# Update on US Exascale Activities

Paul Messina

IESP Meeting

October 6, 2011

# DOE/NNSA activities

- R&D projects funded
- More exascale workshops
  - Programming challenges, architectures
- Explicit exascale funding in ***proposed*** Federal budget for FY 2012
- Request for information from technology providers issued, responses being analyzed
- Three co-design centers launched

# Scope of DOE Exascale Activities

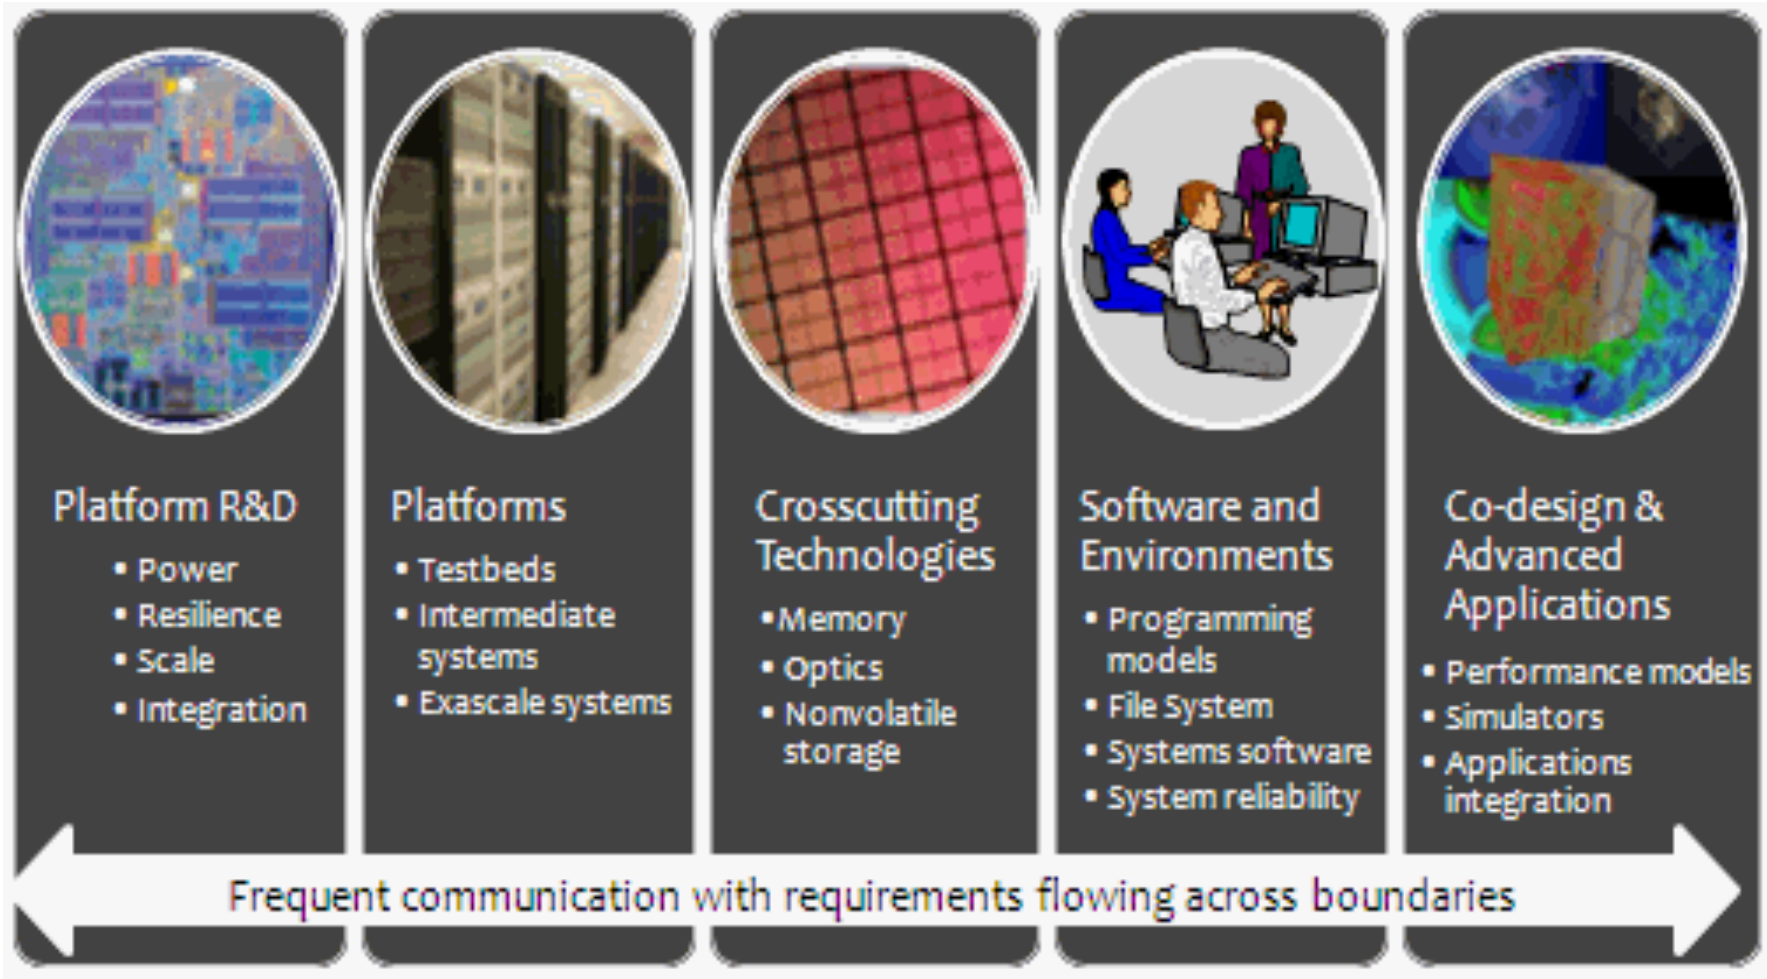

# DOE Exascale Activities

- 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)
  - Computer Science: Scientific Data Management and Analysis (37 proposals requesting ~$22M/year; 11 projects funded at $5M/year)
  - Computational Partnerships: Co-Design (21 Proposals requesting ~$160M/year, 3 funded)
- Exascale Coordination meetings with other Federal Departments and Agencies
- Formal Partnership within DOE between Office of Science and National Nuclear Security Administration (NNSA)

# DOE-ASCR Exascale Research Programs

- Advanced Architectures and Critical Technologies for Exascale
  - 6 projects focused on power management, memory management, and reducing the cost of data movement
- X-Stack Software Research
  - 10 projects focused on operating systems, fault tolerance, programming challenges, performance optimization, etc.
- Scientific Data Management and Analysis at Extreme Scale
  - 10 projects spanning file systems and I/O, data triage, feature detection and data analysis, and visualization

# Not funded (yet)

- The Exascale Software Center plan that was presented at the IESP meeting in April was not funded
  - However there is recognition that we need an integrated software plan and DOE would like it to be inclusive of international participants
- Four of the seven co-design centers that had planning grants and gave presentations at last IESP meeting

# US Senate Subcommittee Funding Position for Exascale

- The report contains a separate section on exascale computing, and specifies funding at the requested level of $126 M
