Exascale Computing Initiative

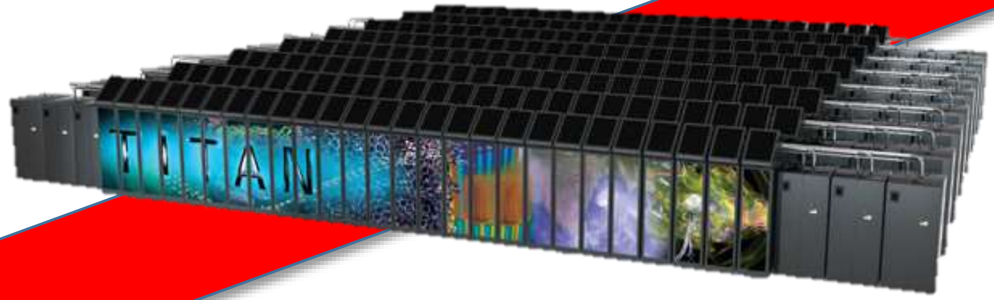

July 2015

ASCAC Review
Draft Report

ASCAC exascale charge

I am charging the ASCAC to review the Department's draft preliminary conceptual design for the Exascale Computing Initiative. Specifically, we are looking for input from the community as to whether there are significant gaps in our plans or areas that need to be given priority or extra management attention. ASCAC should gather, to the extent possible, input from a broad cross-section of the stakeholders.

Patricia Dehmer

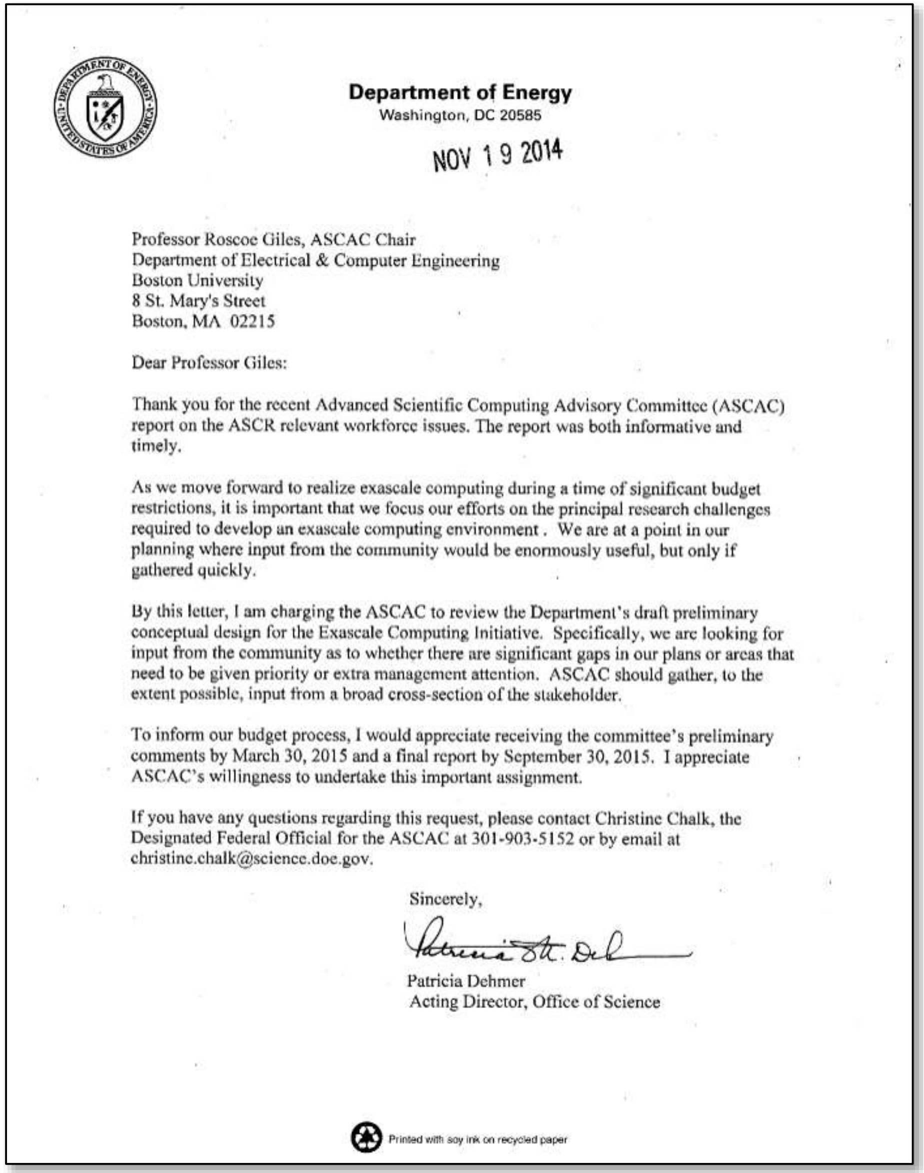

Department of Energy
Washington, DC 20585

NOV 19 2014

Professor Roscoe Giles, ASCAC Chair
Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Dear Professor Giles:

Thank you for the recent Advanced Scientific Computing Advisory Committee (ASCAC) report on the ASCR relevant workforce issues. The report was both informative and timely.

