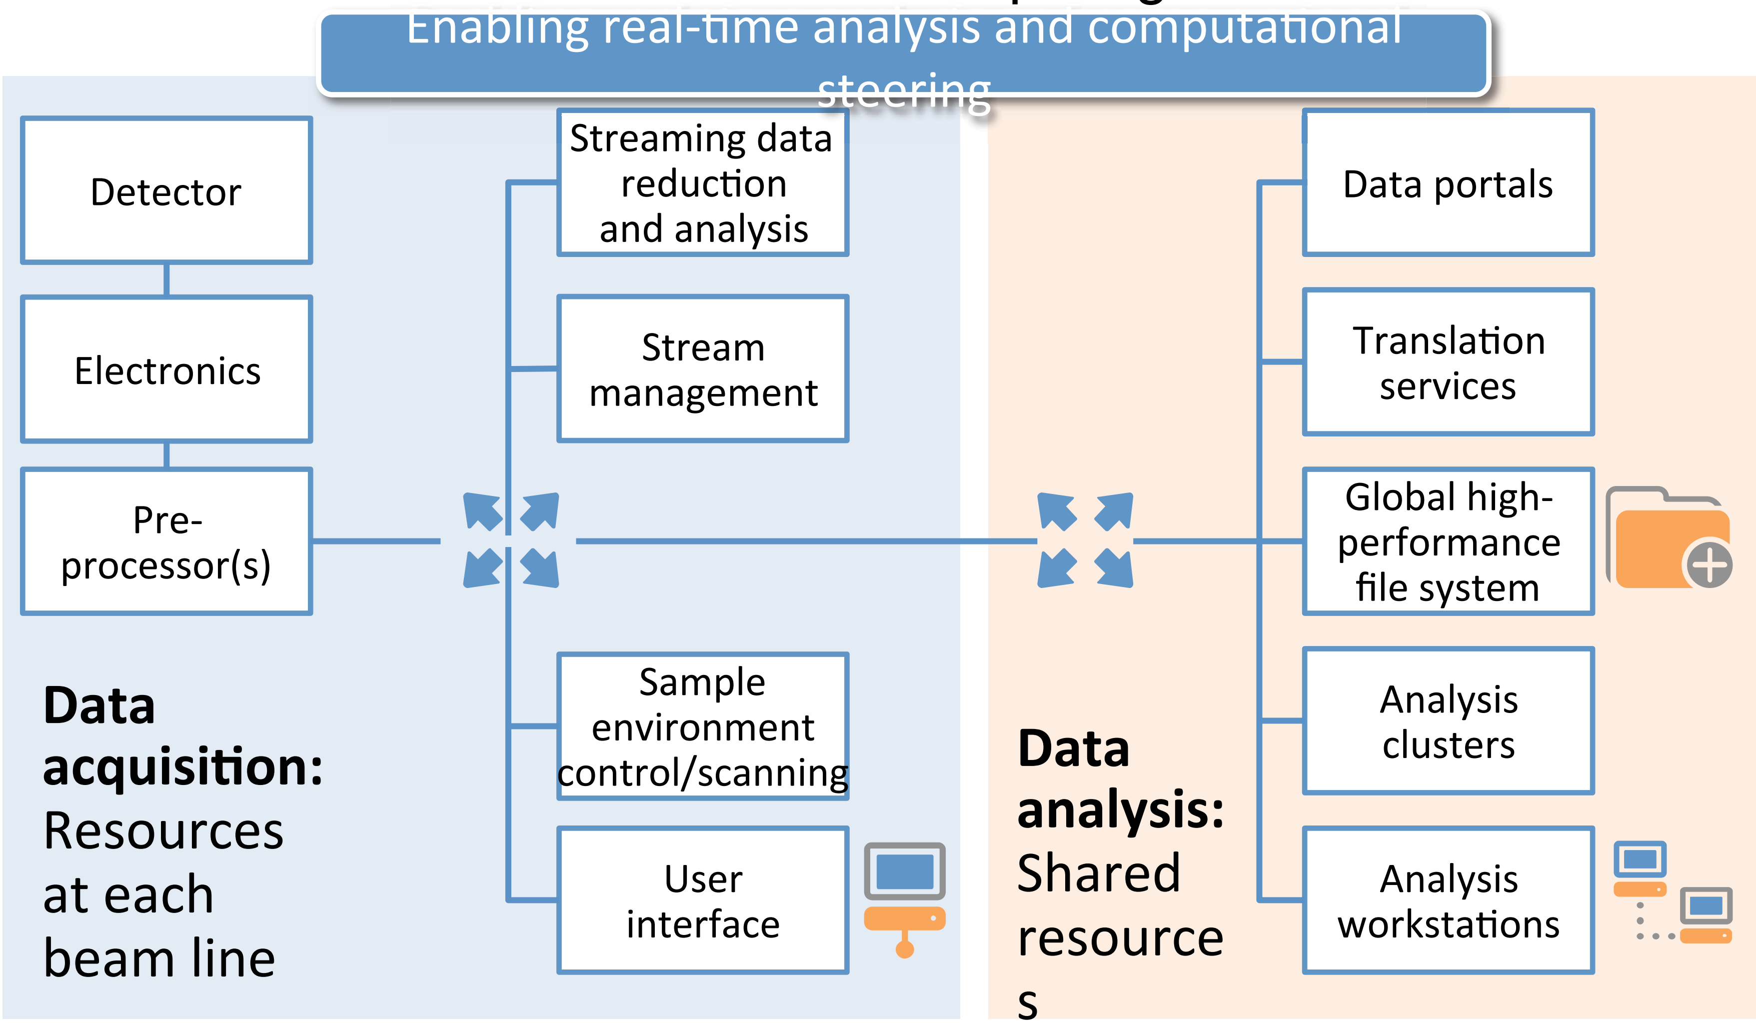

Center for Accelerating Materials Modeling

- BES Proposal Call in *Predictive Theory and Modeling*
- Partnership between Neutron Sciences and Computational Sciences
- Bring materials modeling/simulation directly into the chain for neutron scattering data analysis

