Exascale Update

SC14 - ASCAC

Bill Harrod

Office of Science

November 21, 2014

Mission: Extreme Scale Science

Next Generation of Scientific Innovation

- **DOE's mission is to push the frontiers of science and technology to:**
  - Enable scientific discovery
  - Provide state-of-the-art scientific tools
  - Plan, implement, and operate user facilities
- **The next generation of advancements will require Extreme Scale Computing**
  - 1,000X capabilities of today's Petaflop computers with a similar size and power footprint
- **Extreme Scale Computing, however, cannot be achieved by a "business-as-usual" evolutionary approach**

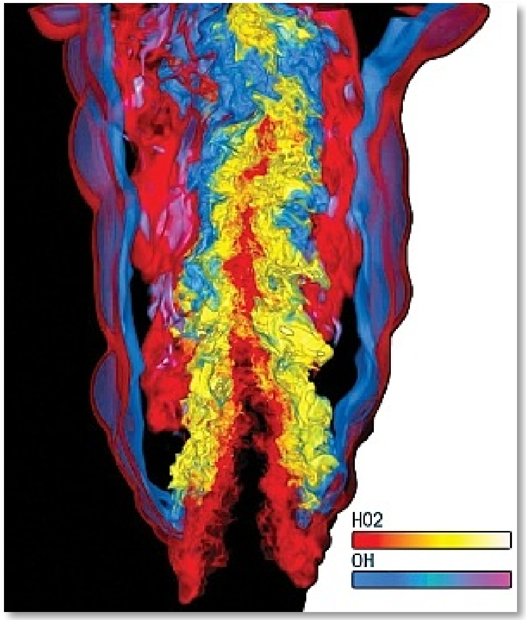

- **Extreme Scale Computing will require major novel advances in computing technology – Exascale Computing**

Exascale Computing Will Underpin Future Scientific Innovations

Exascale Applications Respond to DOE Missions in Discovery, Design, and National Prosperity

Scientific Discovery
- Mesoscale materials and chemical sciences
- Improved climate models with reduced uncertainty
- Plasma physics for fusion energy systems

Engineering Design
- Nuclear power reactors
- Advanced energy technologies
- Resilient power grid

National Prosperity
- Energy Technologies
- Advanced manufacturing
- Health Care
- Transportation