As we move forward to realize exascale computing during a time of significant budget restrictions, it is important that we focus our efforts on the principal research challenges required to develop an exascale computing environment. We are at a point in our planning where input from the community would be enormously useful, but only if gathered quickly.

By this letter, I am charging the ASCAC to review the Department's draft preliminary conceptual design for the Exascale Computing Initiative. Specifically, we are looking for input from the community as to whether there are significant gaps in our plans or areas that need to be given priority or extra management attention. ASCAC should gather, to the extent possible, input from a broad cross-section of the stakeholders.

To inform our budget process, I would appreciate receiving the committee's preliminary comments by March 30, 2015 and a final report by September 30, 2015. I appreciate ASCAC's willingness to undertake this important assignment.

If you have any questions regarding this request, please contact Christine Chalk, the Designated Federal Official for the ASCAC at 301-903-5152 or by email at christine.chalk@science.doe.gov.

Sincerely,

Patricia Dehmer
Acting Director, Office of Science

Printed with soy ink on recycled paper

Review subcommittee membership

- Dan Reed, University of Iowa (chair and ASCAC member)
- Martin Berzins, University of Utah (ASCAC member)
- Bob Lucas, Livermore Software Technology Corporation
- Satoshi Matsuoka, Tokyo Institute of Technology
- Rob Pennington, University of Illinois, retired
- Vivek Sarkar, Rice University (ASCAC member)
- Valerie Taylor, Texas A&M University







Subcommittee review process

- Planning
  - Kickoff planning teleconference
  - Question formulation
  - Work stream development
- DOE headquarters review
  - Office of Science and NNSA
- Individual interviews
  - DOE laboratories
  - U.S. research agency leaders
  - HPC community members
- Document reviews
  - DOE exascale plans
    - Preliminary Conceptual Design for an Exascale Computing Initiative (November 2014)
    - ExaRD Detailed Technical Descriptions (November 2014)
    - Workshop reports, presentations and research papers
  - European and Asian exascale plans
- Weekly subcommittee teleconferences
  - Assessment, writing and review

Review perspective

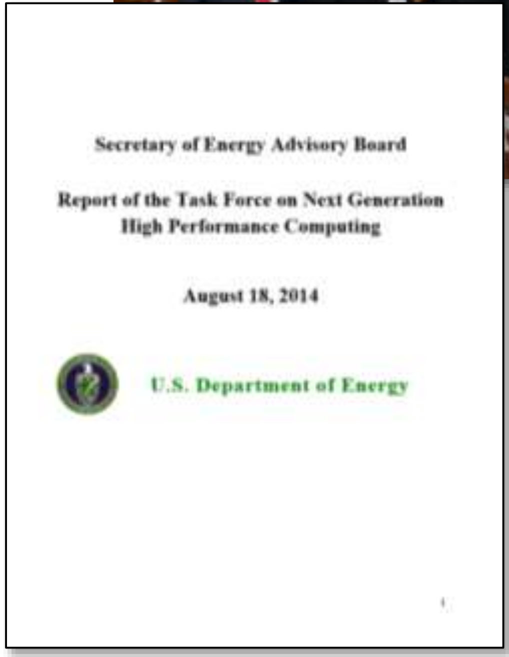

Secretary of Energy Advisory Board

Report of the Task Force on Next Generation High Performance Computing

August 18, 2014

U.S. Department of Energy

- DOE’s exascale plans have been reviewed extensively
  - Community workshops
  - Technical studies
  - Strategic assessments
  - Congressional hearings

over a period of more than *seven years*

- Subcommittee focused primarily on *organization and management*
  - Technical issues and previous studies informed the assessment