  - $90 M in SC/ASCR
  - $36 M in NNSA/ASCI) for the initiative
- The actual funding level will be negotiated with the House
- Most of the proposed funding is **not** new, additional funding for DOE and NNSA in FY 2012

# US Congress: House of Representatives Request

- The Department [DOE] is directed to provide to the Committee, not later than February 10, 2012, a report including its current target date for developing an operational exascale platform, interim milestones towards reaching that target, estimated total ranges of Department investment likely needed to hit those targets, and a complete listing of exascale activities included in the budget request

# Exascale Request for Information (RFI)

- Seven DOE national laboratories have formed a consortium (referred to as E7) to manage the Request for Information (RFI) process to deliver exascale computing capability to the nation including development of prototypes and testbeds and acquisition of exascale technology systems (capable of up to sustained exaflops) in 2019–2020.
  - Argonne National Laboratory
  - Lawrence Berkeley National Laboratory
  - Lawrence Livermore National Laboratory
  - Los Alamos National Laboratory
  - Oak Ridge National Laboratory
  - Sandia National Laboratories, and
  - Pacific Northwest National Laboratory
- E7 is performing a market survey to evaluate obtaining products and services that are the subject of this RFI

# RFI purpose

- Provide the DOE Office of Science and the DOE National Nuclear Security Administration Office of Defense Programs with information for responding to a request from the House Energy and Water Development subcommittee of the House Appropriations Committee and for planning the DOE exascale program
- Among these activities will be the formation of partnerships between laboratories and industry to perform platform and crosscutting co-design and critical technologies research and development (R&D) targeted at delivering exascale computers by 2019–2020.

