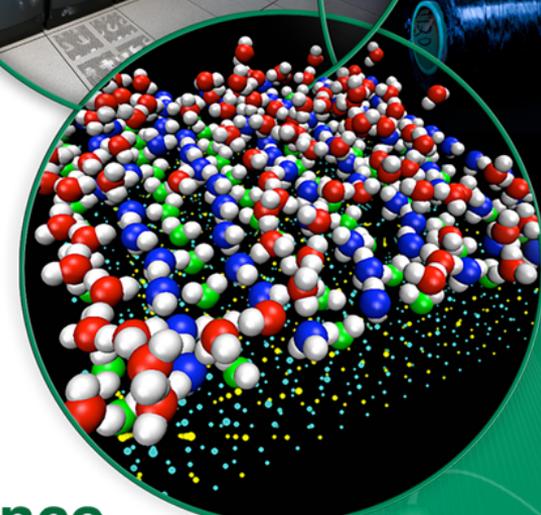
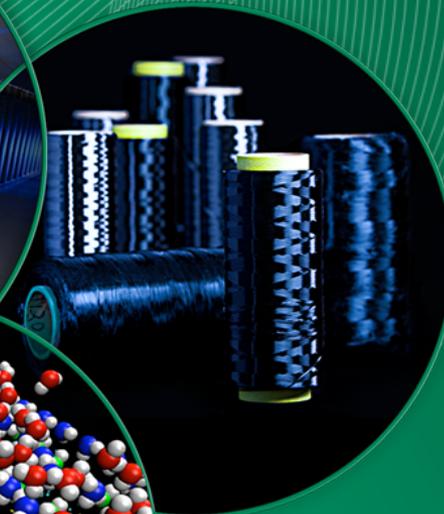


Session 2: Vendor Ecosystems to Deliver Science

Galen Shipman



**Smoky Mountains
Computational Sciences
and Engineering Conference**

Gatlinburg, TN September 4, 2013

Advanced computing infrastructure – a core component of S&T across DOE



National Dialog on Scientific Data

- America COMPETES Reauthorization Act of 2010 - P.L. 111-358
 - Interagency public access committee would coordinate Federal science agency research and policies related to the dissemination and long-term stewardship of the results of federally supported unclassified research
- Office of Science and Technology Policy – Big Data Across the Federal Government March 29, 2012
 - “Today, the Obama Administration is announcing the Big Data Research and Development Initiative. By improving our ability to extract knowledge and insights from large and complex collections of digital data, the initiative promises to help accelerate the pace of discovery in science and engineering, strengthen our national security, and transform teaching and learning.”
- Office of Science and Technology Policy – Memo on data access February 22, 2013
 - Increasing Access to the Results of Federally Funded Scientific Research
 - The Office of Science and Technology Policy (OSTP) hereby directs each Federal agency with over \$100 million in annual conduct of research and development expenditures to develop a plan to support increased public access to the results of research funded by the Federal Government. This includes any results published in peer-reviewed scholarly publications that are based on research that directly arises from Federal funds, as defined in relevant OMB circulars (e.g., A-21 and A-11).
 - Each agency shall submit a draft plan by August 2013.
- DOE Office of Science
 - Requests all facilities to document their data policy
 - Requires all new proposals to develop a data management plan

Foundational infrastructure for S&T

- From large-scale simulations to data capture, analytics and visualization these infrastructures are a critical component for our S&T pursuits
- How are system infrastructures evolving as science and engineering pursuits increasingly rely upon these diverse workloads?
- What are the challenges we must address at various levels of the system stack (Hardware, OS, System Software, Middleware, Algorithms) to support a broad range of workloads from multi-petaflop towards exaflop simulations to data capture, analysis, and visualization?
- What lessons can we learn from the broader Big Data space and what lessons can they learn from HPC?
- Finally, is there convergence or divergence in the future of HPC and Big Data?
 - If not universal convergence, are their points of convergence at the "high end" of HPC and Big Data?

Lineup for our session

The Fusion of Supercomputing and Data Analytics To Drive Scientific Discovery

Pete Ungaro
Cray

The Path to High Efficiency Computing

Bill Dally
Nvidia

New Dimensions in Configurable Computing at runtime simultaneously allows Big Data and fine Grain HPC

Al Gara
Intel

Enabling Integrated HPC, Analytics & Visualization to solve the Big Data Challenge

Jorge Tittinger
SGI

The Fusion of Supercomputing and Data Analytics To Drive Scientific Discovery



- Changing the rules of the game in HPC and Big Data
 - Supercomputing is about getting the highest performance and most realistic simulations possible on “basketball court-sized” tightly integrated systems for scalability and usability. Large Scale data analytics is about getting knowledge out of immense amounts of data on “warehouse-sized”, low-cost commodity systems.... **changes in system architecture and a more holistic approach can change the game for scientific discovery in the future.**
- **Pete Ungaro - President and CEO of Cray**
 - Leads an organization of over 800 employees worldwide delivering highly advanced systems & solutions to demanding science and engineering challenges
 - Named CEO of the year by Seattle Business Monthly, “40 under 40” by Corporate Leader Magazine, and member of the Department of Commerce’s Manufacturing Council
 - Prior includes SVP sales and marketing at Cray and VP of Sales for Worldwide Deep Computing at IBM

The Path to High Efficiency Computing



- Performance & programmability at the end of Dennard scaling
 - HPC and data analytics share challenges of power, programmability, and scalability.... ***The large-scale parallelism and deep storage hierarchy of future machines poses programming challenges.*** This talk will discuss these challenges in more detail and introduce some of the technologies being developed to address them.
- **Bill Dally - Chief Scientist & SVP - Research at NVIDIA Corporation, Stanford Professor**
 - Bill and his group have developed system architecture, network architecture, signaling, routing, and synchronization technology that can be found in most large parallel computers today.
 - Member of the National Academy of Engineering, a Fellow of the IEEE, a Fellow of the ACM, and a Fellow of the American Academy of Arts and Sciences.
 - Prior includes Bell Labs, Caltech, MIT, and Chair of CS at Stanford

New Dimensions in Configurable Computing at runtime simultaneously allows Big Data and fine Grain HPC



- Technological advancement provides a path towards convergence
 - New technologies and concepts allow for future machine architectures to be configurable in new ways.... ***Disaggregation of configurable compute plus memory blocks are fundamental to this new direction.*** .. The hardware and software implications of this approach will be discussed.
- **Al Gara - Exascale Chief Architect at Intel**
 - Currently driving the system architecture direction for Intel's exascale efforts
 - Recipient of 2 Gordon Bell awards, the Seymour Cray award and the Blue Gene effort received the National Medal of Technology
 - Prior to joining Intel he was an IBM fellow and Chief Architect for the Blue Gene platform for three generations from 2000 till leaving IBM in 2011.

Enabling Integrated HPC, Analytics & Visualization to solve the Big Data Challenge



- HPC, analytics, and cloud converge
 - Leveraging four key strengths to drive this convergence:
 - experience in designing data intensive and low latency systems for volume Cloud providers, velocity traders, the US government and for fraud mitigation;
 - co-design partnerships with the DOE, DOD and industry;
 - track record of contributing to and leveraging open source software;
 - strong heritage in computer graphics and visualization.
- **Jorge Titingler - President and CEO of SGI**
 - Leads an organization of over 1,500 employees worldwide focused on helping customers solve their most demanding scientific and technical computing challenges
 - Prior includes President and CEO of Verigy Ltd, executive positions at FormFactor Inc., KLA-Tencor Corp., and Applied Materials Inc.

Galen Shipman
gshipman@ornl.gov