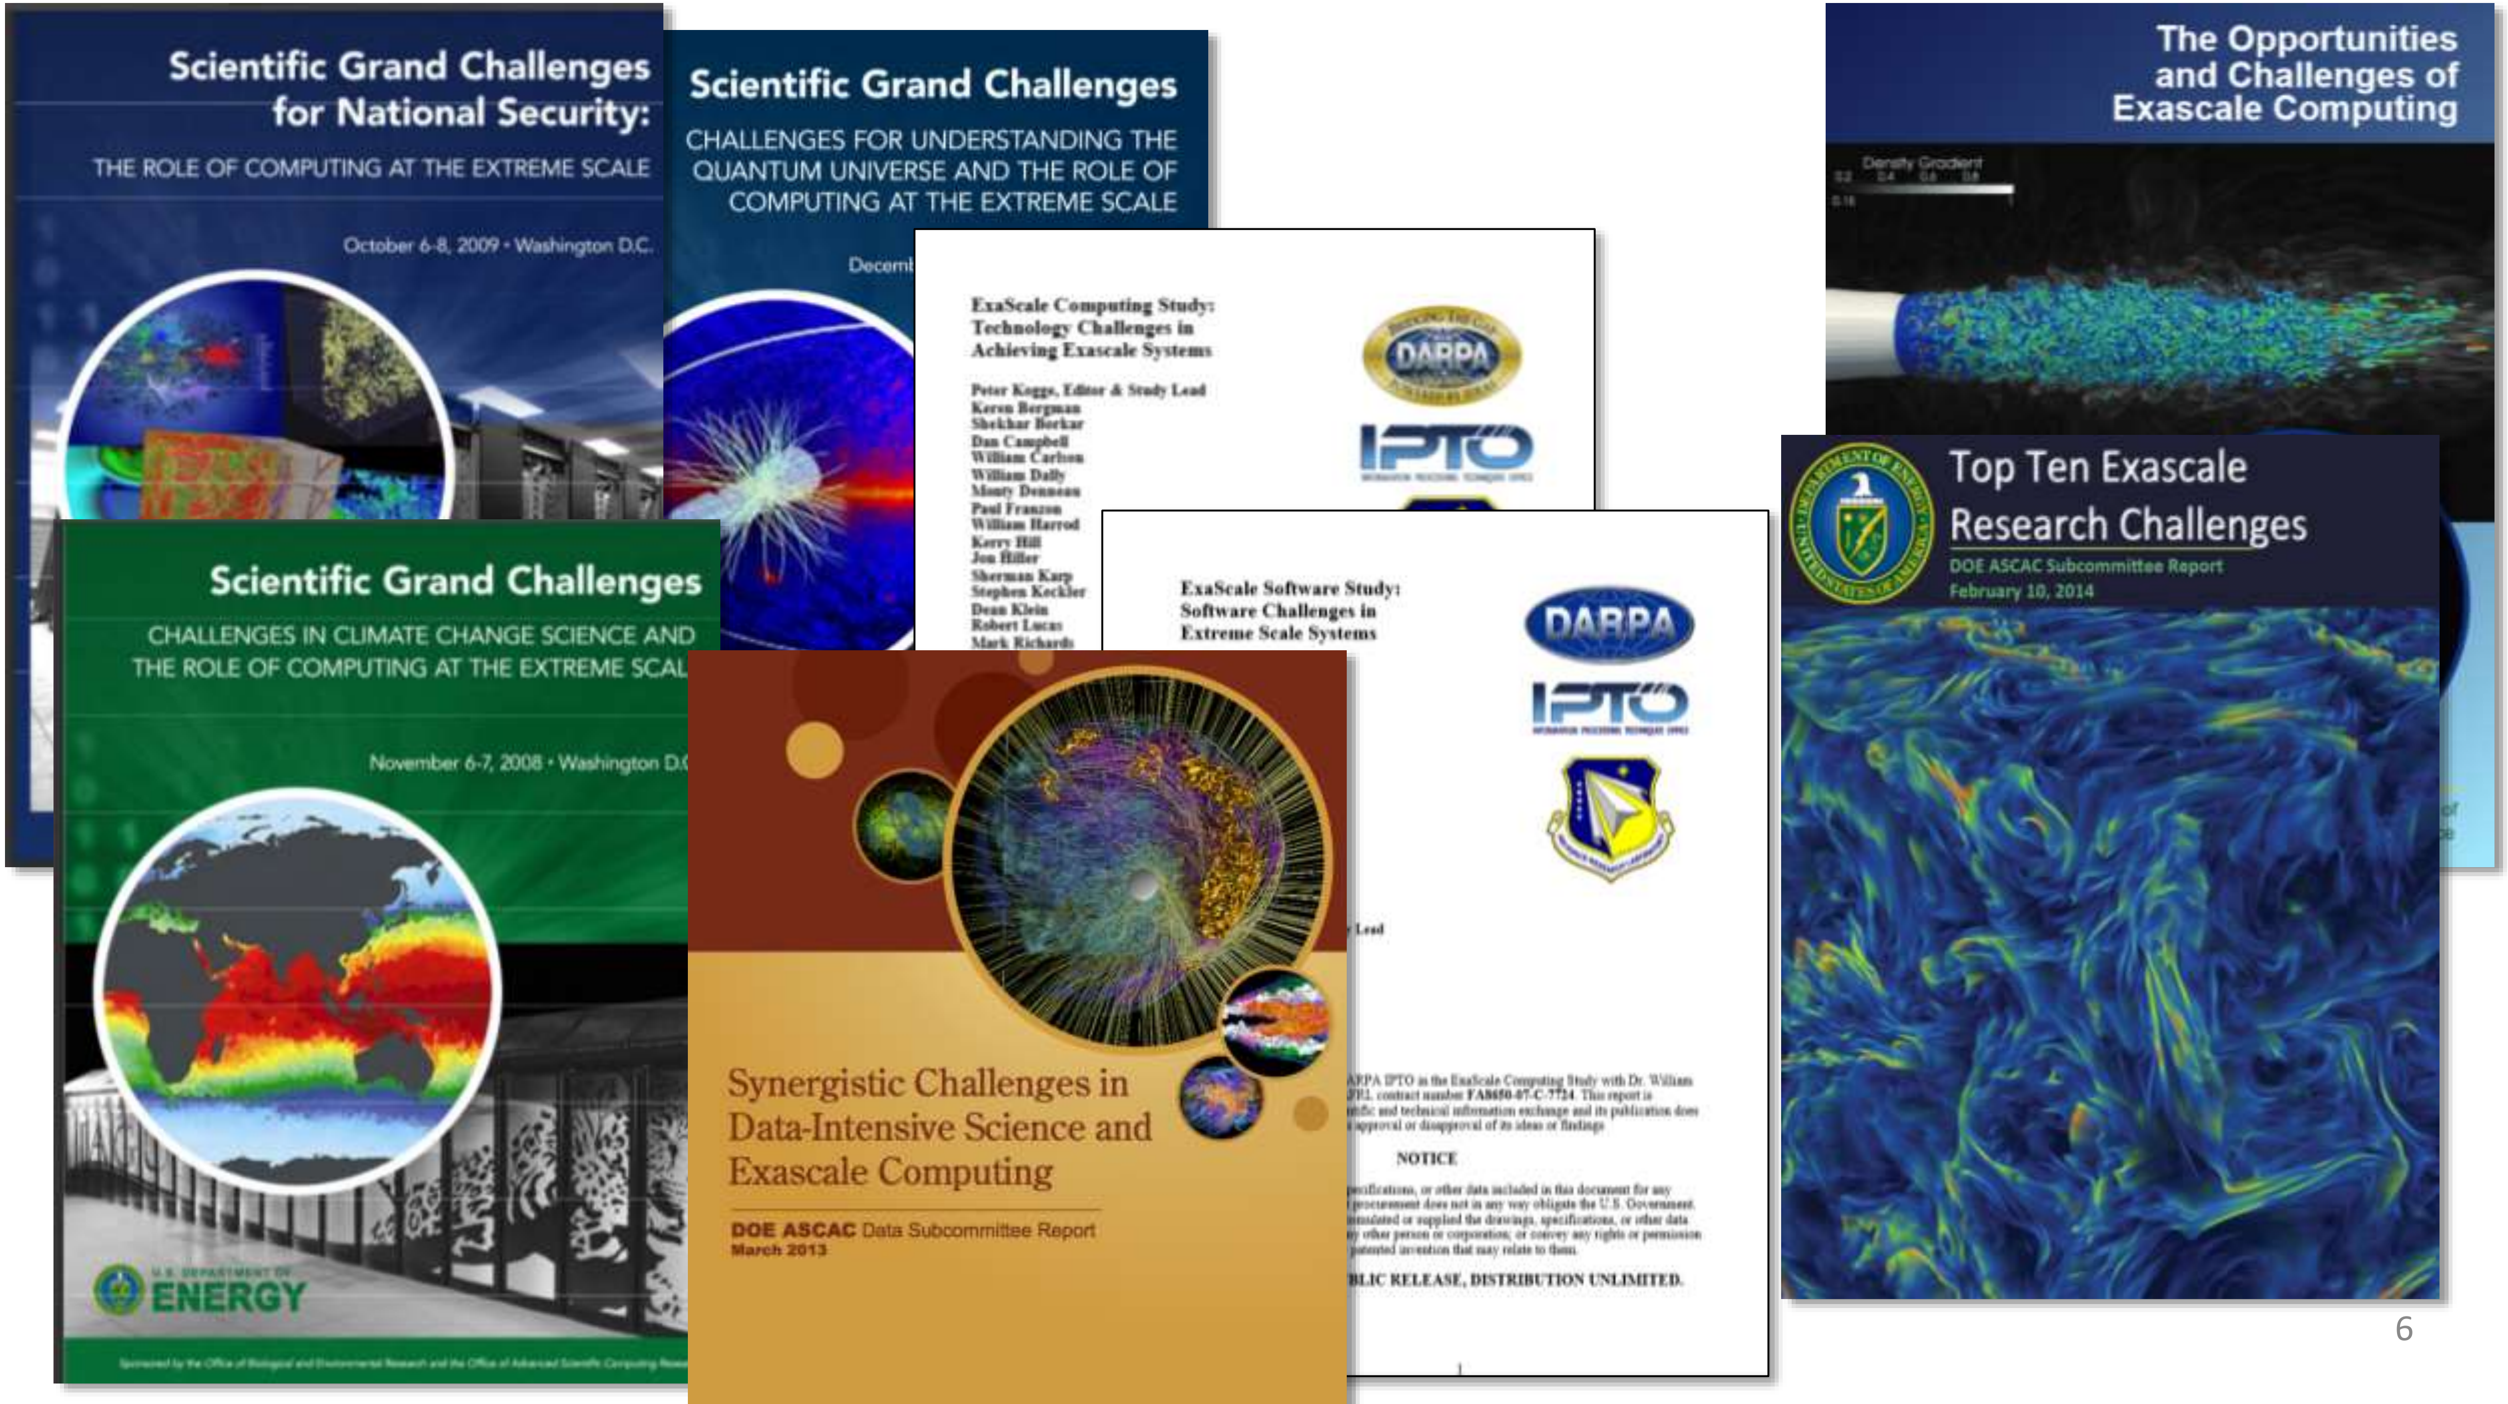
Scientific Grand Challenges for National Security:
THE ROLE OF COMPUTING AT THE EXTREME SCALE
October 6-8, 2009 • Washington D.C.
Scientific Grand Challenges
CHALLENGES FOR UNDERSTANDING THE QUANTUM UNIVERSE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE
ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems
Peter Kogge, Editor & Study Lead
Keren Bergman
Shekhar Borkar
Dan Campbell
William Carlson
William Dally
Monty Denneau
Paul Franzon
William Harrod
Kerry Hill
Jon Hiller
Sherman Karp
Stephen Keckler
Dean Klein
Robert Lucas
Mark Richards
DARPA
IPTO
The Opportunities and Challenges of Exascale Computing
Density Gradient
Top Ten Exascale Research Challenges
DOE ASCAC Subcommittee Report
February 10, 2014
Scientific Grand Challenges
CHALLENGES IN CLIMATE CHANGE SCIENCE AND THE ROLE OF COMPUTING AT THE EXTREME SCAL
November 6-7, 2008 • Washington D.C
U.S. DEPARTMENT OF ENERGY
ExaScale Software Study: Software Challenges in Extreme Scale Systems
DARPA
IPTO
NOTICE
BLIC RELEASE, DISTRIBUTION UNLIMITED.
Synergistic Challenges in Data-Intensive Science and Exascale Computing
DOE ASCAC Data Subcommittee Report
March 2013

