CHALLENGES OF EXASCALE COMPUTING, REDUX

PAUL MESSINA
Argonne National Laboratory

May 17, 2018
5th ENES HPC Workshop
Lecce, Italy

OUTLINE

- Why does exascale computing matter?
- The U.S. DOE Exascale Computing Project
- Challenges: then and now

Argonne NATIONAL LABORATORY

SCIENTIFIC DISCOVERY: SIMULATION OF REALISTIC MATERIALS WITH DEFECTS AND DOPANTS FROM ACROSS THE PERIODIC TABLE

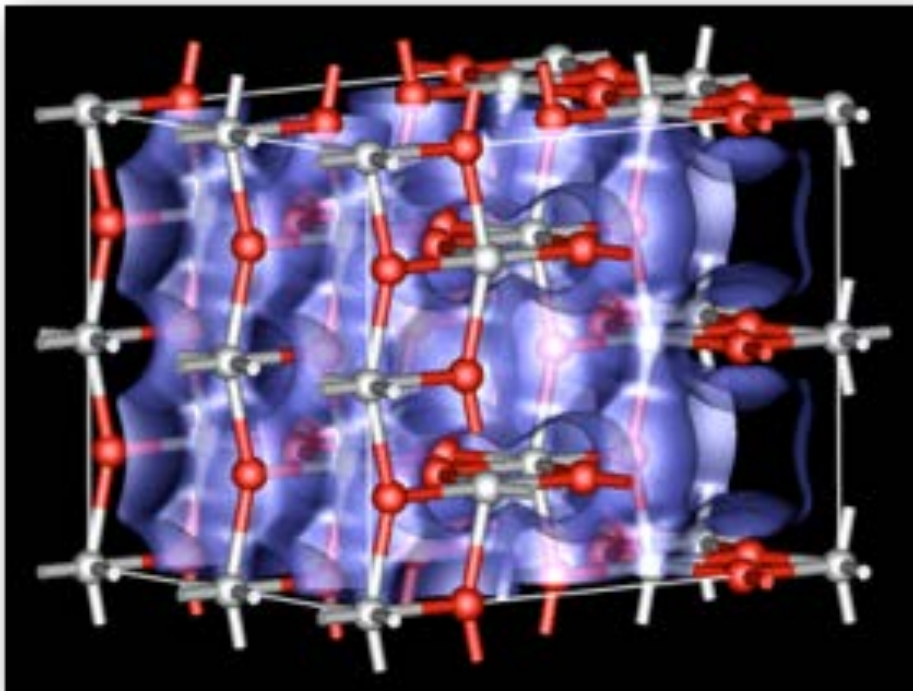

Argonne NATIONAL LABORATORY

ENERGY SECURITY: COMBUSTION MODELING

