BLUE WATERS

SUSTAINED PETASCALE COMPUTING

# Blue Waters and Petascale Science

William Kramer
National Center for Supercomputing Applications,
University of Illinois at Urbana-Champaign


GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY

# Science & Engineering on Blue Waters

*Blue Waters will enable advances in a broad range of science and engineering disciplines. Examples include:*

**Molecular Sc**

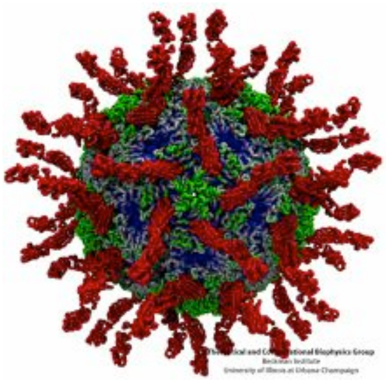

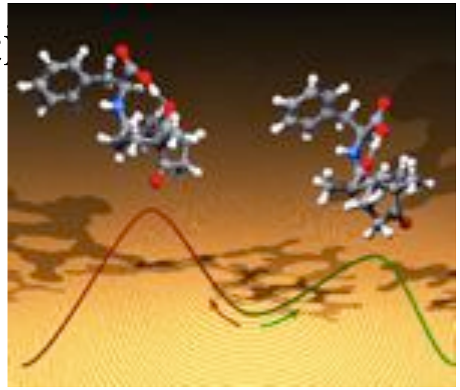

**Weather & Climate**

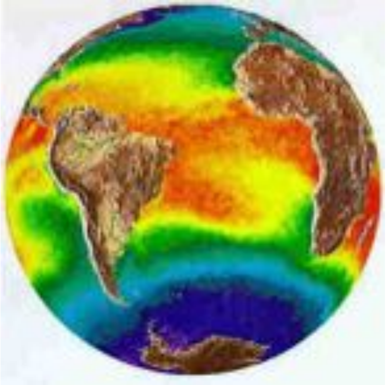

**Astrono**

**Earth**

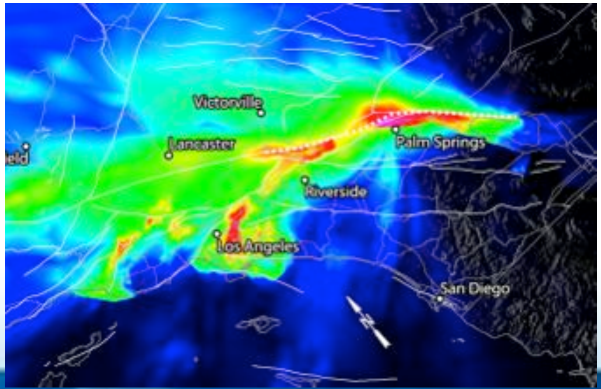


**Health**

# Blue Waters Background

- NSF is focus is science and engineering impact, not technology in its own right.
- 2006 – NSF released the solicitation for Track 1
  - It is not just a system
  - Well balanced – diverse, unspecified workload
- 2007 – NCSA and other submitted a proposal
  - 2008 – NCSA selected and the project began
- NSF funds about 60% of the effort
- Illinois and University funds about 40%

# NSF's Strategy for High-end Computing

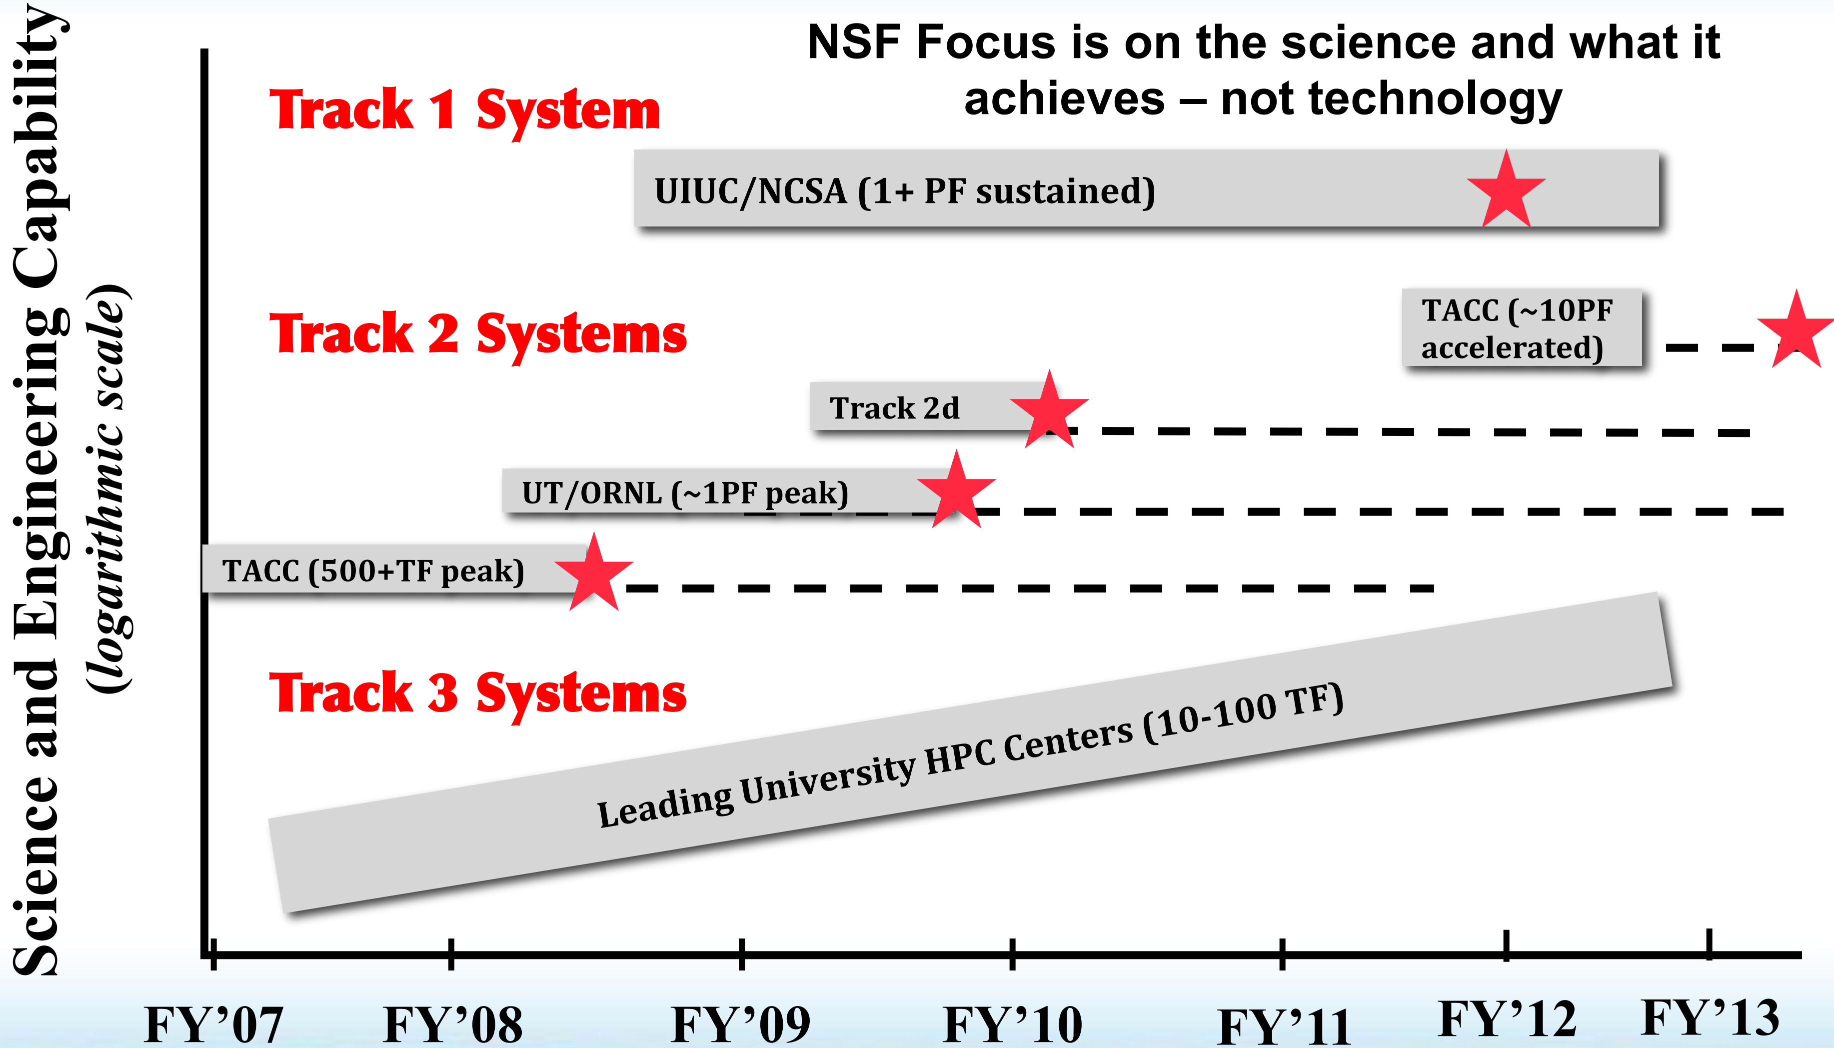

# What These Scientists Told Us They Needed

- Maximum Core Performance
  … to minimize number of cores needed for a given level of performance as well as lessen the impact of sections of code with limited scalability
- Low Latency, High Bandwidth Communications Fabric
  … to maximize scalability of science and engineering applications
- Large, Low Latency, High Bandwidth Memory Subsystem
  … to enable the solution of memory-intensive problems
- Large Capacity, High Bandwidth I/O Subsystem
  … to enable the solution of data-intensive problems
- Reliable Operation
  … to enable long-running simulations

