



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

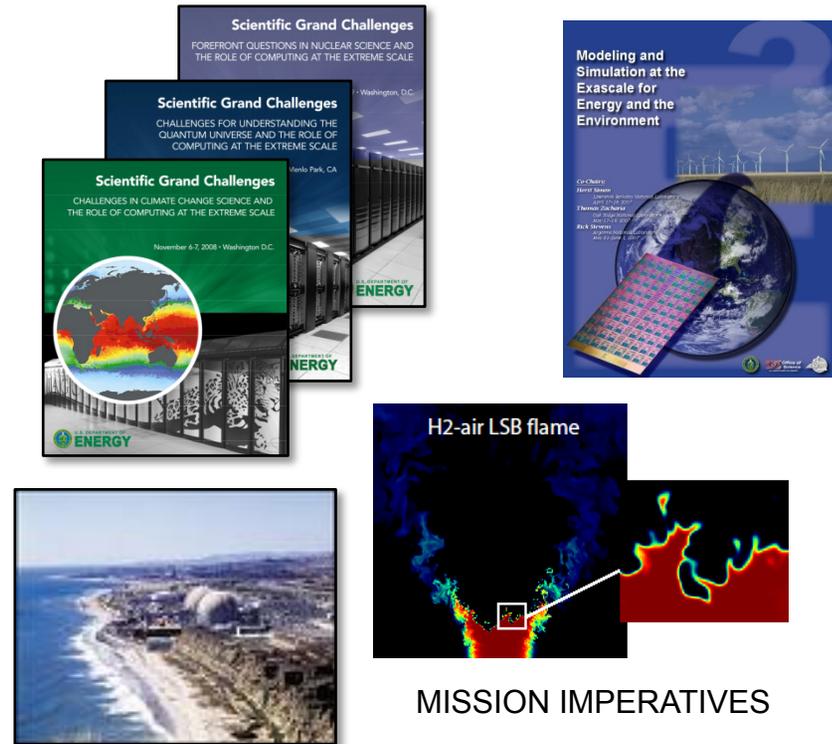
---

# Department of Energy's Exascale Efforts

Barbara Helland  
Advanced Scientific Computing Research  
Office of Science

# Exascale Applications and Technology

- **Town Hall Meetings April-June 2007**
- **Scientific Grand Challenges Workshops November 2008 – October 2009**
  - Climate Science (11/08),
  - High Energy Physics (12/08),
  - Nuclear Physics (1/09),
  - Fusion Energy (3/09),
  - Nuclear Energy (5/09) (with NE)
  - Biology (8/09),
  - Material Science and Chemistry (8/09),
  - National Security (10/09) (with NNSA)
- **Cross-cutting workshops**
  - Architecture and Technology (12/09)
  - Architecture, Applied Mathematics and Computer Science (2/10)
- **Meetings with industry (8/09, 11/09)**
- **External Panels**
  - ASCAC Exascale Charge (FACA)
  - **Trivelpiece Panel**



“The key finding of the Panel is that there are compelling needs for exascale computing capability to support the DOE’s missions in energy, national security, fundamental sciences, and the environment. The DOE has the necessary assets to initiate a program that would accelerate the development of such capability to meet its own needs and by so doing benefit other national interests. Failure to initiate an exascale program could lead to a loss of U. S. competitiveness in several critical technologies.”

Trivelpiece Panel Report, January, 2010



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

# DOE Progress toward Exascale

---

- **Proposals processed in Exascale related topic areas:**
  - **Applied Math: Uncertainty Quantification** (90 proposals requesting ~ \$45M/year; 6 funded at \$3M/yr)
  - **Computer Science: Advanced Architectures**(28 proposals requesting ~ \$28M/year, 6 funded at \$5M/yr)
  - **Computer Science: X-Stack** ( 55 proposals requesting ~\$40M/year, 11 funded at \$8.5M/yr)
  - **Computational Partnerships: Co-Design** (21 Proposals requesting ~ \$160M/year)
- **Exascale Coordination meetings with other Federal Departments and Agencies**
- **Formalizing Partnership with National Nuclear Security Administration (NNSA) within DOE**



