Exascale Computing Project –
Driving a HUGE Change in a Changing World

Al Geist
Chief Technology Officer
Exascale Computing Project

OpenFabrics Conference
March 27, 2017
Austin, TX

EXASCALE COMPUTING PROJECT

Exascale Computing Project - Lift the entire HPC ecosystem and enable continued U.S. leadership in HPC

Leverages the planned DOE facilities acquisitions
- ◊ 2017 CORAL (collaboration of ORNL, ANL, LLNL)
- ◊ 2020 APEX (collaboration of LANL/SNL, NERSC)
- ◊ 2022 CORAL
- ◊ 2024 APEX

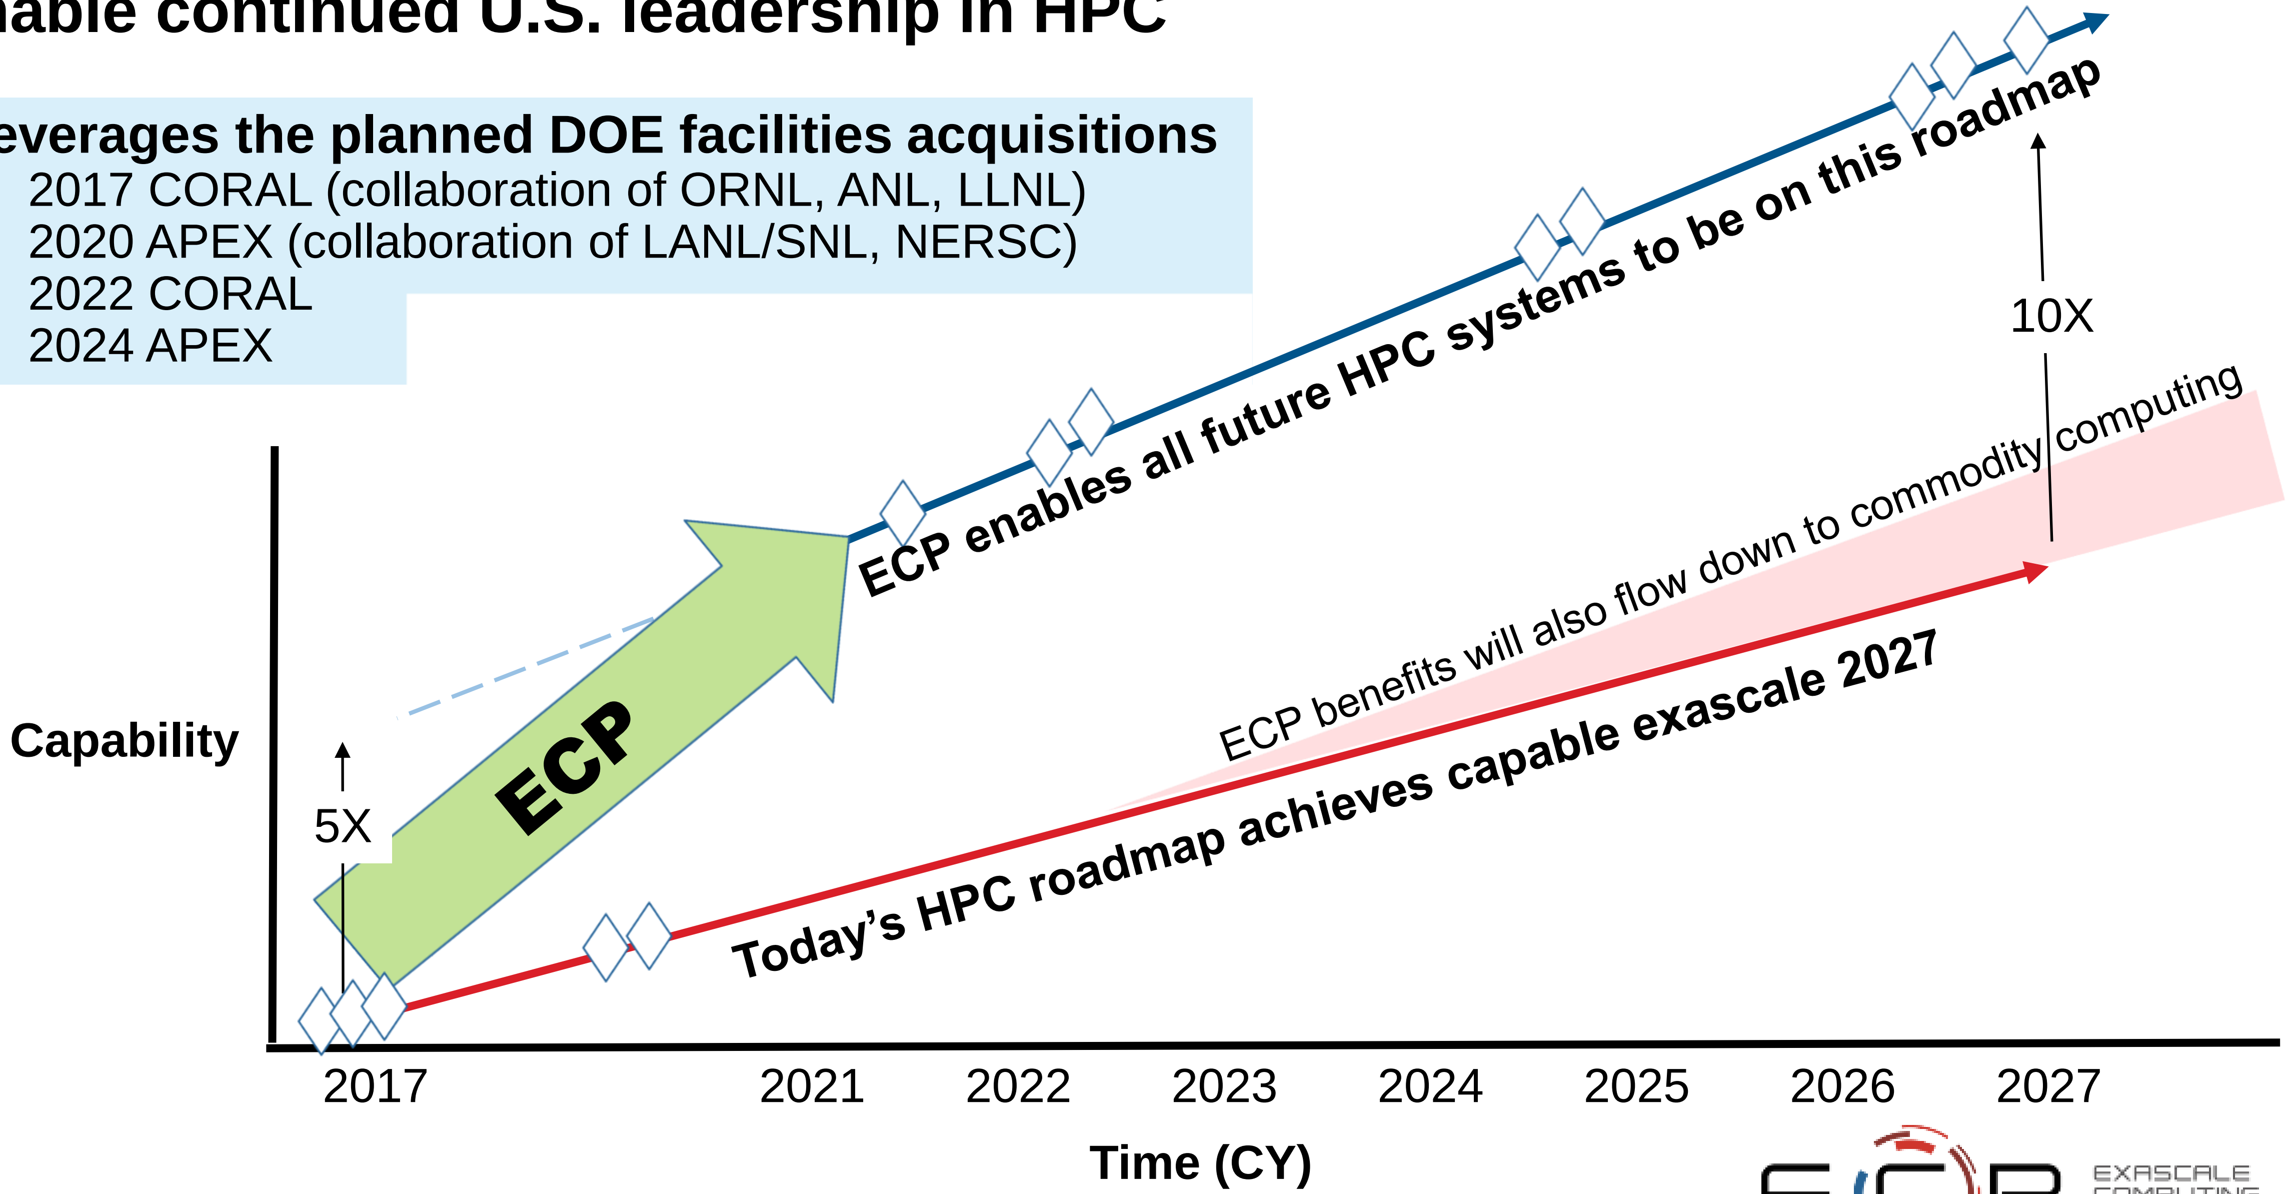

Reaching the Elevated Trajectory will require solving key exascale challenges

- **Extreme Parallelism**
  - For example, an Exaflop @ 1 GHz requires a billion threads executing
- **Memory and Storage**
  - BW, latency, and capacity are not scaling with flops
- **Reliability**
  - Energy saving techniques and number of components drive MTBF down
- **Energy Consumption**
  - 20MW per Exaflop has been a target since 2009