# Blue Waters Goals

- **Deploy a computing system capable of <u>sustaining</u> more than one petaflops or more for a <u>broad</u> range of applications**
  - Cray system achieves this goal using a well defined metrics, including most intense I/O sub-system available
- **Enable the Science Teams to take full advantage of the sustained petascale computing system**
  - Blue Waters Team has established strong partnership with Science Teams, helping them to improve the performance and scalability of their applications
- **Enhance the operation and use of the sustained petascale system**
  - Blue Waters Team is developing tools, libraries and other system software to aid in operation of the system and to help scientists and engineers make effective use of the system
- **Provide a world-class computing environment for the petascale computing system**
  - The NPCF is a modern, energy-efficient data center with a rich WAN environment (100-400 usable Gbps) and data archive (>300 usable PB)
- **Exploit advances in innovative computing technology**
  - Proposal anticipated the rise of heterogeneous computing and planned to help the computational community transition to new modes for computational and data-driven science and engineering

# Focus on Sustained Performance

- **Blue Waters and NSF focus on *sustained* performance in a way few have before.**
- *Sustained* is the computer's performance on a broad range of applications that scientists and engineers use every day.
  - Time to solution is the metric – not Ops/s
  - Tests include time to read data and write the results
- NSF's call emphasized sustained performance, demonstrated on a collection of application benchmarks (application + problem set)
  - Not just simplistic metrics (e.g. HP Linpack)
  - Applications include both Petascale applications (effectively use the full machine, solving scalability problems for both compute and I/O) and applications that use a large fraction of the system
    - Metric is the time to solution
- Blue Waters project focus is on delivering sustained PetaFLOPS performance to all applications
  - Develop tools, techniques, samples, that exploit all parts of the system
  - Explore new tools, programming models, and libraries to help applications get the most from the system
- Introduced the Sustained Petascale Performance (**SPP**) Metrics

# Illinois Centers of Computing

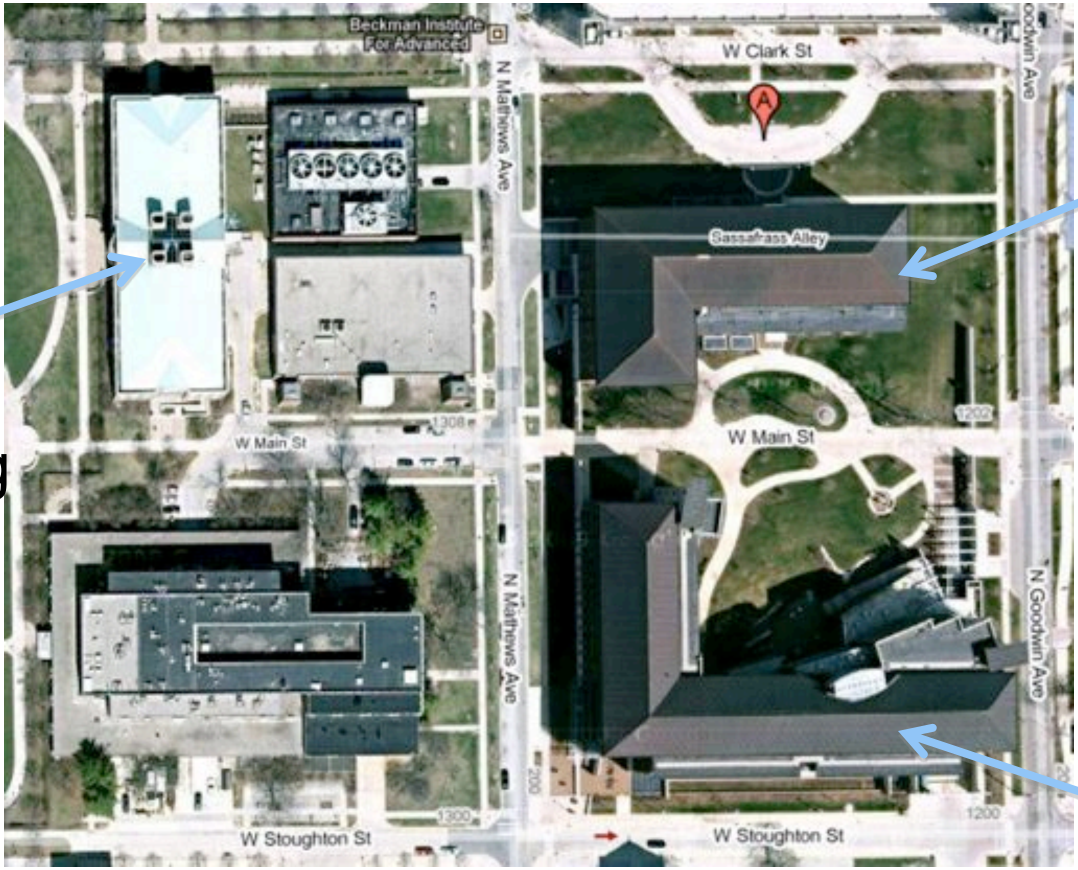


NCSA

Computer Engineering

Computer Science

# Illinois Parallel Computing Events

In sum, Illinois is connected to 23 Nobel Prizes and the entire history of High Performance Computing

1948 Oppenheimer approves Illinois to build a copy of ENIAC
1949 US Army and Illinois jointly Fund ORDVAC and ILLIAC
1951 ORDVAC completed
1951 John Bardeen – co-inventor of the Transistor – joins Illinois
1952 ORDVAC moves to ARL
**1952 ILLIAC – first computer built and owned by an academic Institution become operational**
**1955 4-bit transistor computer completed**
1957 Illinois demonstrates 10x Improved flip-flop
1958 J Robinson demonstrates Efficient binary division
1961 demonstrate 1 nanosecond circuit
**1961 PLATO** – first computer assisted Program for instruction, first time-shared education system
**1962 ILLIAC-II, transistorized computer** 100 times ILLIAC I
**1963 ILLIAC-III** – pattern recognition
1964 Dept of Computer Science formed
**1965 ILLIAC IV**, DARPA, Illinois and Burroughs Collaborate – **SIMD parallel** computation, circuit Card automations, 12 layer ECL, **Semiconductor memory**

1970 **Kuck+ translate conventional code to parallel code**
1970 C.W. Gear develops first program to solve stiff DEs
1972 J **Bardeen wins 2nd Nobel for superconductivity**
1973 Liu and Layland develop basis for timing behavior of multi-program real time computers
**1973 PLATO social networking** pioneering email (notes), group notes, etc. Dorner and Ozzie later found Eudora
**1974 ILLIAC IV operational at IAC**
**1975 ILLIAC IV operational at NASA Ames**
1975 Illinois **gets first UNIX license** from Bell Labs
1976 Illinois uses computer to **solve 4 color problem**
1977 **Kuck+ intro dependency graphs enabling efficient optimization**
1977 Campbell develops PATH Pascal
1978 UIUC Library starts on-line access
1981 NSF funds major computational facility
**1983 Patel Cache Coherent Protocol becomes standard**
1985 NASA funds ICLASS
**1985 NSF funds NCSA**
**1985 NSF funds CSRD**
**1986 NCSA Telnet**
**1988 TAPESTRY heterogeneous parallel system**
**1988 Cedar system completed**
1989 Beckman Institute

1990 Whole Internet Guide
1993 – PERTS system
**1993 Moasic released – starts wide use of the web**
1995 – Vosaic – for video and audio
1995 Reed starts Scalable I/O Initiative
1995 UI Library one of 7 to start Museum Licensing Project
1995 Illinois one of 6 sites for NSF Digital Library Project
1995 NCSA helps White House, Chicago, and others set up web sites
**1997 – PACTI – start of Teragrid**
**1997 "2001" – HAL**
2000 Levchin (CS alum) founds Paypal
2004 Seibel Center Opens
2006 NCSA Building Opens
2005 CS alums found YouTube
2008 UPCRC founded
2008 Blue Waters Award
**2012 Blue Waters Full Service**

**1940 1950 1960 1970 1980 1990 2000 2010 2020**

# National Center for Supercomputing Applications - NCSA

- Established in 1986 as one of the original sites of the National Science Foundation's Supercomputer Centers Program
- Supported by the state of Illinois, the University of Illinois, the National Science Foundation, and grants from other federal agencies.
- A leader in deploying robust high-performance computing resources and in working with research communities to develop new computing and software technologies.
  - Fielded 30 major computational systems for the national community
- NCSA focuses on:
  - Developing and supporting powerful, reliable computing, data, and networking resources that enable researchers to solve demanding science and engineering problems.
    - Develop and explore innovative architectures and techniques to accelerate scientific computing.
  - Working with research communities to help them fully exploit the extraordinary resources available on the Internet (computing systems, data sources and stores, and tools) with cyberenvironments.
  - Developing software, techniques, and tools to improve national cybersecurity and to help law enforcement better respond to cyberattacks.
  - Providing insights into complex systems and sharing the thrill of scientific discovery with the broadest possible audience through artful visualizations of scientific phenomena.
  - Preparing the next generation of scientists and engineers to effectively use computational tools and techniques.
- Two major projects (Blue Waters and XSEDE) and many other activities
- Currently evolving a strategic focus on Data Driven Science

# Parallel Computing Institute and the Center for Extreme-scale Computation

- PCI – Parallel Computing Institute
  - Help meet challenges in new programming technologies, power consumption and performance for scientific applications will require major interdisciplinary engineering efforts.
  - Designed to enable Illinois researchers from across campus to come together in new, application-focused research centers and achieve their scientific goals using the latest and most efficient parallel computing technologies.
- CEC – Center for Extreme Scale Computing
  - Develop applications and technologies to ensure the success of the Blue Waters Project
  - Explore issues critical for the realization of exascale computation in science and engineering

# GLCPC

## Great Lakes Consortium for Petascale Computing

GLCPC Mission:

*"…facilitate and coordinate multi-institutional efforts to advance computational and computer science engineering, and technology research and education and their applications to outstanding problems of regional or national importance…"*

- 2% allocation
- 501c3 Organization*
- 28 Charter members**
- Executive Committee
- Allocations Committee
- Education Committee