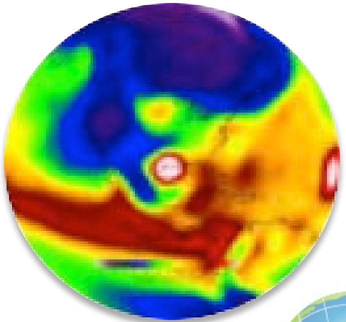

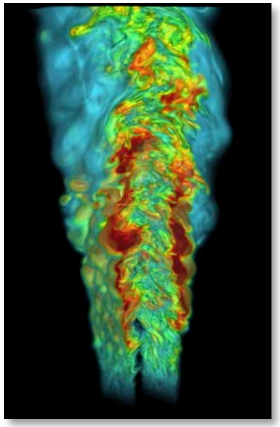

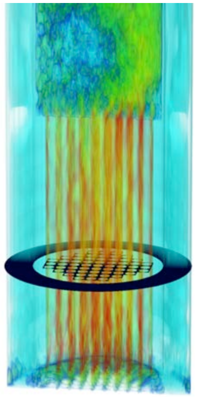

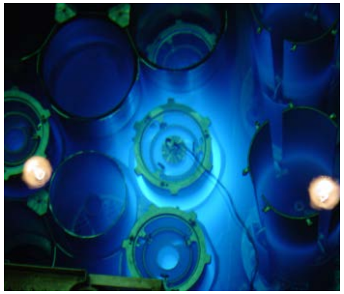

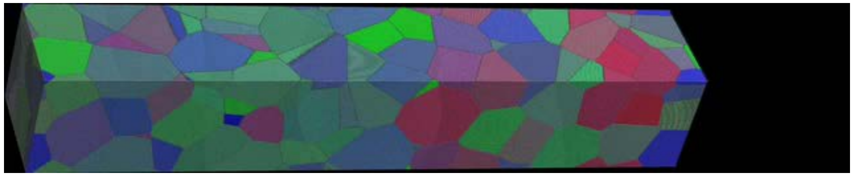

Exascale Challenges and Issues

- **Four primary challenges must be overcome**
  - Parallelism / concurrency
  - Reliability / resiliency
  - Energy efficiency
  - Memory / Storage
- **Productivity issues**
  - Managing system complexity
  - Portability
  - Generality
- **System design issues**
  - Scalability
  - Efficiency
  - Time to solution
  - Dependability (security and reliability) must be integrated at all levels of the design

Exascale Strategy

- Exploit co-design process, driven by the full application workflow
- Develop exascale software stacks
- Partner with and fund vendors to transition research to product space
- Collaborate with other government agencies and other countries, as appropriate

