Los Alamos
NATIONAL LABORATORY
EST.1943

The Role of HPC in Stockpile Stewardship

John Sarrao

ASCAC, 09/21/16

NNSA
National Nuclear Security Administration

Advanced Simulation and Computing (ASC) is a mission driven program

ASC simulations are central to U.S. national security, as they provide a computational surrogate for nuclear testing.

- **Supports the shift in emphasis from test-based confidence to simulation-based confidence**

"Under ASC, computer simulation capabilities are developed to analyze and predict the performance, safety, and reliability of nuclear weapons and to certify their functionality."

While ASC is a relatively new construct, this perspective characterizes the past 70+ years

Los Alamos has been at the forefront of computing since 1943

- The Nuclear Weapons Program has both driven, and taken advantage of, increased computing capability
- **A 16 order-of-magnitude increase in capability in 70 years!**

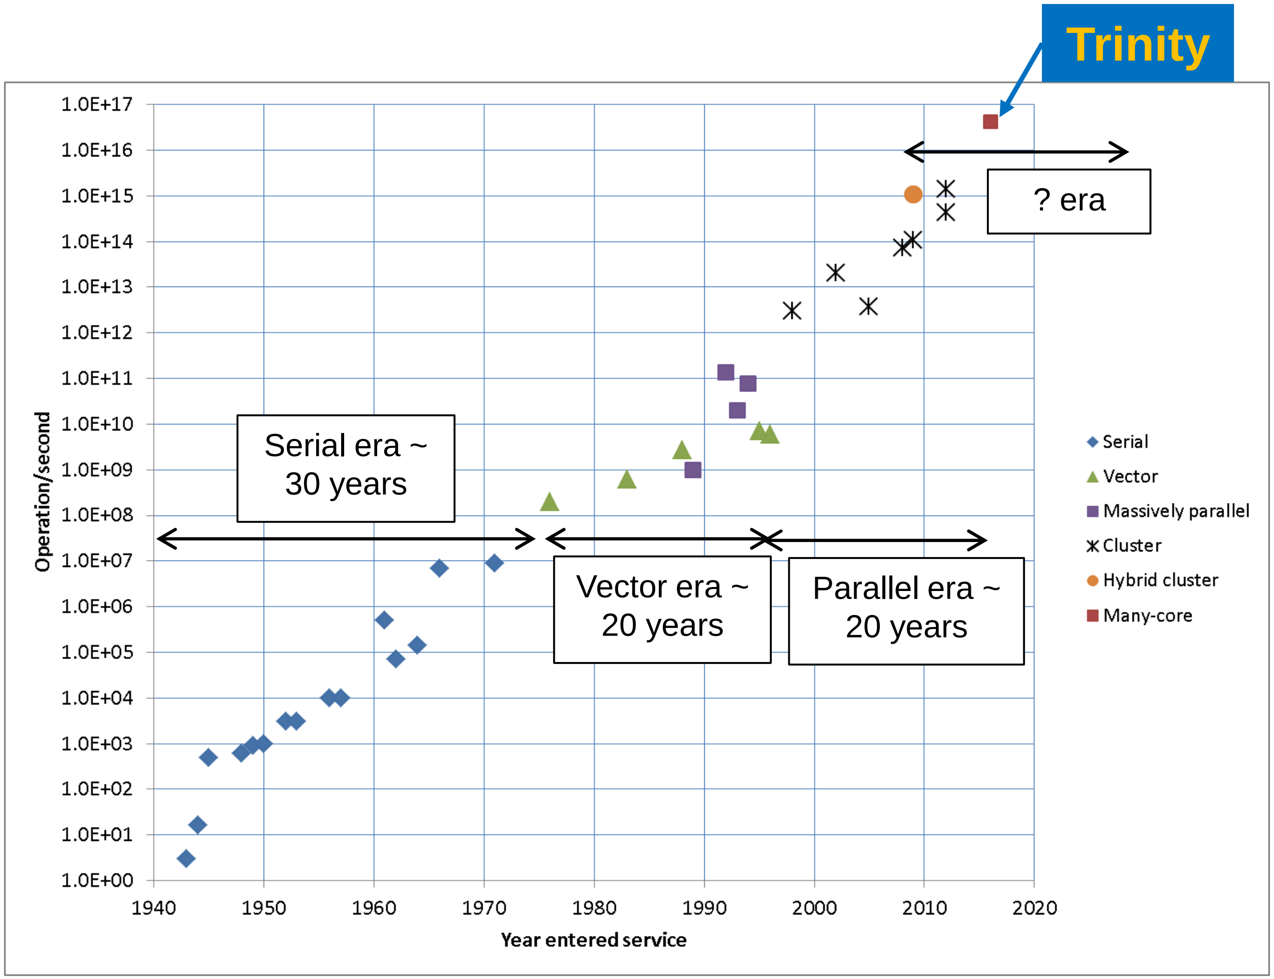

Los Alamos NATIONAL LABORATORY

Technological diversity characterizes the past 70 years (and the future)