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 advancements will require Extreme Scale Computing**
  - 1,000X 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**

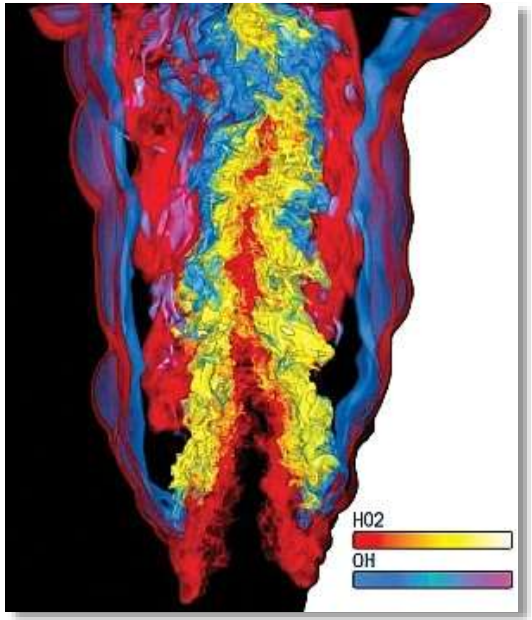

Exascale Computing Will Underpin Future Scientific Innovations

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

Exascale Applications Respond to DOE/NNSA Missions in Discovery, Design, and National Security

Scientific Discovery

- **Mesoscale materials and chemical sciences**
- **Improved climate models with reduced uncertainty**

Engineering Design

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

National Security

- **Stockpile stewardship**
- Real-time cybersecurity and incident response
- Advanced manufacturing

Blue Bold Text indicates planned or existing exascale application projects

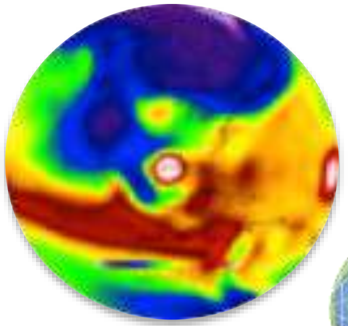

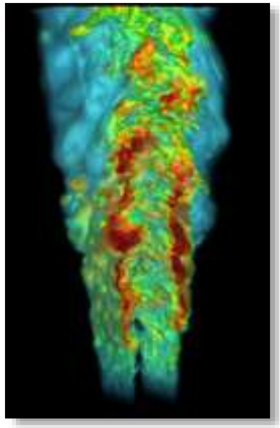

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

Top ten exascale challenges

1. **Energy efficiency**: Creating more energy-efficient circuit, power, and cooling technologies.
2. **Interconnect technology**: Increasing the performance and energy efficiency of data movement.
3. **Memory technology**: Integrating advanced memory technologies to improve both capacity and bandwidth.
4. **Scalable system software**: Developing scalable system software that is power- and resilience-aware.
5. **Programming systems**: Inventing new programming environments that express massive parallelism, data locality, and resilience
6. **Data management**: Creating data management software that can handle the volume, velocity and diversity of data that is anticipated.
7. **Exascale algorithms**: Reformulating science problems and redesigning, or reinventing, their solution algorithms for exascale systems.
8. **Algorithms for discovery, design, and decision**: Facilitating mathematical optimization and uncertainty quantification for exascale discovery, design, and decision making.
9. **Resilience and correctness**: Ensuring correct scientific computation in face of faults, reproducibility, and algorithm verification challenges.
10. **Scientific productivity**: Increasing the productivity of computational scientists with new software engineering tools and environment