* State 501c3 filing complete, federal filing in progress

** 28 Charter members represent over 80 universities, national laboratories, and other education agencies

# NSF Cyber Infrastructure for the 21st Century CIF-21

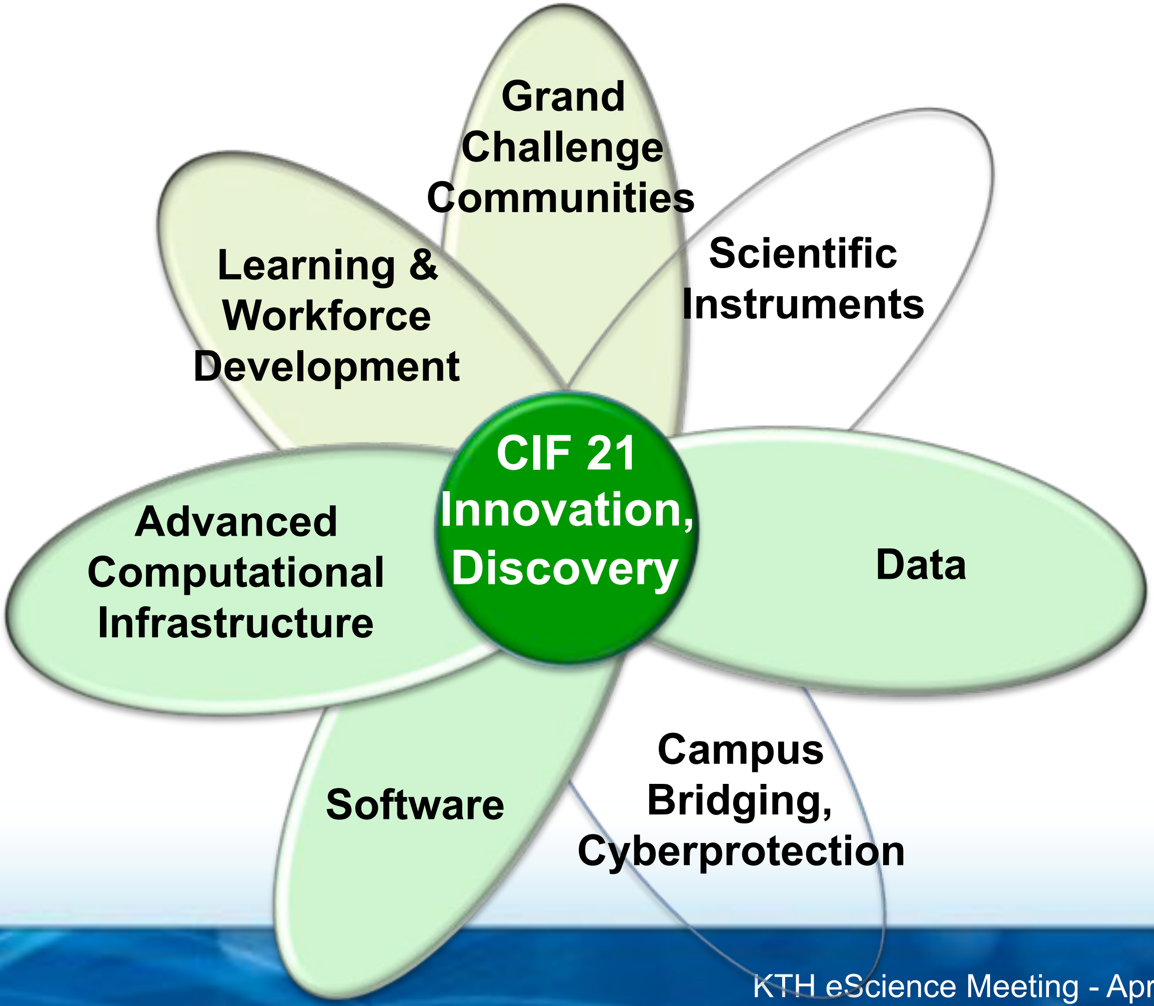

# OCI Advanced Computing Infrastructure Vision

1. Foundational research to fully exploit parallelism and concurrency
2. Applications research and development
3. Sustainable and innovative resources into a collaborative ecosystem
4. Comprehensive education and workforce programs
5. Transformational and grand challenge community programs that support complex problem solving

# XSEDE

- Extreme Science and Engineering Discovery Environment
- XSEDE is the follow-on to TeraGrid and is the virtual organization of 17 institutions/120 FTE providing services and support for the users of NSF-OCI's computational science resources.
  - CMS – coordination and management services
    - architecture, software integration, systems engineering, operations, allocations, project management
  - ECSS – extended collaborative support services
  - TEOS – training, education, and outreach services
  - 5-year award started July 1, 2011

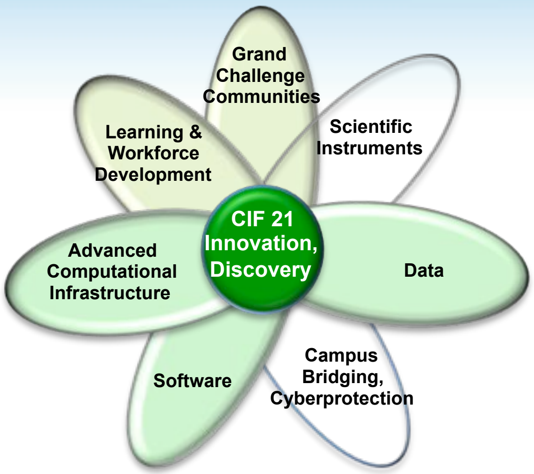
Grand Challenge Communities
Scientific Instruments
Learning & Workforce Development
CIF 21 Innovation, Discovery
Advanced Computational Infrastructure
Data
Software
Campus Bridging, Cyberprotection

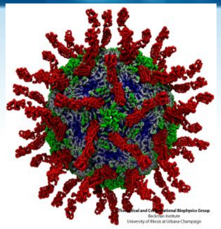
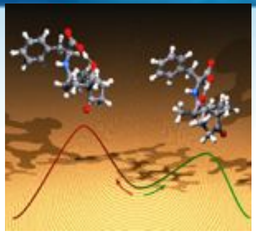

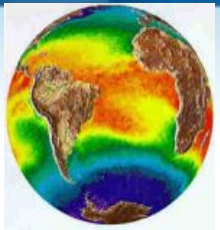

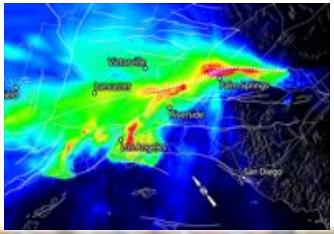


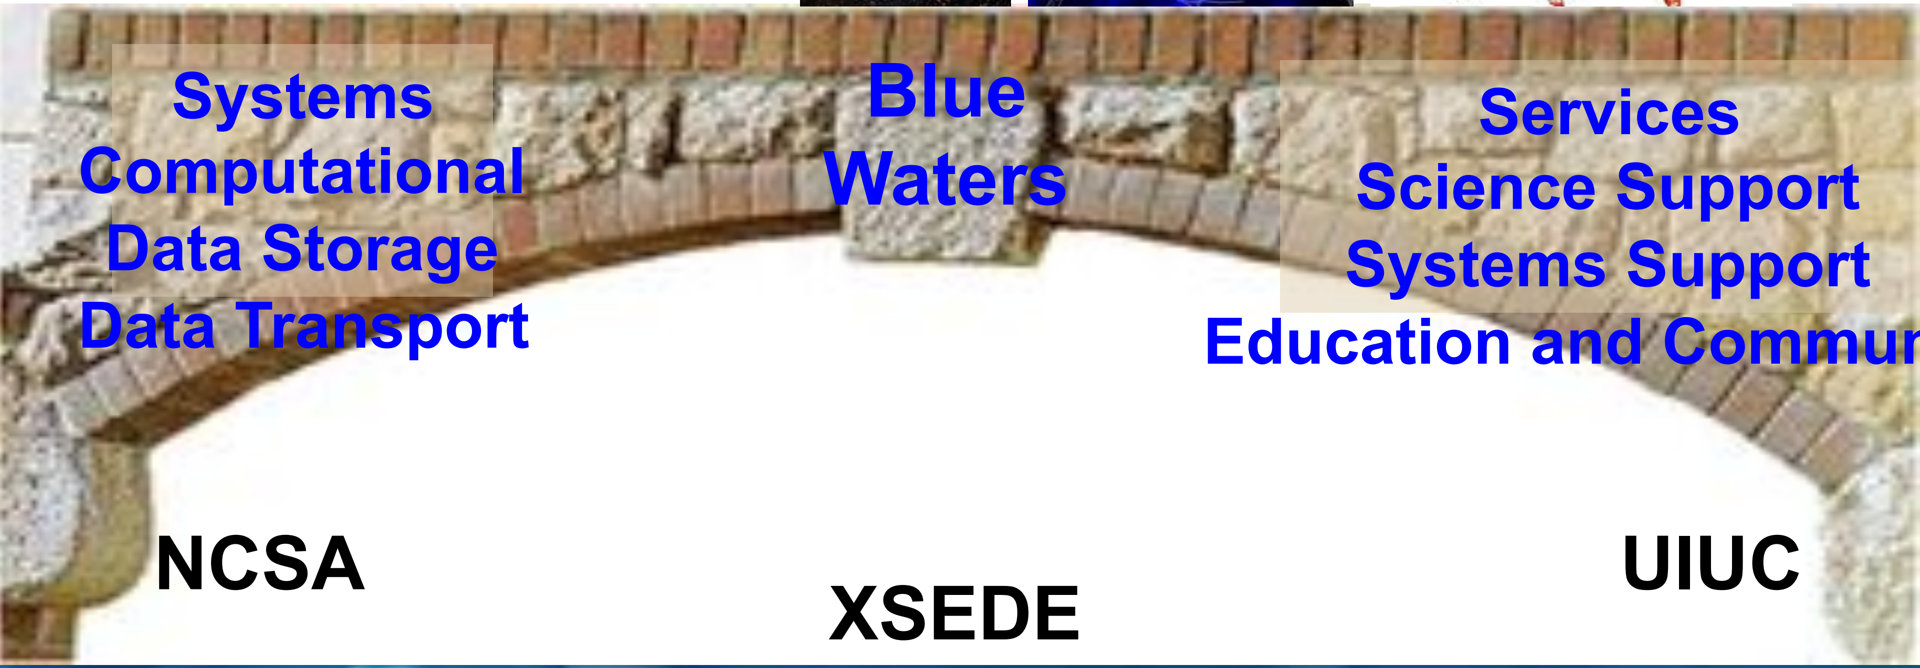
Systems
Computational
Data Storage
Data Transport
Blue Waters
Services
Science Support
Systems Support
Education and Commun
NCSA
XSEDE
UIUC

