



U.S. DEPARTMENT OF
ENERGY

Office of
Science

ASCR

March 17, 2011

CASC

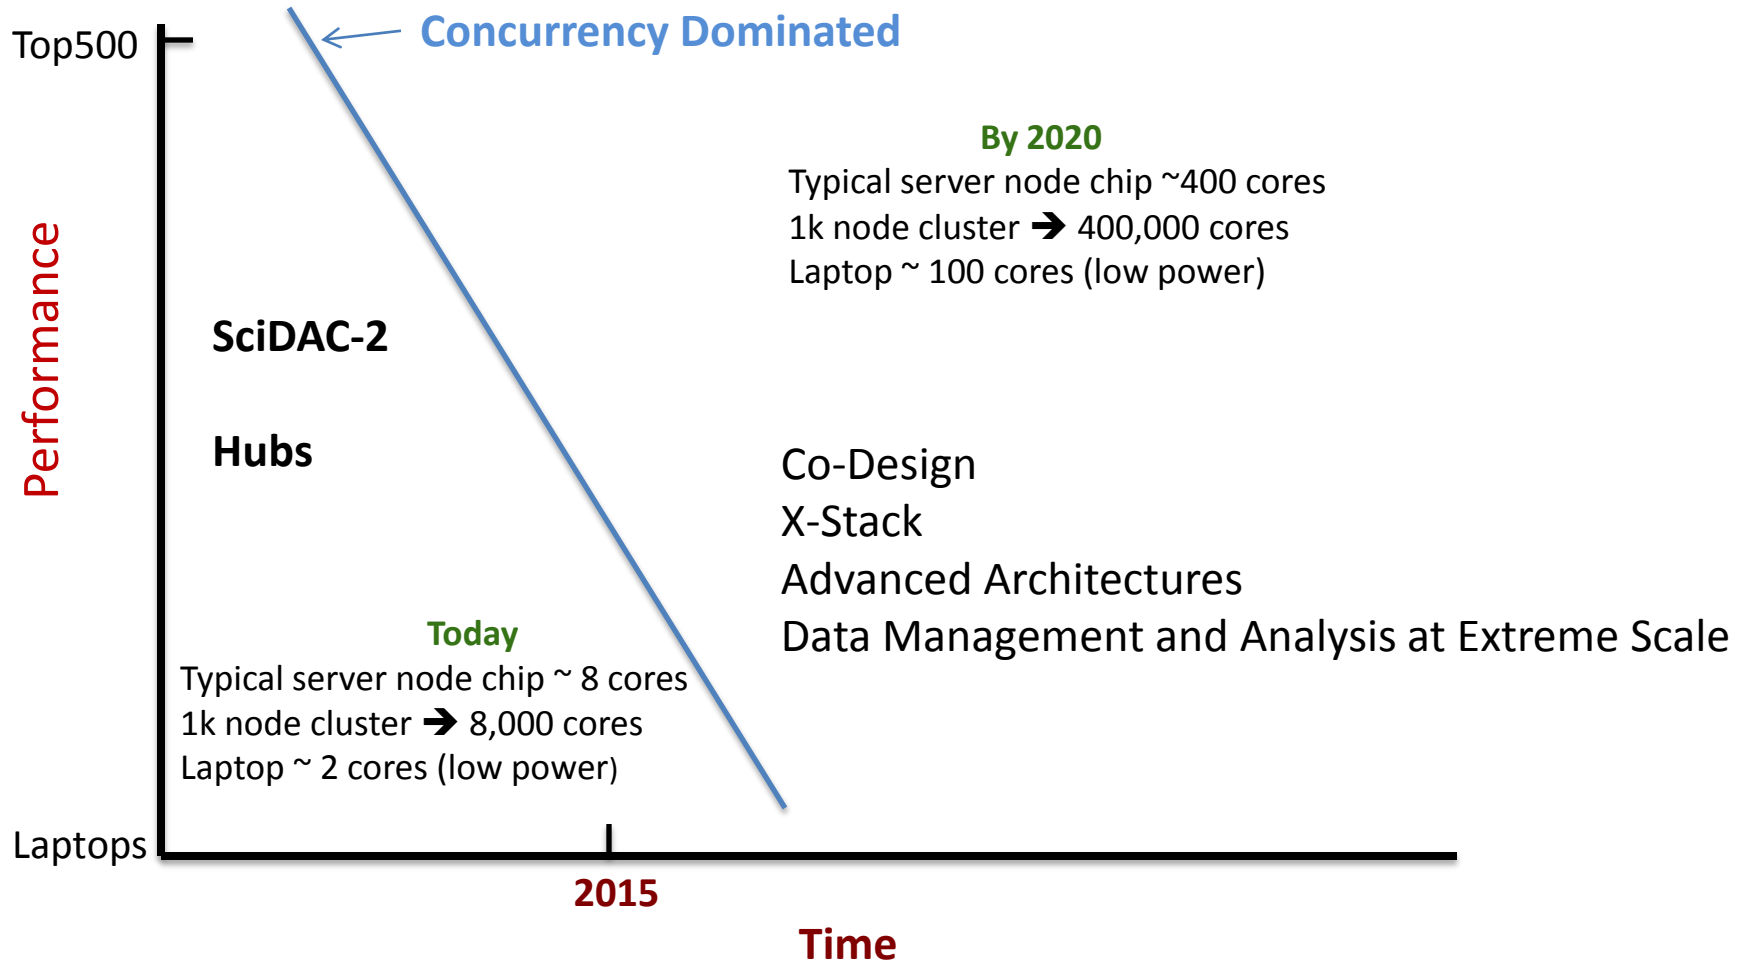
The Challenge of Exascale and Exabyte

Daniel Hitchcock

Acting Associate Director


Advanced Scientific Computing Research

Phase Change in Computing

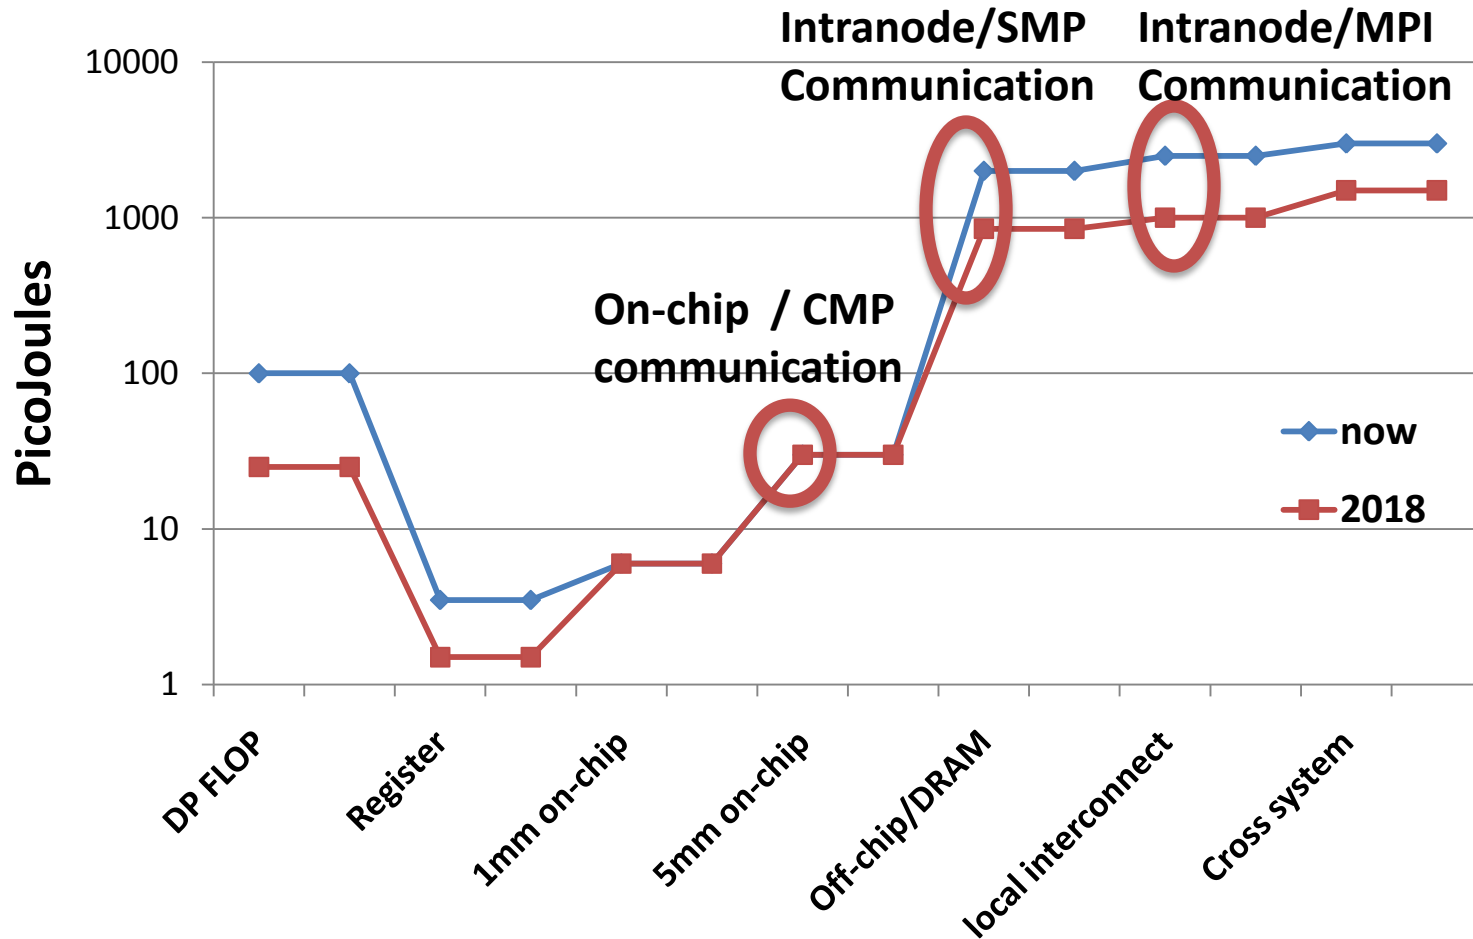


Exascale is about Energy Efficient Computing

- At \$1M per MW, energy costs are substantial
- 1 petaflop in 2010 will use 3 MW
- 1 exaflop in 2018 at 200 MW with “usual” scaling
- 1 exaflop in 2018 at 20 MW is target



The Fundamental Issue: Where does the Energy (and Time) Go?



National Academies Report on Computing Performance

PREPUBLICATION COPY—SUBJECT TO FURTHER EDITORIAL CORRECTION