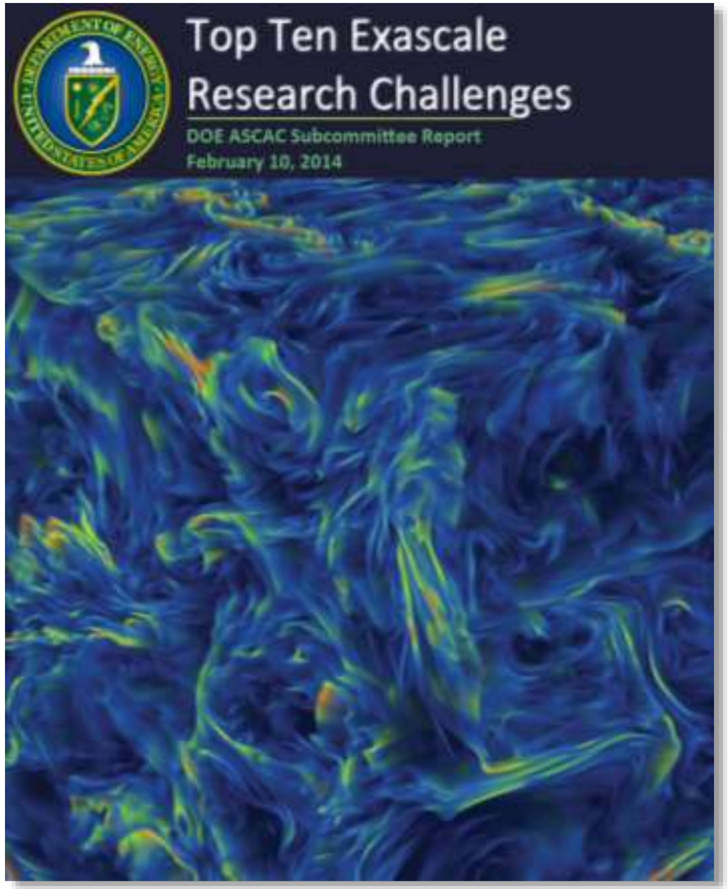

February 2014

Exascale Computing Initiative

High Level Program Structure

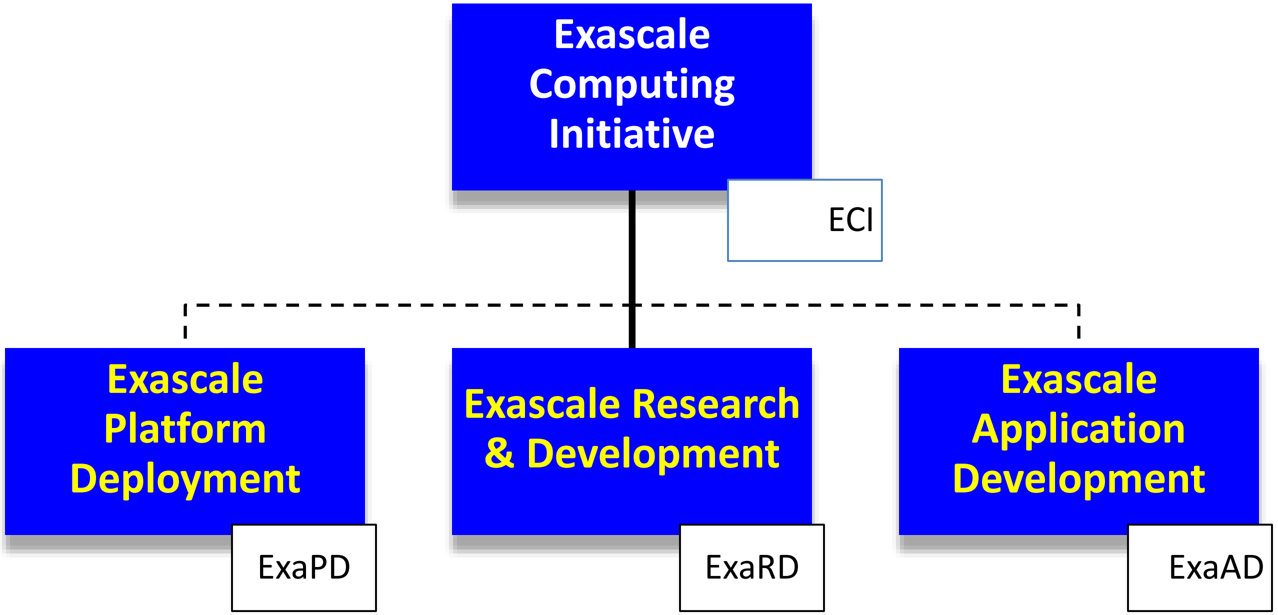

- Exascale technologies will offer new opportunities for pioneering scientific progress, and must be closely coordinated with exascale application development and acquisitions.
- Base funding for applications and acquisitions are provided by the appropriate DOE office.