# Uncertainty Quantification (UQ)

## Topics of Interest

---

- **Mathematical, statistical and hybrid approaches for quantifying and describing the effects and interactions of uncertainty and errors, potentially from multiple sources and with multiple representations**
- **Mathematical and computational frameworks for integrating statistical and deterministic analysis**
- **Mathematical theory and/or implementation of algorithms that demonstrably circumvent the “curse of dimensionality” in UQ analysis for complex system simulations**
- **Mathematical theory and/or algorithms for reduced-order modeling, inference, and inverse problems**
- **Scalable algorithms for numerical solutions of stochastic differential equations**
- **Tractable UQ treatment (intrusive or non-intrusive) for high-concurrency architectures**
- **Memory-access-efficient algorithms that match current and emerging computer architectures and allow for efficient and tractable sampling-based approaches.**

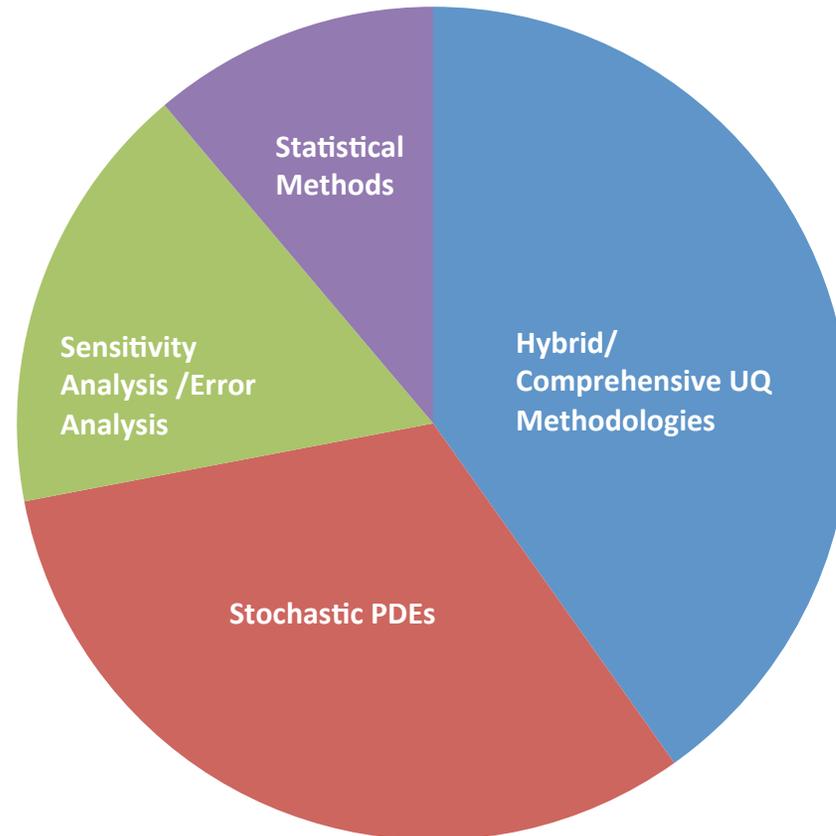
Received: 90 proposals requesting ~\$45M/yr

Funded: 6 projects funded at a total of \$3M/yr



# Uncertainty Quantification Portfolio

---



60% National Laboratory Funding  
40% University Funding



# UQ Portfolio

---

- **Modeling and Simulation of High-Dimensional Stochastic Multiscale PDE Systems at the Exascale, *Guang Lin (PNNL), Nicholas Zabaras (Cornell), and Ioannis Kevrekidis, (Princeton)***
  - Develop a rigorous mathematical framework and scalable uncertainty quantification algorithms on exascale computers to efficiently construct realistic flow-dimensional input models and surrogate low-complexity systems for the analysis, design and control of physical systems represented by multiscale stochastic PDEs (SPDEs). Demonstrate on CO<sub>2</sub> sequestration and nuclear reactor fuel cladding modeling.
- **Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales, *Dongbin Xiu (Purdue), Richard Archibald, Ralf Deiterding, and Cory Hauck (ORNL)***
  - Tackling dimension problem with non-intrusive UQ using three approaches -- “edge detection”, adaptive mesh refinement, adaptive stochastic collocation– with climate and magnetically confined fusion as targeted applications.
- **A High-Performance Embedded Hybrid Methodology for Uncertainty Quantification with Applications, *Charles Tong (LLNL), Barry Lee (PNNL), Gianluca Iaccarino (Stanford)***
  - Mathematical framework that analyzes flow of uncertainties, characterizes the coupling between modules, and includes strategies for model reduction, uncertainty propagation, data fusion and parallel implementation in hybrid UQ; application to flow/thermal transport and neutronics and thermal hydraulics