# The Blue Waters Eco-System

Petascale Education, Industry and Outreach

Petascale Applications
(Computing Resource Allocations)

Petascale Application Collaboration Team Support

Outstanding Full Services
WAN Connections, Science and Engineering Application Support, System Management, Cyber-protection, Operations, …

Great Lakes Consortium for Petascale Computing

Value added hardware (external networking, IDS, nearline storage, import/export, etc) and software (Eclipse, HPSS, Visualization, computational libraries, etc)

Blue Waters Base System – Processors, Memory, Interconnect, Online Storage, System Software, Programming Environment

National Petascale Computing Facility

# Blue Waters Computing Super-system

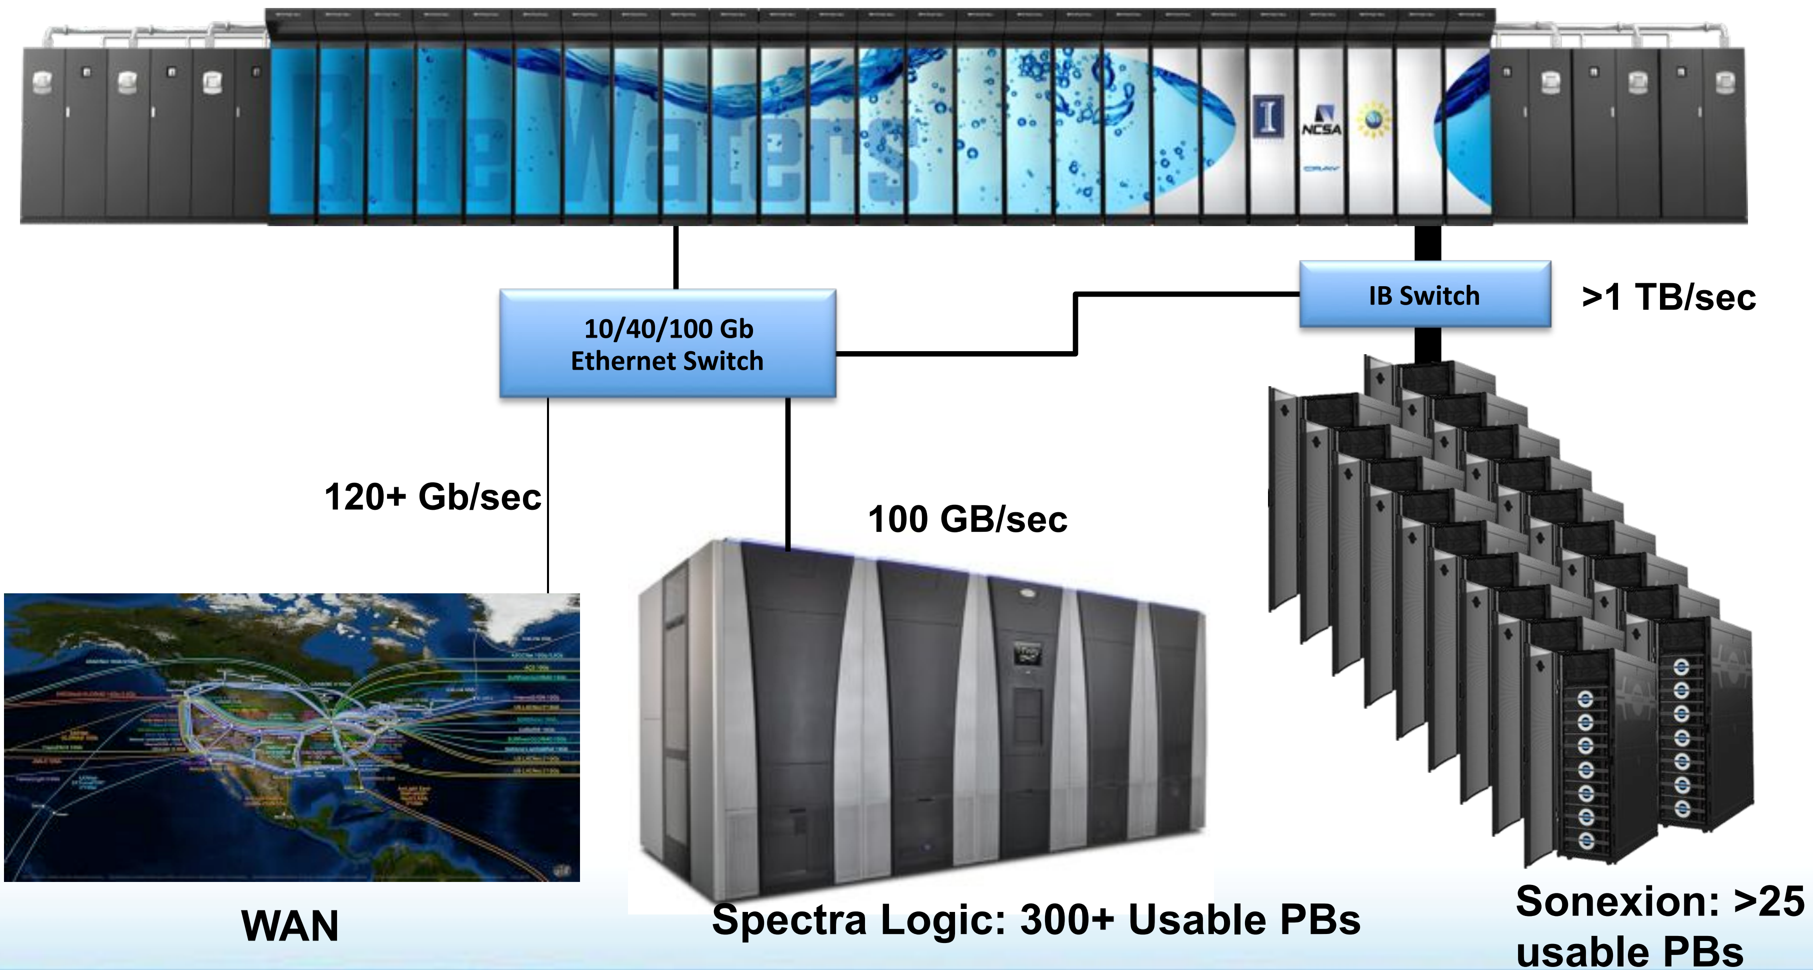

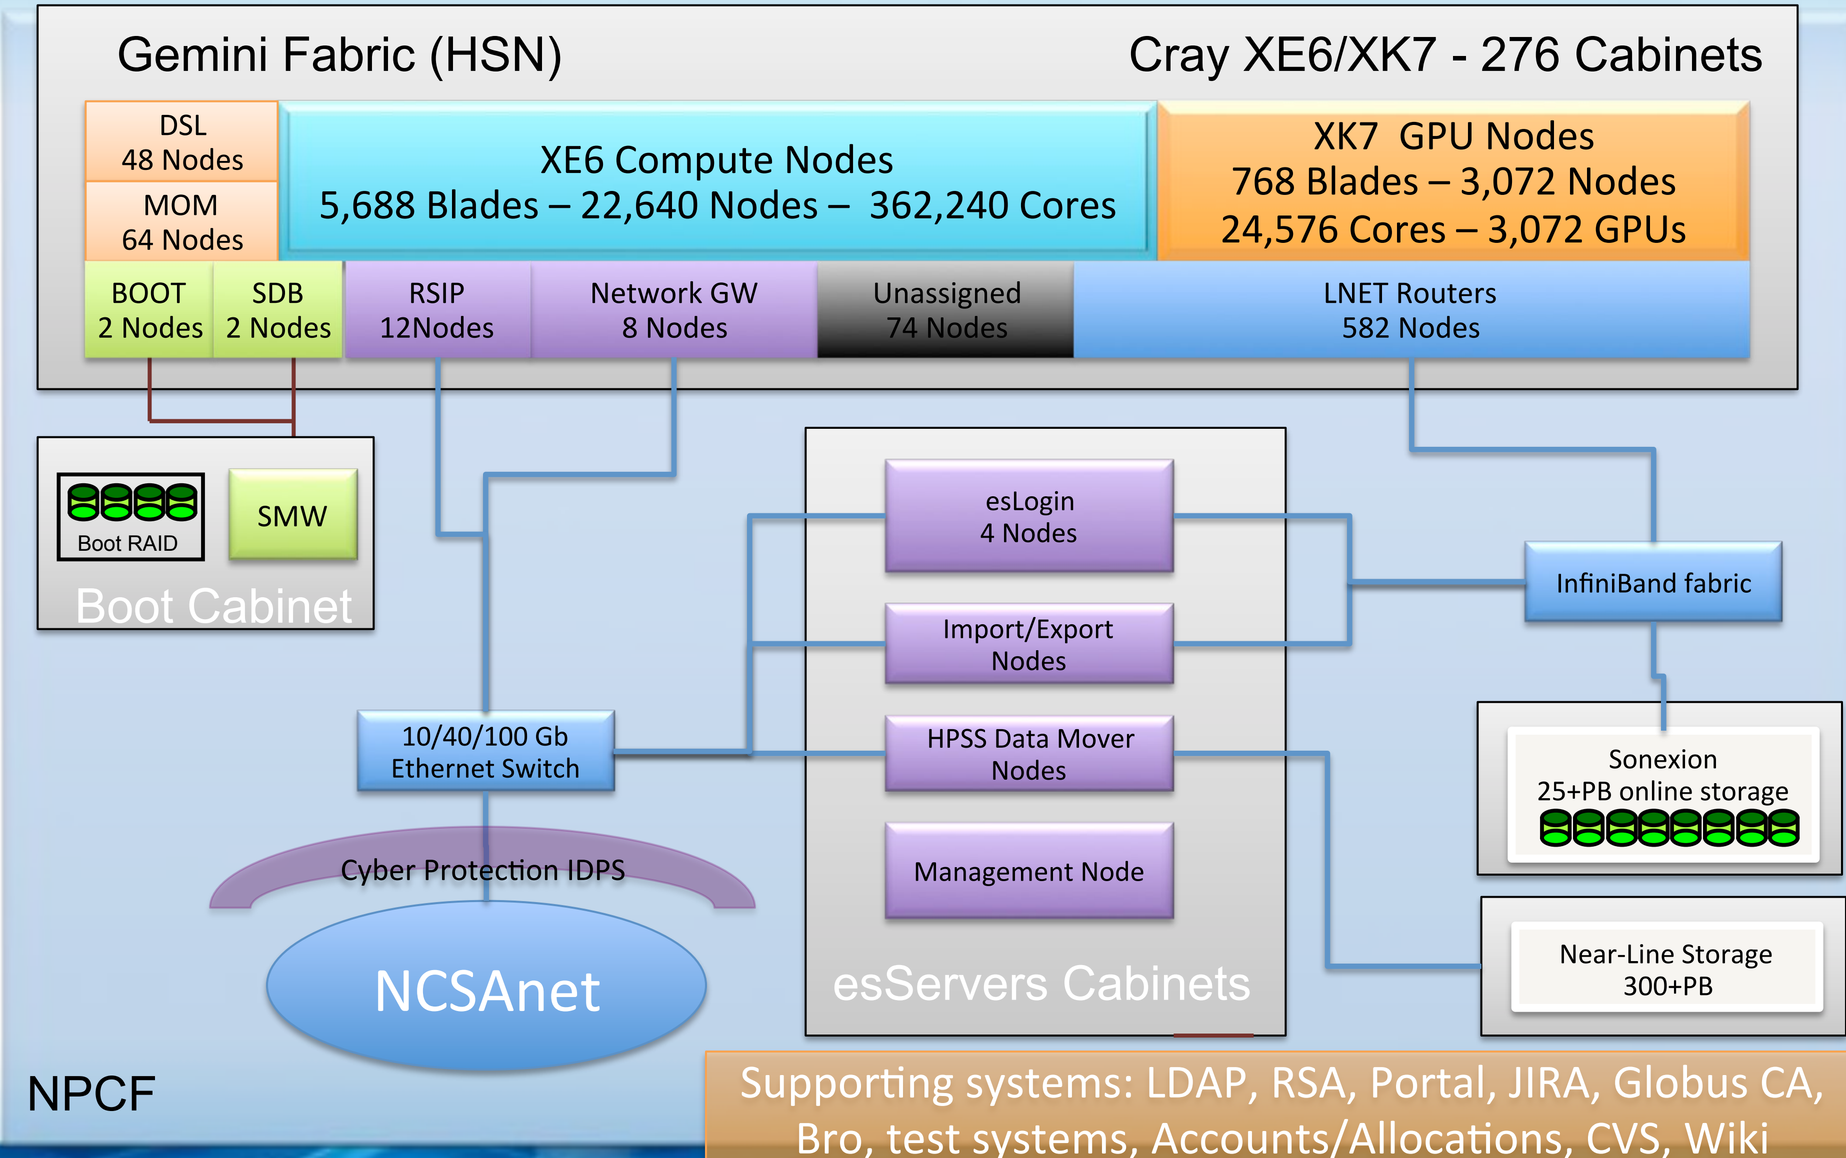
Gemini Fabric (HSN)
Cray XE6/XK7 - 276 Cabinets
DSL
48 Nodes
MOM
64 Nodes
XE6 Compute Nodes
5,688 Blades – 22,640 Nodes – 362,240 Cores
XK7 GPU Nodes
768 Blades – 3,072 Nodes
24,576 Cores – 3,072 GPUs
BOOT
2 Nodes
SDB
2 Nodes
RSIP
12Nodes
Network GW
8 Nodes
Unassigned
74 Nodes
LNET Routers
582 Nodes
Boot RAID
SMW
Boot Cabinet
esLogin
4 Nodes
Import/Export
Nodes
HPSS Data Mover
Nodes
Management Node
esServers Cabinets
InfiniBand fabric
Sonexion
25+PB online storage
Near-Line Storage
300+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

# Defining a Core - AMD Wide AVX mode

- In this mode, only one integer scheduler unit is used
  - Most common mode for S&E applications
    - Code is Floating Point dominated and makes use of AVX instructions
    - Code needs more memory per MPI rank
- Implications
  - This core has *exclusive* access to the 256-bit FP unit and is capable of 8 FP results per clock cycle
  - The core has twice the memory capacity
  - The core has twice the memory bandwidth
  - The L2 cache is effectively twice as large
  - The peak performance of the chip is not reduced
- AMD refers to this as a "Core Module"