In addition, the exascale advanced architecture will need to solve emerging data science and machine learning problems in addition to the traditional modeling and simulations applications.

Radical, Novel, Advanced solutions are not a Requirement but may be needed

We want the vendors to propose what they see as being needed to meet performance, reliability, programmability, data science convergence, and power requirements.

- If vendors can meet the requirements without needing new radical solutions that is fine and likely preferred.
- If it involves radical new concepts, we are interested in hearing about these solutions.
- We want to encourage vendors to propose new ideas where they provide a path for addressing our requirements but we don’t need novelty or “advancedness” just so we can claim things are “advanced”.

Goals of the Exascale Computing Project

Foster application development	Ease of use	Rich exascale ecosystem	US HPC leadership
Develop scientific, engineering, and large-data applications that exploit the emerging, exascale-era computational trends caused by the end of Dennard scaling and Moore's law	Create software that makes exascale systems usable by a wide variety of scientists and engineers across a range of applications	Enable exascale by 2021 and by 2023 at least two diverse computing platforms with up to 50× more computational capability than today's 20 PF systems, within a similar size, cost, and power footprint	Help ensure continued U.S. leadership in architecture, software and applications to support scientific discovery, energy assurance, stockpile stewardship, and nonproliferation programs and policies

The ECP Plan

- Use a *holistic/co-design* approach across four focus areas:
  - Application Development
  - Software Technology R&D
  - Hardware Technology R&D
  - Exascale Systems Development
- Enable an initial exascale system to be delivered in 2021 (power consumption and reliability requirements may be relaxed)
- Enable *capable exascale* systems to be delivered in 2022 as part of the CORAL DOE facility upgrades
- System acquisitions and costs are **outside of the ECP plan**, and will be carried out by DOE facilities