[Small/large core memory]


Cray 1 1976
[Vector machine]

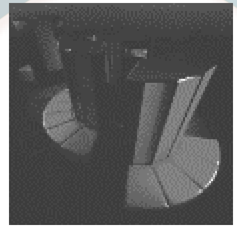
Cray X-MP 1983


TMC CM-5 1992 [fat tree]


Blue Mountain 1998
[Massively parallel]

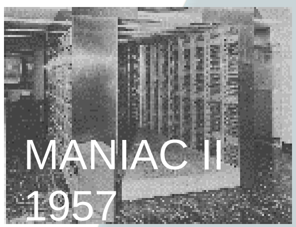

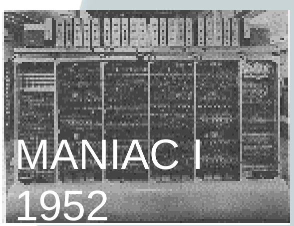

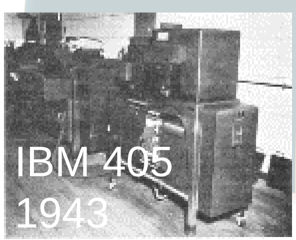

- Core NW mission needs have been major industry driver, but that has changed.
- Significant changes in architecture have accompanied the increasing power
- Resulted in rich capability of coupling scientific algorithms to varied architectures (i.e. scaling, messaging, and vectorization)


Lightning (LNXI) 2004
[commodity computing]


Roadrunner 2005-2008
[Hybrid architecture]

IBM Punch Card Accounting Machines (PCAM) were used for the Fat Man implosion hydrodynamics

16 Ops

IBM 601 Multiplier

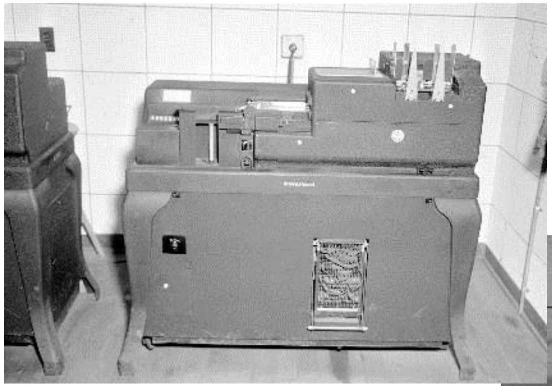

IBM Punch Card

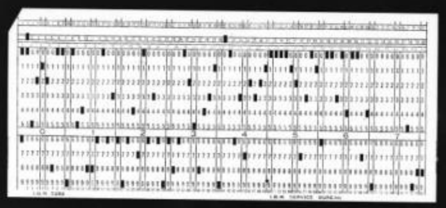

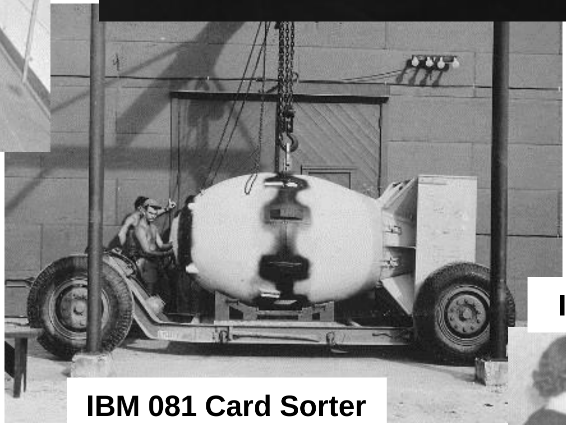

IBM 405 Accounting Machine

IBM 513 Reproducing / Summary

IBM 081 Card Sorter

IBM 031 Key Punch

IBM 077 Collator

Serial, Electro-mechanical, 1944-1950

The first calculation on ENIAC, one of the first electronic computers, was by Los Alamos

- **The first calculation on ENIAC, 1945**
  - First thermonuclear calculation, for the Super bomb
  - John von Neumann, Stan Frankel, Anthony Turkevich
- **Monte Carlo method developed at Los Alamos, 1947**
  - Uses random particle method to solve nuclear problems
  - By Stan Ulam and John von Neumann
- **First ever Monte Carlo calculation, on ENIAC in 1948**
  - John von Neumann, Klara von Neumann, Herman Goldstine, Adele Goldstine, Nick Metropolis

500 Ops

18,000 vacuum tubes

From Columbia U.