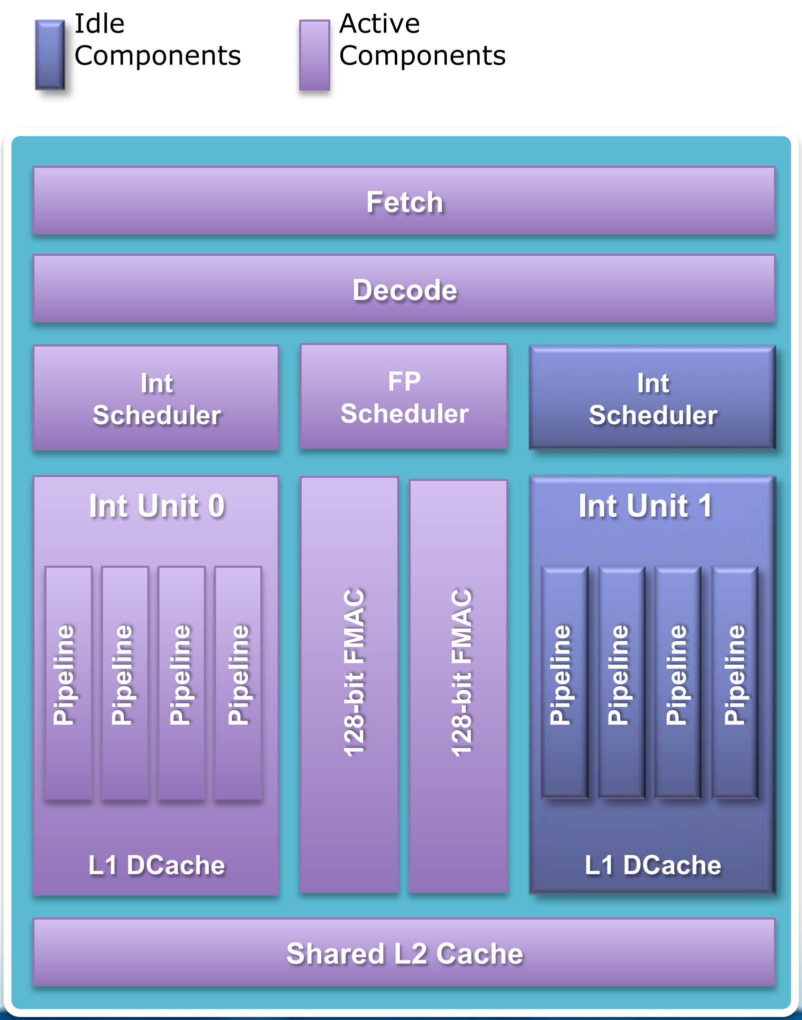

# Blue Waters HighSpeed Network

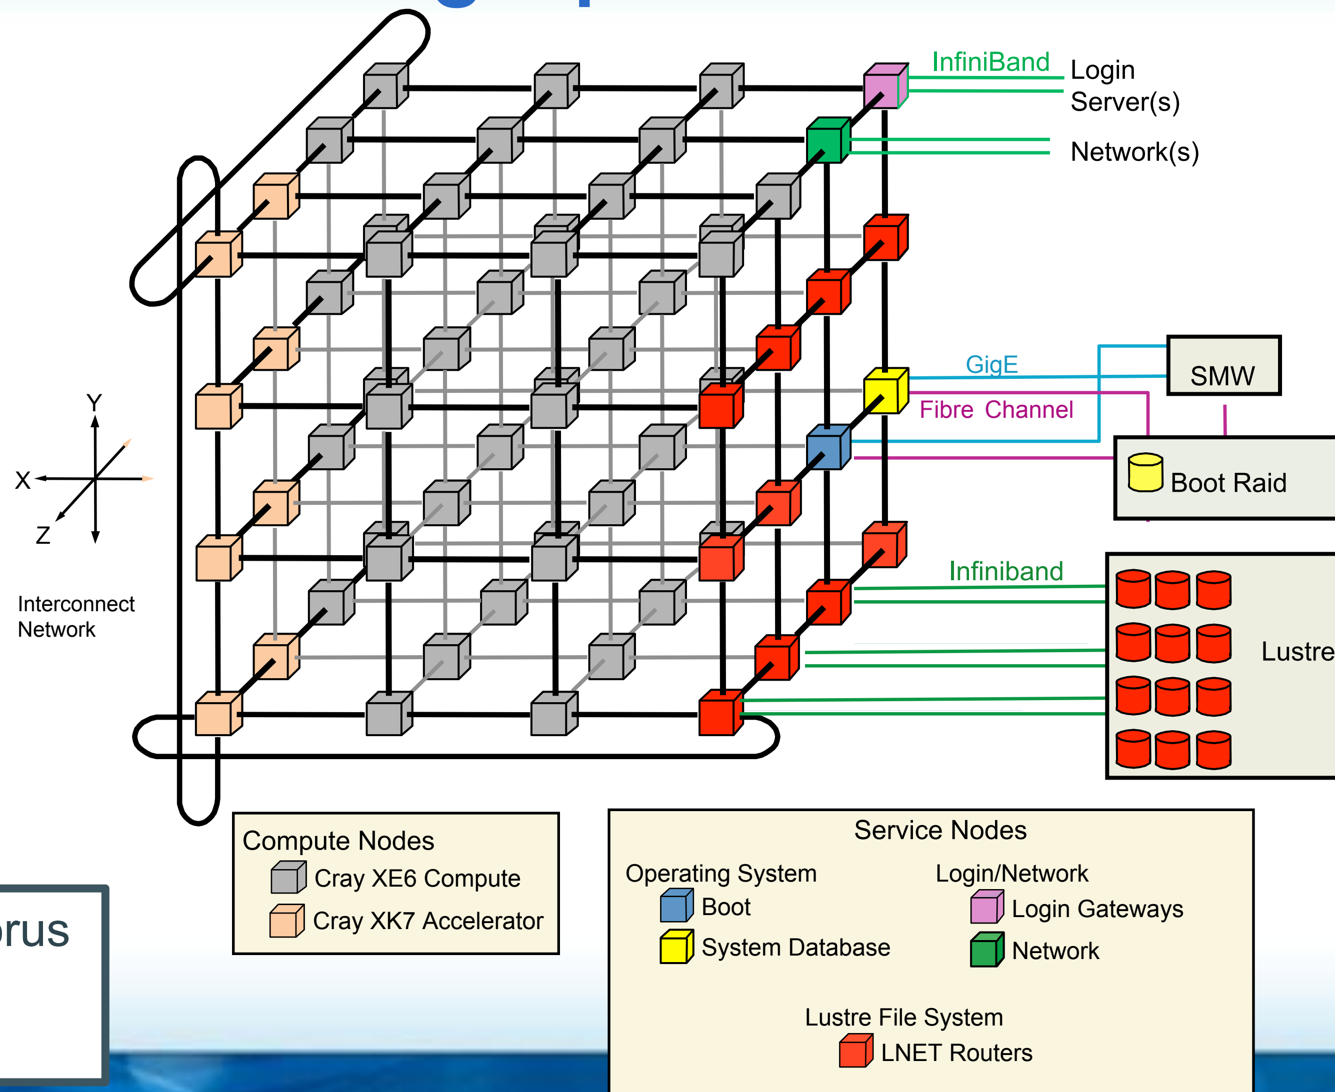


Blue Waters 3D Torus Size
23 x 24 x 24

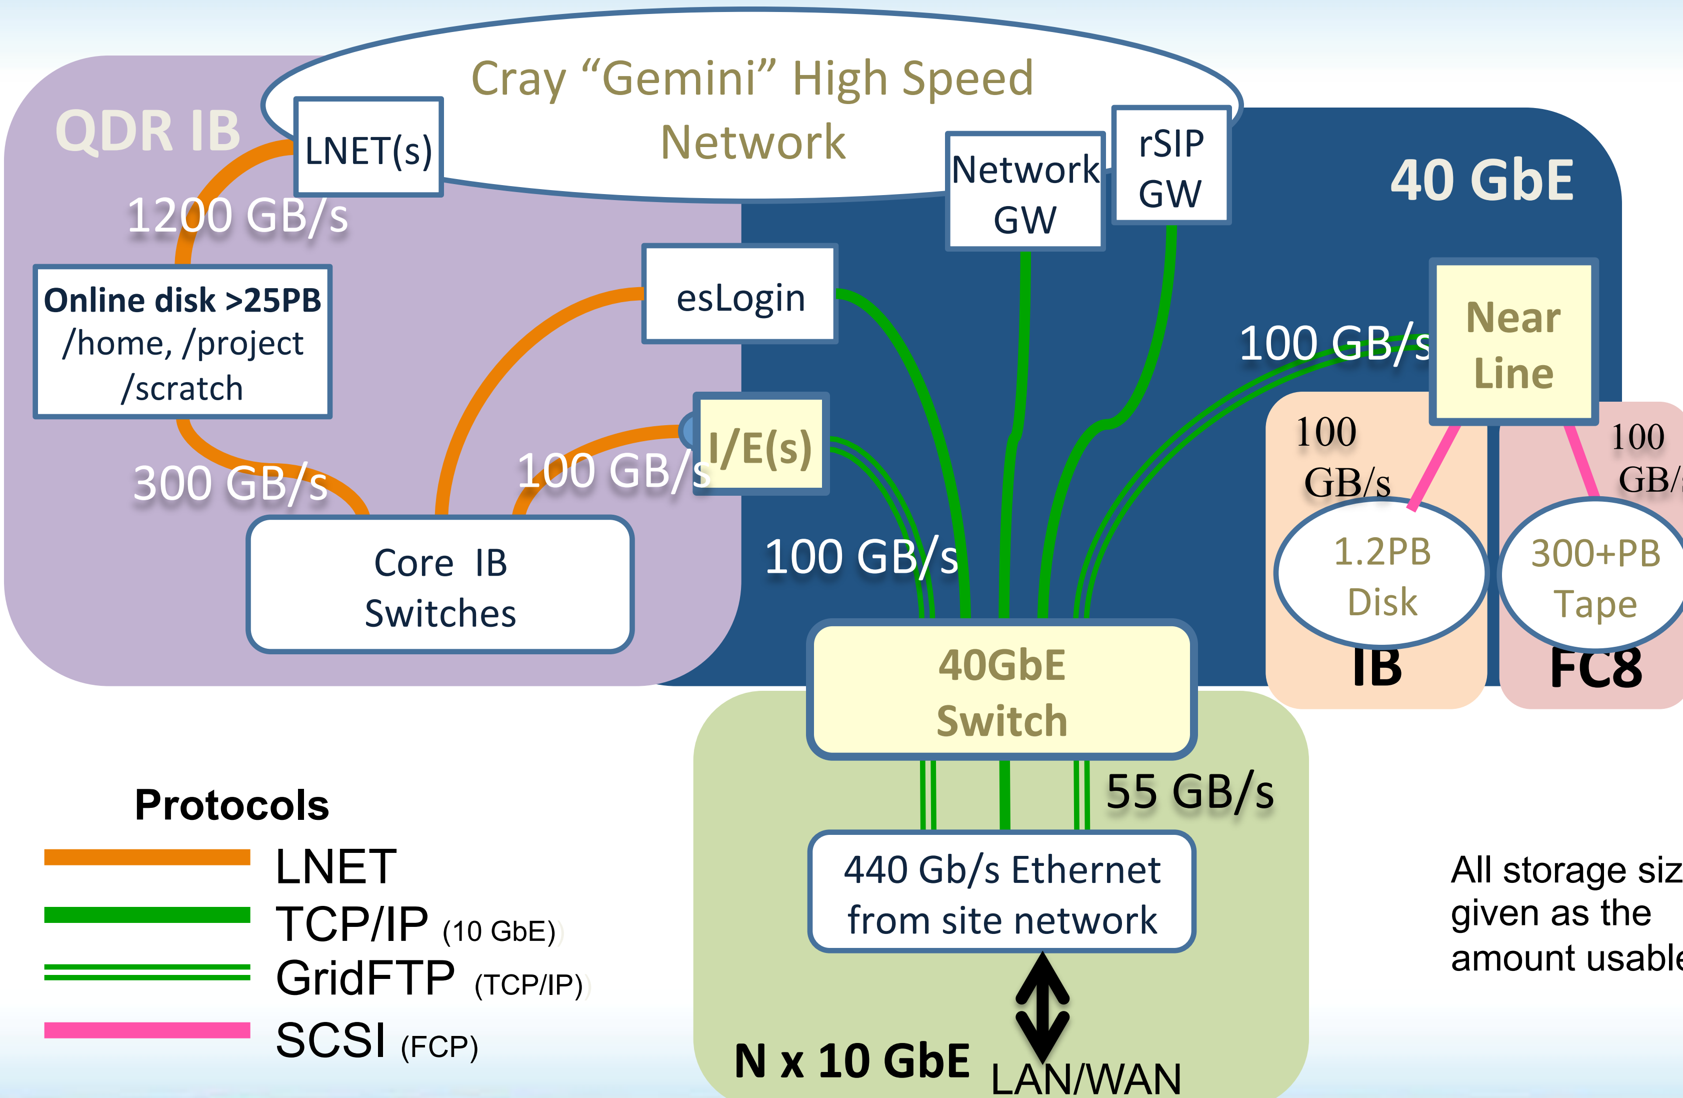


All storage sizes given as the amount usable

# Computational Microscopy

Klaus Schulten (UIUC)

- Simulations of 4 different biological systems
  - Protein elongation in the ribosome
  - Structural transitions in poliovirus entry
  - Sculpting cellular membranes by BAR domains
  - Energy conversion by the chromatophore organelle
- NAMD molecular dynamics code
  - System size will exceed 100M atoms for the first time

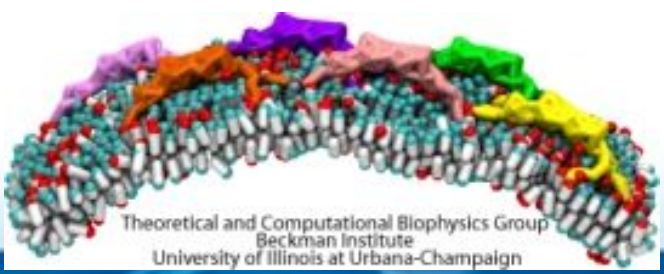


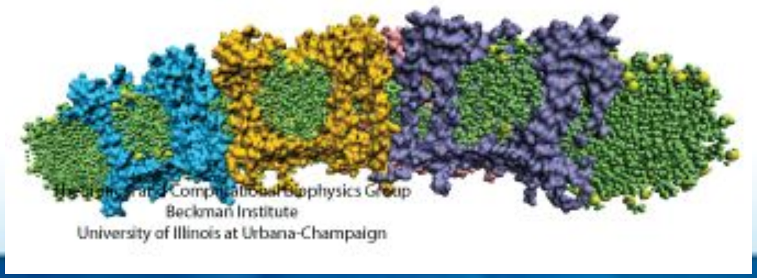

# Breakthrough Petascale Quantum Monte Carlo Calcs.

Shiwei Zhang, David Ceperley (UIUC)

- Computational Goals
  - Properties of hydrogen and hydrogen-helium mixtures in astrophysically relevant conditions
  - Electronic & magnetic structures of transition-metal compounds
  - Phases and properties of fermionic atoms in optical lattices
  - Properties of materials for energy conversion
  - Dynamical QMC simulations of liquid water and other disordered materials

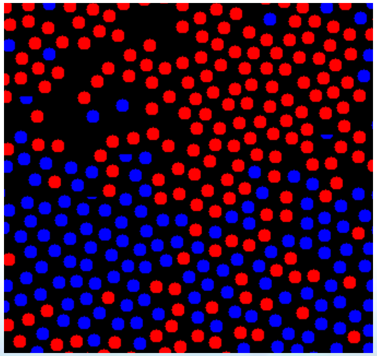

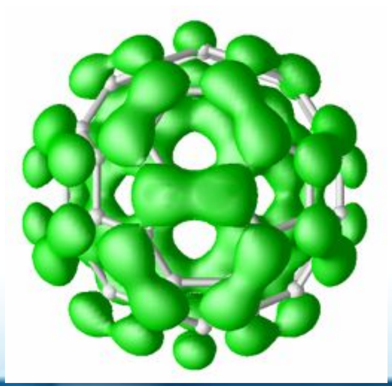

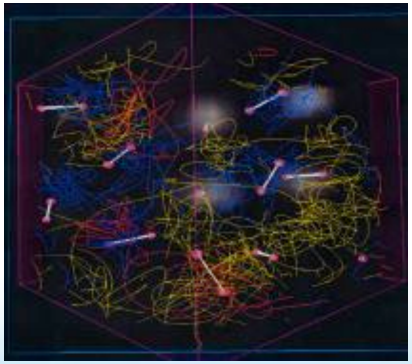

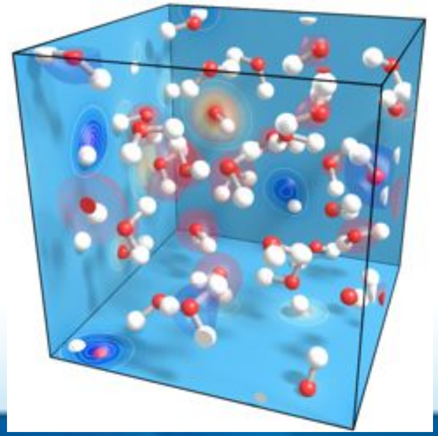

# Understanding Tornadoes and Their Parent Supercells Through Ultra-High Resolution Simulation/Analysis

Bob Wilhelmson (UIUC), et al

- Goal
  - Simulate development, structure, & demise of large damaging tornadoes in supercells at resolution sufficient to capture low-level tornado inflow, thin precipitation shaft that forms “hook echo” adjacent to tornado, and other smaller scale structures
- Approach
  - Simulate storms with model that solves (using finite differences) the equations of motion for air and water substances (droplets, rain, ice, …)

# Earthquake System Science (PressOn)

Tom Jordan, Phil Maechling

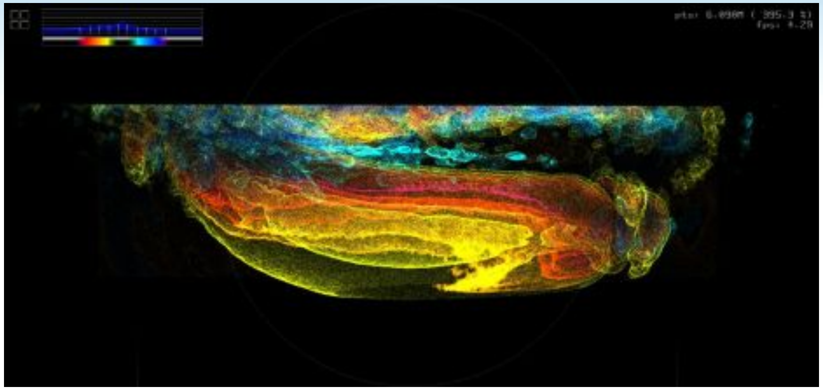

- Prepare 3 seismic & engineering codes for Blue Waters
  1. AWP-Olsen: finite difference dynamic rupture & wave propagation
  2. Hercules: finite element wave propagation
  3. OpenSees: structure modeling for earthquake engineering
- Science Goals
  - Improve resolution of dynamic rupture simulations by ~10X to investigate realistic friction laws, near-fault stress states, & off-fault plasticity
  - Investigate upper frequency limit of deterministic ground-motion prediction by simulating strong motions up to 3 Hz using realistic 3D structural models for Southern California
  - Improved understanding of engineering response to earthquakes by modeling response of buildings to broadband ground motions

# Testing Hypotheses About Climate Prediction

Stan, Kirtman, Large, Randall

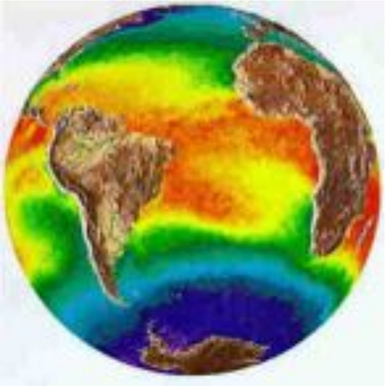

- Hypotheses
  - The large errors in current-generation climate models are associated with fundamentally flawed assumptions in the parameterizations of cloud processes and ocean eddy mixing processes
  - Climate predictability, a synthetic quality entirely associated with a given model, increases with increasing model resolution by virtue of the changing representation of atmospheric and oceanic noise
- Target Problems
  - Annual Cycle Experiment using the Co. St. U. Global Cloud-Resolving Atmospheric General Circulation Model
    - Test if annual cycle of physical climate system quantities such as the precipitation and surface temperatures are more accurately reproduced when both cloud processes and the ocean meso-scale are resolved and coupled

# Industry S&E Teams

- Industry can participate in the NSF PRAC process
- 5+% allocation can dedicated to industrial use
  - Specialized support by the NCSA Private Sector Program (PSP) staff
  - Blue Waters staff will support the PSP staff as needed
  - Potential to provide specialized services within Service Level Agreements parameters
    - E.g. throughput, extra expertise, specialized storage provisions, etc.

High interest shared by partner companies in the following:

- Scaling capability of a well-known and validated CFD code
- Temporal and transient modeling techniques and understanding.

Two example cases under discussion:

- NASA OVERFLOW at scale for CFD flows
- Temporal modeling techniques using the freezing of $H_2O$ molecules as a use case and as a reason to conduct both large-scale single runs and to gain significant insight by reducing uncertainty.

# Engagement with Science Teams

*Deep partnership with Science Teams focused on sustained performance for their applications*

- NCSA PoC for each major Science Team
  - In-depth engagement on science & engineering applications, e.g.
    - Optimize performance on chip, node and system
    - IO optimization for data-intensive applications
- New capabilities for Science Teams, e.g.
  - Petascale application development environment
  - Optimized libraries and methods (numerical, communications, IO)
  - New algorithms, for both CPUs and GPUs
- Computational models of applications
  - In-depth understanding of performance shortfalls/bottlenecks

| Science Area | Number of Teams | Codes | Structured Grids | Unstructured Grids | Dense Matrix | Sparse Matrix | N-Body | Monte Carlo | FFT | Significant I/O |
|---|---|---|---|---|---|---|---|---|---|---|
| Climate and Weather | 3 | CESM, GCRM, CM1, HOMME | X | X | | X | | X | | |
| Plasmas/ Magnetosphere | 2 | H3D(M), OSIRIS, Magtail/UPIC | X | | | | X | | X | X |
| Stellar Atmospheres and Supernovae | 2 | PPM, MAESTRO, CASTRO, SEDONA | X | | | X | | X | | X |
| Cosmology | 2 | Enzo, pGADGET | X | | | X | X | | | |
| Combustion/ Turbulence | 1 | PSDNS | X | | | | | | X | |
| General Relativity | 2 | Cactus, Harm3D, LazEV | X | | | X | | | | |
| Molecular Dynamics | 4 | AMBER, Gromacs, NAMD, LAMMPS | | | X | | X | | X | |
| Quantum Chemistry | 2 | SIAL, GAMESS, NWChem | | | X | X | X | X | | X |
| Material Science | 3 | NEMOS, OMEN, GW, QMCPACK | | | X | X | X | X | | |
| Earthquakes/ Seismology | 2 | AWP-ODC, HERCULES, PLSQR, SPECFEM3D | X | X | | | X | | | X |
| Quantum Chromo Dynamics | 1 | Chroma, MILD, USQCD | X | | X | X | X | | X | |
| Social Networks | 1 | EPISIMDEMICS | | | | | | | | |
| Evolution | 1 | Eve | | | | | | | | |
| Computer Science | 1 | | | X | X | X | | | X | X |

# Motivation – Effective Use of Systems

- Additional effort by the science and engineering teams needed to achieve full potential of Blue Waters system, with all of its advanced technology
- Larger computational science and engineering (S&E) community needs assistance to take full advantage of the capability provided by today's computing systems, from increasingly powerful campus clusters to high-end systems (Blue Waters, Kraken, Stampede)
- Increasing performance requires dramatic increases in parallelism that then generates complexity challenges for science and engineering teams~~.~~
  - **Scaling applications to large core counts on general-purpose CPU nodes**
  - **Effectively using accelerators and "many-core" processors**
  - **Using homogeneous general purpose and heterogeneous (accelerated) nodes in a single, coordinated simulation**
  - **Effectively using parallel IO systems for data-intensive applications**
  - **Effectively using limited bandwidth of interprocessor network**
  - **Enhancing application flexibility to increasing effective, efficient use of systems**

# Scalability

- Developing better process-to-node mapping using for graph analysis to determine MPI behavior and usage patterns.
- Topology Awareness in Applications and in Resource Management
- Improve use of the available bandwidth (MPI implementations, lower level communication, etc.).
  - For example, the DNS analysis assumes only a relatively low fraction of available bandwidth will be achieved – can this be improved? Most likely.
- Considering alternative programming methods that improve efficiency of calculations (e.g., one-sided access can reduce memory bandwidth requirements).
- UI Staff and  other NCSA collaborators and partners, working closely with the Science Teams, will explore the above approaches.
  - Most of the above approaches will provide an increase of a factor of 2-6 in effective bandwidth.

# Many Core and Accelerated Units

- Help the science teams to make more effective use of GPUs consists of two major components.
  - Introduce compiler and library capabilities into the science team workflow to significantly reduce the programming effort and impact on code maintainability. Examples:
  - Compiler based directives
  - Tools
    - GMAC - a library that provides global shared memory and automates data transfer/coherence between the CPUs and the GPUs in a node
    - DL is a compiler-based memory layout transformation tool that uses a combination of compiler and runtime support to easy the task of adjusting memory layout to satisfy conflicting needs between the CPU and the GPU
    - TC is a compiler based tool for thread coarsening and data tiling.
- Provide expert support to the science teams through hand-on workshops, courses, and individualized collaboration programs.

# Approach to Accelerator Programming

- Most important hurdle for widespread adoption of accelerated computing is programming difficulty
- The community needs a single programming model that is portable across machine types, is forward scalable in time AND **allows users to maintain a single code base**
  - Portable expression of heterogeneity and multi-level parallelism
  - Programming model and optimization should not be significantly difference for “accelerated” nodes and multi-core x86 processors
- Approach:
  - Use accelerator specific coding + OpenACC compiler support
  - Support 3rd party GPU/Accelerator tools and languages for compatibility
    - CUDA and OpenCL
    - PGI Fortran compiler
    - Allinea debugger
    - UIUC tools for data layout, granularity adjustment and regulaization
  - Optimized scientific libraries for Accelerators
  - Cray compiler with native support for Accelerators
    - C, C++ and Fortran;  MPI and OpenMP
    - Directives based on OpenMP for identifying parallel work
  - Whole program scoping tools

# Using Both General Purpose And Accelerated Nodes In Single Application.

- For multi-physics applications that provide a natural decomposition into modules is to deploy the most appropriate module(s) different computational units.
  - NCSA will assist in identifying appropriate modules, and in the mechanics of heterogeneous partitioning.
- For applications, such as NAMD, Episimdemics, and possibly ENZO, that use the Charm++ adaptive runtime system, heterogeneity can be handled without significant changes to the application itself.
- MPI applications may be able to leverage the Charm++ runtime system by converting them to adaptive MPI (AMPI) first - EVE and CM1.
- Some applications naturally involve assigning multiple blocks to individual processors include multiblock codes (typically in fluid dynamics), and the codes based on structured adaptive mesh refinement.
  - The application-level load balancing algorithms can be modified to deal with the performance heterogeneity created by the mix of nodes. The NCSA/Illinois staff will assist in such modification.
- Some applications use frameworks for accomplishing their load-balancing (Zoltan, UNITAH, Paramesh and Chombo, etc.) that already address the issue of differential performance of different processors.

# Application Based Resiliency

- Multiple layers of Software and Hardware have to coordinate information and reaction
- Analysis and understanding is needed before action
- Correct and actionable messages need to flow up and down the stack to the applications so they can take the proper action with correct information
- Applications need to understand circumstances and take action
- Flexible resource provisioning needed in real time
- Interaction with other constraints so sub-optimization does not adversely impact overall system optimization

# Enhanced Intellectual Services

- Direct, short-term funding for science teams to implement at least one new method to address scaling or heterogeneity improvements
- Funds to encourage science team and new method providers to work together to make mid-term transformative improvements
- Education
  - Blue Waters Fellowships
  - Year round Virtual School for Computational Science and Engineering
- Community Building
  - Fund Workshops

| PI | Award Date | Project Title |
|---|---|---|
| Sugar | 04/15/2009 | Lattice QCD on Blue Waters |
| Bartlett | 04/15/2009 | Super instruction architecture for petascale computing |
| Nagamine | 04/15/2009 | Peta-Cosmology: galaxy formation and virtual astronomy |
| Bissett | 05/01/2009 | Simulation of contagion on very large social networks with Blue Waters |
| O'Shea | 05/01/2009 | Formation of the First Galaxies: Predictions for the Next Generation of Observatories |
| Schulten | 05/15/2009 | The computational microscope |
| Stan | 09/01/2009 | Testing hypotheses about climate prediction at unprecedented resolutions on the NSF Blue Waters system |
| Campanelli | 09/15/2009 | Computational relativity and gravitation at petascale: Simulating and visualizing astrophysically realistic compact binaries |
| Yeung | 09/15/2009 | Petascale computations for complex turbulent flows |
| Schnetter | 09/15/2009 | Enabling science at the petascale: From binary systems and stellar core collapse To gamma-ray bursts |

| PI | Award Date | Project Title |
|---|---|---|
| Woodward | 10/01/2009 | Petascale simulation of turbulent stellar hydrodynamics |
| Tagkopoulos | 10/01/2009 | Petascale simulations of Complex Biological Behavior in Fluctuating Environments |
| Wilhelmson | 10/01/2009 | Understanding tornadoes and their parent supercells through ultra-high resolution simulation/analysis |
| Wang | 10/01/2009 | Enabling large-scale, high-resolution, and real-time earthquake simulations on petascale parallel computers |
| Jordan | 10/01/2009 | Petascale research in earthquake system science on Blue Waters |
| Zhang | 10/01/2009 | Breakthrough peta-scale quantum Monte Carlo calculations |
| Haule | 10/01/2009 | Electronic properties of strongly correlated systems using petascale computing |
| Lamm | 10/01/2009 | Computational chemistry at the petascale |
| Karimabadi | 11/01/2010 | Enabling Breakthrough Kinetic Simulations of the Magnetosphere via Petascale Computing |
| Mori | 01/15/2011 | Petascale plasma physics simulations using PIC codes |

| PI | Award Date | Project Title |
| --- | --- | --- |
| Voth | 02/01/2011 | Petascale multiscale simulations of biomolecular systems |
| Woosley | 02/01/2011 | Type Ia supernovae |
| Cheatham | 02/01/2011 | Hierarchical molecular dynamics sampling for assessing pathways and free energies of RNA catalysis, ligand binding, and conformational change |
| Wuebbles | 04/15/2011 | Using petascale computing capabilities to address climate change uncertainties |
| Gropp | 06/01/2011 | System software for scalable applications |
| Klimeck | 09/15/2011 | Accelerating nano-scale transistor innovation |
| Pande | 09/15/2011 | Simulating vesicle fusion on Blue Waters |

# Blue Waters Early Science System

- BW-ESS Configuration
  - 48 cabinets, 4,512 XE6 compute nodes, 96 service nodes
  - 2 PBs Sonexion Lustre storage appliance
- Access through Blue Waters Portal
  - https://bluewaters.ncsa.illinois.edu/

- Current Projects
  - **Biomolecular Physics**—K. Schulten, University of Illinois at Urbana-Champaign
  - **Cosmology**—B. O'Shea, Michigan State University
  - **Climate Change**—D. Wuebbles, University of Illinois at Urbana-Champaign
  - **Lattice QCD**—R. Sugar, University of California, Santa Barbara
  - **Plasma Physics**—H. Karimabadi, University of California, San Diego
  - **Supernovae**—S. Woosley, University of California Observatories
  - *Three more projects held in reserve*

# Virtual School of Comp. Science & Engineering

- Goal
  - Prepare the current & next generation of scientists and engineers to utilize leading-edge computer systems
- Strategy
  - Provide courses unavailable in graduate schools, e.g.,
    - Petascale Computing Environments and Tools
    - Proven Algorithmic Techniques for Many-core Processors
- Implementation
  - Eight Week-long Summer Schools (2008-11)
    - HD video broadcast to multiple sites
    - Over 1,300 students attended Summer Schools
    - Videos, slides and exercises are publicly available

# Undergraduate Education Program

- Goals
  - Develop awareness and knowledge of petascale computing and petascale-enabled science and engineering
  - Promote participation of faculty and students from diverse institutions and fields
    - 2- and 4-year primarily undergraduate colleges and universities
    - Minority Serving Institutions
    - EPSCoR institutions
- Program Elements
  - Professional development workshops for undergraduate faculty (>700)
  - Undergraduate course materials development (32 modules)
  - Undergraduate student internships (42)

# The Blue Waters Team (partially)

# Blue Waters Sustaining Goals

- Deploy a capable, balanced system of sustaining more than one petaflops or more for a broad range of applications
- Enable Science Teams to take full advantage of sustained petascale systems
- Enhance the operation and use of the sustained petascale system
- Provide a world-class computing environment for the petascale system
- Exploit advances in innovative computing technology
- Provide National Leadership

# Acknowledgements

This work is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (award number OCI 07-25070) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign, its National Center for Supercomputing Applications, Cray, and the Great Lakes Consortium for Petascale Computation.

The work described is achievable through the efforts of the Blue Waters Project.

Individual Contributions From

- Thom Dunning, Marc Snir, Wen-mei Hwu, Bill Gropp
- Cristina Beldica, Brett Bode, Michelle Butler, Greg Bauer, Mike Showerman, John Melchi, Scott Lathrop, Merle Giles
- Sanjay Kale, Steve Lumetta, Ravi Iyer, David Padua
- The Blue Waters Project Team and our partners
- NSF/OCI
- Cray, Inc, AMD, NVIDIA, Xyratex, Adaptive, Allinea