# RFI purpose (continued)

- Provide information that would guide the formulation of a Request for Proposals
  - Cost
  - Schedule
  - Likely technologies that could be requested

# Conceptual exascale program roadmap that can serve as a guide to the timing constraints on responses

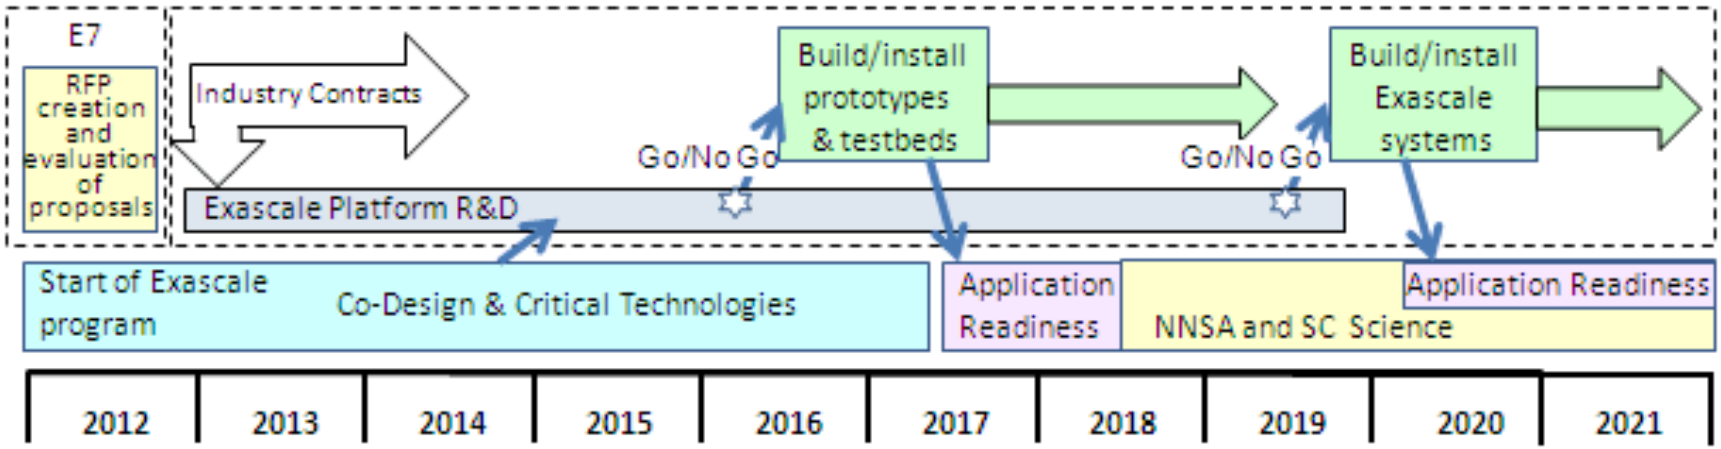

**Table 1. Exascale System Goals**

| Exascale System | Goal |
|---|---|
| Delivery Date | 2019 |
| Performance | 1000 PF LINPACK and 300 PF on to-be-specified applications |
| Power Consumption* | 20 MW |
| MTBAI** | 6 days |
| Memory including NVRAM | 128 PB |
| Node Memory Bandwidth | 4 TB/s |
| Node Interconnect Bandwidth | 400 GB/s |

*Power consumption includes only power to the compute system, not associated storage or cooling systems.

**The mean time to application failure requiring any user or administrator action must be greater than 24 hours, and the asymptotic target is improvement to 6 days over time. The system overhead to handle automatic fault recovery must not reduce application efficiency by more than half.

PF = petaflop/s, MW = megawatts, PB = petabytes, TB/s = terabytes per second, GB/s = gigabytes per second, NVRAM = non-volatile memory.