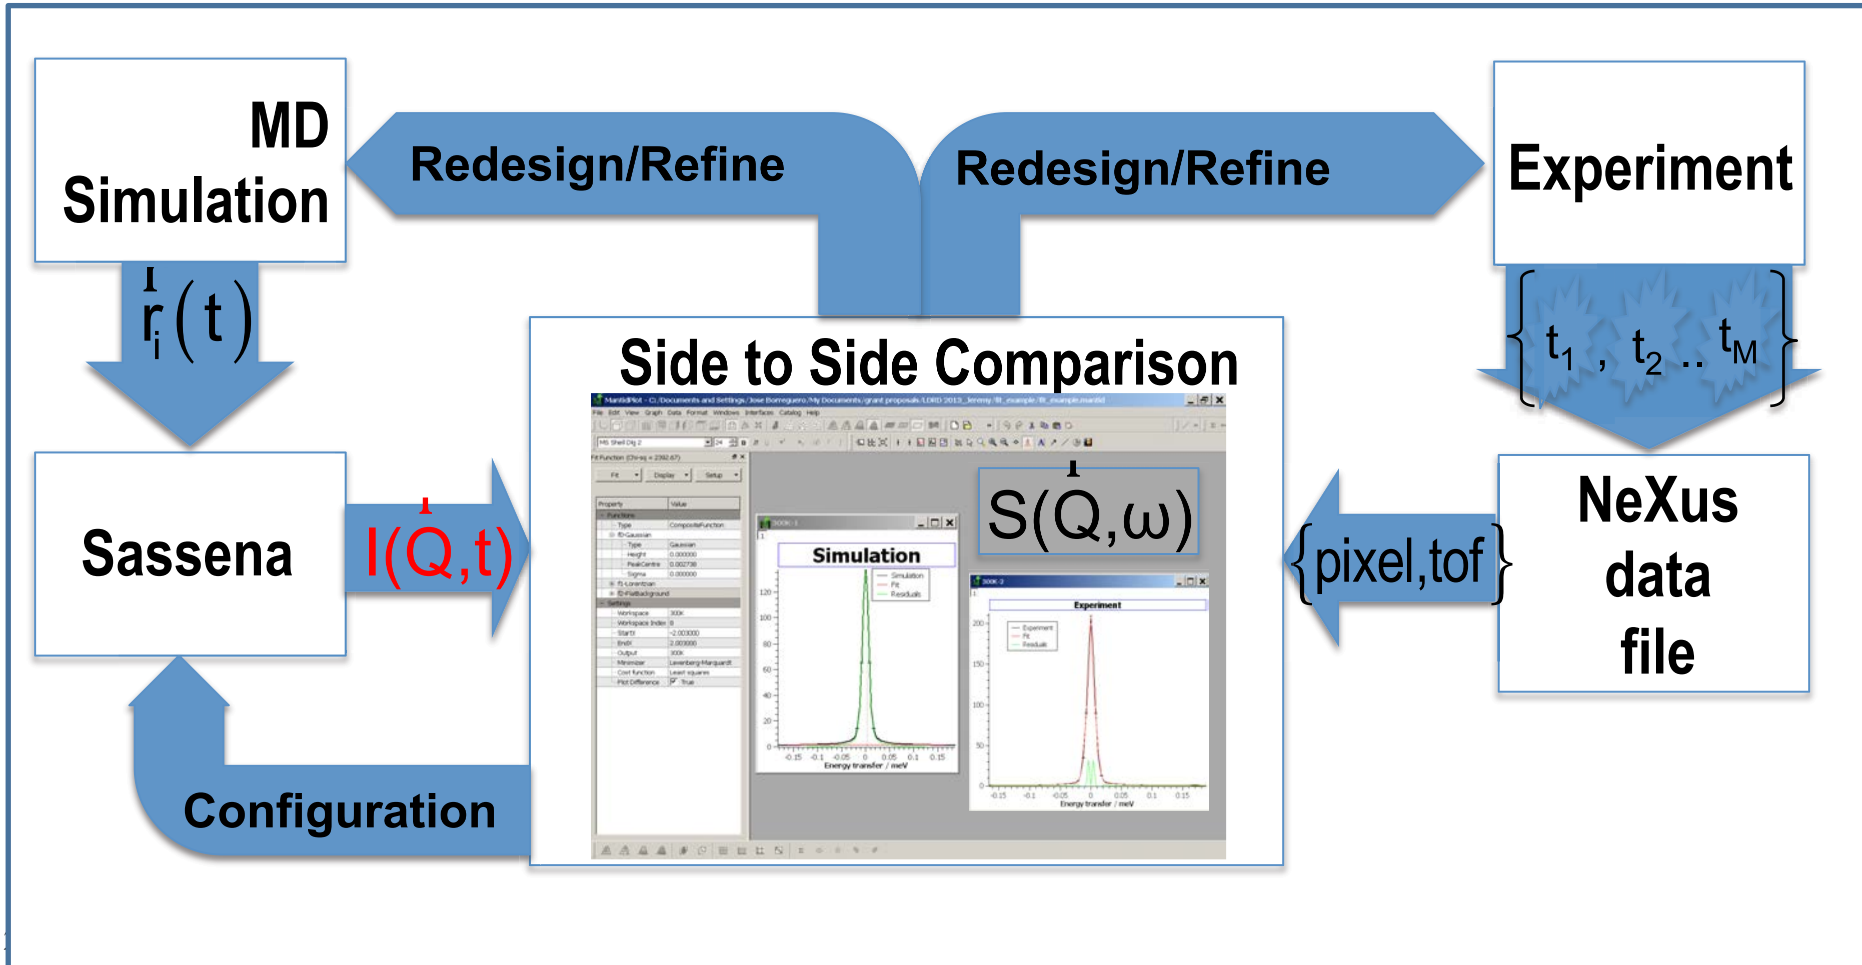

CAMM: Refinement Workflow Components

- Kepler workflow for overall management of job sequencing, submission, and management.
- Dakota to predict new parameters for the next iteration of the refinement.
- Python script to update the force-field.
- NAMD to perform the MD simulations using the new parameters.
- Sassena to calculate the dynamics structure factor from the MD trajectory.
- Various python modules to account for instrumental artifacts in the data; convolution with the instrument resolution function, background scattering, elastic line scattering from the sample can and small offsets in the energy scale due to sample size/positioning.
- Mantid python application programming interface (API) for manipulation of experimental data and conversion of simulated data to match the experimental data format.
- ActiveMQ to enable message passing between Dakota and Kepler. In the future, the messaging brokerage will allow communication with other services

Modeling LiCl data with workflow

- Kepler workflow submits simulations to HPC resource
- Dakota optimizes model parameters
- Process automated with Kepler GUI for input