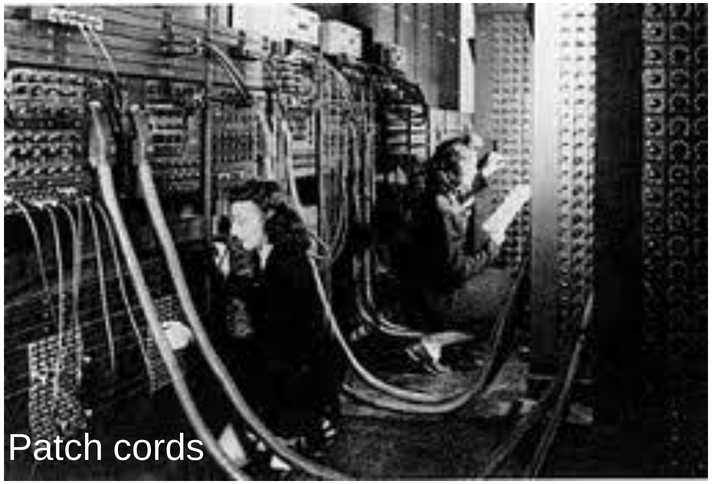

Patch cords

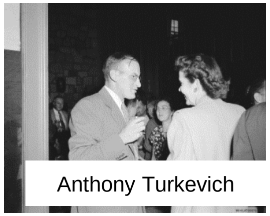

Anthony Turkevich

Co-design got an early start

The Cray vector systems were used by the weapons program for about 30 years

- **Cray 1 was co-designed with Los Alamos over a 6 year period**
  - Los Alamos had serial #1
  - Fastest machine, used integrated circuits
  - Bare iron, LLL wrote operating system and LASL wrote compilers, math and graphics libraries
- Cray XMP arrived in 1983, 4 CPU
- **Cray YMP arrived in 1988, 8 CPU**
- **LANL kept a T94 until 2003 and a J90 until 2004**

6 GOps

Cray J90

200 MOps

Cray 1

600 MOps

Cray XMP 4800

3 GOps

Cray YMP

7 GOps

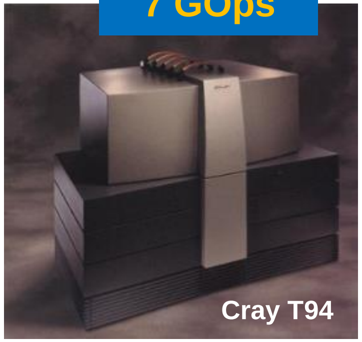

Cray T94

Computing has always been a core component of the weapons program; landscape changed in 1992

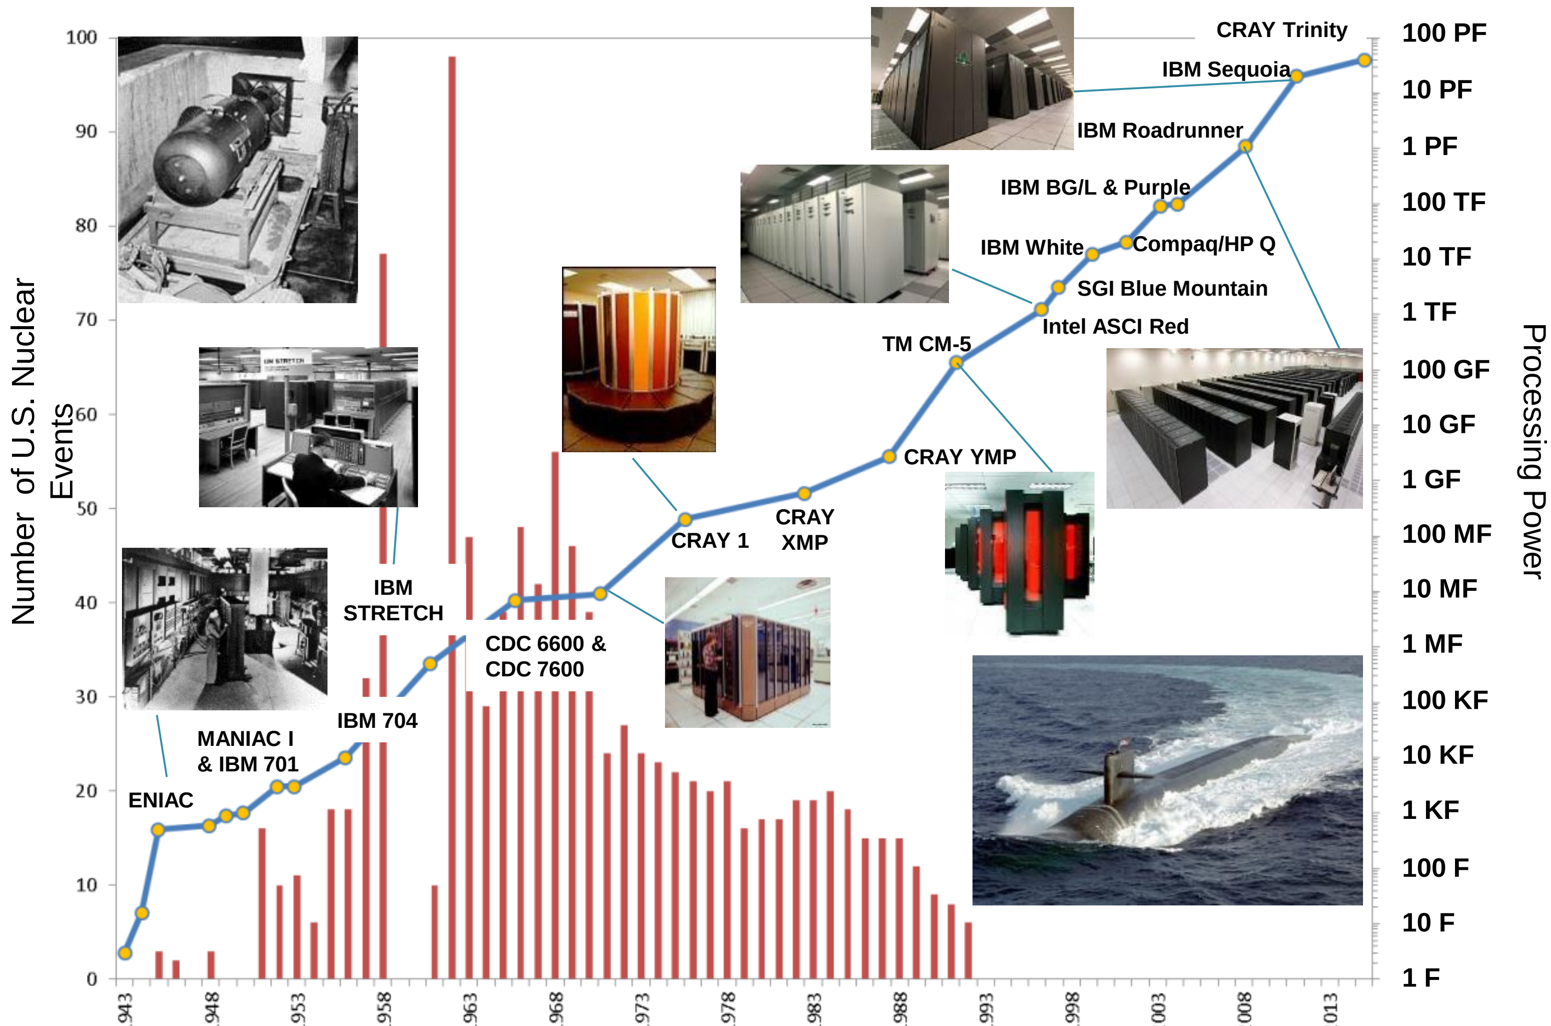

The end of testing drove the adoption of the Accelerated Strategic Computing Initiative (at a transition point in HPC)

- ASCI is essential to analyze and assess performance of nuclear weapons, predict their safety and reliability, and certify their functionality in the absence of nuclear testing
- ASCI's objective is to support high-confidence assessments and stockpile certification through higher fidelity simulations
- ASCI will accelerate development of HPC far beyond what might be achieved in the absence of nuclear testing

Content circa 1998

Courtesy: M. McCoy, LLNL

In summary, hard ASCI tri-lab/Headquarters work resulted in an *integrated plan for development of infrastructure and software*

