Mission: Extreme Scale Science

Next Generation of Scientific Innovation

- **DOE's mission is to push the frontiers of science and technology to:**
  - **En**able scientific discovery
  - Provide state-of-the-art scientific tools
  - Plan, implement, and operate user facilities
- **Causing a data explosion – a natural component of exascale computing**
  - Experimental facilities face exponentially burgeoning data caused by technology advances
- **Extreme Scale Computing, however, cannot be achieved by a "business-as-usual" evolutionary approach**

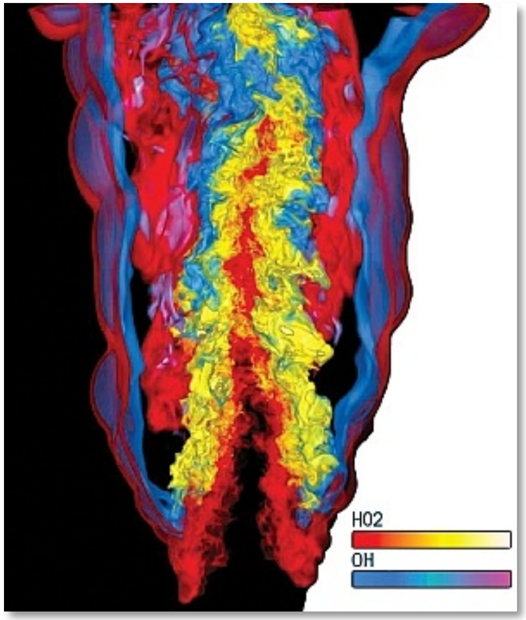

- **Extreme Scale Computing will require major novel advances in computing technology – Exascale Computing**

Exascale Computing Will Underpin Future Scientific Innovations

Mission: Extreme Scale Science

Data Explosion

Genomics

Data Volume increases to 10 PB in FY21

High Energy Physics (Large Hadron Collider)

15 PB of data/year

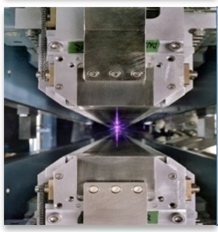

Light Sources

Approximately 300 TB/day

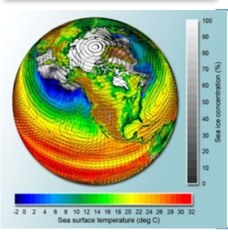

Climate

Data expected to be 100 EB

Driven by exponential technology advances

Data sources

- Scientific Instruments
- Scientific Computing Facilities
- Simulation Results
- Observational data

Big Data and Big Compute

- Analyzing Big Data requires processing (e.g., search, transform, analyze, …)
- Extreme scale computing will enable timely and more complex processing of increasingly large Big Data sets

1 EB = 10^{18} bytes of storage
1 PB = 10^{15} bytes of storage
1 TB = 10^{12} bytes of storage

"Very few large scale applications of practical importance are NOT data intensive." – Alok Choudhary, IESP, Kobe, Japan, April 2012

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
  - Time to solution
  - Efficiency
- **Extensive Exascale Studies**
  - US(DOE, DARPA, …), Europe, Japan, …

Key Performance Goals
for an exascale computer (ECI)

Parameter	
Performance	Sustained 1 – 10 ExaOPS
Power	20 MW
Cabinets	200 - 300
System Memory	128 PB – 256 PB
Reliability	Consistent with current platforms
Productivity	Better than or consistent with current platforms
Scalable benchmarks	*Target speedup over "current" systems TBD*
Throughput benchmarks	*Target speedup over "current" systems TBD*

ExaOPS = 10^{18} Operations / sec

Exascale Target System Characteristics

- 20 pJ per average operation
- Billion-way concurrency (current systems have Million-way)
- Ecosystem to support new application development and collaborative work, enable transparent portability, accommodate legacy applications
- High reliability and resilience through self-diagnostics and self-healing
- Programming environments (high-level languages, tools, …) to increase scientific productivity

FY2011:

- **Computer Science:** Execution Models
- **Computational Partnerships:** 3 Exascale Co-Design Centers Funded
- **Networking:** Terabit Networking for Extreme-Scale Science
- **Request for Information:** Critical and Platform Technologies