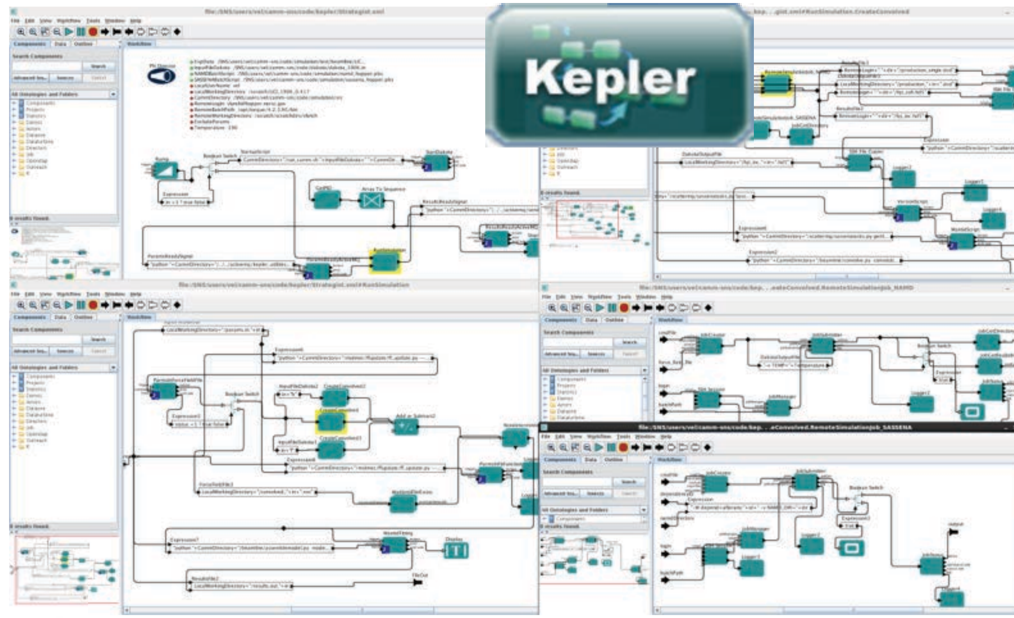

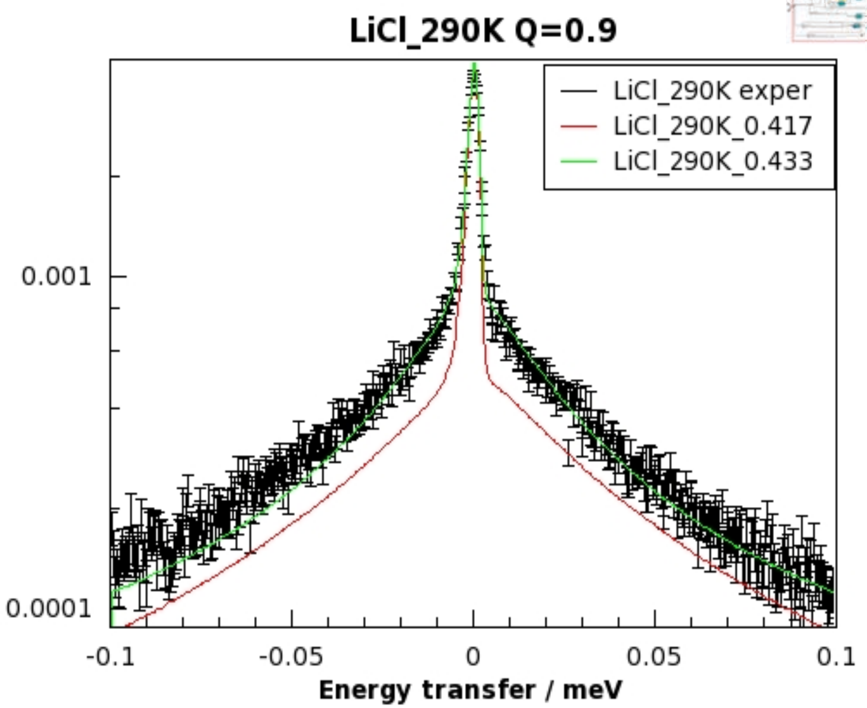

<<<<< Best data captured at function evaluation 2
Confidence Interval for c0 is [7.7181005978e-01, 7.9133911535e-01]
Confidence Interval for b0 is [3.5683719031e-05, 4.1095168905e-05]
Confidence Interval for e0.0 is [4.7649690032e-02, 5.0604743707e-02]
Confidence Interval for e0.1 is [3.8658107943e-02, 4.0925369124e-02]
Confidence Interval for e0.2 is [3.2393364793e-02, 3.4198221349e-02]
Confidence Interval for eshift is [2.2784599458e-04, 2.7276832227e-04]
Confidence Interval for FF1 is [4.3254547020e-01, 4.3403369489e-01]