# A ***potential*** partnership model

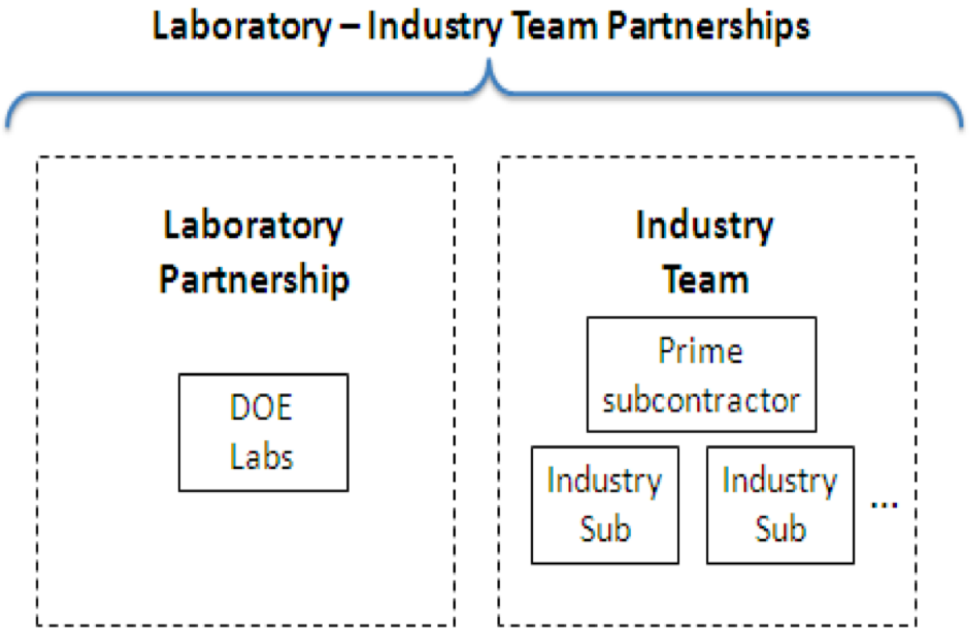


**Figure 3. Industry teams led by a prime will form partnerships with DOE laboratories.**

# 22 responses to the RFI

- Integrated Systems Providers
- Microprocessor Providers
- Networking and Interconnect
- Memory and Solid State Storage
- File system and Storage
- Software

# Preliminary findings from responses to RFI

- Responders all mentioned that the power target would be difficult to hit without investment.
- Many challenge areas (system software, resilience, s/w tools) require a coordinated/ integrated approach.
- Responses were very light on the data challenges

# Preliminary findings from responses to RFI

- As a group the responses depicted an optimism towards meeting the scope and schedule set forth in the RFI
- Responses had little discussion on risks and risk mitigation strategies
  - Are some companies underestimating the serious integration challenges?
- Open source was not embraced by many responders
- Did not see much clarity on the subject of cross-cuts
  - RFI stressed need for cross-cutting technologies in support of all exascale systems, co-design centers help identify
- E7 is still processing RFI responses and working on fleshing out a plan

# Three co-design centers established

- Advanced Nuclear Reactors
- Combustion
- Materials in Extreme Environments

# The Center for Exascale Simulation of Advanced Reactors (CESAR)

- The need for energy technologies that both avoid further contributions to global warming and serve as reliable (base) energy sources has led to renewed interest in nuclear power.
- But current codes – though highly tuned and calibrated for commercial light-water reactors – lack the physics fidelity to seamlessly carry over to new reactor classes with significantly different design characteristics.
  - The new generation of modeling tools need to have more predictive power, thus be more science-based
  - Such tools – at the requisite level of physical fidelity – will require exascale-level of computational power