Overarching findings and observations

- *The subcommittee strongly endorses the DOE plan for exascale computing development and deployment.*

- Like any ambitious undertaking, DOE's proposed exascale computing initiative (ECI) involves some risks. Despite the risks, the benefits of the initiative to scientific discovery, national security and U.S. economic competitiveness are clear and compelling.
- The subcommittee believes the ECI is a well-crafted plan designed to meet DOE mission needs while also advancing broader national security and competitiveness goals.
- DOE has a successful record of managing complex projects of this type.

Detailed recommendations (1-4 of 7)

1. Develop a detailed management and execution plan that defines clear responsibilities and decision-making authority to manage resources, risks, and dependencies appropriately across vendors, DOE laboratories, and other participants.
2. As part of the execution plan, clearly distinguish essential system attributes (e.g., sustained performance levels) from aspirational ones (e.g., specific energy consumption goals) and focus effort accordingly.
3. Given the scope, complexity, and potential impact of the ECI, conduct periodic external reviews by a carefully constituted advisory board.
4. Mitigate software risks by developing evolutionary alternatives to more innovative, but risky alternatives.

Detailed recommendations (5-7 of 7)

5. Unlike other elements of the hardware/software ecosystem, application performance and stability are mission critical, necessitating continued focus on hardware/software co-design to meet application needs.
6. Remain cognizant of the need for the ECI to support for data intensive and computation intensive workloads.
7. Where appropriate, work with other federal research agencies and international partners on workforce development and long-term research needs, while not creating dependences that could delay or imperil the execution plan.

1. Management and execution plan

- *Develop a detailed management and execution plan that defines clear responsibilities and decision-making authority to manage resources, risks, and dependencies appropriately across vendors, DOE laboratories, and other participants.*

- Establish a leadership structure that operates below and in concert with the present, high-level leadership at DOE headquarters. This leadership structure's sole focus should be the exascale program.
  - It could be either a small management group or a single, overall leader

- Develop a formal risk and assessment plan in concert with detailed execution planning

Key Performance Goals

Parameter	
Performance	Sustained 1 – 10 ExaOPS for applications
Power	20 MW
Cabinets	200 - 300
System Memory	128 PB
Mean Time Between Application Failure	6 days
Productivity	*TBD: programmability and code portability*
Scalable benchmarks	*Will utilize an approach that is similar to the CORAL RFP benchmarks*
Throughput benchmarks	*Will utilize an approach that is similar to the CORAL RFP benchmarks*

ExaOPS = 10^{18} Operations / sec

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

2. Irreducible goals versus aspirations

- *As part of the execution plan, clearly distinguish essential system attributes (e.g., sustained performance levels) from aspirational ones (e.g., specific energy consumption goals) and focus effort accordingly.*

- At a point when many new technologies components are still maturing, the ECI must not commit prematurely. If target numbers are publicized and shape activities prematurely, there is a danger that the ECI could be perceived as a failure for not reaching initial objectives.

- Failure to create a broad ecosystem will very likely disincentive both the users and the vendors, and as a result, will fail to leverage their mainstream research and development efforts, ultimately resulting in fewer technological advances and lower overall performance.

3. External review board assessment and validation

- *Given the scope, complexity, and potential impact of the ECI, conduct periodic external reviews by a carefully constituted advisory board.*

- There must be a well-defined process to monitor technology developments, potential risks and benefits; careful co-ordination across stakeholders; and rigorous assessment of project priorities and directions.

- The primary rationale for the ECI is the new scientific discoveries and technical capabilities it will enable. DOE must quantify what that means, ensuring there are credible application and discovery measures for the success, or failure, of the ECI.

- An external review board may also be useful in resolving complex resource issues and in assessing research and development risks when technology is changing quickly. It can also help define success metrics.

4. Software and risk mitigation

- *Mitigate software risks by developing evolutionary alternatives to more innovative, but risky alternatives.*

- Exascale software development has two distinct goals. The first is allowing applications to execute at scale as quickly as possible, with minimal change. The second is shifting the software base to post-petascale architectures and ensuring broader uptake and use of exascale systems.

- Ensure applications have both an evolutionary and a revolutionary path to exascale execution.

Exascale Application Challenges
Key Issues

- **Parallelism**
  - Today's parallelism is "weak scaling".
  - In the future we must be able to exploit "strong scaling"
- **Reliability**
  - Application code developers and users must have techniques for overcoming faults and heterogeneity.
- **Scaling Application Codes to Exploit Large-Scale, Advanced Architectures**
  - Lack of scalability will limit the throughput of many scientific applications
- **Data-Intensive Science**
  - Must address the challenges of large-scale simulations and data analysis.