The Future of Computing Performance

Game Over or Next Level?

Samuel H. Fuller and Lynette I. Millett, Editors

Committee on Sustaining Growth in Computing Performance

Computer Science and Telecommunications Board

Division on Engineering and Physical Sciences

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, DC
www.nap.edu

- **Report makes same points**
 - Past performance increases have driven innovation
 - Processor speeds stalled
 - Energy is limitation
- **Impacts everything**
 - Data
 - Simulations
 - Control



Performance Growth

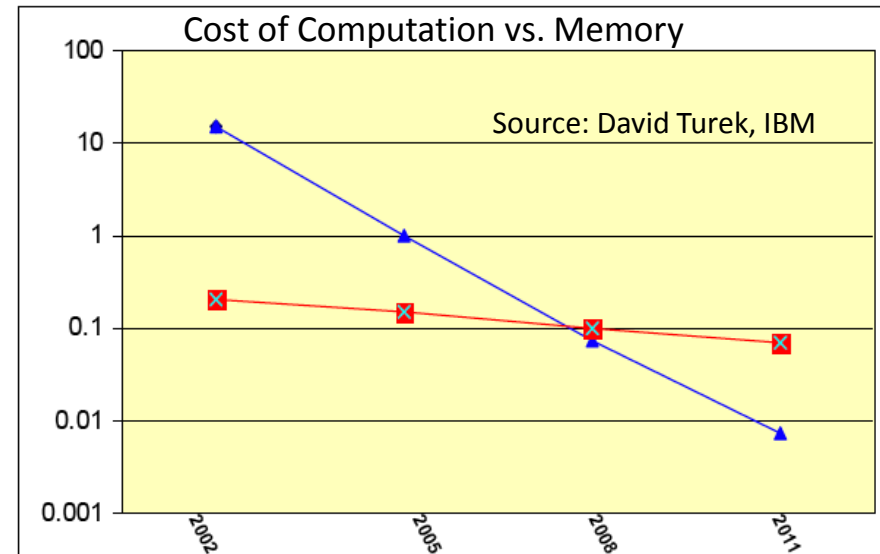

- 1) **System power** is the primary constraint
- 2) **Memory** bandwidth and capacity are not keeping pace
- 3) **Concurrency** (1000x today)
- 4) **Processor** architecture is an open question
- 5) **Programming model** heroic compilers will not hide this
- 6) **Algorithms** need to minimize data movement, not flops
- 7) **I/O bandwidth** unlikely to keep pace with machine speed
- 8) **Reliability and resiliency** will be critical at this scale
- 9) **Bisection bandwidth** limited by cost and energy

Unlike the last 20 years most of these (1-7) are equally important across scales, e.g., 100 10-PF machines

Memory is Not Keeping Pace

Technology trends against a constant or increasing memory per core

- Memory density is doubling every three years; processor logic is every two
- Memory costs are dropping gradually compared to logic costs

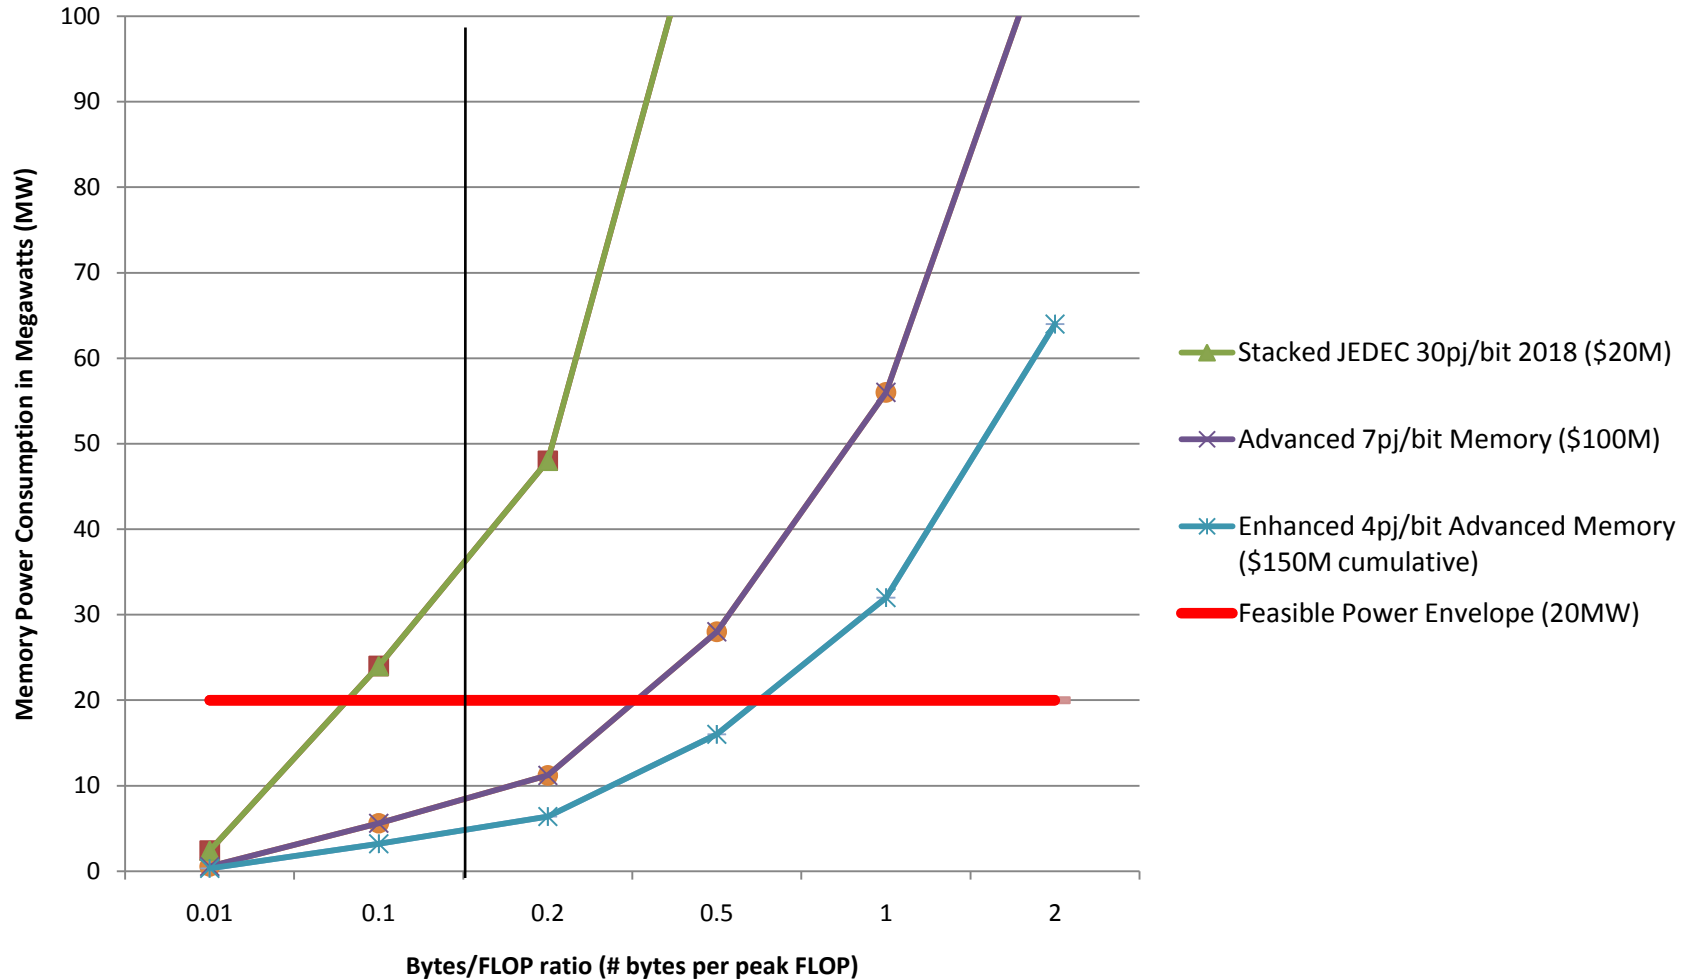


The cost to sense, collect, generate and calculate data is declining much faster than the cost to access, manage and store it

Question: Can you double concurrency without doubling memory?



Memory Technology: Bandwidth costs power



Approaches

- **Locality, Locality, Locality!**
- **Billion Way Concurrency;**
- **Uncertainty Quantification (UQ) including hardware variability;**
- **Flops free data movement expensive so:**
 - Remap multi-physics to put as much work per location on same die;
 - Include embedded UQ to increase concurrency;
 - Include data analysis if you can for more concurrency
 - Trigger output to only move important data off machine;
 - Reformulate to trade flops for memory use.
- **Wise use of silicon area**