# UQ Portfolio (cont)

---

- **Enabling Predictive Simulation and UQ of Complex Multiphysics PDE Systems by the Development of Goal-Oriented Variational Sensitivity Analysis and a-Posteriori Error Estimation Methods, *John Shadid (SNL), Don Estep (CSU), Victor Ginting (UWyoming)***
  - Develop variational adjoint-based methods for sensitivity analysis and error estimation for specific quantities of interest (QoI) in complex multiphysics PDE systems with application to magnetohydrodynamics. Represents an important extension to DOE complex problems and DOE applications communities.
- **Bayesian Uncertainty Quantification in Predictions of Flows in Highly Heterogeneous Media and its Application to CO2 Sequestration, *Yalchin Efendiev (Texas A&M), Panayot Vassilevski (LLNL)***
  - Develop framework, modeling techniques and algorithms for characterizing uncertainty in highly heterogeneous porous media. Statistical characterization involves the use of level-set methods, Markov chains, sampling based methods along with efficient multiscale methods and preconditioners for flow simulations and includes a validation component.
- **Large-Scale Uncertainty and Error Analysis for Time-Dependent Fluid/Structure interactions in Wind Turbine Applications, *Juan Alonso (Stanford) and Michael Eldred (SNL)***
  - Comprehensive methodology for uncertainty quantification by combining stochastic expansions—both nonintrusive polynomial chaos and stochastic collocation—, the adjoint approach, and fusion with experimental data to account for aleatory and epistemic uncertainties from random variable, random field, and model form sources.



# Advanced Architecture FOA

## Topics of Interest

---

- Approaches for reducing and/or managing power requirements for high performance computing systems, including the memory and storage hierarchy
- Approaches for reducing and/or managing heat in high performance computing systems
- Methods for improving system resilience and managing the component failure rate, including approaches for shared information and responsibility among the OS, runtime system, and applications
- Co-design of systems that support advanced computational science at the extreme scale
- Scalable I/O systems, which may include alternatives to file systems
- Approaches to information hiding that reduce the need for users to be aware of system complexity, including heterogeneous cores, the memory hierarchy, etc.

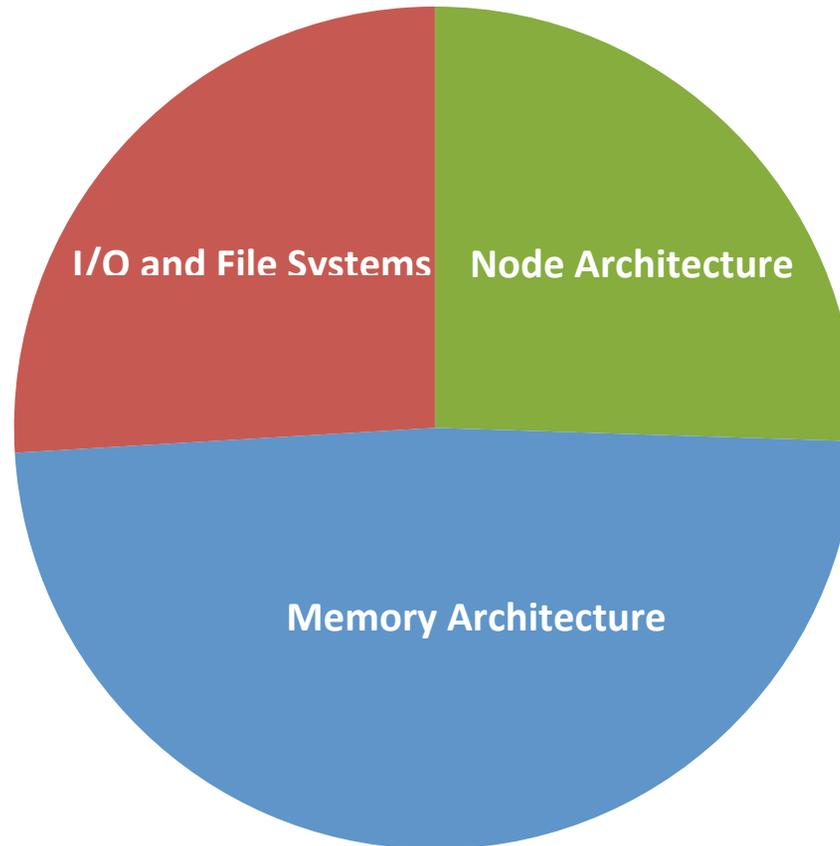
Received: 28 proposals requesting ~\$28M/yr