ECP Timeline has Two Phases – and ends 2022

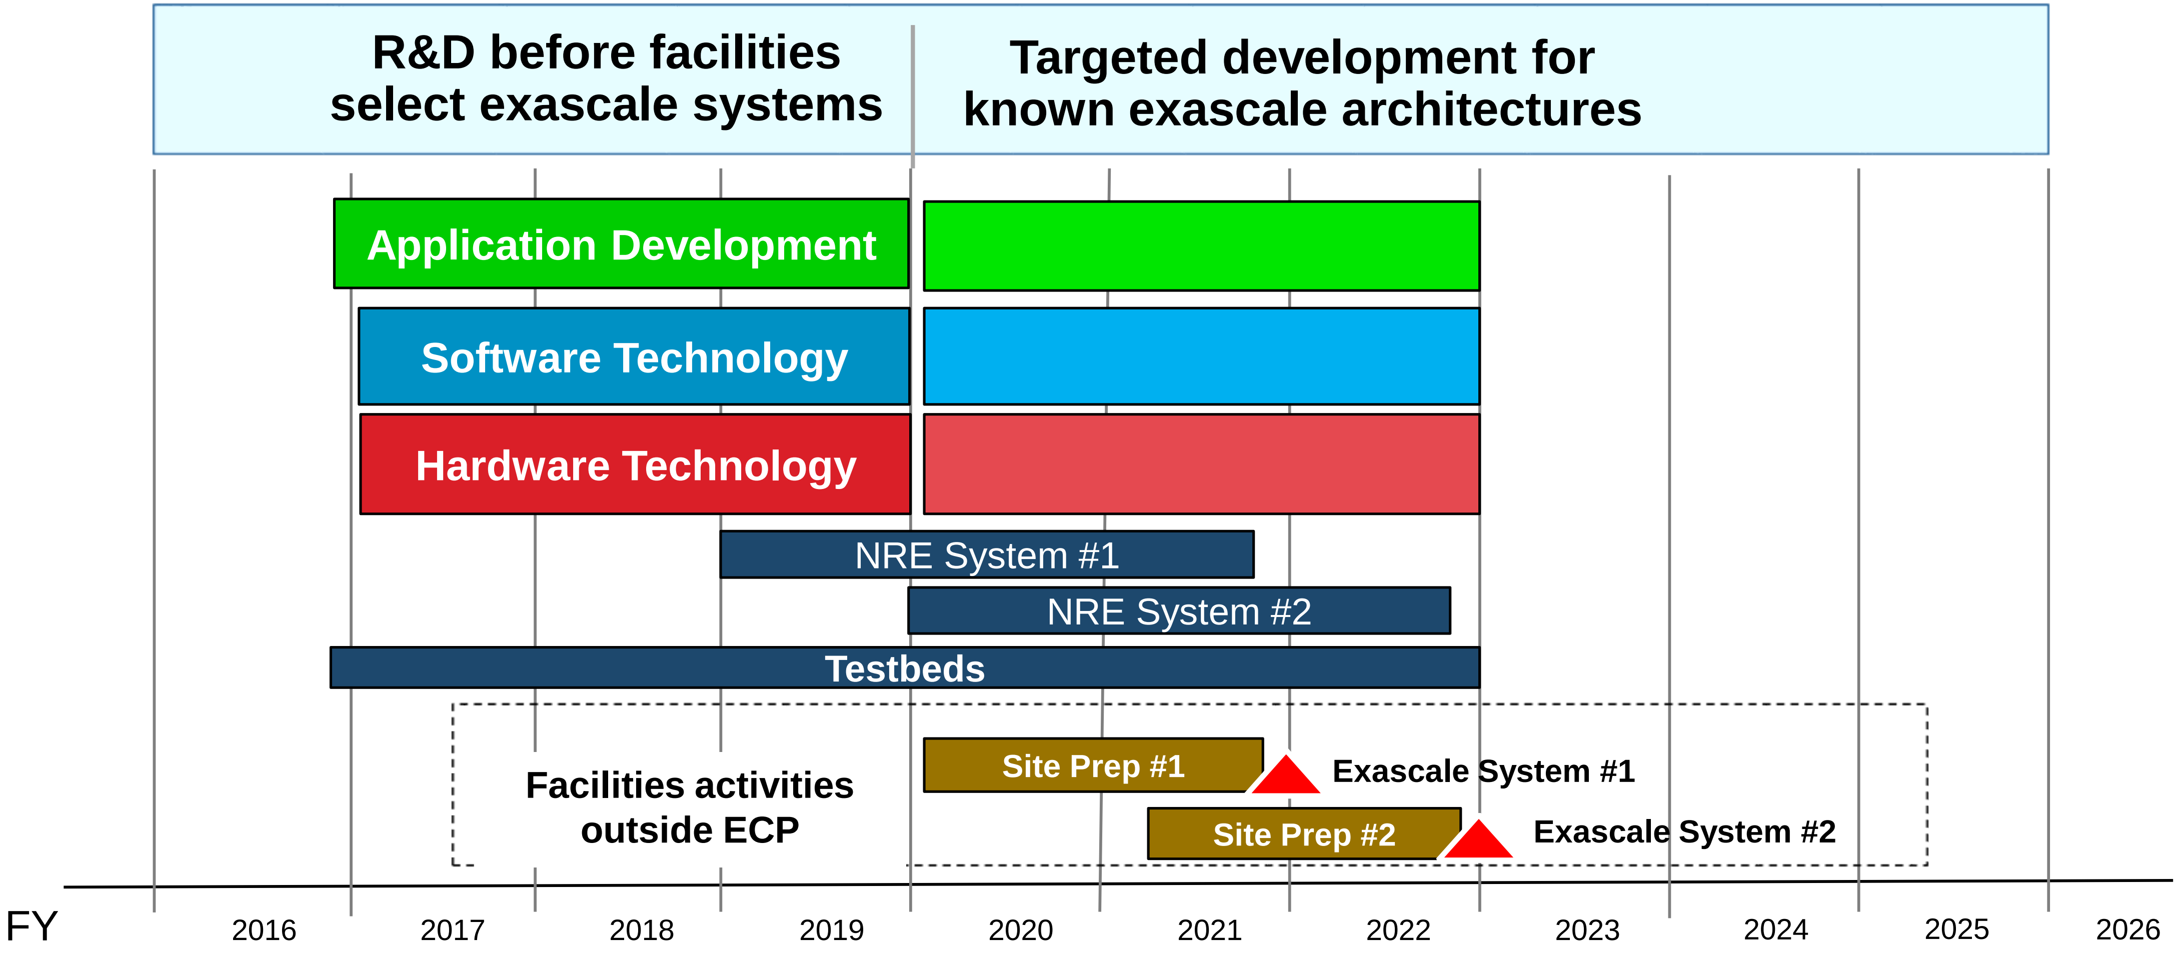

What about the 2021 System?