Exascale Progress since May, 2010

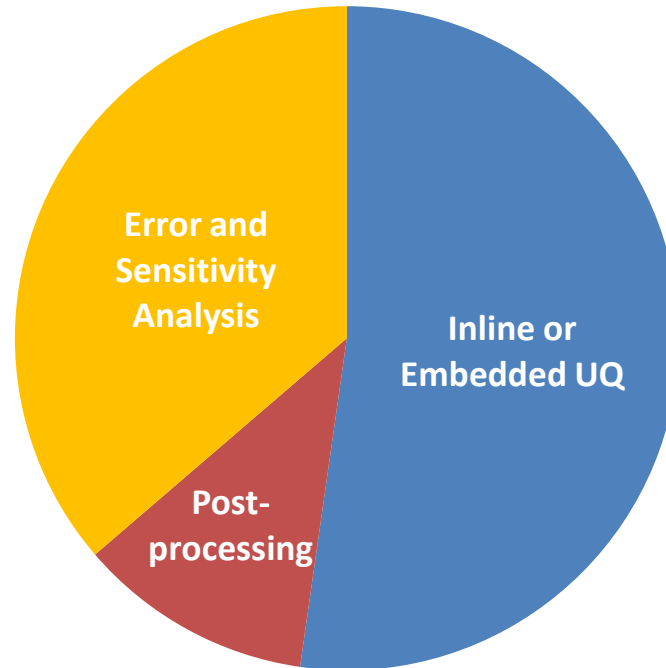
- **Proposals processed in Exascale related topic areas:**
 - Applied Math: Uncertainty Quantification
 - Computer Science: Advanced Architectures
 - Computer Science: X-Stack
 - Computer Science: Data Management and Analysis at Extreme Scales
 - Computational Partnerships: Co-Design (21 Proposals requesting ~ \$160M/year)
- **Exascale Coordination meeting with DOD and DARPA – June 15 with decisions to --**
 - Have yearly coordination/planning meetings in accordance with current MOU
 - Follow-on meeting of Program Managers to identify gaps and overlaps in computer science and applied math areas currently funded
 - Address critical technologies issues, initially in memory
- **Ongoing meetings with NNSA**



ASCR Exascale Research Kick-off Meeting

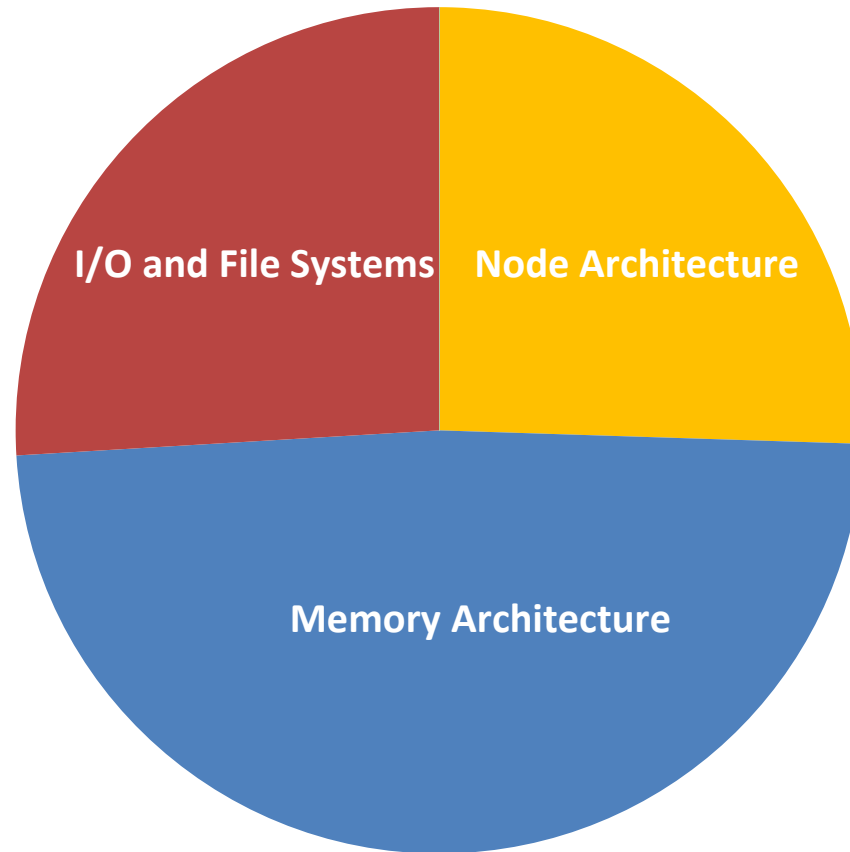
- **PI meeting for**
 - Advanced Architecture, X-Stack, Scientific Data Management and Analysis at Extreme Scales awardees
 - Co-Design planning grant recipients
- **Expected Outcomes:**
 - Awareness **within** each solicitation communities and ASCR what members are doing and areas where they can leverage and supplement their work
 - Awareness **across** solicitation communities of what is going on and where each project fits in relation to the broad spectrum
 - Identification of gaps in ASCR exascale research potfolio
 - Lay groundwork for collaboration/cooperation with NNSA Exascale Roadmapping activities
- **March 7-11, 2011**

Uncertainty Quantification Portfolio



60% National Laboratory Funding
40% University Funding

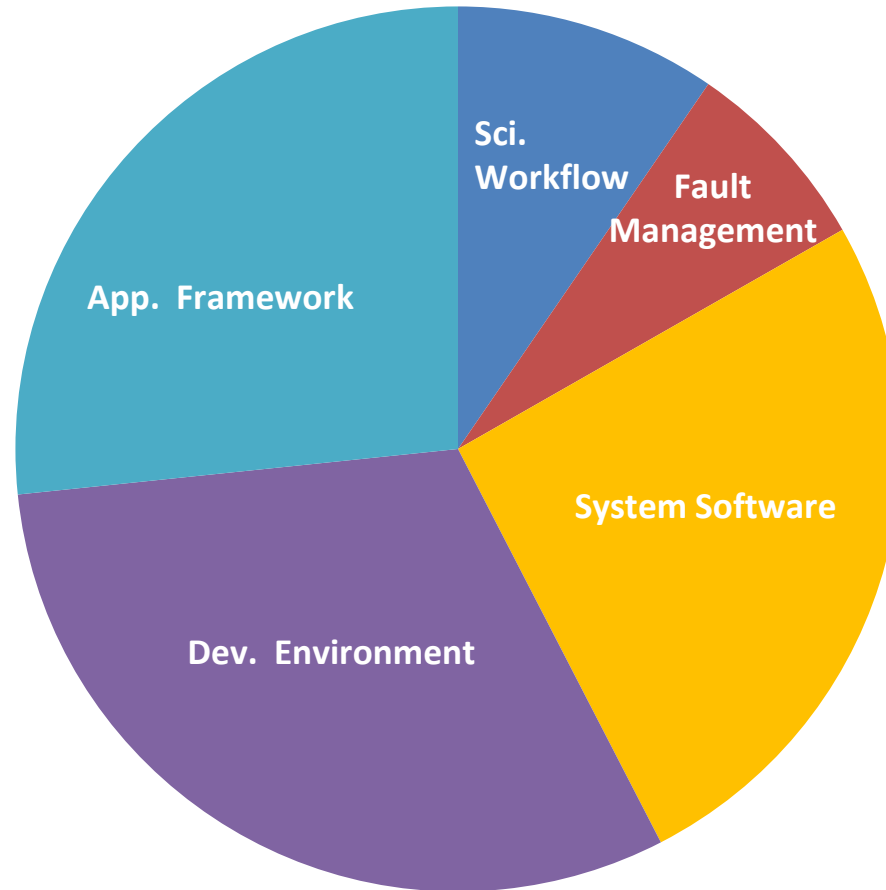
Advanced Architecture Portfolio



60% DOE Laboratory
40% University and Industry



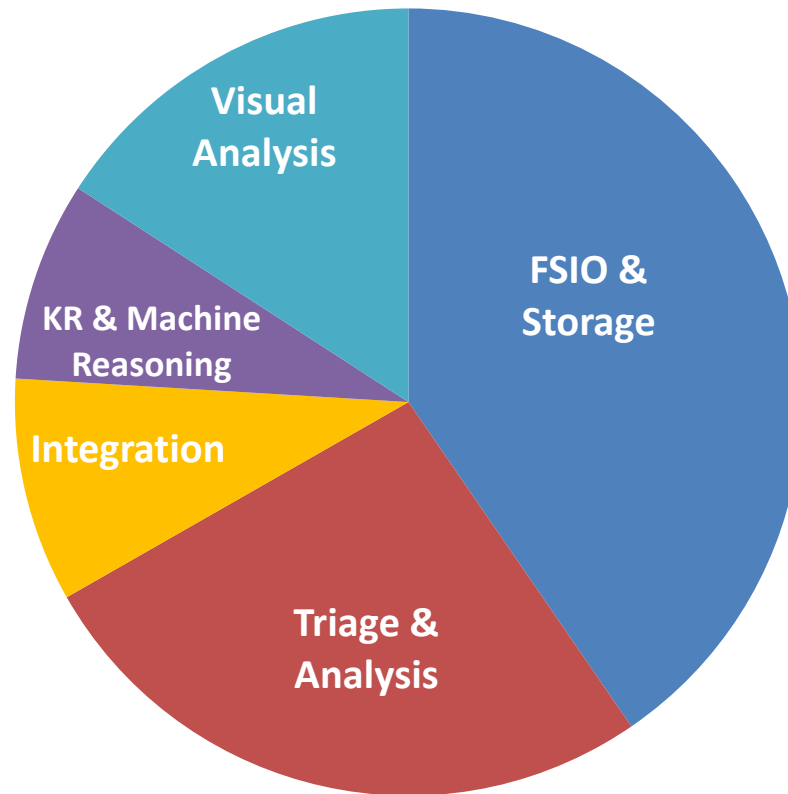
X-Stack Portfolio



50% National Laboratory
50% University and Industry



Scientific Data Management & Analysis

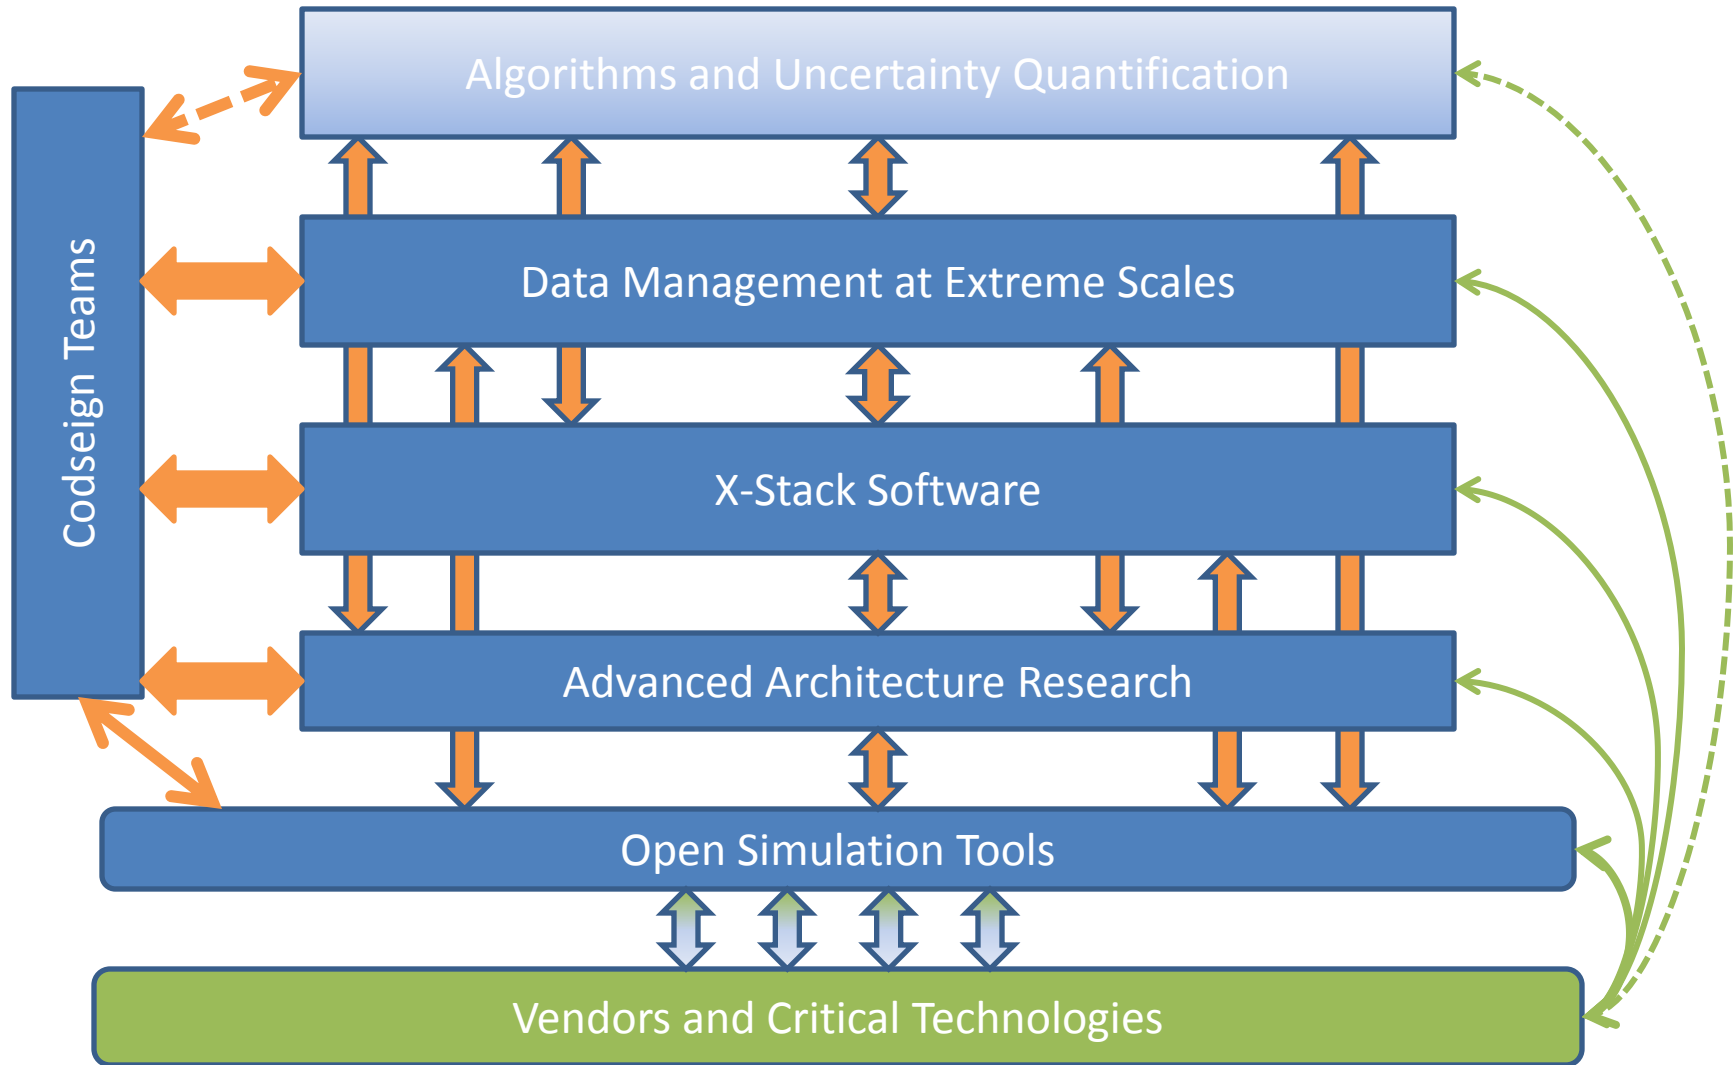


60 % Lab

40% University and Industry



Kickoff Meeting Focus:



ASCR at a Glance

Office of Advanced Scientific Computing Research

Associate Director – Daniel Hitchcock (acting)

Phone: 301-903-7486

E-mail: Daniel.Hitchcock@science.doe.gov



Research

Division Director – William Harrod

Phone: 301-903-5800

E-mail: William.Harrod@science.doe.gov

Facilities

Division Director – Daniel Hitchcock

Phone: 301-903-9958

E-mail: Daniel.Hitchcock@science.doe.gov

Relevant Websites

ASCR: science.doe.gov/ascr/

ASCR Workshops and Conferences:

science.doe.gov/ascr/WorkshopsConferences/WorkshopsConferences.html

SciDAC: www.scidac.gov

INCITE: www.science.doe.gov/ascr/incite/

Exascale Software: www.exascale.org

DOE Grants and Contracts info: science.doe.gov/grants/