Funded: 6 projects funded at a total of \$5M/yr



# Advanced Architecture Portfolio

---



60% DOE Laboratory  
40% University and Industry



# Advanced Architecture Portfolio

---

- **Blackcomb: Hardware-Software Co-design for Non-Volatile Memory in Exascale Systems, Jeffrey S. Vetter (ORNL), Robert Schreiber (Hewlett Packard Labs), Trevor Mudge (University of Michigan), Yuan Xie (Penn State University)**
  - Identify and evaluate the most promising hardware and software technologies centered on non-volatile memory that can address the resilience, energy and performance requirements in future DOE exascale systems
- **Enabling Co-Design of Multi-Layer Exascale Storage Architectures, Robert B. Ross (ANL), Samuel Lang (RPI)**
  - Exploration and co-design of exascale storage system by providing a detailed accurate and highly parallel simulation toolkit for exascale storage capturing the complexity, scale and multi-layer nature of exascale hardware and software
- **NoLoSS: Enabling Exascale Science through Node Local Storage Systems, Robert B. Ross (ANL), Shane Canon (LBNL), Bronis R. de Supinski (LLNL)**
  - Conduct a detailed assessment of potential roles and benefits of node local storage in exascale computational science by exploring existing hardware options and assessing software mechanisms that best exploit them



# Advanced Architecture Portfolio (cont)

---

- **CoDEx: A Hardware/Software Co-Design Environment for the Exascale Era, *John Shalf (LBNL), Curtis Janssen (SNL, California), Dan Quinlan (LLNL), Sudhakar Yalamanchili (GATECH)***
  - Development of hardware/software co-design environment that combines a highly configurable, cycle accurate simulation of node architectures with novel automatic extraction and exascale extrapolation of memory and interconnect traces and scalable simulation of massive interconnection networks
- **Data Movement Dominates: Advanced Memory Technology to Address the Real Exascale Power Problem, *Arun Rodrigues (SNL), Paul Hargrove (LBNL), David Resnick (Micron), Keren Bergman (Columbia), Bruce Jacob (UMD)***
  - Examines architectural possibilities enabled by 3D integrated memory systems, consisting of both DRAM and logic on the same “stack”, and inexpensive, high bandwidth chip-to-chip communication enabled by optical interconnects
- **Thrifty: An Exascale Architecture for Energy-Proportional Computing, *Josep Torrellas (UIUC), Daniel Quinlan (LLNL), Allan Snaveley (UCSD), Wilfred Pinfold (Intel)***
  - Development and evaluation of detailed architecture simulator and implementation and evaluation of a power-aware optimizing compiler and development of libraries for application tuning and characterization



# X- Stack

## Topics of Interest

---

- **System software, including operating systems, runtime systems, adaptable operating and runtime systems, I/O systems, system management/administration, resource management and means of exposing resources, and external environments**
- **Fault management, both by the operating & runtime systems & by applications**
- **Development environments, including programming models, frameworks, compiler, & debugging tools**
- **Application frameworks**
- **Cross-cutting dimensions, including resilience, power management, performance optimization, & programmability**
- **Design and/or development of high-performance scientific workflow systems that incorporate data management & analysis capabilities**