Exascale Project Schedule

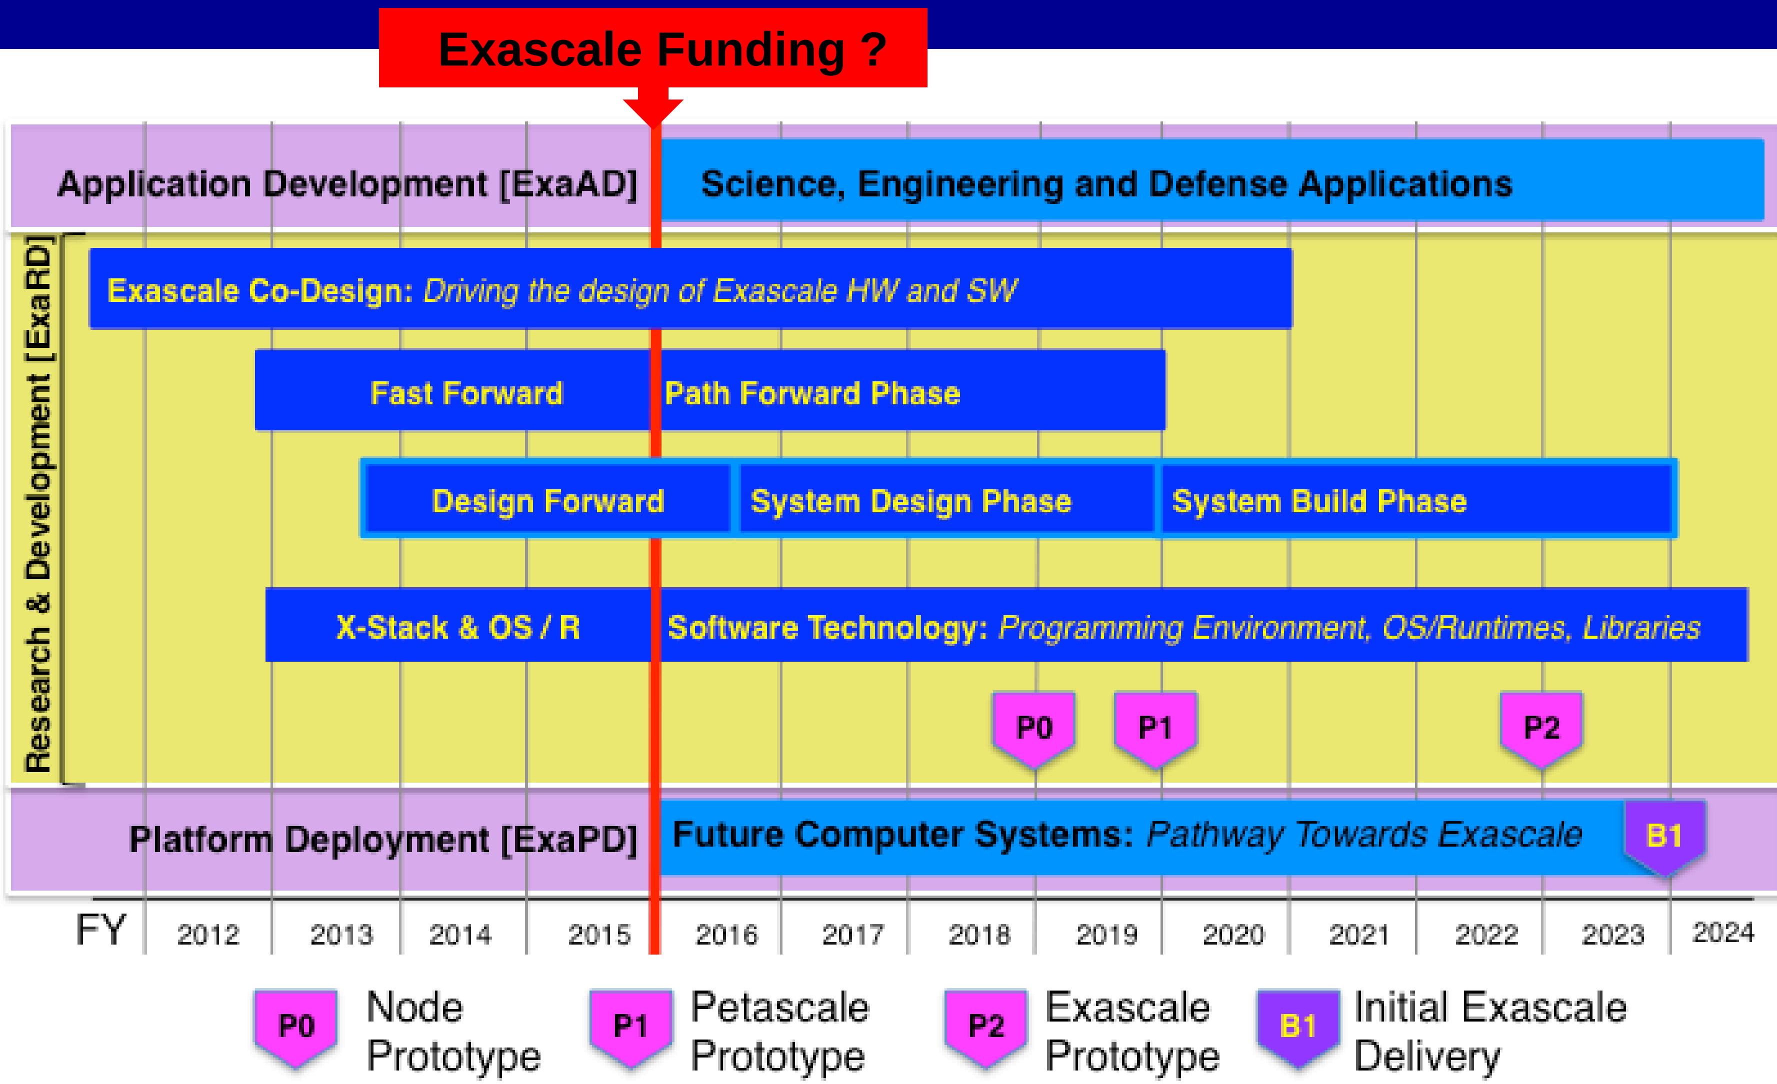

Fast Forward and Design Forward Programs

Fast Forward Program

- Jointly funded by SC & NNSA
- Phase 1: Two year contracts, started July 1, 2012, Phase 2: Two year contracts, starting Fall 2014

Project Goals & Objectives

- Initiate partnerships with multiple companies to accelerate the R&D of critical node technologies and designs needed for extreme-scale computing.
- Fund technologies targeted for productization in the 5–10 year timeframe.

Design Forward Program

- Jointly funded by SC & NNSA
- Phase 1: Two year contracts, started Fall 2013, Phase 2: Two year contracts. Starting Fall 2014

Project Goals & Objectives

- Initiate partnerships with multiple companies to accelerate the R&D of interconnect architectures and conceptual designs for future extreme-scale computers.
- Fund technologies targeted for productization in the 5–10 year timeframe.