- The site of the 2021 system is TBD and will be decided by DOE around June 2017.

- If the site is one of the CORAL labs, then the CORAL RFP will state:

  "Within the goal of having three capable exascale systems by 2022-2023, if an early exascale system can be delivered in 2021 and upgraded to a capable exascale system by 2023, then provide the upgrade as an option."

- If the site of the 2021 system is outside the CORAL labs then (in addition to the CORAL RFP for three 2022 systems) a separate RFP for a single 2021 system will be released in 2018 by the chosen lab.

What is a *capable* exascale computing system?

ECP defines a capable exascale system as a supercomputer that

- Can solve science problems 50x faster (or more complex, for example, more physics, higher fidelity) than the 20 PF systems of today can solve comparable problems.
- Must use a software stack that meets the needs of a broad spectrum of applications and workloads
- Have a power envelope of 20-30 MW
- Must be sufficiently resilient such that user intervention due to hardware or system faults is required on the order of a week.

EXASCALE COMPUTING PROJECT

Diversity is Very Important to DOE

- In 2018 a single CORAL RFP will be released for delivery of three capable exascale systems by the 2022-2023 timeframe. The RFP will also include NRE for the systems.
- These systems will have to be designed to solve emerging data science and machine learning problems in addition to the traditional modeling and simulations applications.
- The DOE Leadership Computing Facility has a requirement that the ANL and ORNL systems must have diverse architectures.
- Given the ECP goal of fostering a rich exascale ecosystem, LLNL has the option to choose a system that is diverse from both the ANL and ORNL systems.

There are Many Types of System Diversity

- **Systems can vary from one another in many different dimensions**
  - System (architecture, interconnect, IO subsystem, density, resilience, etc.)
  - Node (heterogeneous, homogeneous, memory and processor architectures, etc.)
  - Software (HPC stack, OS, IO, file system, prog environment, admin tools, etc.)
  - Hardware e.g.

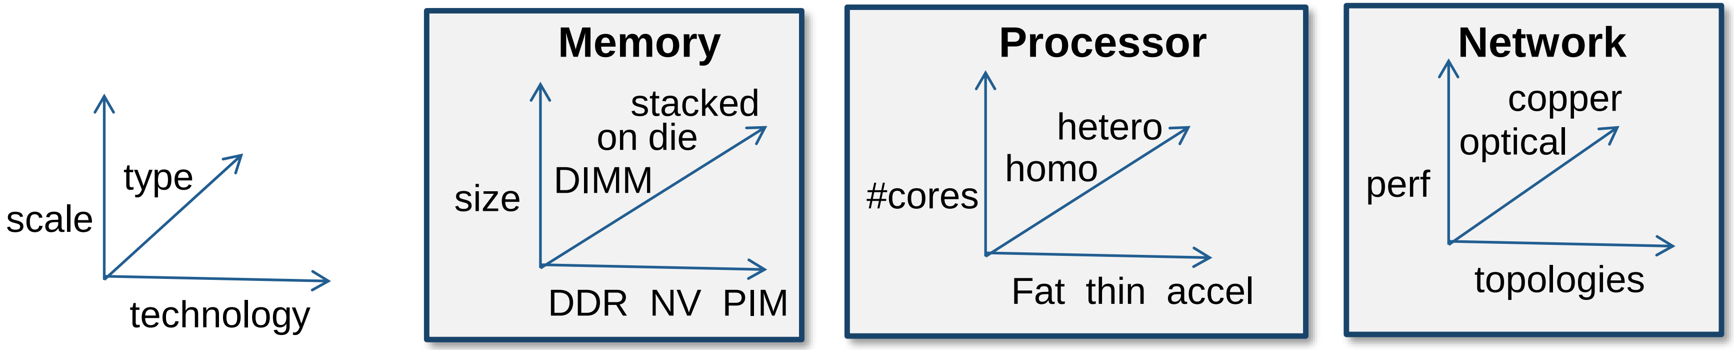

- **Ways Systems can be diverse**
  - Few big differences
  - Many little differences
  - Different technologies
  - Different ecosystems, i.e., vendors involved

How Diverse is Enough?

How diverse is enough? There is no hard metric, Labs will evaluate diversity by how much it will benefit the exascale ecosystem

Having system diversity provides many advantages.