FY2012:

- **Computer Science:** Programming Environments (X-Stack), Performance Modeling (BMS), HWArch, e.g. CAL
- **Applied Math:** Resilient Extreme-Scale Solvers (RX-Solvers)
- **Networking:** Scientific Collaborations at Extreme-Scale
- **FastForward**: Critical / Cross Cutting technologies (joint with NNSA)

FY2013:

- **Exascale Strategy Plan to Congress**
- **Computer Science:** Operating System / Runtime (OS/R)
- **Applied Math:** Uncertainty Quantification
- **DesignForward**: Critical / Cross Cutting technologies (joint with NNSA)
- **FastForward:** Critical/Cross Cutting technologies (joint with NNSA)

FY2014:

- **CORAL:** The joint Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL)
- **Computer Science:** Scientific Data Management, Analysis and Visualization at Extreme Scale
- **Computer Science:** Software Productivity
- **Exploratory Research for Extreme-Scale Science (EXPRESS)**
- **Networking:** Analytical Modeling for Extreme-Scale Computing Environments
- **FastForward 2:** Critical/Cross Cutting technologies (joint with NNSA)
- **DesignForward 2:** Critical/Cross Cutting technologies (joint with NNSA)

FY2015:

- **Preliminary Conceptual Design for an Exascale Computing Initiative:** Developed jointly with NNSA
- **Computer Science:** Resilience for Extreme-Scale Supercomputing Systems