FastForward 2 Projects

- **AMD** will conduct research for an integrated exascale node architecture. Particular areas of emphasis include near-threshold-voltage logic and other low-power computing technologies.   AMD will investigate a new standardized memory interface
- **Cray Inc.** will explore alternative processor design points, including ARM microprocessor designs.
- **Intel** will use this award to continue to advance research in energy efficient node and system architectures, including software targeted at developing extreme-scale systems.
- **NVIDIA** will build on its work in FastForward 1, with a strong focus on energy efficiency, programmability, and resilience.
- **IBM** will investigate next-generation standardized memory interface.

First steps on a Path Towards Exascale:

Three Exascale Co-Design Centers

Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)

- Director: Timothy Germann (LANL)
- http://www.exmatex.org

Center for Exascale Simulation of Advanced Reactors (CESAR)

- Director: Andrew Siegel (ANL)
- https://cesar.mcs.anl.gov

Center for Exascale Simulation of Combustion in Turbulance (ExaCT)

- Director: Jacqueline Chen (SNL)
- http://exactcodesign.org

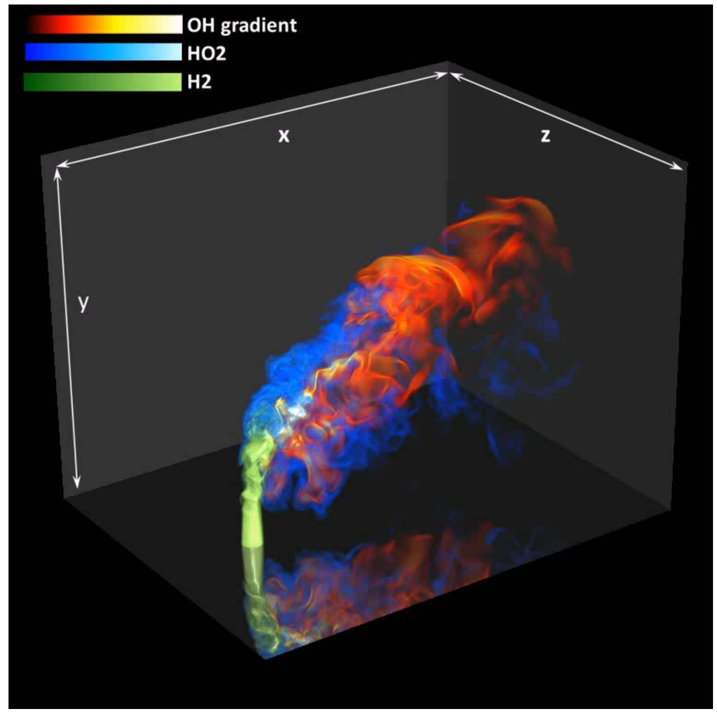

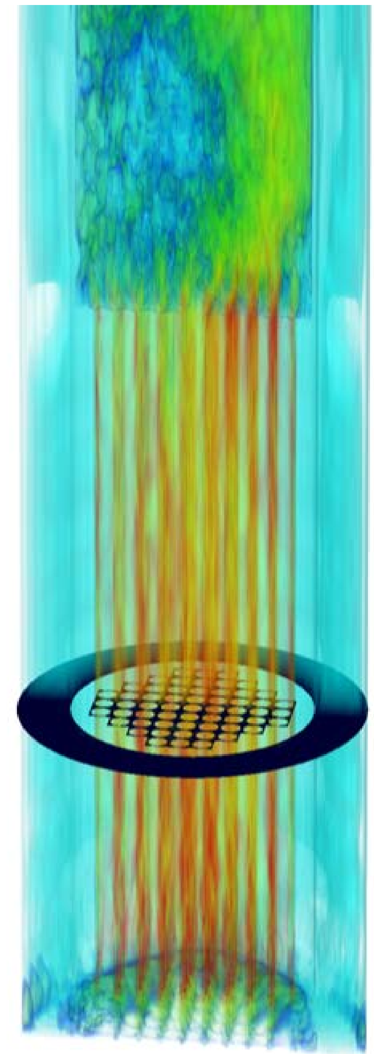

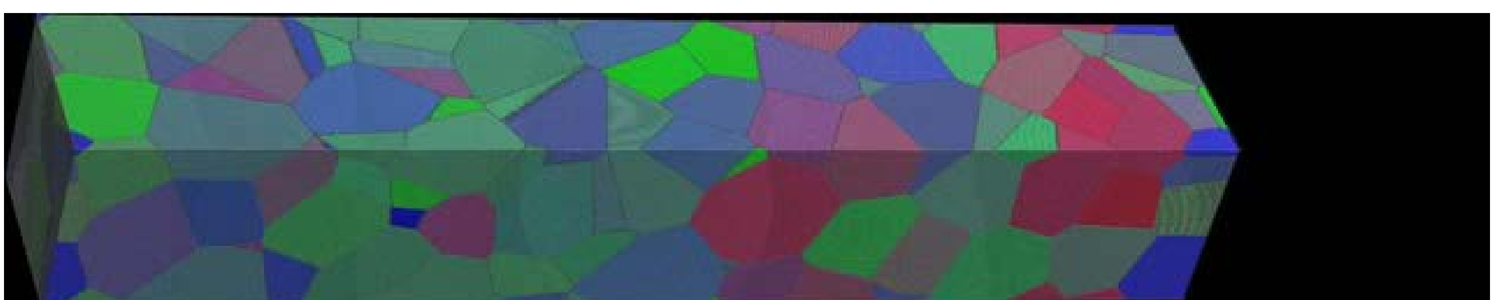

ASCR Exascale Co-Design

Status

Notable Successes

- Informed & influenced vendors node designs by providing "proxy apps", codes derived from actual applications, that go beyond the traditional math kernels
- Re-examined current algorithms and made changes so that they take advantage of anticipated exascale architectures
- Explored programming models and environments that are suitable for the data layout of the applications