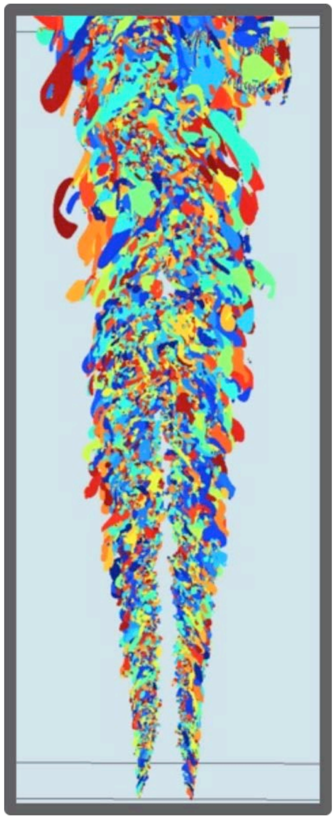

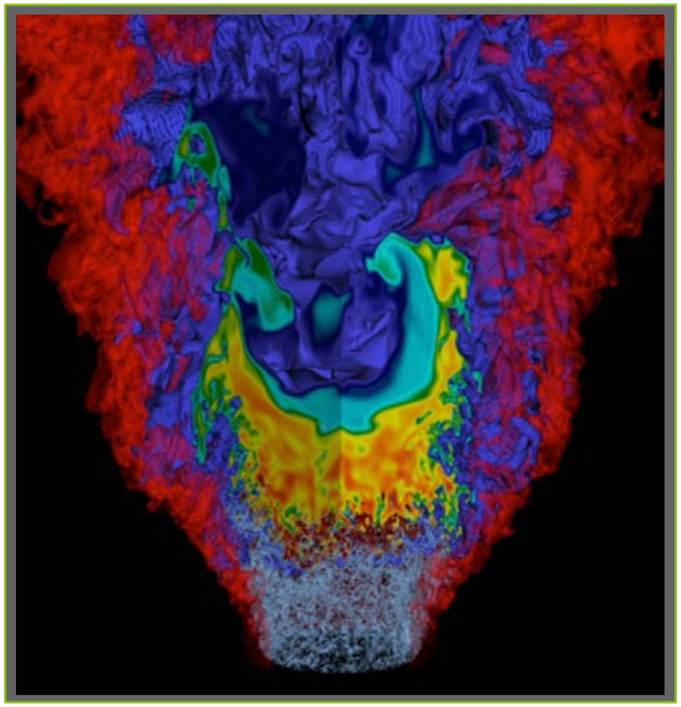

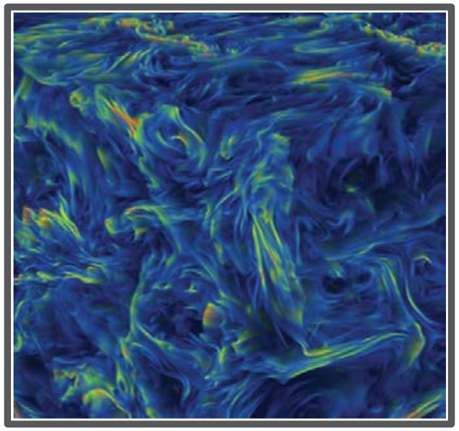

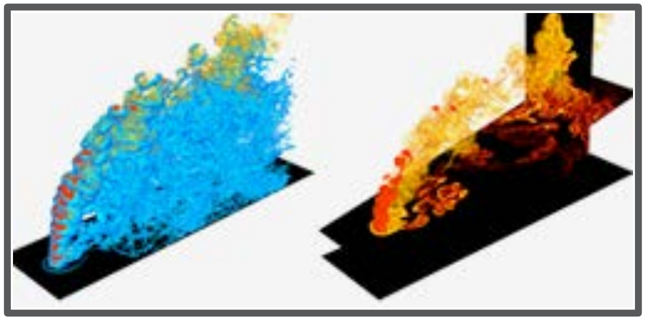

Argonne NATIONAL LABORATORY

HEALTH CARE:
PRECISION MEDICINE FOR CANCER

ECONOMIC SECURITY: ADDITIVE MANUFACTURING

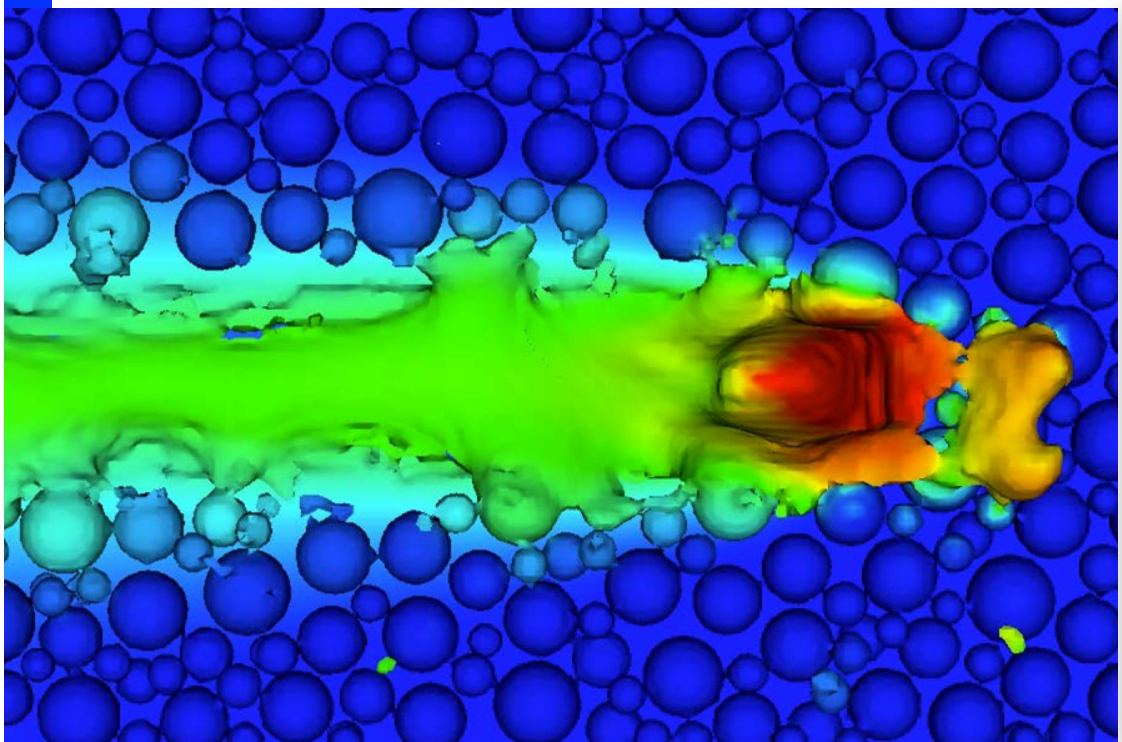

EARTH SYSTEM: CLOUD-RESOLVING CLIMATE MODELING

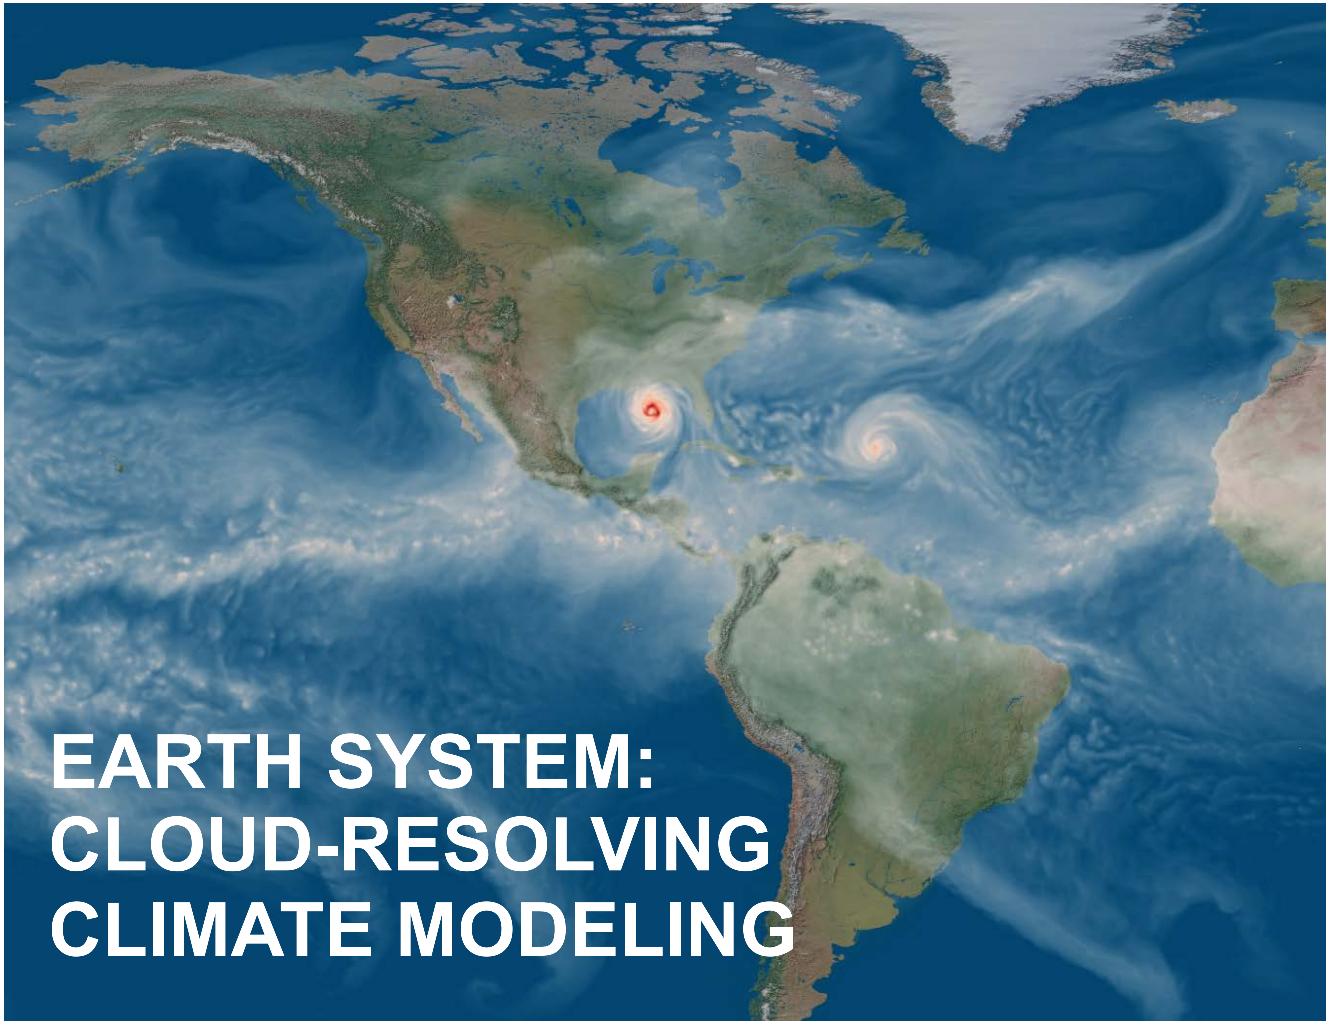

ENERGY SECURITY

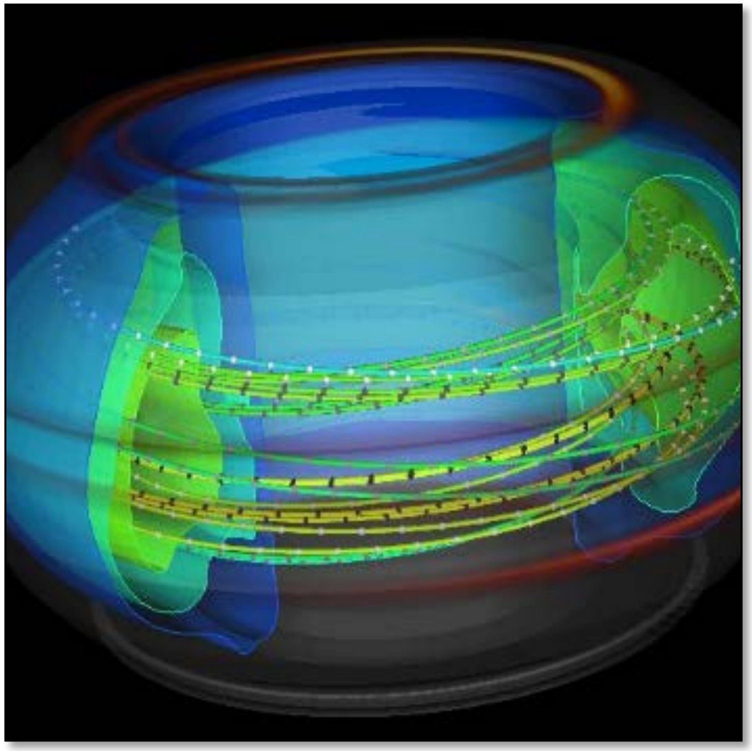

Magnetic Fusion Energy