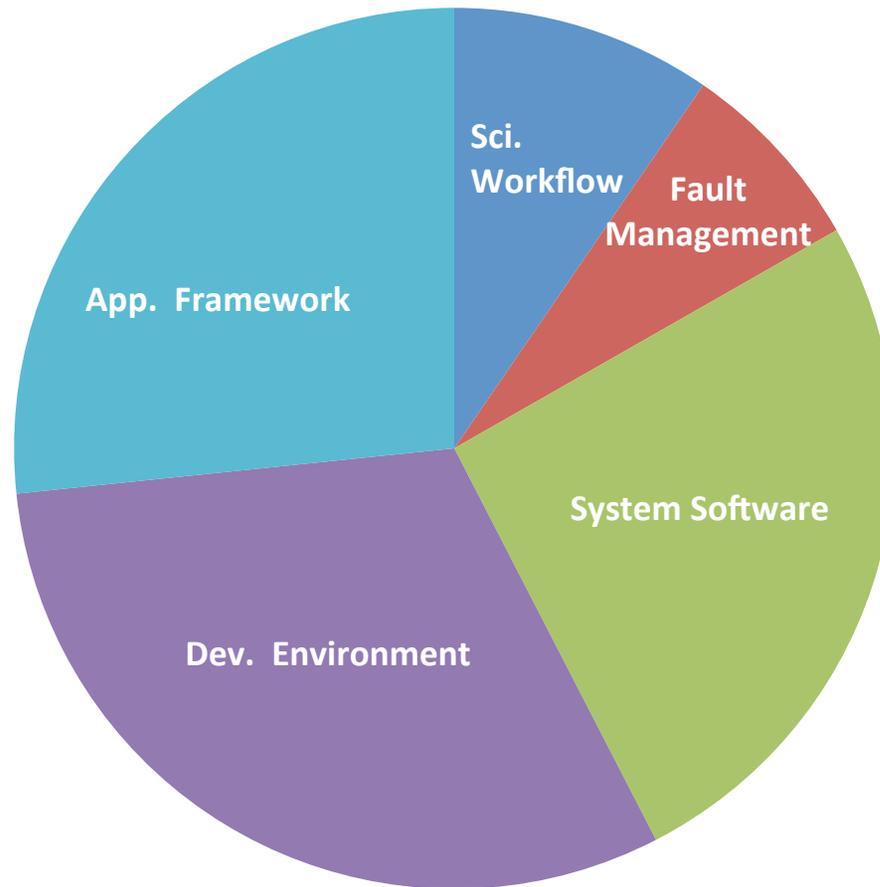
Received: 55 proposals requesting ~\$40M/yr

Funded: 11 projects funded at a total of \$8.5M/yr



# X-Stack Portfolio

---



50% National Laboratory  
50% University and Industry



# X-stack Portfolio

---

- **A Fault-oblivious Extreme-scale Execution Environment, *Ronald G Minnich (SNL CA), Sriram Krishnamoorthy (PNNL), Maya Gokhale (LLNL), P. Sadayappan (OSU), Eric Van Hensbergen (IBM), Jim McKie (Alcatel-Lucent Bell Labs), Jonathan Appavoo (BU)***
  - Develop a hierarchical data model that provides simple, familiar visibility and access to data structures through the file system hierarchy, while providing fault tolerance through selective redundancy. The hierarchical task model features work queues whose form and organization are represented as file system objects. By exposing the relationships between data and work to the runtime system, information is available to optimize execution time and provide fault tolerance.
- **Auto-tuning for Performance and Productivity on Extreme-scale Computations, *Samuel W. Williams (LBNL), John Gilbert (UCSB), Stephane Ethier (PPPL)***
  - Build a series of broadly-applicable, auto-tuned efficiency layer components to address auto-tuning's two principal limitations: an interface ill-suited to the forthcoming hybrid SPMD programming model and its scope limited to fixed function numerical routines.
- **Software Synthesis for High Productivity ExaScale Computing, *Armando Solar-Lezama (MIT), Rastislav Bodik & James Demmel(UC Berkeley)***
  - Develop a programming model that integrates program validation, synthesis and auto-tuning built upon the SKETCH domain-independent synthesizer . The goal is to simplify the challenges of programming heterogeneous architectures, such as implementing application-specific data distribution and communication.