Lessons Learned

- Interdisciplinary teams discover key issues that could not have been discovered within the boundary of any one disciplinary - and find solutions that are "co-designed"
- Bandwidth for interactions: personnel in Co-Design Centers in high-demand, as the Centers provide the only "application use case" for testing research results from computer science, applied math, and hardware architectures
- Horizontal integration: need to create opportunities for the Centers to leverage each other

Exascale Programming Models

X-Stack Goals

- **New Programming Models:** approaches to managing parallelism and data movement through innovations in interfaces
- **Dynamic Runtime Systems:** adapt to changing application goals and system conditions
- **Locality-aware and Energy-efficient Strategies:** manage locality and minimize energy consumption
- **Interoperability:** facilitate interoperability across different HPC languages and interfaces, as well as cross-cutting execution models

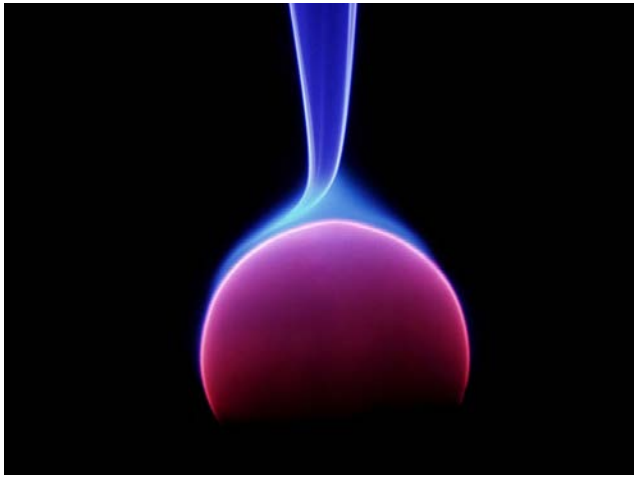

Exascale Programming Models
Observations & Lessons Learned

- **Computations and communications will be irregular due to heterogeneous processors characteristics, deep memory hierarchies, and much higher rate of faults than current platforms.**
- **Exascale system will have deep memory hierarchy: generating optimized code will be very challenging.**
- **Using DSLs, application code size can be reduced by 100x, and the code generated and optimized has about 3x faster execution time.**
- **Unprecedented parallelism, asynchrony, and irregular computations and communications are well matched to task/thread-oriented dataflow programming models.**

Status

Meetings with

- Secretary of Energy
- SEAB
- OMB
- OSTP
- Congressional Meetings
- Other agencies

Project Design Planning

- Develop preliminary version of project design document, involving lab researchers and DOE/NNSA program managers.
- Held "Red Team" review of document
- Held external review of document (federal employees from other agencies)
- Updated project design document

Initiated New Programs

- FastForward 2 and DesignForward 2, Resilience, Analytical Modeling for Extreme-Scale Computing Environments, Scientific Data Management, Analysis and Visualization at Extreme Scale

Preliminary Conceptual Design for the Exascale Computing Initiative Document

- **This document presents a preliminary strategy and contains limited descriptions of project management details.**
- **A preliminary list of proposed projects is provided**

Table of Contents

U.S. DEPARTMENT OF ENERGY | Office of Science

Technical Approach
Project Components

1. ***Exascale Research, Development and Deployment (ExaRD):***
   - ***Exascale Co-Design Centers and Beyond Exascale****,* exploratory research to co-design hardware and software architecture
   - ***Software Technology Research and Development****,* aimed at specific hardware and software technologies;
   - ***Vendor Research and Development***, aimed at developing exascale node and system architectures;
2. ***Exascale Application Development (ExaAD):***
   - Readiness to Use Capable Exascale Systems, initiating the development of a suite of exascale applications software packages that will be operational in 2023 to ensure maximal scientific and engineering impact of the exascale systems
3. ***Exascale Platform Deployment (ExaPD)****:*
   - Coordinated Acquisition Strategy for exascale platforms, including long-lead site preparations and system platforms

Risk	Impact (H/M/L)	Likelihood (H/M/L)	Mitigation
Insufficient funding within either NNSA or SC	H	M	Prioritize elements within the ECI and, if necessary, de-scope ECI and/or spread the acquisition costs over a longer time period
Full system is unreliable	H	M	Invest in robust, multilayered approaches to manage or resolve faults
Failure to achieve critical integration of DOE-developed software into vendor software	H	M	Early development of XBUS and integration with funded vendors software environments
Software environments do not satisfy DOE application needs	H	M	Determine the workload requirements of critical applications as early as possible
Key algorithms that do not scale may not have timely, suitable alternatives	H	L	Early investments in exploratory algorithms research
Exascale computer architecture departs significantly from expected designs, after significant R&D in software, tools, and algorithms were invested based on the expected designs	H	M	A clearly defined communication channel with vendors and effective communication so that the ECI participants are well-informed regarding evolving architectures and new directions

Software Technology Projects
Appendix 1

1 Application Foundations
1.1 Co-Design (CD)
1.2 Applied Mathematics (AM)
1.3 Data Analytics and Visualization (DV)
2 User Experiences (UE)
2.1 Productivity (PR)
2.2 Collaborative Environment (CE)
2.3 Resilience (RE)
2.4 Application Integrity
2.5 Cybersecurity (CS)
3 Software Stack (SS)
4 Performance Execution (PE)
5 Data Management (DM)
6 Hardware Architecture (HA)
7 System Engineering and Integration

Delivering on the Exascale Promise

Exascale computing systems are essential for the scientific fields that will transform the 21st Century global economy

- Energy
- Biotechnology
- Nanotechnology
- Materials science

Goal: Lead computational sciences and HPC to continue to develop and deploy HPC hardware and software systems through exascale platforms

DOE ASCAC subcommittee identifies top 10 technology advancements

- Critical to making productive, economically viable exascale systems
- Innovations beyond anticipated conventional practices and their incremental extensions

Technical Approaches

The Top Ten Exascale Challenges, with Technical Approaches

1. Energy efficiency
2. Interconnect technology
3. Memory Technology
4. Scalable System Software
5. Programming systems
6. Data management
7. Exascale Algorithms
8. Algorithms for discovery, design, and decision
9. Resilience and correctness
10. Scientific productivity