- It promotes price competition, which increases the value to DOE.
- It promotes a competition of ideas and technologies, which helps provide more capable systems for DOE's mission needs.
- It reduces risk that may be caused by delays or failure of a particular technology or shifts in vendor business focus, staff or financial health.
- It helps promote a rich and healthy high performance computing ecosystem, which is important for national competitiveness and DOE's strategic plan.

The ECP holistic approach uses co-design and integration to achieve capable exascale

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable and productive software stack	Hardware technology elements	Integrated exascale supercomputers

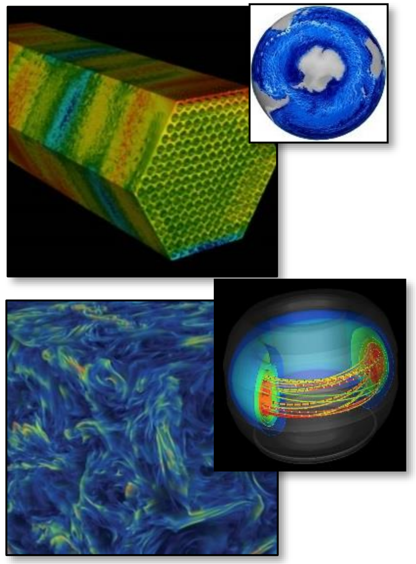

Co-design centers
Proxy apps

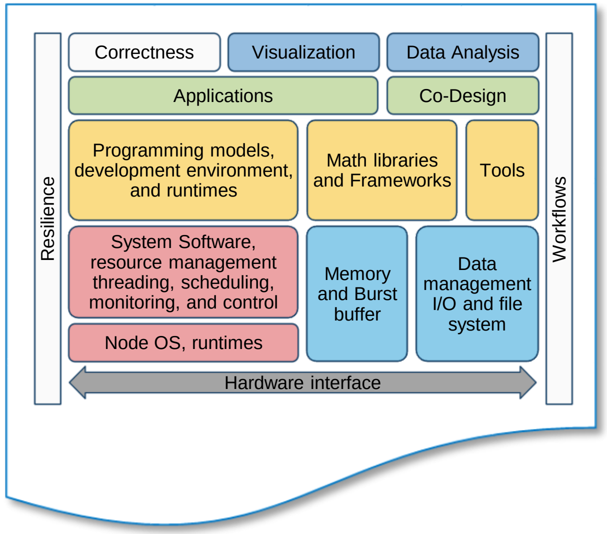

Integration of NNSA and Office of Science SW efforts

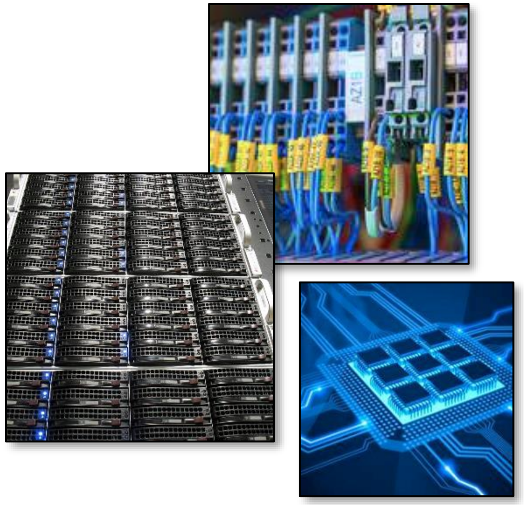

PathForward
Design Space Evaluation

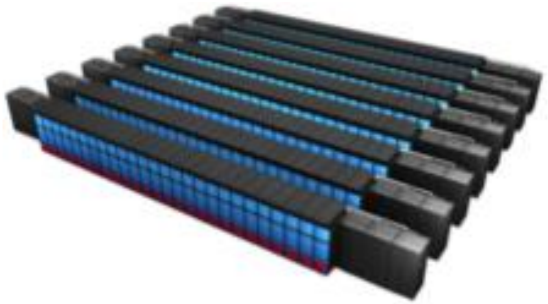

Testbeds
NRE

EXASCALE COMPUTING PROJECT

Application Scope determined by Mission Needs

Support DOE science *and* energy missions

- Discover and characterize next-generation **materials**
- Systematically understand and improve **chemical processes**
- Analyze the extremely large datasets resulting from the next generation of **particle physics** experiments
- Extract knowledge from **systems-biology** studies of the microbiome
- Advance **applied energy** technologies (e.g., whole-device models of plasma-based fusion systems)

Meet national security needs

- Stockpile Stewardship Annual Assessment and Significant Finding Investigations
- Robust uncertainty quantification (UQ) techniques in support of lifetime extension programs
- Understanding evolving nuclear threats posed by adversaries and in developing policies to mitigate these threats

Key science and technology challenges to be addressed with exascale

- Materials discovery and design
- Climate science
- Nuclear energy
- Combustion science
- Large-data applications
- Fusion energy
- National security
- Additive manufacturing
- Many others!

ECP EXASCALE COMPUTING PROJECT

ECP Application Development – (1/3)

Applications chosen based on National impact and DOE Offices priorities

Fundamental Laws (NP)	Climate (BER)	Materials Science (BES)	Chemical Science (BES, BER)	Combustion (BES)
QCD-based elucidation of fundamental laws of nature: Standard Model validation and beyond SM discoveries	Accurate regional impact assessment of climate change*	Find, predict, and control materials and properties:	Biofuel catalysts design; stress-resistant crops	Design high-efficiency, low-emission combustion engines and gas turbines*

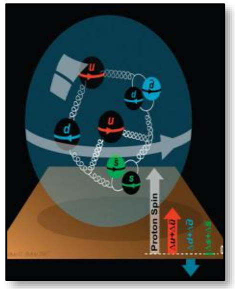

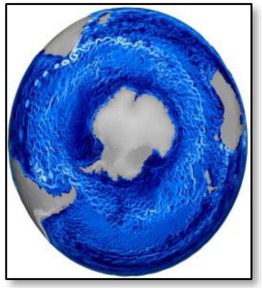

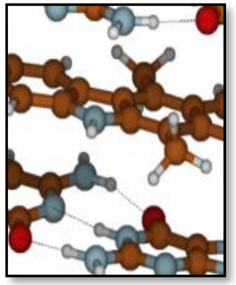

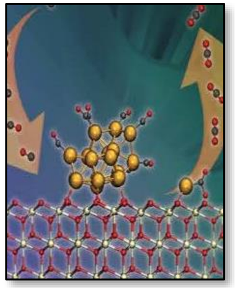

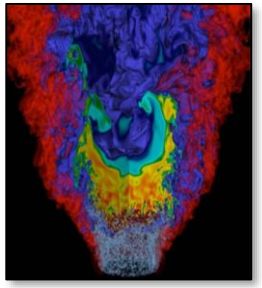

EXASCALE COMPUTING PROJECT

ECP Application Development – (2/3)

Some applications also include a significant machine learning component *

Precision Medicine for Cancer (NIH)	Seismic (EERE, NE, NNSA)	Genomics (BES)	Metagenomic (BER)	Chemical Science (BES)
Accelerate and translate cancer research in RAS pathways, drug responses, and treatment strategies*	Reliable earthquake hazard and risk assessment in relevant frequency ranges* treaty verification	Protein structure and dynamics; 3D molecular structure design of engineering functional properties*	Leveraging microbial diversity in metagenomic datasets for new products and life forms*	Design catalysts for conversion of cellulosic-based chemicals into fuels, bioproducts

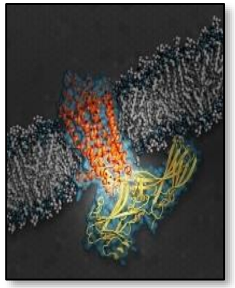

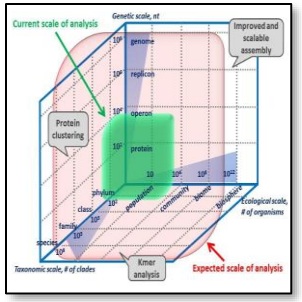

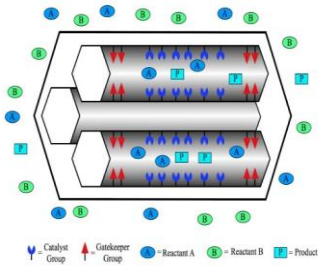

ECP Applications Development – (3/3)

Some applications also include a significant data science component *

Magnetic Fusion Energy (FES)	Nuclear Energy (NE)	Wind Energy (EERE)	Cosmology (HEP)	Astrophysics (NP)
Predict and guide stable ITER operational performance with an integrated whole device model*	Accelerate design and commercialization of next-generation small modular reactors*	Increase efficiency and reduce cost of turbine wind plants sited in complex terrains*	Cosmological probe of standard model (SM) of particle physics: Inflation, dark matter, dark energy*	Demystify origin of universe and nuclear matter in universe*