DOE Progress Towards Exascale

Schedule Baseline

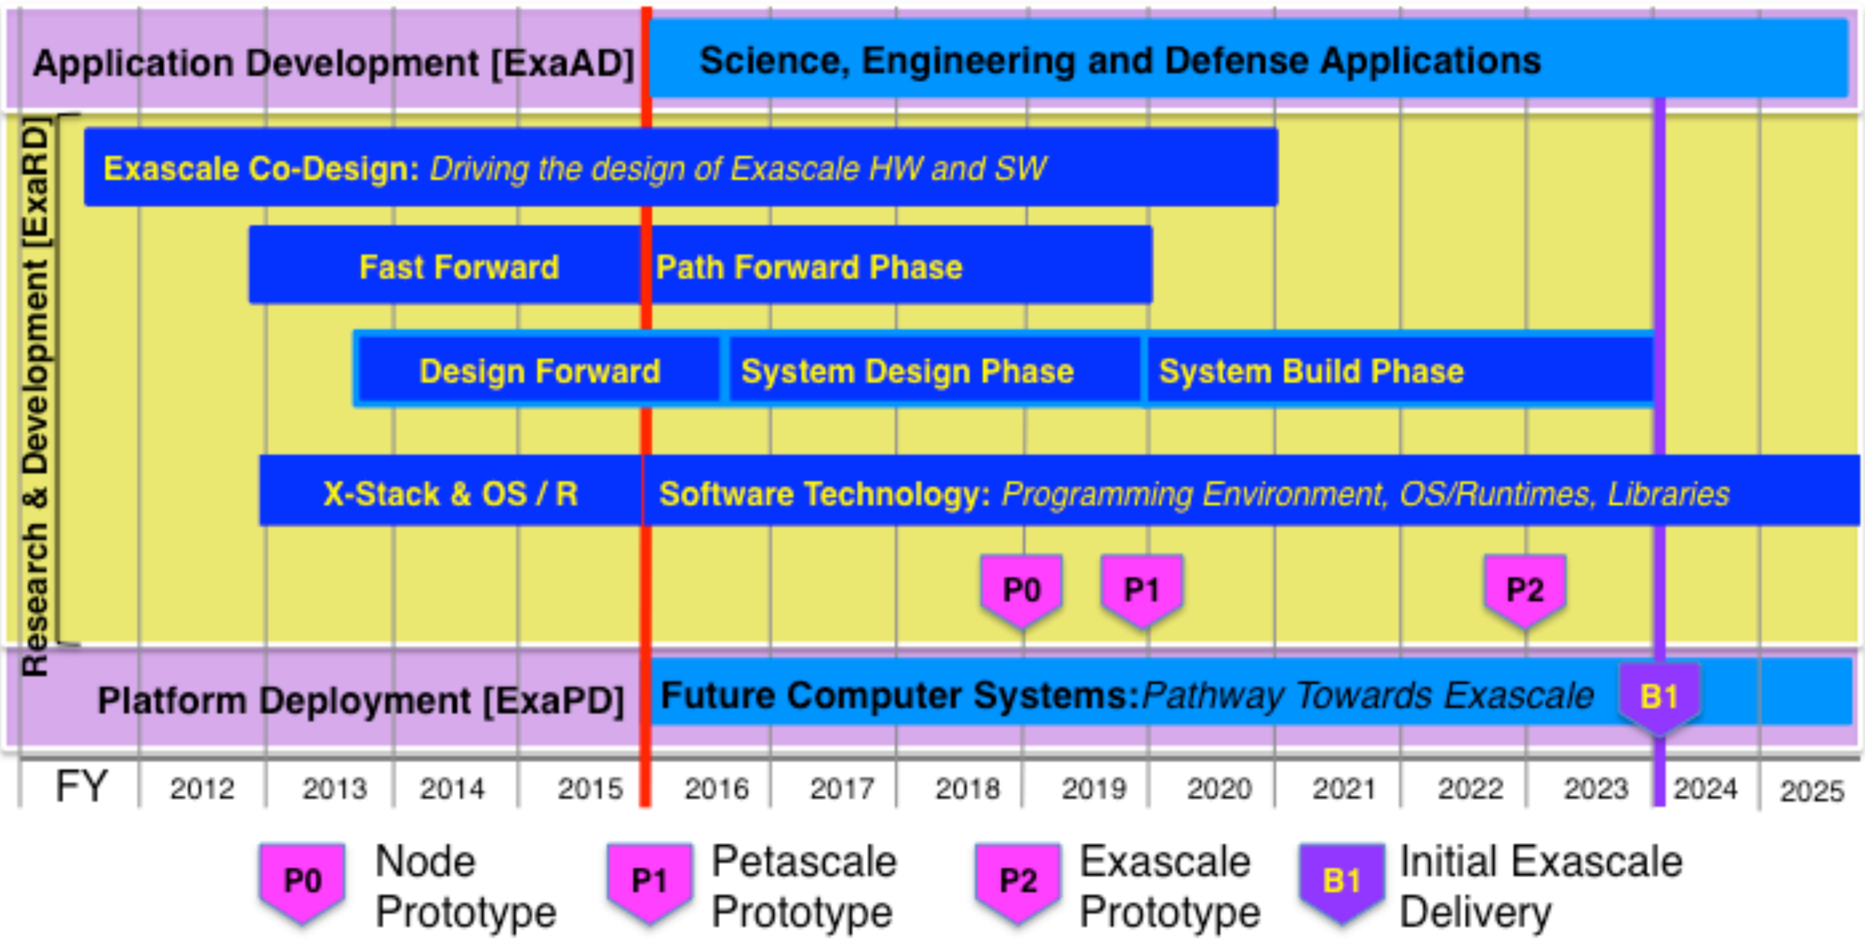

Current partnerships with vendors
Fast and Design Forward Programs

Fast Forward Program – *node technologies*

- Jointly funded by SC & NNSA
- **Phase 1:** Two year contracts, started July 1, 2012**, Phase 2:** Two year contracts, starting Fall 2014: IBM, Cray, AMD, NVIDIA, Intel ($64M / $100M)

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 – *system technologies*

- Jointly funded by SC & NNSA
- **Phase 1:** Two year contracts, started Fall 2013**, Phase 2:** Two year contracts. Starting Winter 2015: Cray, AMD, IBM, Intel ($23M / $10M)

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.

Summary

- **High-performance computing (HPC) and large-scale data analysis will advance national competitiveness** in a wide array of strategic sectors, including basic science, national security, energy technology, and economic prosperity.
- The **U.S. semiconductor and HPC industries have the ability to develop the necessary technologies** for an exascale computing capability early in the next decade.
- **An integrated approach to the development** of hardware, software, and applications is required for the development of exascale computers.
- ECI's goal is to **deploy, by FY-2023, two capable exascale computing systems.**

BACK-UP