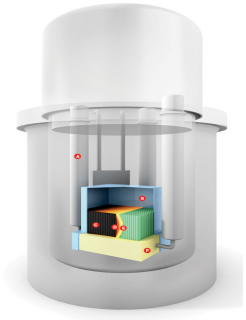

TerraPower Traveling Wave reactor

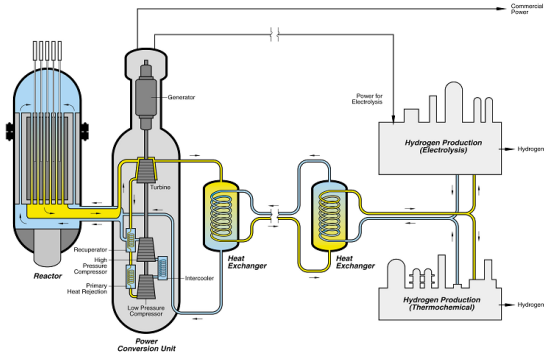


GA High temperature gas-cooled reactor

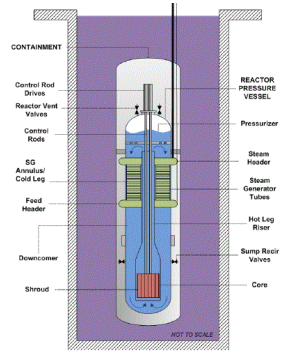


NuScale Small Modular reactor

# CESAR's Approach

- Start with existing petascale codes for thermo-hydraulics (NEK) and neutronics (UNIC) – and, if future budgets allow, structures (DIABLO)
- Couple codes and evolve into TRIDENT, an exascale code capable of high-fidelity modeling of real states of advanced reactors
- Be guided by advanced reactor vendors, solving problems they pose
- Co-design with computer vendors and system software developers

# – and its Benefits …

- **The benefit(s):** Simulating a complete nuclear power system in fine detail will fundamentally change the paradigm of how nuclear reactors are built, tested and operated.
  - Every step of the nuclear regulatory timeline can be compressed by guiding expensive experiment efforts.
  - New designs can be rapidly prototyped, accident scenarios can be studied in detail, material properties can be discovered, and design margins can be dramatically narrowed.
  - Scientists can analyze problems for a wide range of novel reactor systems

# Status

- Funded at $4M/yr for 5 years, starting September, 2011
- Progress
  - Kickoff meeting September 16, 2011
  - Identification of 1st-year targets in multiple areas:
    - Performance Modeling of a Mode Carlo neutronics code
    - Porting of some kernels (to be identified) to GPGPU nodes
    - Identification and execution of one flagship thermal hydraulics computation
    - Custom visualization interface for MOAB coupling code
    - Initial uncertainty quantification experiments at large scale
    - Starting postdoc at IBM on specific codesign project to be determined

# Concerns

- Need to find support for coupling to a structures code (removed from original proposal during funding downsizing).
- The Exascale Software Center was intended to supply exascale software specifically needed by CESAR in areas of OS, file structures, scalable MPI-3 implementation, and tools. Depending on future ASCR plans, we may need to adjust our deliverables if we ourselves end up responsible for general-purpose exascale software.

# The CESAR Team

- The management team:
  - Bob Rosner (Director), Bob Hill (Reactors), Rusty Lusk (Computer Science), Andrew Siegel (Deputy Director), Kord Smith (Chief Scientist/Engineer)
- Labs:
  - Argonne National Laboratory (lead)
  - Lawrence Livermore National Laboratory
  - Los Alamos National Laboratory
  - Oak Ridge National Laboratory
  - Pacific Northwest National Laboratory
- Universities:
  - Massachusetts Institute of Technology
  - Texas A&M University
  - Rice University
- Computer vendors:
  - IBM
  - The other vendors: Cray, HP, Intel, …
- Reactor vendors: AREVA, TerraPower, General Atomics, NuScale

# NSF HPC Task Force Report
## *March 2011*

- The HPC Task Force considers the issues confronting the NSF in the area of high performance computing to fall into the following 3 categories:
  - 1. Cyberinfrastructure Sustainability
  - 2. **Exascale Computing**
  - 3. Broader Engagement

# NSF HPC Task Force Report
## Strategic Exascale recommendations

- Develop a sustainable model to provide the academic research community with access, by 2015–2016, to a rich mix of HPC systems that:
  - deliver sustained performance of 20–100 petaflops on a broad range of science and engineering applications;
  - are integrated into a comprehensive national CI environment; and are supported at national, regional, and/or campus levels.
- Invest now to prepare for exascale systems that will be available by 2018–2020, including
  - **co-design partnerships** to provide the HPC systems
  - **data CyberInfrastructure** needed to enable data-driven science.

# Selected general recommendations to NSF

- Given the major challenges involved in the transition to HPC at the exascale, NSF should consider **new models for partnerships**, such as expanded collaborations with industry, academia, and other agencies
- Given the opportunities and challenges presented by the generation of exabytes of digital data, NSF should provide funding for a **digital data framework** designed to address the issues of knowledge discovery in the exascale ecosystem

# Summary

- There has been progress in some areas
- Not as much as I had hoped in the last year
  - Especially in the software issues that IESP has identified
- There are still many issues to be addressed in project planning, management, technical options to be pursued
  - But there is intense activity to address them