# X-stack Portfolio (cont)

---

- **COMPOSE-HPC: Software Composition for Extreme Scale Computational Science and Engineering, *David E. Bernholdt (ORNL), Matt Sottile (Galois, Inc.), Tom Epperly (LLNL), Manojkumar Krishnan (PNNL), Rob Armstrong (SNL)***
  - Will address a range of significant technical challenges related software composition facing applications developers on exaflop machines by developing practical, high-performance tools and supporting infrastructure. Will build capabilities to simplify the composition of threaded software and creating concurrent multi-tasking applications on top of infrastructure.
- **Damsel: A Data Model Storage Library for Exascale Science, *Alok Choudhary (Northwestern University), Robert Ross (ANL), Nagiza Samatova (NCSU), Quincey Koziol (The HDF Group)***
  - Enable exascale computational science applications to interact conveniently and efficiently with storage through abstractions that match their data models by 1) identifying major data model motifs and developing benchmarks; and developing and hardening a data model storage library, that provides efficient storage data layouts, optimizes I/O operations and is fault tolerant.
- **Vancouver: Designing a Next-Generation Software Infrastructure for Productive Heterogeneous Exascale Computing, *Jeffrey S. Vetter (ORNL), Wen-mei Hwu (UIUC), Allen D. Malony (University of Oregon), Rich Vuduc (GATECH)***
  - Will address performance and productivity challenges of scalable heterogeneous computer systems through development of libraries and infrastructure for benchmarking, auto-tuning, data movement and task scheduling; tools for code analysis, inspection and transformation with integrated performance analysis and prediction; and new programming models to support global address spaces and multiple levels of parallelism



# X-stack Portfolio (cont)

---

- **Enabling Exascale Hardware and Software Design through Scalable System Virtualization**, *Patrick G. Bridges (UNM), Peter Dinda (Northwestern University), Kevin Pedretti (SNL), Stephen Scott (ORNL)*
  - Will enhance the Palacios virtual machine monitor to enable and support research across the breadth of the exascale system hardware/software stack. Proposed enhancements include support for: large-scale system software testing through the normal job allocation process; evaluating new HEC architecture innovations; and vertical profiling features; and configurability for a wide range of customizations.
- **ZettaBricks: A Language, Compiler and Runtime Environment for Anyscale Computing**, *Saman Amarasinghe, Alan Edelman, Una-May O'Reilly, Martin Rinard, Chris Hill (MIT)*
  - Plan to build ZettaBricks, a language, compiler, and runtime environment that will provide a *write once run anywhere programming model* and reduce burden of programming exscale systems by eliminating need for architecture-specific performance tuning and simplifying error handling.
- **Compiled MPI: Cost-Effective Exascale Application Development**, *Torsten Hoefler (UIUC), Daniel J. Quinlan (LLNL), Andrew Lumsdaine (Indiana)*
  - Development of CoMPI, a novel compiler-driven approach to enable existing MPI applications to scale to exascale systems with incremental modifications over the application's lifetime. CoMPI is based on ROSE and will include a new set of source code annotations; compile transformation framework that leverages annotations; MPI runtime implementation with a rich set of functionality extensions and a novel compiler analysis
- **ExM: System support for extreme-scale, many-task applications** *Michael Wilde (ANL), Daniel S. Katz (Chicago), Matei Ripeanu (University of British Columbia)*
  - Design, develop and evaluation two system software components– *ExM data store* to allow communication between concurrent and asynchronous application tasks and persistent storage and *ExM task manager* to allow for rapid, data-aware and efficient dispatch of many tasks to exascale computing systems with fault tolerant execution.



# Planning Grants

---

- **Because of Development nature of projects, grants for further planning provided for**
  - From Xstack: An Open Integrated Software Stack for Extreme Scale Computing , *Pete Beckman (ANL), Kathy Yelick (LBNL), Al Geist (ORNL), Mike Heroux (SNL), Jack Dongarra (UTK)*
  - 5 proposed Co-Design Centers
- **Plans must include**
  - refined goals and objectives with milestones and deliverables
  - management plan
  - risk plan
- **Submitted plans reviewed in Spring, 2011**