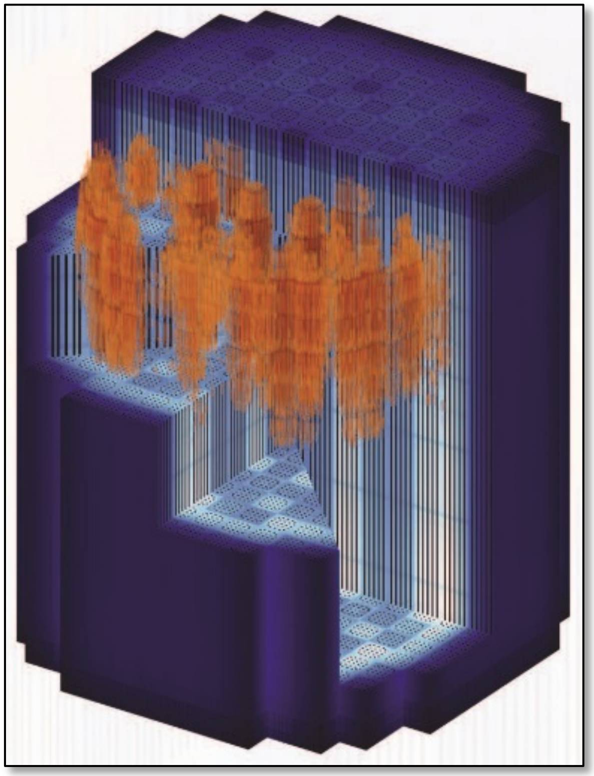

Nuclear Energy

Argonne NATIONAL LABORATORY

SCIENTIFIC DISCOVERY: ASTROPHYSICS

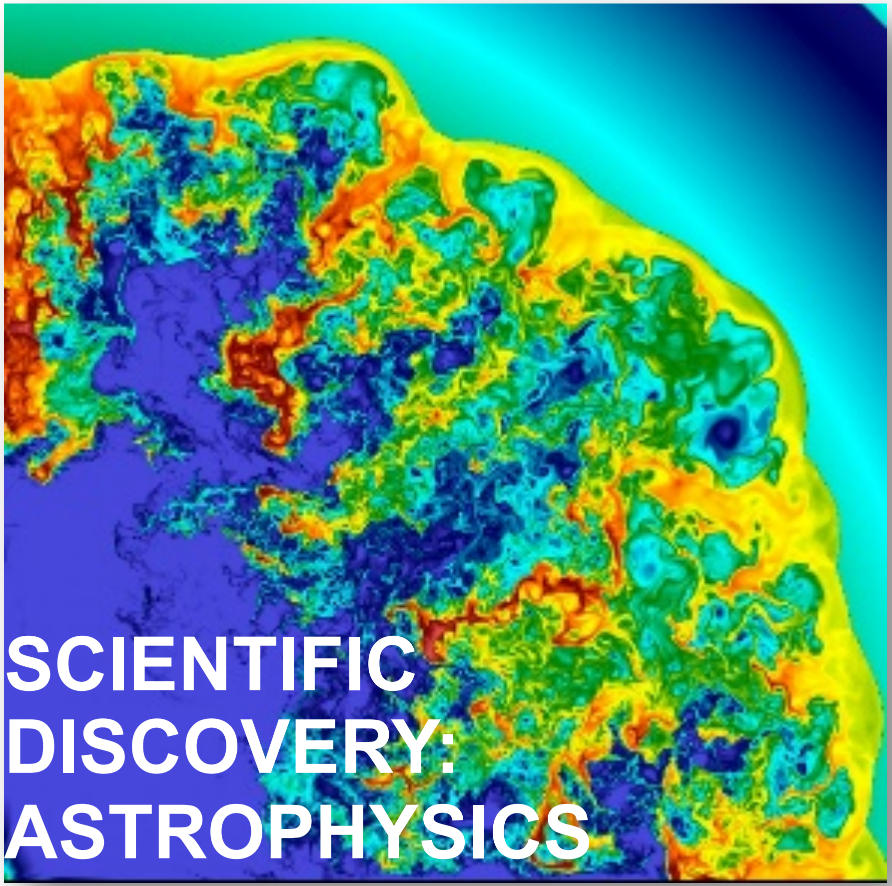

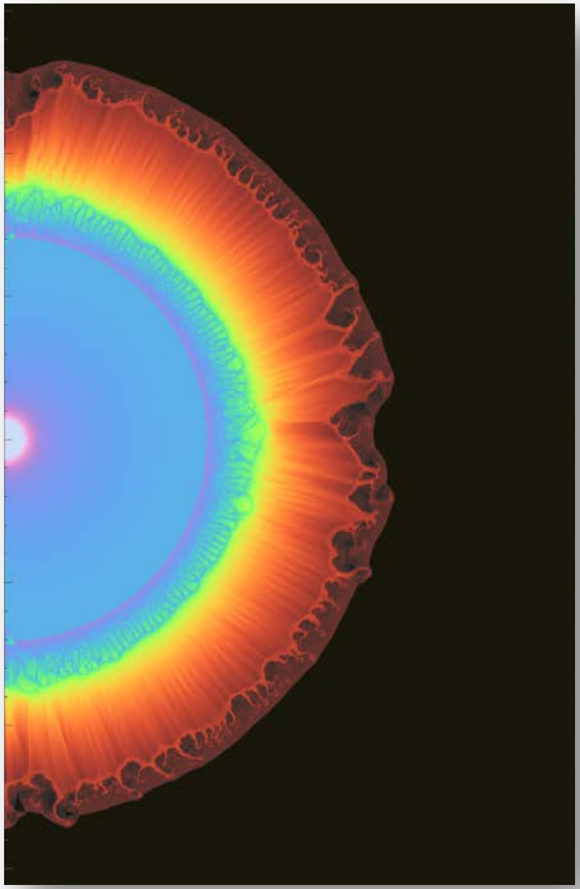

Argonne NATIONAL LABORATORY

ENERGY SECURITY: WIND ENERGY

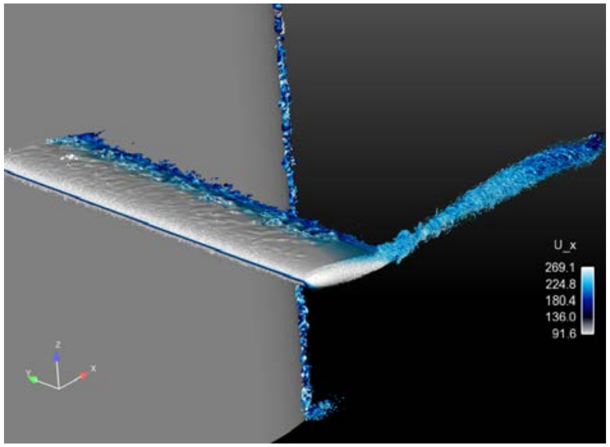

Argonne NATIONAL LABORATORY

THE U.S. EXASCALE COMPUTING PROJECT (ECP)

THE ECP AIMS TO CREATE AN EXASCALE ECOSYSTEM THAT WILL

- Enable classical simulation and modeling applications to tackle problems that are currently out of reach,
- Enable new types of applications to utilize exascale systems, including ones that use machine learning, deep learning, and large-scale data analytics,
- Support widely used programming models as well as new ones that promise to be more effective on exascale architectures or for applications with new computational patterns, and
- Be suitable for applications that have lower performance requirements currently, thus providing an on ramp to exascale should their future problems require it.

ECP AREAS OF TECHNICAL FOCUS

Application Development

The Application Development effort develops and enhances the predictive capability of applications critical to the DOE, including the science, energy, and national security mission space. The scope of the AD focus area includes

- targeted development of requirements-based models, algorithms, and methods,
- integration of appropriate software and hardware via co-design methodologies,
- systematic improvement of exascale system readiness and utilization, and
- demonstration and assessment of effective software integration.

Software Technology

Software Technology spans low-level operational software to high-level applications software development environments, including the software infrastructure to support large data management and data science for the DOE SC and NNSA computational science and national security activities at exascale. Projects will have:

- line of sight to application's efforts
- inclusion of a Software Development Kit to enhance the drive for collaboration, and
- delivery of specific software products across this focus area.

Hardware and Integration

This focus area is centered on the integrated delivery of specific outcomes (ECP Key Performance Parameters, or KPPs) and products (e.g., science as enabled by applications, software, and hardware innovations) on targeted systems at leading DOE computing facilities. Areas include:

- PathForward
- Hardware Evaluation
- Application Integration at Facilities
- Software Deployment at Facilities
- Facility Resource Utilization
- Training and Productivity

Argonne NATIONAL LABORATORY

ECP IS A COLLABORATION AMONG SIX US DOE NATIONAL LABORATORIES

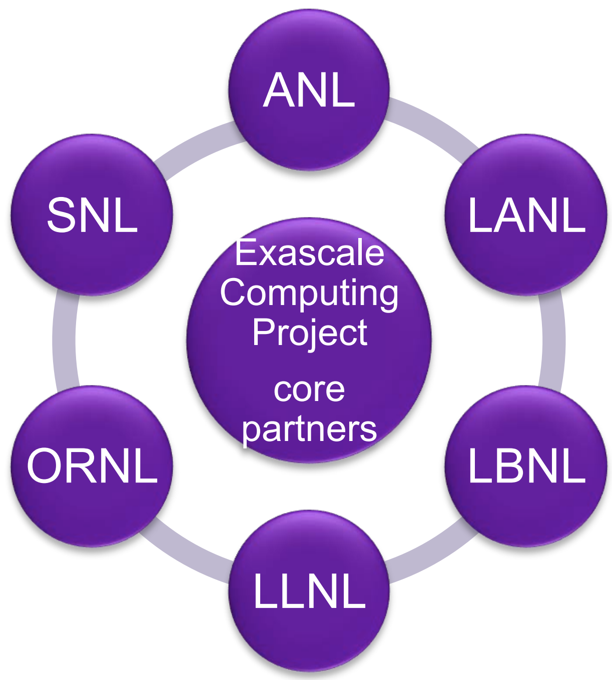

But the work is carried out at many institutions

ECP
EXASCALE COMPUTING PROJECT

DOE LABORATORIES & AGENCY PARTNERS
22

PRIVATE SECTOR PARTNERS
10

UNIVERSITY RESEARCH PARTNERS
39

INDUSTRY COUNCIL MEMBERS
18

THE ECP ECOSYSTEM

- 800+ Researchers
- 25 Application Development Projects
- 66 Software Development Projects
- 5 Co-Design Centers

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

MEMBERS OF THE ECP INDUSTRY COUNCIL

United Technologies

Westinghouse

Lilly

GM

Cummins

ExxonMobil

CASCADE TECHNOLOGIES

FedEx

GE

ANSYS

MARS

Chevron

Altair

P&G

TRI ALPHA ENERGY
THE POWER OF INGENUITY

Dow

Whirlpool CORPORATION

DREAMWORKS

ECP OUTCOME

- ***Accelerated delivery of a capable exascale computing ecosystem*** to provide breakthrough solutions addressing our most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security
  - ***Capable***: Wide range of applications can effectively use the systems developed through ECP, ensuring that both science and security needs will be addressed **(affordable, usable, useful)**
  - ***Exascale***: Deliver 50× the performance of today's 20 PF systems, supporting applications that produce high-fidelity solutions in less time and address problems of greater complexity
  - ***Ecosystem***: Not just more powerful systems, but all methods and tools needed for effective use of ECP-enabled exascale systems to be acquired by DOE labs

DOE SUPERCOMPUTING FACILITIES WILL ACQUIRE AND OPERATE THE EXASCALE SYSTEMS – NOT THE ECP

- The mission of those facilities is to provide access to the most powerful computing systems at any given time
- Hence the ECP is working closely with the DOE supercomputing facilities to ensure that the systems meet the mission needs
- Current plans call for delivery of the first exascale system to Argonne National Laboratory in 2021, with additional exascale systems to follow at other SC and NNSA laboratories in 2021-2022 and over the next several years to meet identified mission needs

ARGONNE LEADERSHIP COMPUTING FACILITY 2021 EXASCALE SUPERCOMPUTER – A21

Intel Aurora supercomputer planned for 2018 shifted to 2021

Scaled up from **180 PF to over 1000 PF**

Support for three “pillars”

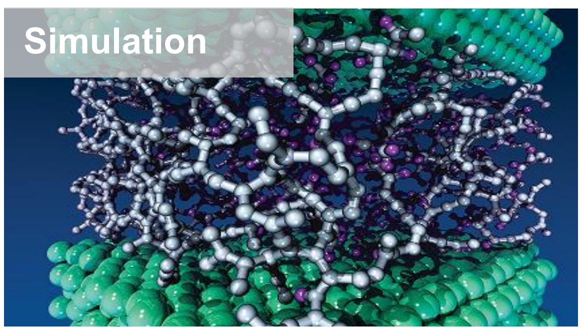

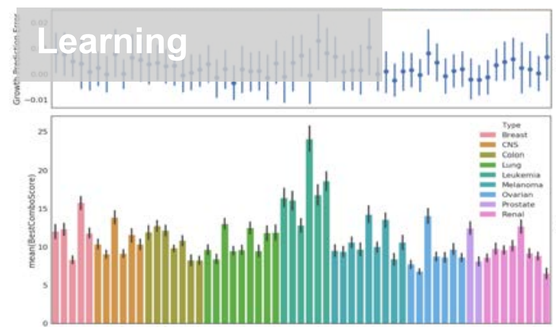

ECP APPLICATION DEVELOPMENT (AD)

Chemistry and Materials Applications	Describes underlying properties of matter needed to optimize and control the design of new materials and energy technologies
Energy Applications	Model and simulation of existing and future technologies for the efficient and responsible production of energy to meet the growing needs of the U.S.
Earth and Space Science Applications	Spans fundamental scientific questions from the origin of the universe and chemical elements to planetary processes and interactions affecting life and longevity
Data Analytics and Optimization Applications	Applications partially based on modern data analysis and machine learning techniques rather than strictly on approximate solutions to mathematical equations
National Security Applications	Stewardship of the US nuclear stockpile, assessment of future threats; related physics and engineering modeling and scientific inquiries consistent with that mission space
Co-Design	Focused on crosscutting algorithmic methods that capture the most common patterns of computation and communication in ECP applications

ECP APPLICATIONS: EXASCALE-CAPABLE MODELING, SIMULATION, DATA

Goal
Develop and enhance the predictive capability of applications critical to DOE across science, energy, and national security mission space

- Targeted development of requirements-based methods
- **Integration of software and hardware via co-design methodologies**
- Systematic improvement of exascale readiness and utilization
- **Demonstration and assessment of effective software integration**

ECP SOFTWARE: PRODUCTIVE, SUSTAINABLE ECOSYSTEM

Goal
Build a comprehensive, coherent software stack that enables application developers to productively write highly parallel applications that effectively target diverse exascale architectures

- Extend current technologies to exascale where possible
- Perform R&D required for new approaches when necessary
- Coordinate with and complement vendor efforts
- Develop and deploy high-quality and robust software products

Programming Models and Runtime

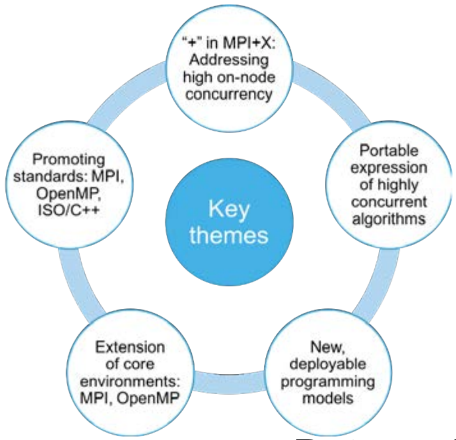

Development Tools

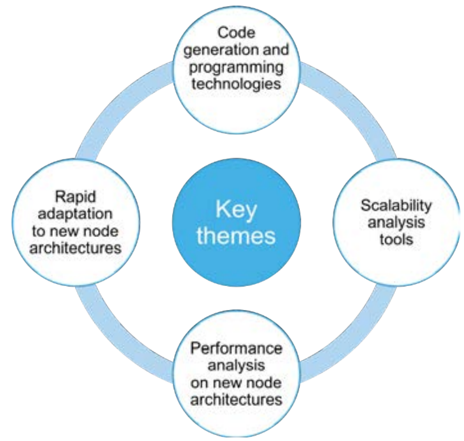

Math Libraries

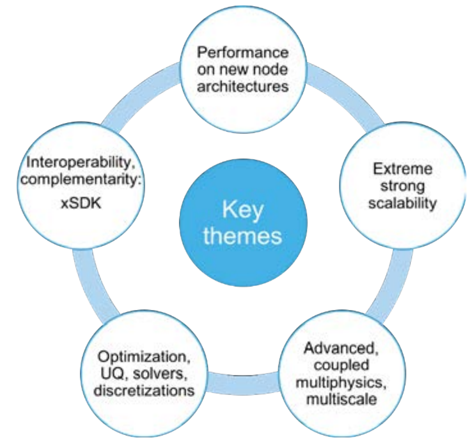

Data and Visualization

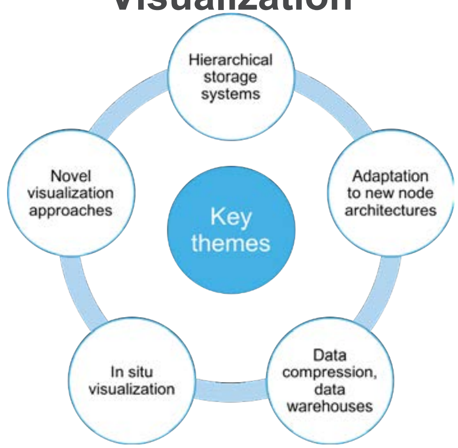

Software Ecosystem and Delivery

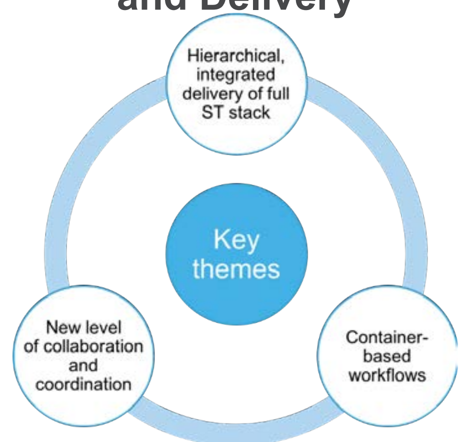

ECP HARDWARE AND INTEGRATION: DELIVERY OF INTEGRATED ECP/DOE FACILITY PRODUCTS

Goal
A capable exascale computing ecosystem made possible by integrating ECP applications, software and hardware innovations within DOE facilities

- Innovative supercomputer architectures for competitive exascale system designs
- Accelerated application readiness through collaboration with the facilities
- Integrated and **continuously tested** exascale software ecosystem deployed at facilities through collaboration with facilities
- **Training on key ECP technologies**, help in accelerating the software development cycle and in optimizing the productivity of application and software developers

PATHFORWARD

Objective

- Three to six supercomputer vendors competing for upcoming DOE exascale system acquisitions

Approach

- Awarded six commercial R&D contracts in early 2017
  - AMD, Cray, IBM, Intel, HPE, and NVIDIA
  - Each includes at least a 40% cost share by the awardee
- Direct R&D to maximize energy efficiency while improving the performance of DOE extreme-scale applications
- Guide the R&D through interaction of lab experts with the contractors to the mutual benefit of DOE and the US commercial supercomputing industry
- Monitor and track vendor progress

CHALLENGES THEN AND NOW

CHALLENGES: THEN AND NOW

- 10 years ago
  - Identify a spectrum of important applications that would benefit from exascale computing
  - Energy efficiency
  - Billion-way concurrency

Argonne NATIONAL LABORATORY

SCIENTIFIC GRAND CHALLENGE WORKSHOPS HELD IN 2008-2010 IDENTIFIED ADVANCES ENABLED BY EXASCALE AND BEYOND

Reports available at

http://extremecomputing.labworks.org/index.stm

CHALLENGES: THEN AND NOW

- 3- 5 years ago
  - Communicating with hardware vendors to develop exascale systems that are programmable, efficient and affordable
  - How to structure the project
  - How much innovation to pursue
  - Top 10 research challenges

TOP 10 EXASCALE RESEARCH CHALLENGES
2014 REPORT

- Energy efficiency
- Interconnect technology
- Memory technology
- Scalable system software
- Programming systems
- Data management
- Exascale Algorithms
- Algorithms for discovery, design, and decision
- Resilience and correctness
- Scientific productivity

KEY HIGH-LEVEL TECHNICAL CHALLENGES THAT MUST BE TACKLED TO ACHIEVE EXASCALE

- *Massive Parallelism* – 100+ times greater than today's largest systems
- *Memory and Storage* – effective use of many levels of hierarchy
  - Memory and storage efficiencies consistent with increased computational rates and data movement requirements
- *Reliability* – system adaptation and recovery from faults in complex system components and designs
- *Energy Consumption* – Energy consumption reduced beyond current industry roadmaps
  - would be prohibitively expensive at this scale
  - hardware and software techniques for minimizing it

SELECTED APPLICATIONS CHALLENGES (1)

- Exploiting on-node memory and compute hierarchies
- Programming models: which to use
- Integrating software components that use disparate approaches (e.g., for on-node parallelism)
- Developing and integrating co-designed motif-based community components
- Achieving good performance on complex new architectures
- Achieving portable performance (without “if-def’ing” 2 different code bases)

SELECTED APPLICATIONS CHALLENGES (2)

- Multi-physics coupling: both algorithms and software
- Integrating sensitivity analysis, data assimilation, uncertainty quantification technologies
- Guiding the vendors with our understanding of our applications and support software requirements while workload is changing
  - Difficult to design proxy applications that represent the key features of the full applications
  - Vendors can't deal with too many full or proxy applications
- Understanding requirements of data analytic and machine learning applications – moving target

Argonne NATIONAL LABORATORY

CHALLENGES FOR SOFTWARE TECHNOLOGY DEVELOPERS (1)

- Ecosystem: Large, complex, evolving SW environment
  - **In particular, coordinating with the vendors' software stack**
- Billion-way concurrency
  - Several novel compute node architectures
- Programming models and languages
  - How many are viable to support?
  - Which ones to support in this project for existing applications?
  - What models will new applications enabled by exascale need?
- Coupled applications: Physics, scales, in situ data, more
- Data-driven: New software HPC environments, containers

CHALLENGES FOR SOFTWARE TECHNOLOGY DEVELOPERS (2)

- How to meet needs of new types of applications or existing applications whose needs evolve due to exascale hardware features
- Getting software building blocks (e.g., math libraries) ready early enough that application teams can take advantage of them
- Some ST projects have so many application teams that want deliverables from them that they can't keep up
  - In contrast to fear that applications would not use ECP software products

Argonne NATIONAL LABORATORY

CHALLENGES: THEN AND NOW

- Now
  - Energy efficiency
  - Memory technology
  - Packaging and testing software during development
  - Co-design at such a large scale
  - Continuous Integration
  - Delivering software that works on existing platforms and that production facilities will install
  - Developing and delivering exascale software that production computing facilities will install and in some cases maintain
  - Sustainability

Argonne NATIONAL LABORATORY

ECP SOFTWARE STACK

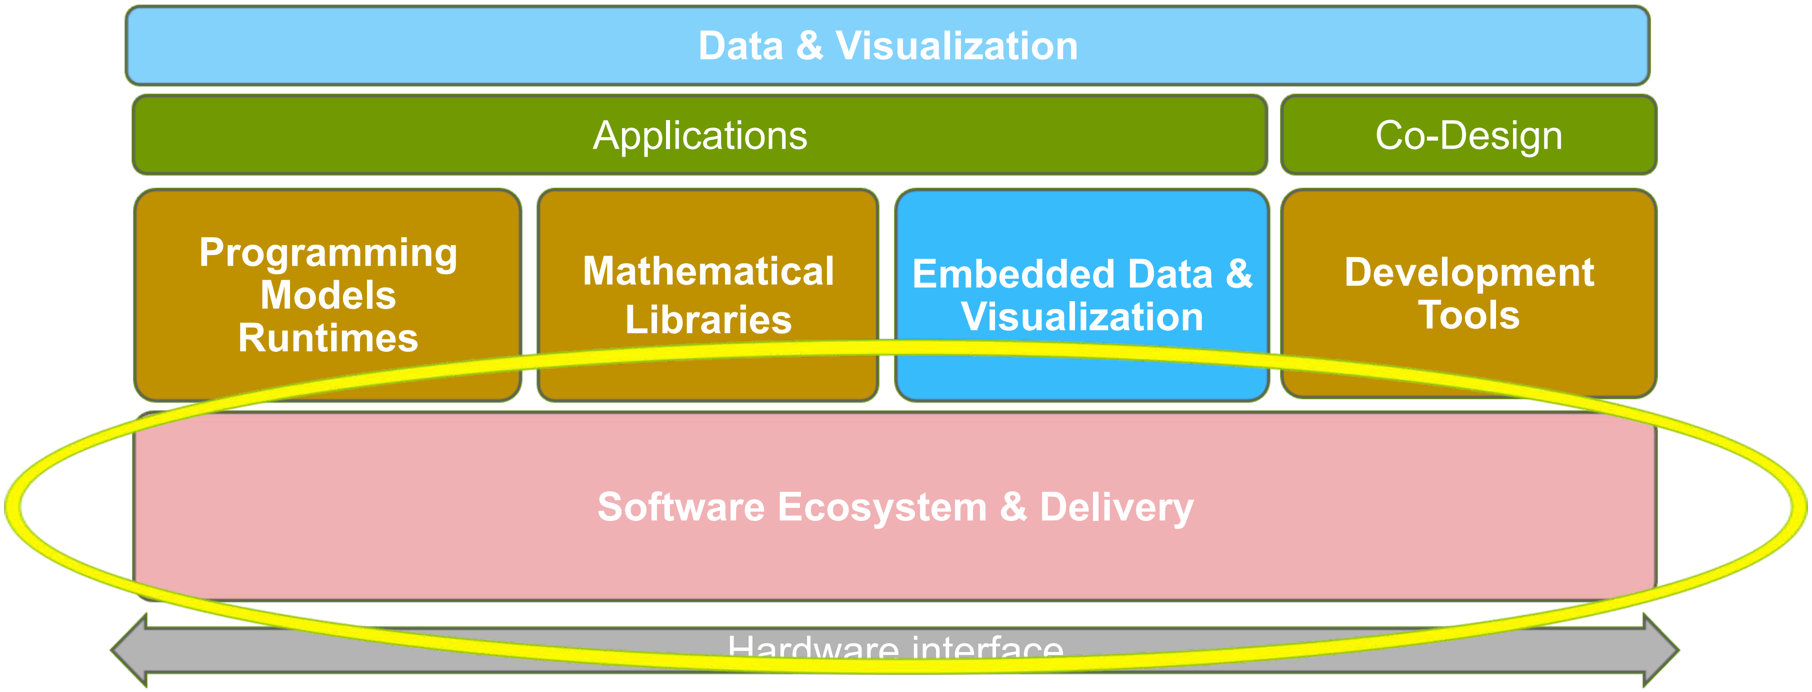

Argonne NATIONAL LABORATORY

BIGGEST CHALLENGE: CREATING AN ENDURING EXASCALE COMPUTING ECOSYSTEM

- Multiple commercial exascale systems
  - Scalable from small to exascale
- The system software for those systems, most of it the same from the user's standpoint
- The software building blocks for applications, including the APIs
- The application software, sustainable, open source (mostly)
- A software stack that can be used from the laptop to the exascale
- **The people with the expertise to develop, maintain, and use the ecosystem**
- All aspects of a particular supercomputing ecosystem, be they hardware, software, algorithms, or people, must be strong if the ecosystem is to function effectively

Argonne NATIONAL LABORATORY

ECP SOFTWARE STACK, DELIVERY AND SUSTAINABILITY

ECP software projects
Each project establishes release channels: SDKs, vendor/3rd party, Direct to Facilities

SDKs*

Reusable software libraries embedded in applications; cohesive/interdependent libraries released as sets modeled on xSDK

- Regular coordinated releases
- Hierarchical collection built on Spack
- Products may belong to >1 SDK based on dependences
- Establish community policies for library development
- Apply Continuous Integration and other robust testing practices

Math SDK
Tools SDK
PM&RT SDK
DataViz SDK
Facility SDK

Vendors/3rd Parties

Potential strategy for binary distributions.

- **Potential increased incorporation into HPC vendor stacks, OpenHPC stack.**
- Develop scalable releases targeting higher end systems

Direct2Facility

Platform-specific software in support of a specified 2021–2023 exascale system

- Software **exclusively** supporting a specific platform
- System software, some tools and runtimes

Assume all releases are delivered as "build from source" via Spack – at least initially

Focus on ensuring that software compiles robustly on all platforms of interest to ECP (including testbeds)

*Extreme-scale Scientific Software Development Kit (xSDK)

Argonne NATIONAL LABORATORY

SUMMARY

- It is essential to focus early in the development cycle on delivery and integration – "The last kilometer"
- Start collaborating as soon as possible with the computer facilities that will buy and operate the exascale systems
- Collaborate as much as possible with the system integrators on the software ecosystem

THANK YOU!

www.anl.gov

Argonne NATIONAL LABORATORY