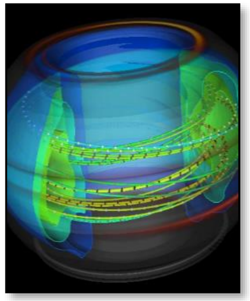

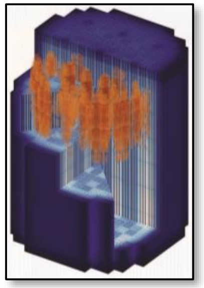

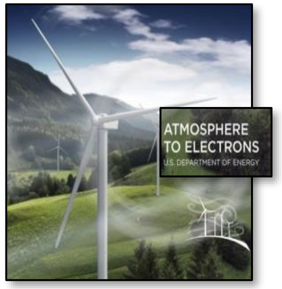

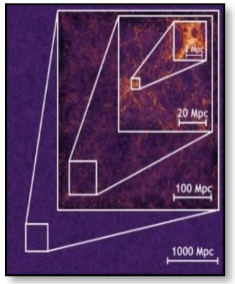

* Scope includes a significant data science component

ECP Application Development Co-Design Centers

- **Center for Online Data Analysis and Reduction at the Exascale (CODAR)**
- **Block-Structured AMR Co-Design Center (AMReX)**
- **Center for Efficient Exascale Discretizations (CEED)**
- **Co-Design Center for Particle Applications (CoPA)**
- **Graph and Combinatorial Methods for Enabling Exascale Applications (GraphEx)**

ECP Software Technology Summary

- ECP will build a comprehensive and coherent software stack that will enable application developers to productively write highly parallel applications that can portably target diverse exascale architectures

- ECP will accomplish this by extending current technologies to exascale where possible, performing R&D required to conceive of new approaches where necessary, coordinating with vendor efforts, and developing and deploying high-quality and robust software products

ECP Hardware Technology Summary

Objective: Fund R&D to design hardware that meets ECP's Targets for application performance, power efficiency, and resilience

- Issue *PathForward and PathForward-II* Hardware Architecture R&D contracts
- Participate in evaluation and review of PathForward and LeapForward deliverables
- Lead Design Space Evaluation through Architectural Analysis, and Abstract Machine Models of PathForward/PathForward-II designs for ECP's holistic co-design

Goals for PathForward (issued last year – vendor awards pending)

- Improve the quality and number of competitive offeror responses to the Capable Exascale Systems RFP

- Improve the offeror's confidence in the value and feasibility of aggressive advanced technology options that would be bid in response to the Capable Exascale Systems RFP

- Improve DOE confidence in technology performance benefit, programmability and ability to integrate into a credible system platform acquisition

Goals of PathForward-II (planned for issue in 2017)

- Support high payoff, innovative hardware technologies and systems technologies that may have higher risk. It is focused on component, node, and system architecture designs that will intersect with the 2021 exascale system.
- Also of interest to the PathForward-II RFP team:
  - Innovations that may enable dramatic acceleration of certain applications, for example, delivering a 100x increase in 2021 on some classes of applications while still being able to solve the full range of DOE applications
  - Developments that promote wider diversity in the exascale ecosystem
  - Innovations in power consumption, performance, programmability, reliability, data science, machine learning, or portability
  - Reducing total cost of ownership

EXASCALE COMPUTING PROJECT

ECP Exascale Systems Summary

- **Funds Non-Recurring Engineering (NRE)**
  - Brings to the product stage promising hardware and software research and integrates it into a system
  - Includes application readiness R&D efforts
  - Must start early enough to impact the system - more than two full years of lead time are necessary to maximize impact
- **Funds Testbeds**
  - ECP ECP testbeds will be deployed each year throughout the project
  - FY17 testbeds will be acquired through options on existing contracts at Argonne and ORNL
  - Testbed architectures will track SC/NNSA system acquisitions and other promising architectures

This is a very exciting time for computing in the US

- Unique opportunity to do something HUGE for the nation in HPC
- The exascale systems in 2021 and 2022 afford the opportunity for
  - More rapid advancement and scaling of mission and science applications
  - More rapid advancement and scaling of an exascale software stack
  - Rapid investments in vendor technologies and software needed for 2021 and 2022 systems
  - More rapid progress in numerical methods and algorithms for advanced architectures
  - Strong leveraging of and broader engagement with US computing capability
- When ECP ends, we will have
  - Prepared industry and critical applications for a more diverse and sophisticated set of computing technologies, carrying US supercomputing well into the future
  - Demonstrated integrated software stack components at exascale
  - Invested in the engineering and development, and participated in acquisition and testing of capable exascale systems

Thank you!

www.ExascaleProject.org

EXASCALE COMPUTING PROJECT