J.M. Borreguero, 2014 TMS Annual Meeting

Modeling LiCl data from BASIS at SNS

- NAMD simulations
- 154 Li+ Cl- pairs
- 1128 water molecules
- Total of 150ns production
- S(Q,ω) obtained with Sassena + Mantid

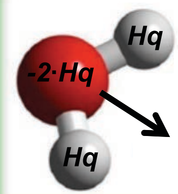

- Standard Hydrogen charge values:
  - *Hq=0.417* (TIP3P)
  - *Hq=0.424* (SPC/E)
- We vary *Hq* from 0.32 to 0.50
- Converged to *Hq=0.433 in 3 iterations*

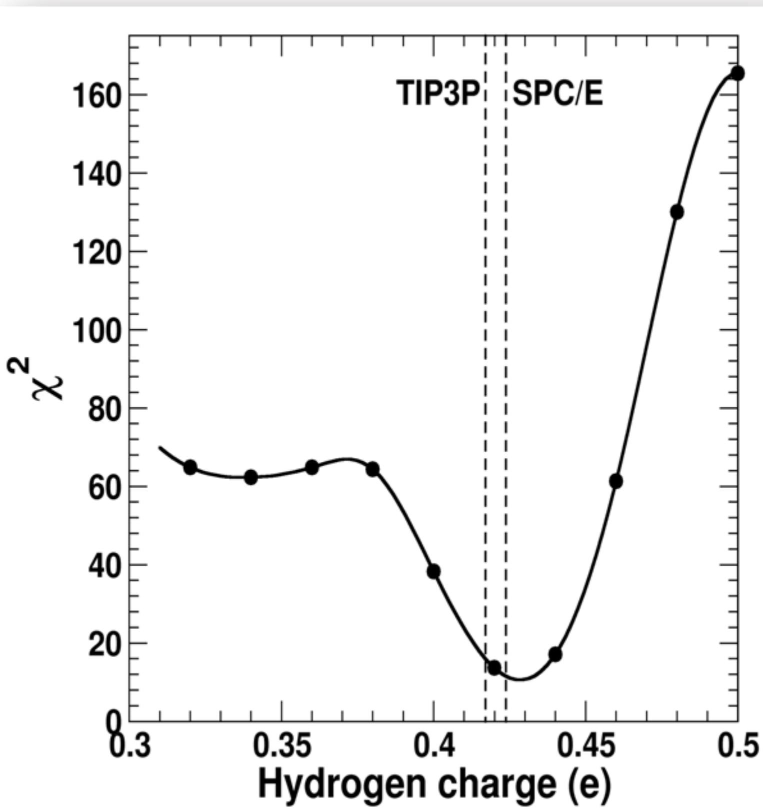

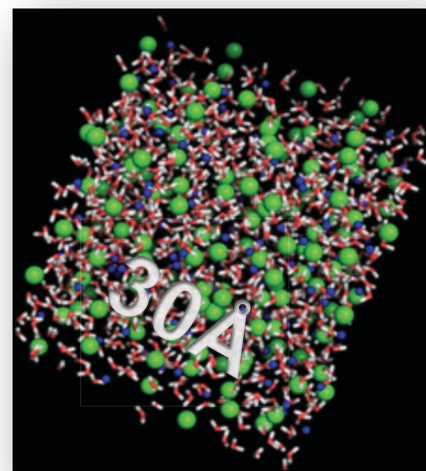

Li Cl in water

J.M. Borreguero,
2014 TMS Annual Meeting

CADES
Compute & Data Environment for Science

Oak Ridge National Laboratory

We have established CADES to provide core compute and data services such as those required by ADARA and CAMM projects to major facilities and programs

CADES is a cross-cutting center: it shares both data infrastructure and compute & data science expertise with and among many projects

A rich set of flexibly composable services coupled with exp in data science partnering with domain scientists on their cl

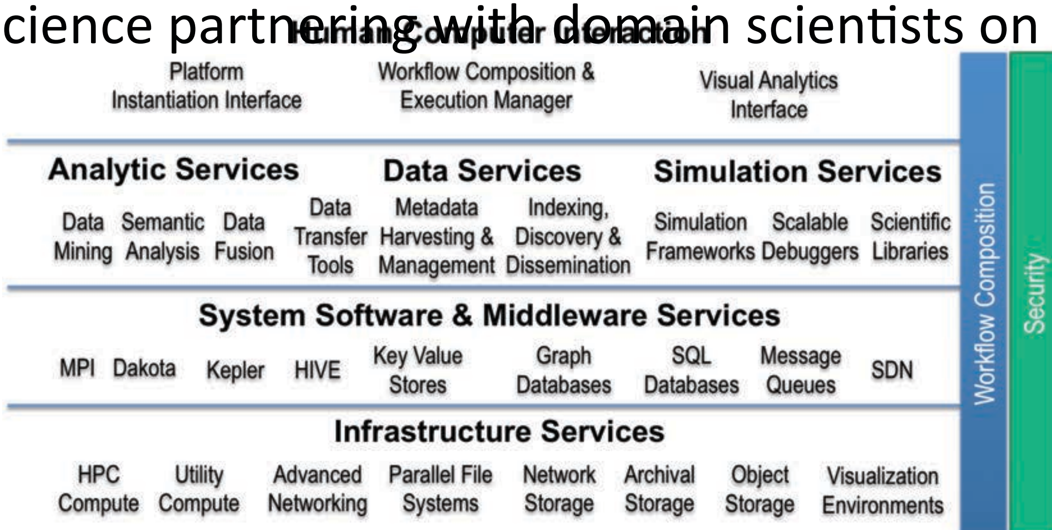

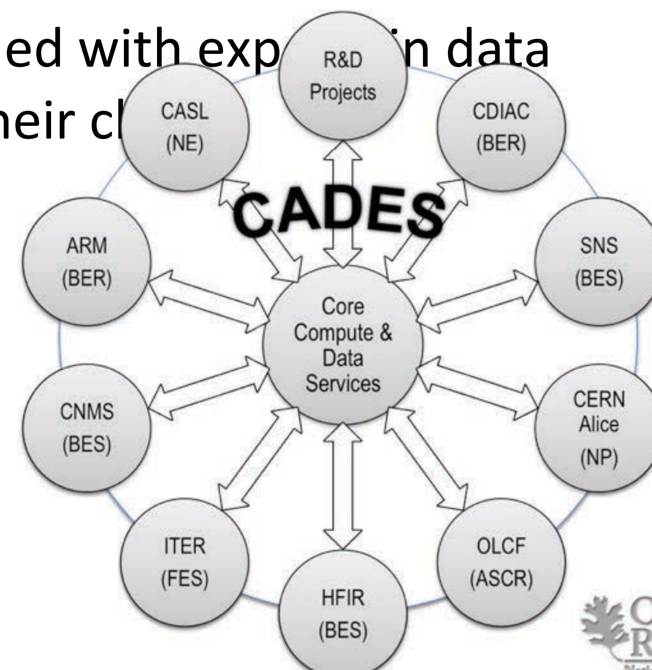

Conclusions

- Leadership computing facilities for the critically important problems that need the most powerful compute and data infrastructure
- User requirements integration is an important element in shaping the outcomes of large supercomputer facilities.
- New and exciting opportunities are availed by integrated large experimental facilities with supercomputing facilities.

Acknowledgements

OLCF Users: CS Chang (PPPL)

Visualization: Mike Matheson and Dave Pugmire (ORNL) for visualizations

OLCF user requirements process: Ricky Kendall & Doug Kothe (ORNL)

CAMM & CADES: G. Shipman, J. Borreguero, and T.Proffin (ORNL)

OLCF-3 Vendor Partners: Cray, AMD, NVIDIA, CAPS, Allinea

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Questions?

WellsJC@ornl.gov

Contact us at
http://olcf.ornl.gov
http://jobs.ornl.gov