Co-Design

Application-driven co-design is the process by which:

– Scientific problem requirements guide computer architecture and system software design
– Computer technology capabilities and constraints inform formulation and design of algorithms and software

Need shared global perspective across the design-space – to establish conceptual framework for co-design and interoperability

– Parallelism
– Latency
– Overhead
– Dependability

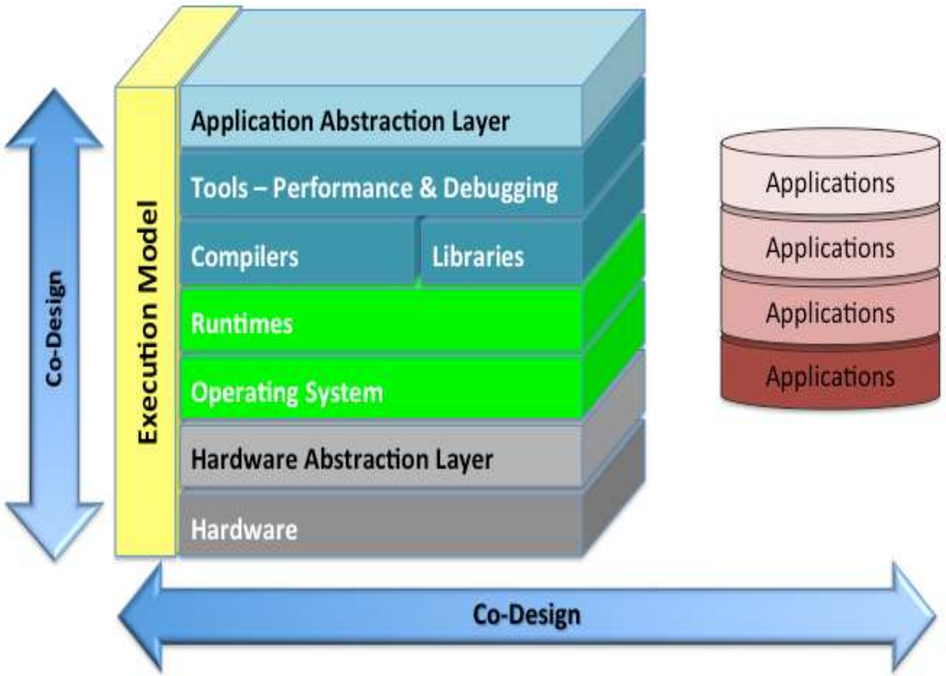

5. Application centrality and co-design

- *Unlike other elements of the hardware/software ecosystem, application performance and stability are mission critical, necessitating continued focus on hardware/software co-design to meet application needs.*

- The ECI should identify a set of these mission-critical applications from its target domains (computational materials science, next generation climate models, stockpile stewardship) and make them yardsticks against which exascale systems are evaluated.

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

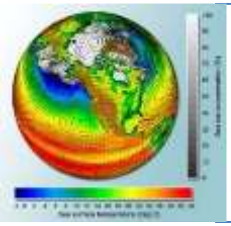

Climate

Data expected to be hundreds of PB

Driven by exponential technology advances

Data sources

- Experimental & Observational data from Scientific Instruments
- Simulation Results from scientific computing facilities

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

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

6. Big data and analytics

- *Remain cognizant of the need for the ECI to support data intensive* and *computation intensive workloads.*

- Modeling and data analysis are inextricably intertwined enablers of innovation and discovery; both draw on the same ecosystem of hardware and software technologies; and both are crucial elements of DOE's ECI.

- A new generation of data analytic tools and libraries are needed to aid in the interpretation and validation of the data generated from exascale applications

7. Interagency and international collaborations

- *Where appropriate, work with other federal research agencies and international partners on workforce development and long-term research needs, while not creating dependences that could delay or imperil the execution plan.*

- The February 2015 ASCAC report identified the need for long-term partnerships among the stakeholders, including government agencies, academia and vendors, to address these fundamental requirements, as derived from science and mission needs.

7. Interagency and international collaborations (continued)

- Develop plans for interagency research collaborations and mechanisms to incorporate salient research results, while not creating dependencies that could delay or imperil its execution plan. More broadly, interagency research collaborations would expand and accelerate development of a highly trained and flexible workforce that is aware of, contributing to and utilizing exascale systems.

- There are additional opportunities for bilateral and multilateral international collaborations to ensure development of consistent and interoperable software ecosystems and applications

Summary

- Advanced computing's benefits are broad and deep

- Exascale computing is not a destination, but the next milestone in a journey
  - Scientific discoveries
  - National security
  - Economic competitiveness

- *The subcommittee strongly endorses the DOE plan for exascale computing development and deployment.*

Next steps

- The subcommittee welcomes comments and feedback

- Remaining timeline and issues
  - Document completion by end of August
  - ASCAC review and approval

Acknowledgments

- Steve Binkley
- Christine Chalk
- Bill Harrod
- Barb Helland
- Thuc Hoang
- Bob Meisner

- … and a host of others

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration