HOW CAN LESSONS LEARNED IN THE PAST FORTY YEARS GUIDE FUTURE HPC RESEARCH STRATEGIES?

PANELISTS

- Buddy Bland, Oak Ridge National Laboratory
  - Computing facilities and their impact on computational science, computer architectures and computer science
- Jackie Chen, Sandia National Laboratories
  - Computational science
- Phil Colella, Lawrence Berkeley National Laboratory
  - Applied mathematics
- Eli Dart, Lawrence Berkeley National Laboratory
  - Network research and facilities
- Jack Dongarra, University of Tennessee and Oak Ridge National Laboratory
  - Software libraries
- Thom Dunning, University of Washington and Pacific Northwest National Laboratory
  - Computational science
- Wendy Huntoon, Keystone Initiative for Network Based Education and Research
  - Network technology, facilities and applications
- Paul Messina, Argonne National Laboratory
  - Moderator
- Dan Reed, University of Utah
  - Computer science, supercomputer facilities

CONTEXT FOR THIS PANEL

- The HPC, computer science and mathematics communities face disruptive changes due to
  - the end of Dennard scaling,
  - the looming end of Moore's Law and
  - the emergence of quantum, neuromorphic and other technologies

- Having successfully weathered previous disruptions, are there lessons learned and insights into effective investment strategies and guidance in research investments in the coming decades?

TOPICS FOR DISCUSSION

- Drawing on the experience of the Advanced Scientific Computing Research program of the U.S. Department of Energy, as well as other agencies, the panelists will address the following questions and seek your opinions and insights:
  - Which programmatic techniques, including partnering with other agencies or industry, were most successful and might translate to the future?
  - What are examples of investment failures that might have been avoided?
  - What lessons have been learned to help adapt to the end of Moore's Law and the likely shift in software architecture that will be needed?
  - What has been the impact of supercomputing facilities and computer networks and will such facilities continue to be relevant?

BACKGROUND: THIS PANEL WAS INSPIRED BY A CHARGE FROM THE DOE OFFICE OF SCIENCE TO WRITE A REPORT THAT ADDRESSES THESE QUESTIONS

- What are the major scientific accomplishments that have shaped the ASCR-supported disciplines in the last 40 years? How has ASCR contributed to those advances?
- What impacts have those accomplishments had on the Department's missions in energy, environment, or security?
- What are the key aspects of the ASCR's investment strategy that have had the greatest impacts?
- Looking to the future, and building on the ASCAC reports, what research areas and funding strategies to pursue those areas could further strengthen ASCR in serving the DOE missions?

In this panel we would like to draw on your experiences and expertise -- you, the audience – to answer these questions broadly

THE HPC ECOSYSTEM IS VERY BROAD

- Computational science
- Applied mathematics
- Numerical libraries
- Computer science
- Computer architecture
- Supercomputing facilities
- Networking facilities
- Industry applications
- Workforce
- Education and training

THE HPC ECOSYSTEM UNDERWENT MAJOR CHANGES IN THE LAST 40 YEARS

- We have an HPC ecosystem that has adapted to many changes in computer architectures and types of applications
- The adaptation required changes in applications, mathematical models and algorithms, software environment and tools, computing and storage facilities, networks, computer architectures and technologies, …
- and the people who carry out the work and how they work: more and more, in multidisciplinary teams

EXAMPLES OF LESSONS LEARNED WE HAVE IDENTIFIED

DO YOU IN THE AUDIENCE AGREE?

- Unsolicited proposals often led to accomplishments with big impact, sometimes over a period of decades
- Block/base funding over long periods likewise has been instrumental in achieving numerous breakthroughs
- Multi-agency funding and collaborations have been fruitful, sometimes essential for involving players with the necessary expertise and facilities

EXAMPLES OF LESSONS LEARNED WE HAVE IDENTIFIED

DO YOU IN THE AUDIENCE AGREE?

- Integration of the ecosystem elements has been essential for successfully transitioning to major technology changes in
  - Computer architectures
  - Multiphysics applications
  - Multi-institution collaborations
  - Workflows, data-centric computing
- Open software policy and type of license have played important roles

SCIENCE RESULTS AND WHAT ENABLES THEM:

AN EXAMPLE BY JACKIE CHEN

Combustion Science and Technology

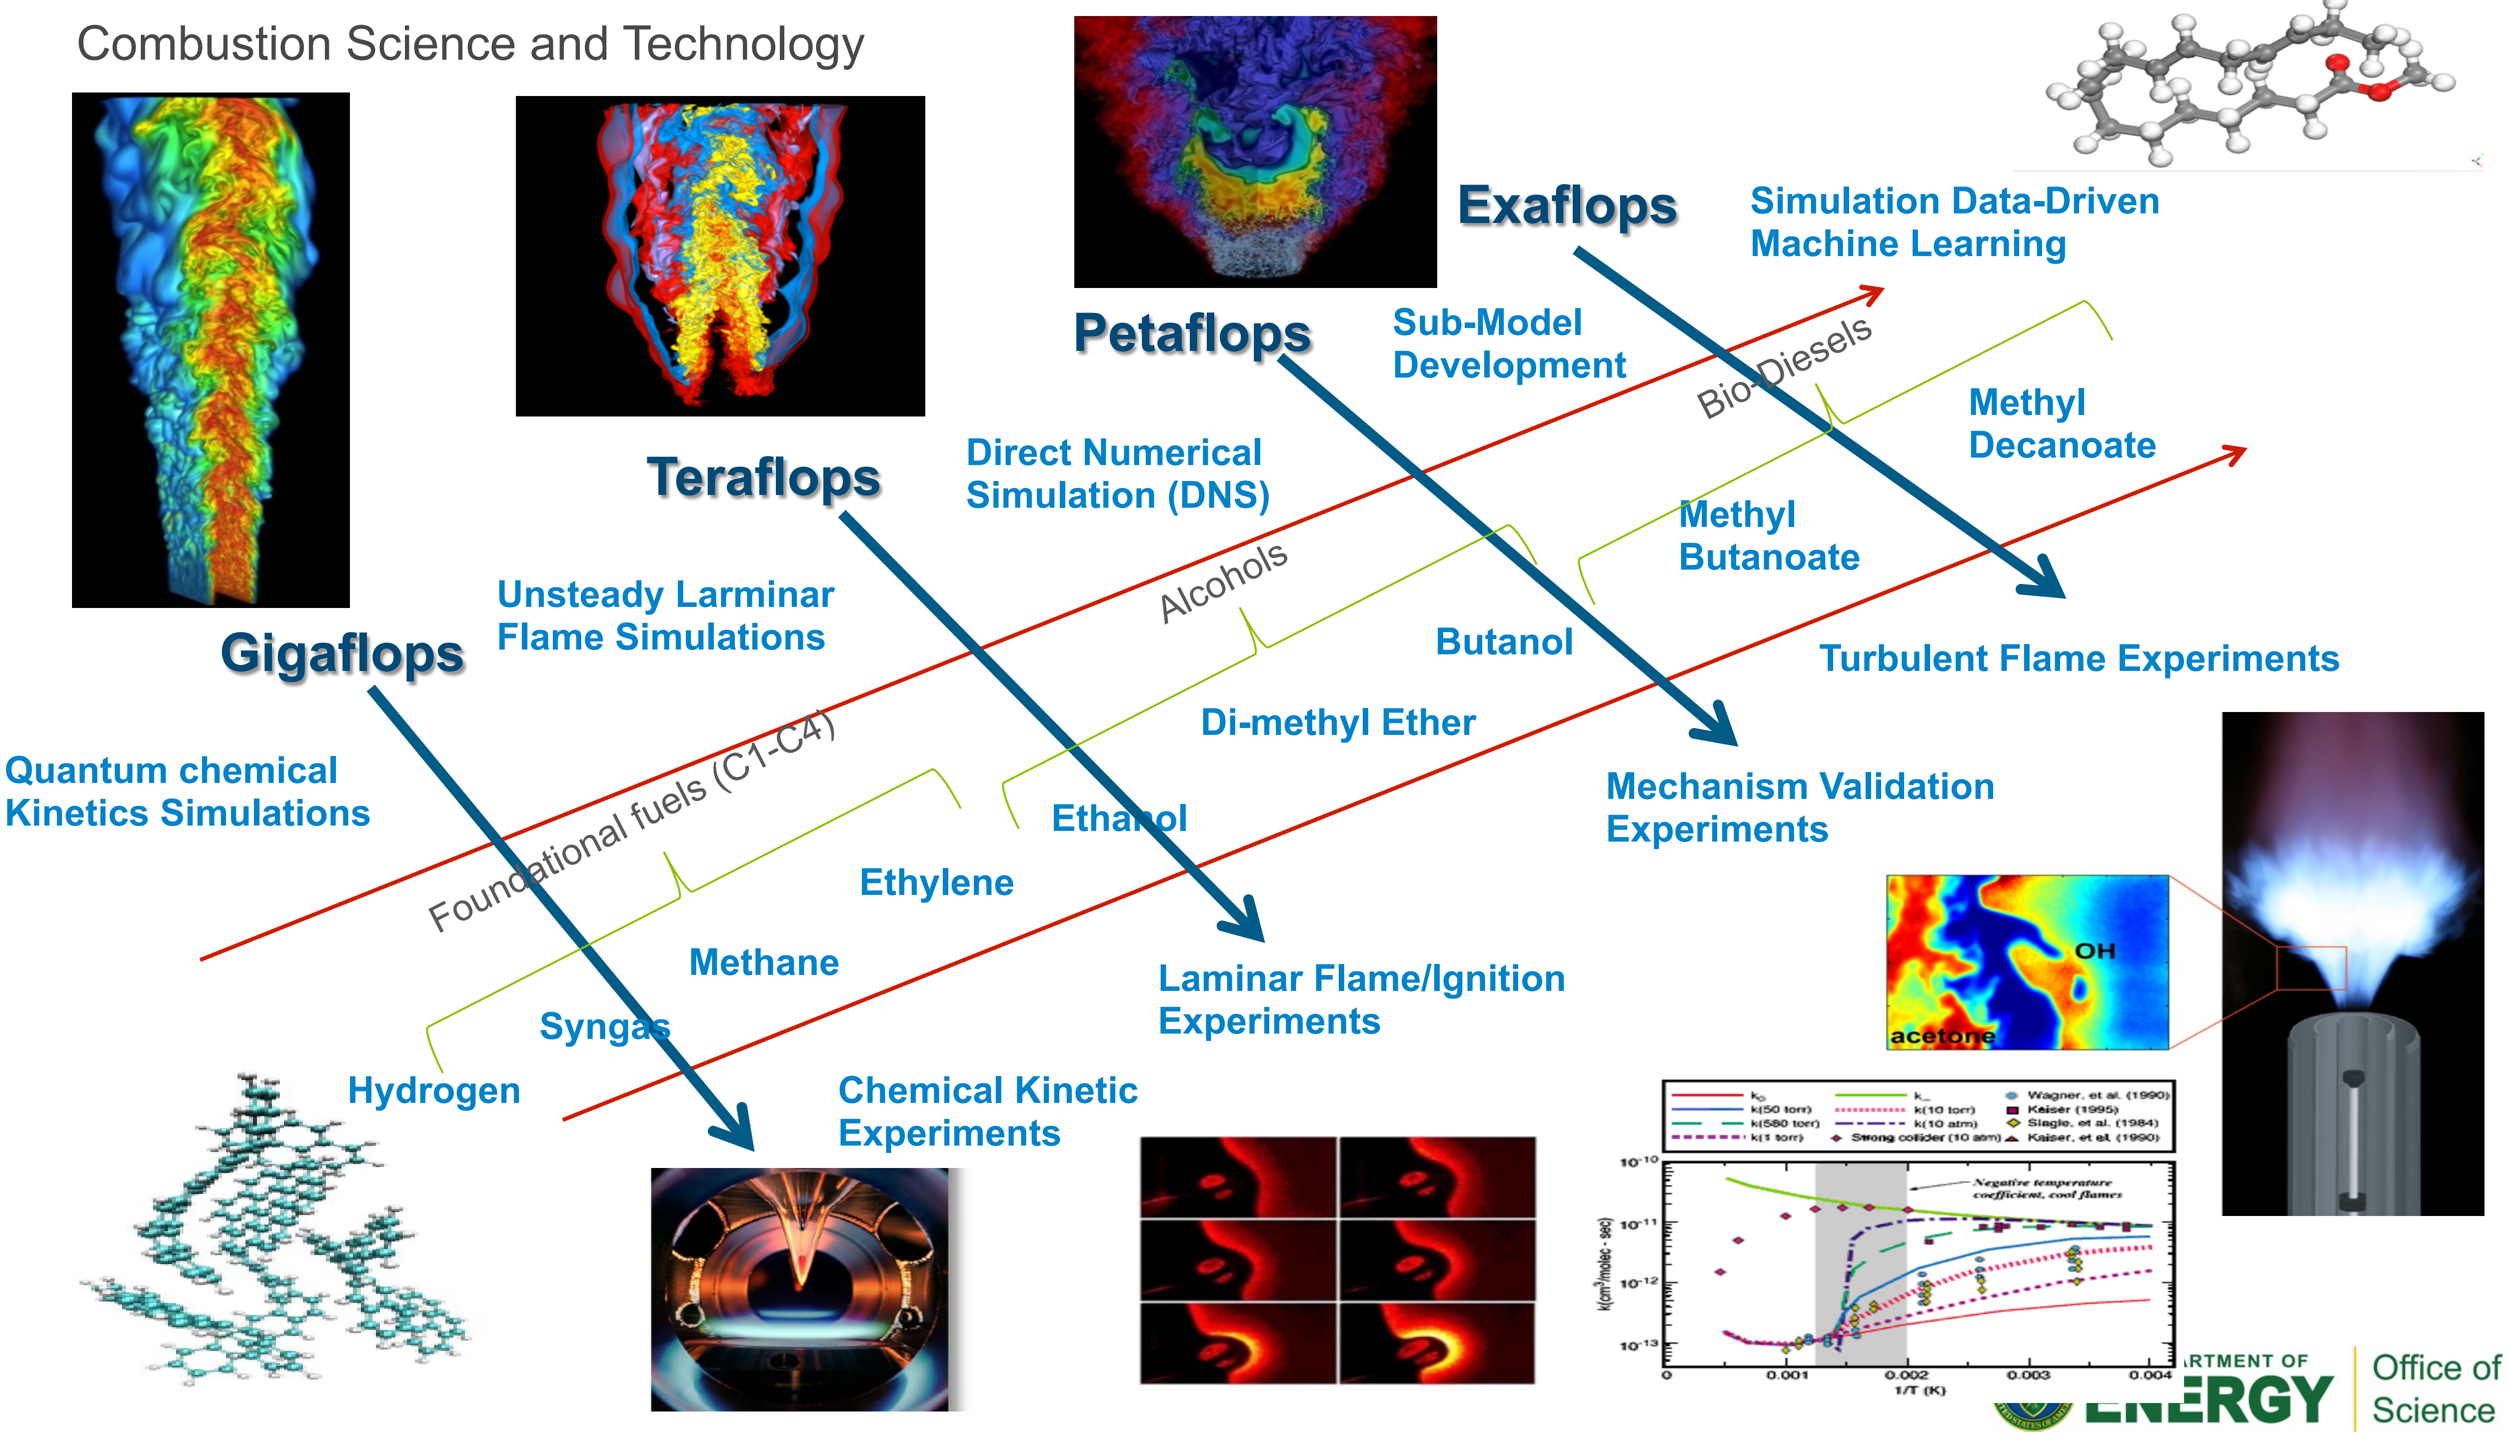

Office of Science

COMBUSTION SIMULATION AT EXTREME SCALE ENABLED BY CODESIGN THROUGH ASCR: SCIDAC, APPLIED MATH, ECI & ECP, INCITE AND LEADERSHIP COMPUTING

SCIDAC

SDAV: FastBit Indexing, Kepler Workflow, ADIOS, Topological Feature

Ultrascale Visualization USV:
In situ Visualization

MMICCS

PHILMS: Physics Informed Learning Machines for Multiscale Prob.

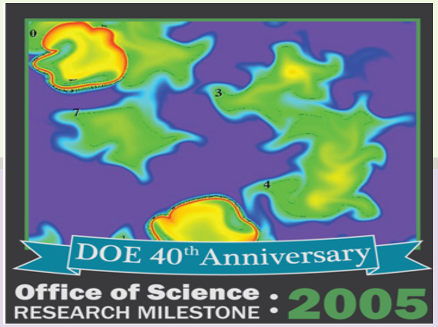

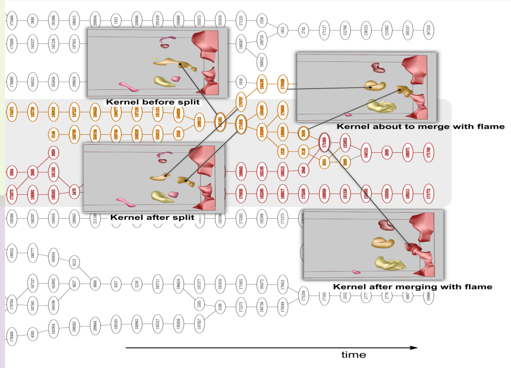

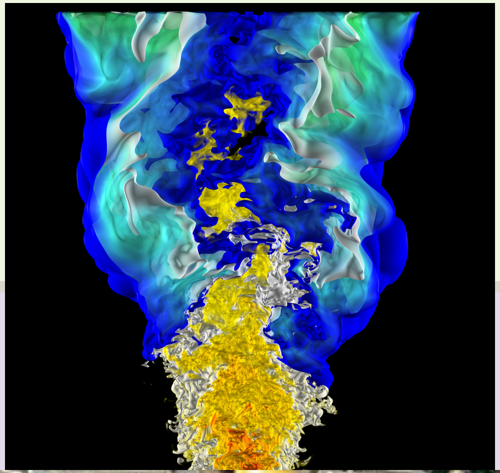

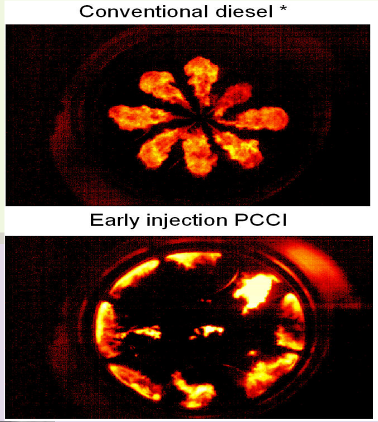

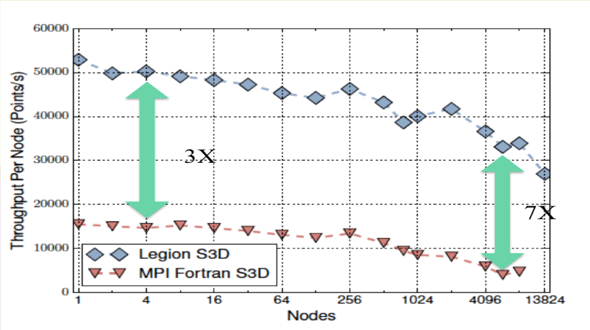

ExaCT Exascale Co-Design: Legion, AMR, Adjoint UQ, Analytics workflow

Exascale Computing Initiative and Project

Exascale Computing Project: Pele AMR codes with multiphysics and embedded boundary

U.S. DEPARTMENT OF ENERGY | Office of Science

MATHEMATICS' ROLES IN HPC

PHIL COLELLA, PRESENTED BY PAUL MESSINA

MATHEMATICS IS CENTRAL TO THE HPC ECOSYSTEM

- Mathematical models: describing science in a way that is amenable to representation of a computer.
- Approximation / discretization: replacing an infinite number of degrees of freedom with a finite number of degrees of freedom.
- Solvers and software: accurate and efficient representation on high-end computers.

Deep mathematical questions are associated with each of these topics, and they strongly interact with one another. A major accomplishment for ASCR has been resolving these questions for a broad range of important applications.

EXAMPLES OF MATHEMATICAL ACCOMPLISHMENTS

- Mathematical models: Low Mach number asymptotics for reacting fluid flow, well-posedness of shock hydrodynamics.
- Approximation / Discretization: high-order methods for PDE, adaptive mesh refinement, level-set methods for front propagation, particle methods for fluids and plasmas, simulations in complex geometries.
- Solvers and software: broadly-used solvers and frameworks for HPC simulation.

Scientific impacts both in DOE mission areas (combustion, fusion, defense applications, …) and beyond (aerodynamics, astrophysics / cosmology, semiconductor process design…).

HOW HAS MATH BEEN SUCCESSFUL ?

- There are tradeoffs up and down the stack of models, discretization, and software, requiring actors that can look beyond their stovepipes.
- Experts on crosscutting mathematical technologies working to solve a specific scientific problem often find it leads to ideas and tools that are applicable to a broad range of problems.
- Long lead times (10 years or more) from first appearance of a new idea to its broad acceptance in the scientific community requires patience from funding sources.

LINEAR ALGEBRA SOFTWARE'S IMPACT
:
JACK DONGARRA

THE PACKS FOR DENSE LINEAR ALGEBRA
EISPACK, LINPACK, THE BLAS, LAPACK, AND SCALAPACK

- EISPACK 1972
  - Effort to translate Algol to Fortran
- LINPACK 1978
  - Designed to use Level 1 BLAS (vector ops)
- Level 2 & 3 BLAS 1988 & 1990
  - Higher level of granularity for matrix operations; GEMV and GEMM
- Sca/LAPACK 1992
  - Refactored to perform block algorithms for performance on “modern” architectures

– PACKs

  - Reliability, Robustness, Structure, Usability, Portability, and Validity

- LAPACK is one of the most successful and influential software packages produced under the DOE‘s ASCR program.
  - Used by 13+ HW vendors in their software libraries
  - Used by 17+ SW companies in their products

- Architectural changes have come every decade or so, thereby creating a need to refactor the software to the emerging architectures.
- PACK's became joint efforts with DOE ASCR and other organizations.
- PACK's widely embraced and helped form standards (e.g., BLAS, IEEE floating point standard, and MPI).
- DOE ASCR provided a stable organization with decades-long lifetime which can maintain and evolve the software.
  - The ASCR DOE laboratories are critical elements in the development and support of HPC software.
  - Funding streams of an academic institution or the NSF supercomputing centers is too unstable.
- Important mechanism in this story was the existence of a block grant to do cutting-edge computer science at a the DOE laboratories.

– Success of PACKs in part based on the choice of the software license used.

RESEARCH COMPUTERS AND COMPUTER SCIENCE:

DAN REED

EXEMPLAR RESEARCH COMPUTERS

- Caltech Cosmic Cube (Chuck Seitz/Geoffrey Fox)
  – Lattice QCD calculations using commodity microprocessors
- Illinois Cedar multiprocessor (David Kuck)
  – Shared memory computing and compiler technology
- NYU Ultracomputer (Jack Schwartz)
  – Shared memory fetch & add intrinsics
- Argonne Advanced Computing Research Facility
  – Parallel computing research experiments plus education
  – Alliant FX/8, Denelcor HEP, Encore Multimax
  – Sequent Balance, Intel iPSC, AMT DAP

- Ecosystems that shaped computing
  – Vendor relations and new computer architecture co-design
  – Software systems and broad uptake and deployment
  – Staff and researcher training and education

EXEMPLAR COMPUTER SCIENCE RESEARCH

- Programming models, languages, and tools
  - PGAS languages and MPI standards
  - Vectorizing and parallelizing compiler technology
  - Scalable debugging and performance tools

- System software and libraries
  - Parallel I/O libraries
  - Autotuning software
  - Scalable visualization tools
  - Operating system scaling

A FUNDING STRATEGY TO PROMOTE SCIENTIFIC DISCOVERY THROUGH ADVANCED COMPUTING

THOM DUNNING, PRESENTED BY JACKIE CHEN

BACKGROUND OF SCIDAC

- Context
  - Major change in computer architectures in 1990s—parallel computers
    - Software technologies needed to use parallel computers still in development
    - New mathematical algorithms required for parallel computers
    - Scientific codes not poised to take advantage of new computing technology
- Scientific Discovery through Advanced Computing
  - Support teams of computational scientists, computer scientists and applied mathematicians to create new generation of codes
  - Support development of software technologies needed by scientific applications to enable efficient and effective use use of parallel computers
  - Support development of new algorithms need to optimize performance of scientific and engineering applications on parallel computers

IMPACT OF SCIDAC

- General Observations
  - COV Report 2014:
    - "SciDAC remains the gold standard nationally and internationally for fostering interaction between disciplinary scientists and HPC."
  - Has become standard model in Office of Science for advancing the state-of-the-art in applications for computational science and engineering
- Advancements Enabled by SciDAC
  - Development of new generation of science and engineering codes along with many accompanying scientific advancements
    - BER, BES, FES, NP, and HEP
  - Development of many new software technologies for parallel computing systems
  - Development of many new mathematical algorithms for parallel computing systems

COMPUTING FACILITIES' EVOLUTION AND THEIR ROLE IN THE HPC ECOSYSTEM:

BUDDY BLAND

ECOSYSTEM OF FACILITIES IS A SCIENTIFIC ENABLER

- Interconnected system of facilities provides unique capabilities
  - Facilities are more than just big computers
    - Software such as schedulers, debuggers, file system, data archives, math libraries, data analysis tools, etc.
    - Buildings, power & cooling infrastructure, operations & maintenance
    - Linking experimental, observational, and computational facilities
  - People who understand tools, computing, science, methods
    - User support, Computational and Data Liaisons, AI specialists
  - Bring all those resources to bear on big problems
    - Energy, Medicine, National Security, Basic & Applied Sciences, Engineering
    - Industrial competitiveness, Global challenges

DOE OFFICE OF SCIENCE COMPUTE FACILITIES

- The first supercomputers were purchased by specific programs or laboratories to address specific programs
  - Defense Programs
  - Fusion
  - High Energy Physics
  - Uranium Enrichment
- Transitioned the National Magnetic Fusion Energy Computing Center into NERSC to provide a center available to all DOE science programs.
- In 1992 created High Performance Computing Research Centers at Oak Ridge and Los Alamos to explore and scale parallel architectures and algorithms.
- In 2004 created Leadership Computing Facilities to address a broader mission of delivering the most powerful systems to government, academia, and industry
- In 2015 worked with NNSA to create the Exascale Computing Project to develop needed applications, software and technology to get to exascale systems.

COLLABORATION: VENDORS, INSTITUTIONS, AGENCIES

- Integration of research and production (e.g. Centers of Excellence)
- Synergistic relationship between the Labs, vendors, agencies
  - Co-design of systems and applications to get the requirements right
    - Government projects to create early parallel system
    - Rapid feedback loop for innovation
    - Worked with application scientists, and Industry to discover the possible
    - Then production versions were deployed
  - Examples:
    - DARPA Touchstone ➢ Intel Paragon, Sandia Red Storm ➢ Cray XT line
- Critically important to maintain a path to funding for important ideas
- Also important to both maintain long-term relationships with vendors and allow for innovation through competition

NETWORKS' GLOBAL IMPACT:

WENDY HUNTOON

ESNET CONNECTS USERS GLOBALLY

- ESnet connects and peers with regional, national and international networks.
- Supports access to DoE resources for thousands of scientists, engineers, faculty, staff and students at leading academic institutions and other research organizations.

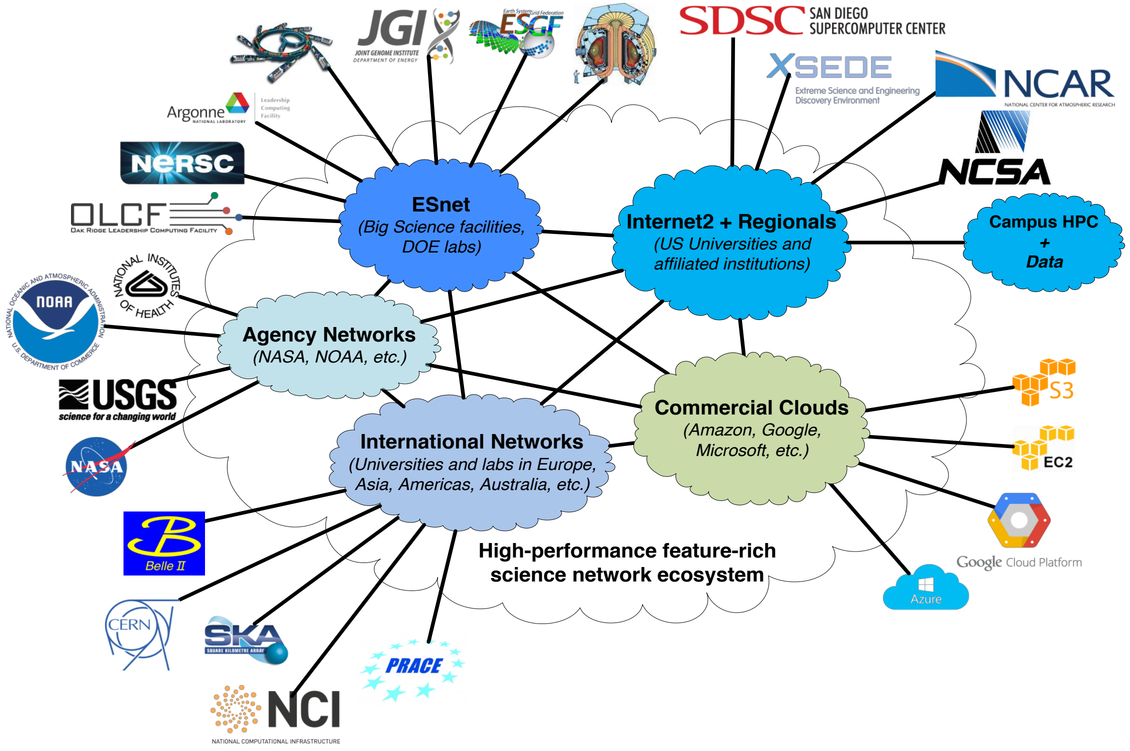

LHC EXPERIMENTS – INTERNATIONALLY CONNECTED

- Distributed, interconnected set of LHC sites bring computing, data and people together
- Ecosystem of connected sites accomplishes things unachievable at a single university or Laboratory
- Higgs Boson discovery

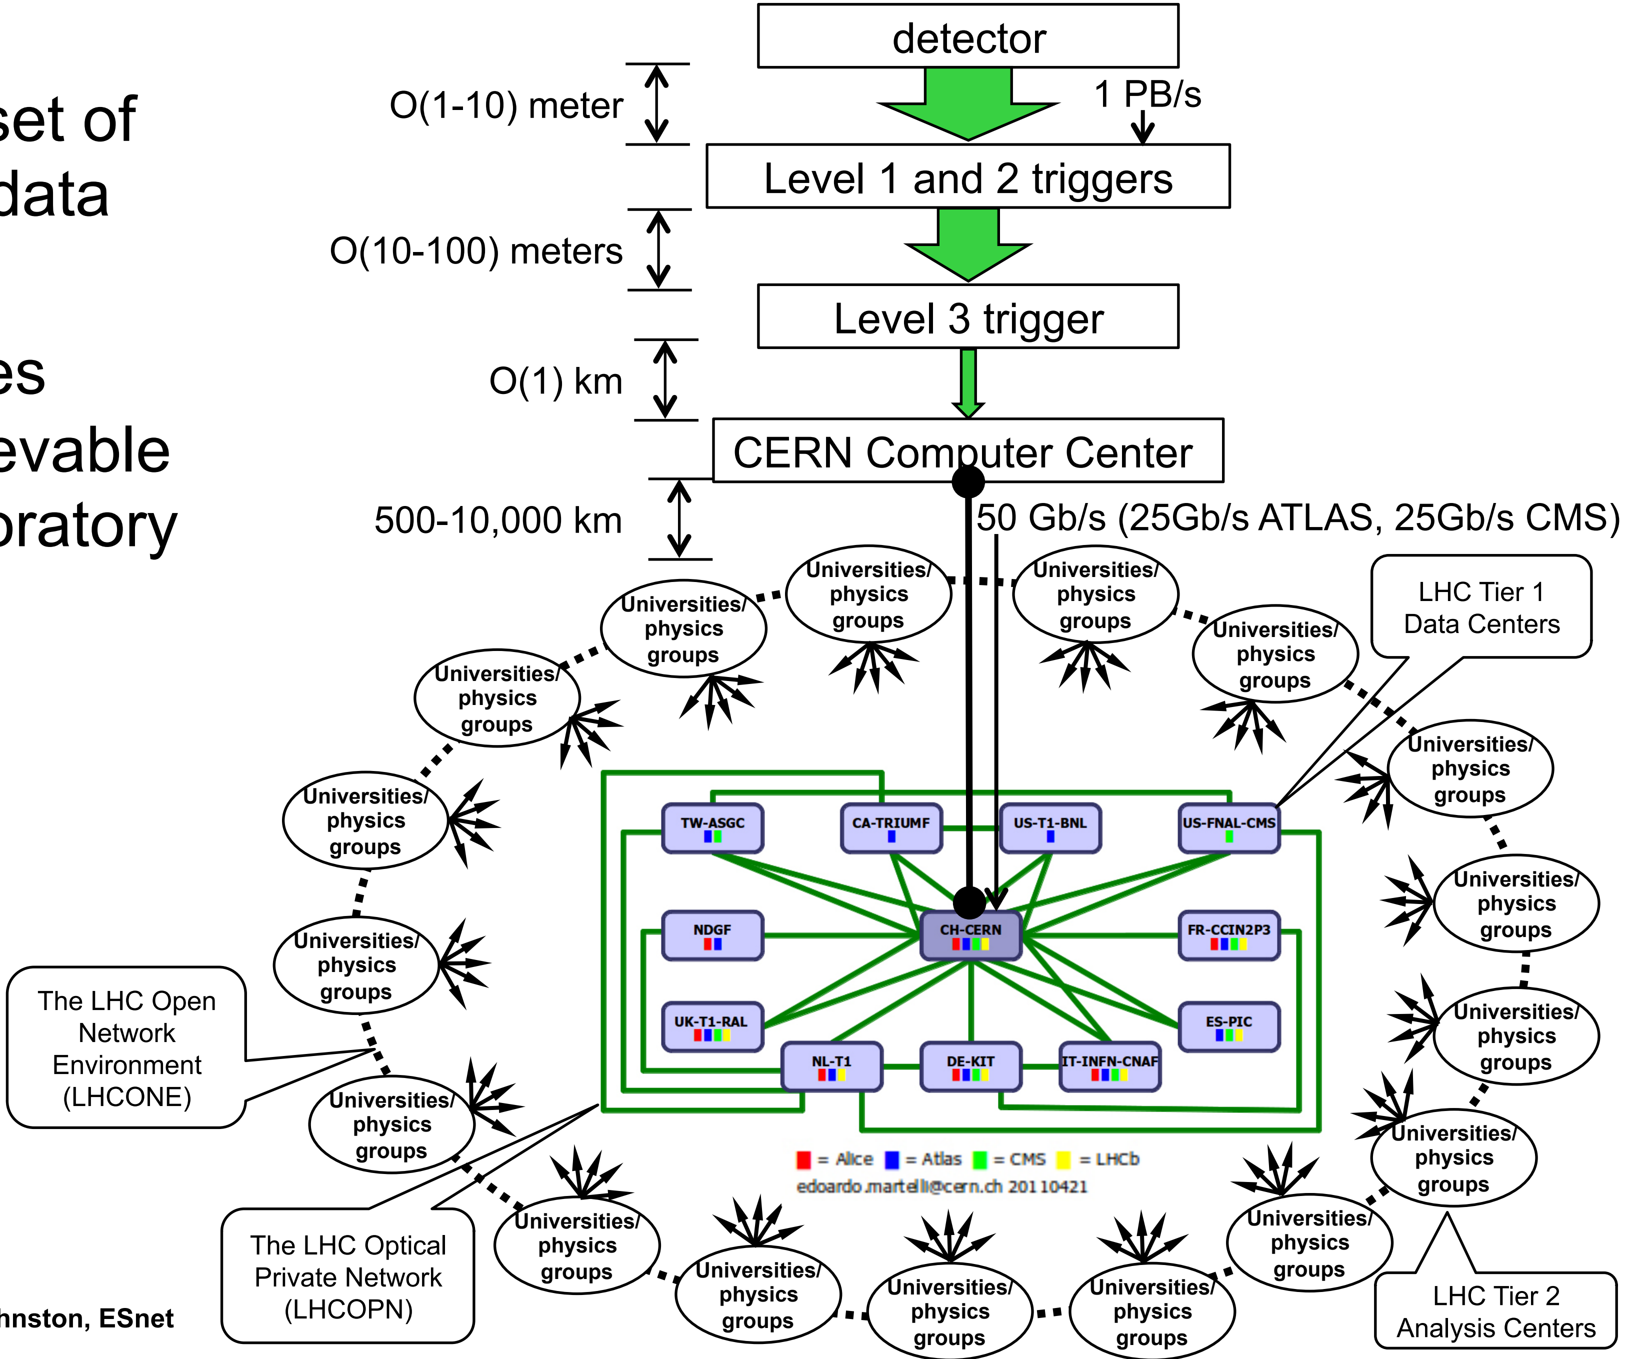

Image source: Bill Johnston, ESnet

NETWORKS AS FACILITIES, RESEACH ON NETWORK TECHNOLOGIES

ELI DART

DOE OFFICE OF SCIENCE AND ESNET – THE ESNET MISSION

- **ESnet - the Energy Sciences Network - is an Office of Science user facility whose <u>primary mission is to enable large-scale science</u> that depends on**:
  - Multi-institution, world-wide collaboration
  - Data mobility: sharing of massive amounts of data
  - Distributed data management and processing
  - Distributed simulation, visualization, and computational steering
  - Collaboration with the US and International Research and Education community
- ESnet traffic reflects this: Growth of about 10x every 4 years, and currently moving about 80 Petabytes/month
- ESnet is a multi-hundred gigabit/sec network whose backbone covers the US and Europe in order to connect the DOE National Laboratories, science instruments and user facilities to each other and to collaborators worldwide

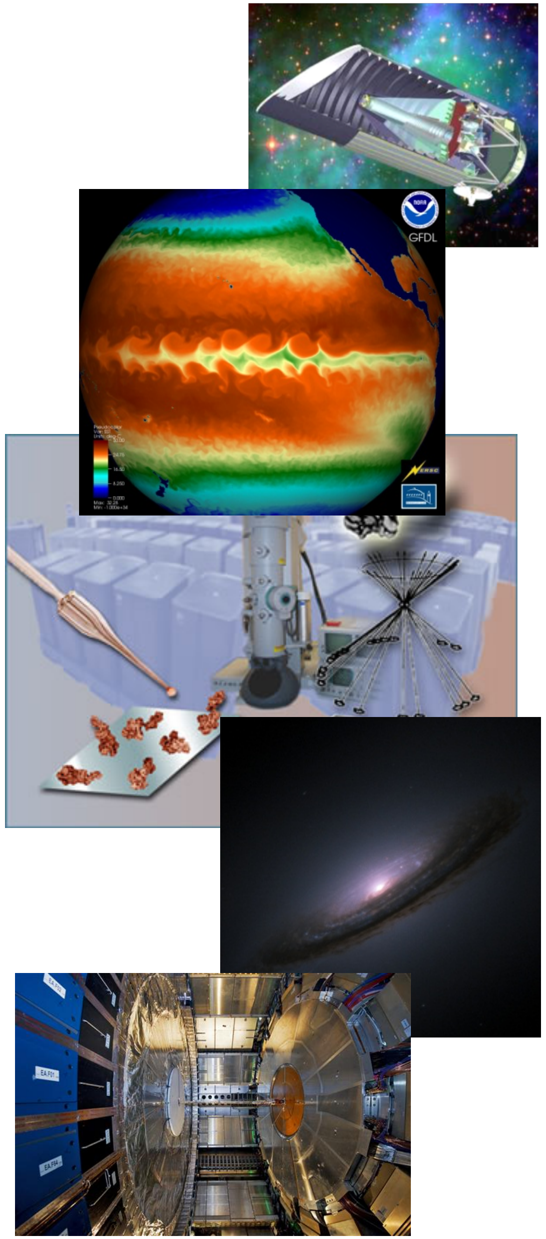

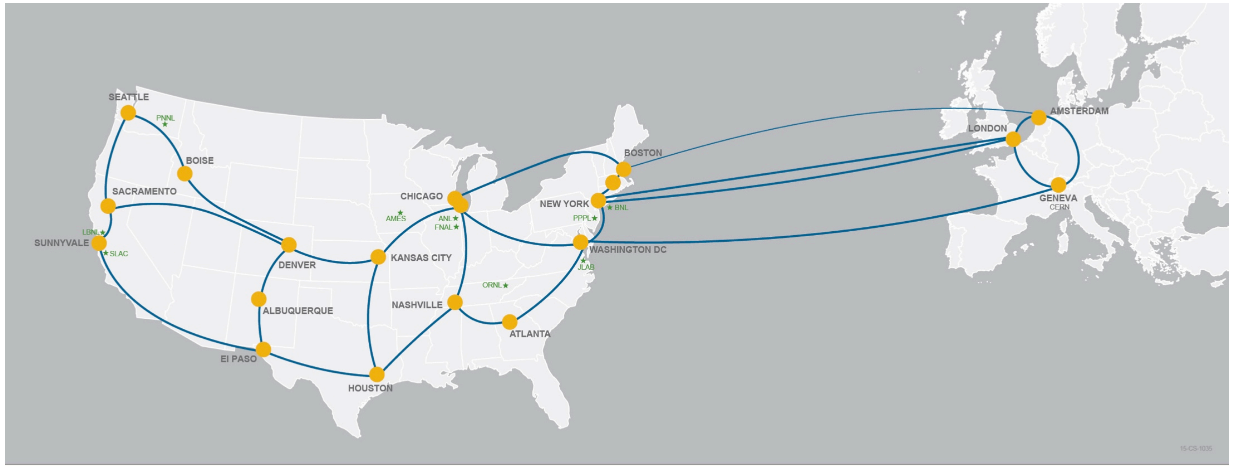

FIXING INTERNET CONGESTION COLLAPSE

- Van Jacobson at LBNL, 1986
  - 1000x decrease in throughput between LBNL and UCB
  - Several key improvements to TCP
    - Slow start
    - Congestion window
    - Others
  - Fixed the throughput problem and laid the foundation for dramatic Internet growth
- Critical architectural improvements served as a platform for multiple technologies

Figure 3: Startup behavior of TCP without Slow-start

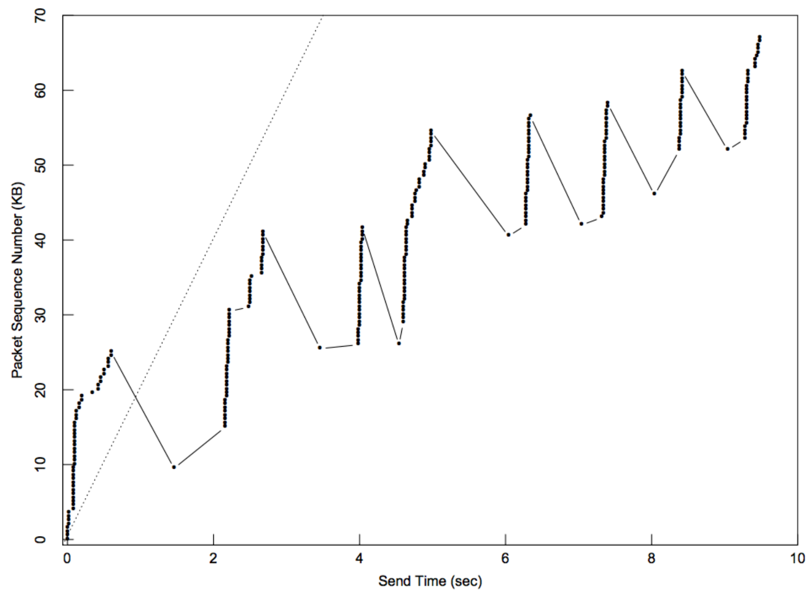

Figure 4: Startup behavior of TCP with Slow-start

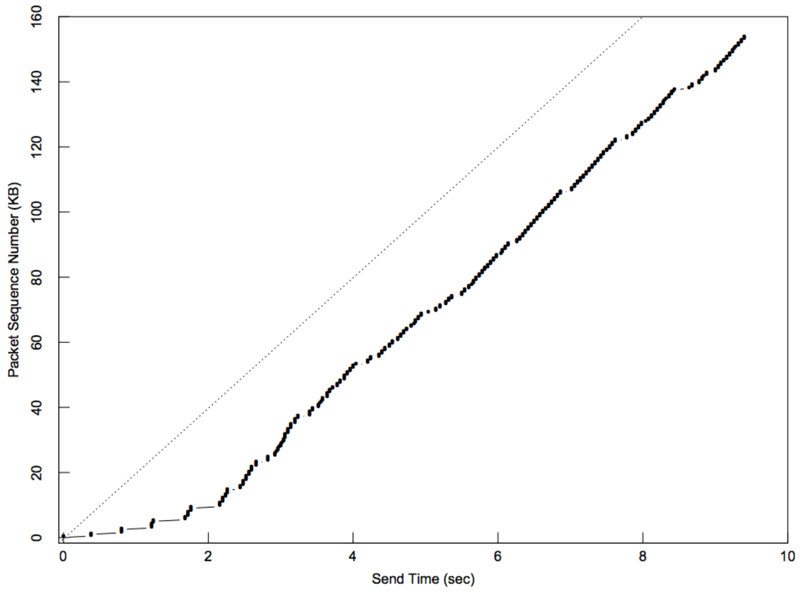

Images: JACOBSON, V. Congestion avoidance and control. In Proceedings of SIGCOMM ’88 (Stanford, CA, Aug. 1988), ACM.

SCIENCE DMZ DESIGN PATTERN

- High-speed interface between site/ campus network and wide area science network
- Allows for easy deployment of scalable data transfer services
- Data placement at scale allows HPC to be brought to bear on previously intractable problems
- Working toward a future where all major scientific facilities have high speed connections to capable science networks

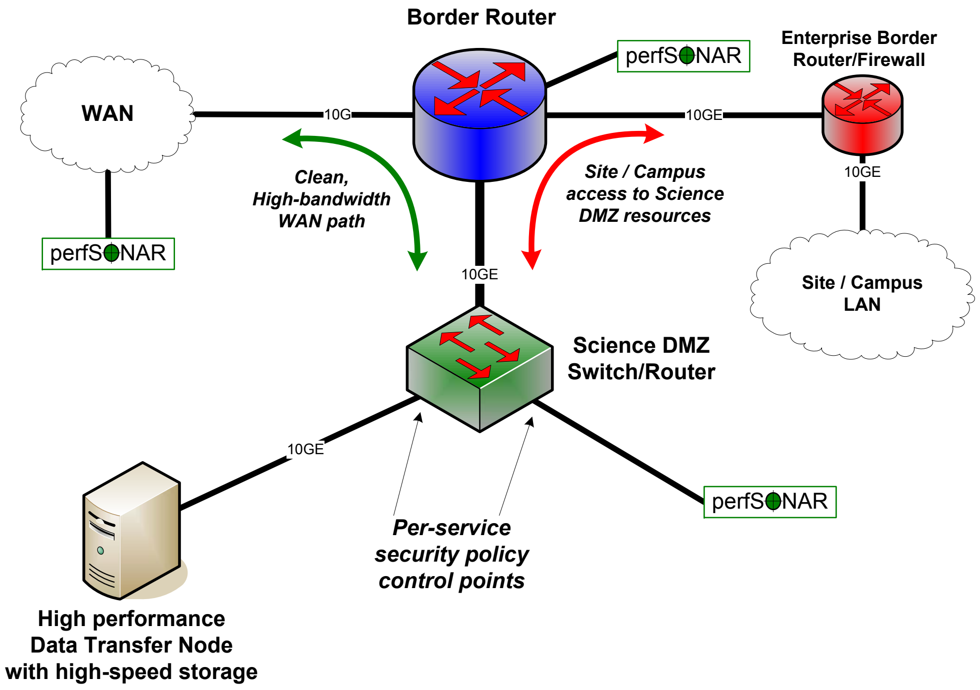

E. Dart, L. Rotman, B. Tierney, M. Hester and J. Zurawski, "The Science DMZ: A network design pattern for data-intensive science," *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Denver, CO, 2013, pp. 1-10. doi: 10.1145/2503210.2503245

ESNET DRIVING SOFTWARE-DEFINED NETWORKING CONCEPTS ACROSS DOE AND INDUSTRY SINCE 2006

- SDN will do for networking in the future what modern orchestration and service constructs have done for computing
  - More responsive, self-healing, more efficient, higher performance
  - The APIs, interfaces, and services are an area of active research
- Getting the architecture right for multiple technology generations with new ESnet6 network

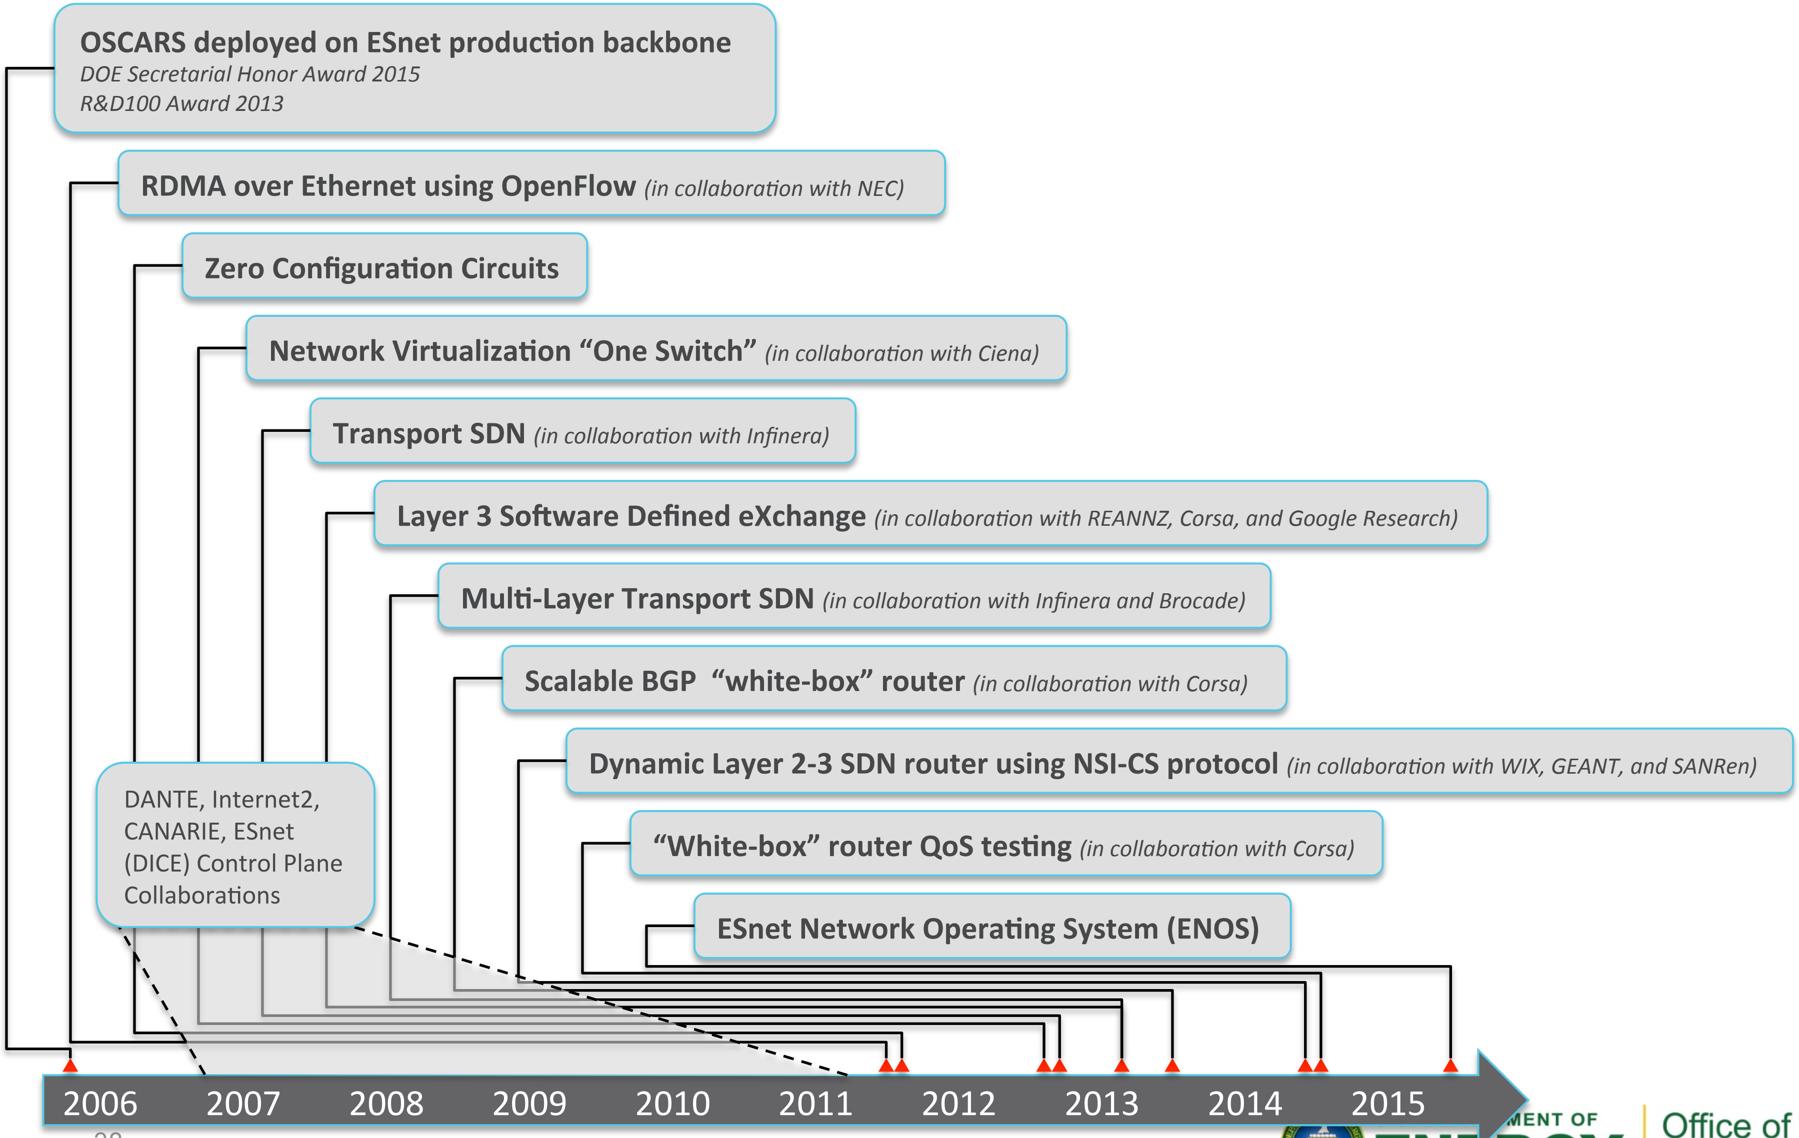

IN THE REMAINING TIME, PLEASE GIVE US YOUR INPUT

- What are the major scientific accomplishments that have shaped the ASCR-supported disciplines in the last 40 years? How has ASCR contributed to those advances?
- What impacts have those accomplishments had on the Department's missions in energy, environment, or security?
- What are the key aspects of the ASCR's investment strategy that have had the greatest impacts?
- Looking to the future, and building on the ASCAC reports, what research areas and funding strategies to pursue those areas could further strengthen ASCR in serving the DOE missions?