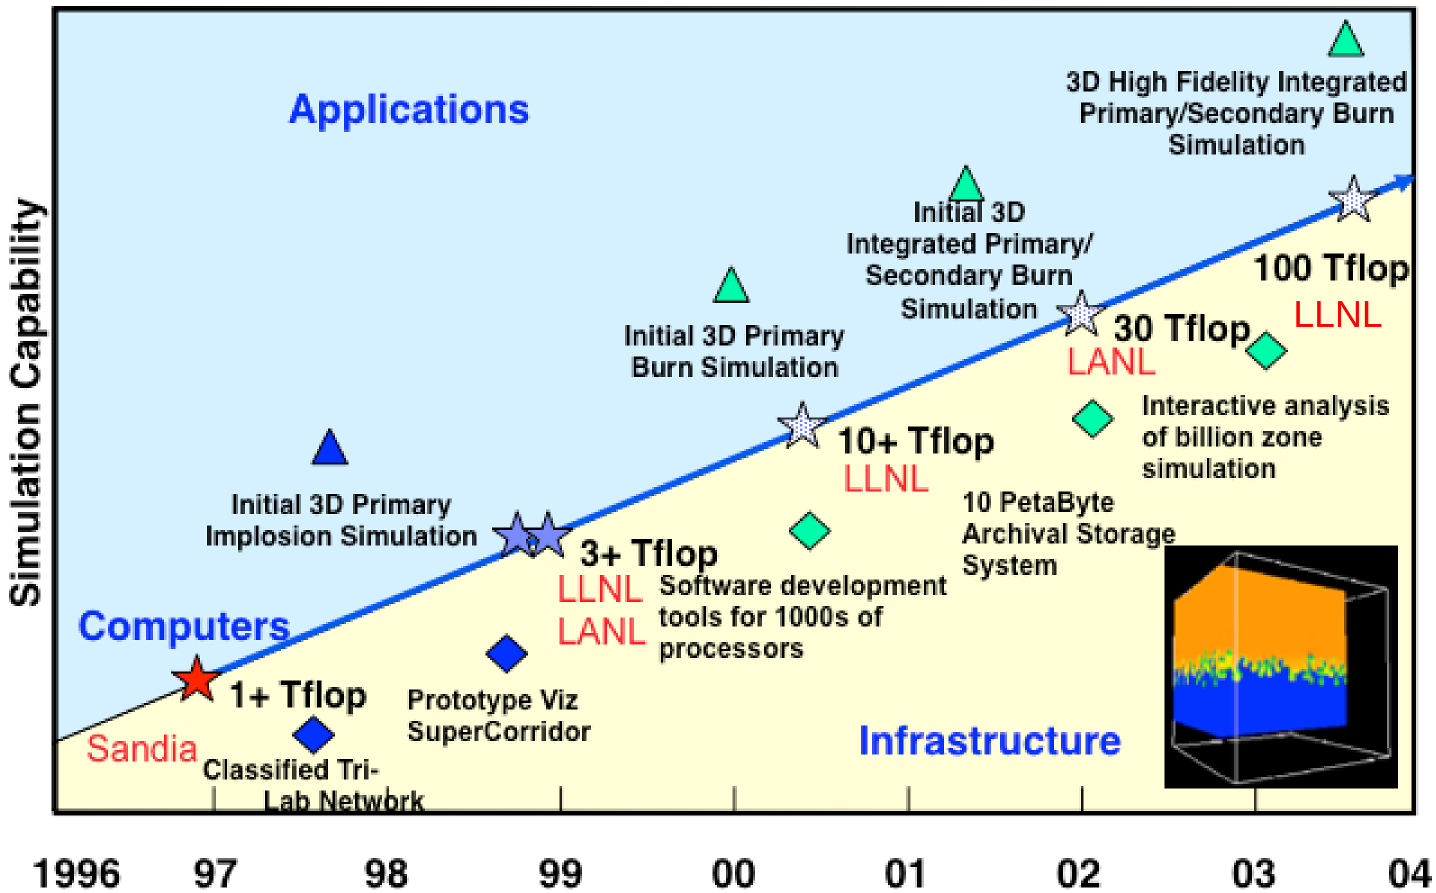

Content circa 1998

Courtesy: M. McCoy, LLNL

Stockpile Stewardship at the Petascale: *Understanding energy balance*

Nuclear Performance

- **High resolution 2D petascale simulations revolutionized NNSA's strategy for doing science**
- **Energy balance:**
  - First success of Predictive Capability Framework
  - Prior to solution provided by petascale (and experiment) simulations were calibrated to nuclear tests
- **UQ performed in 2D at low resolution**

Among the reasons nuclear tests were performed was our lack of understanding of energy balance and boost

Weapons Science

Mission-directed goals requires increases in both simulation scale and model complexity

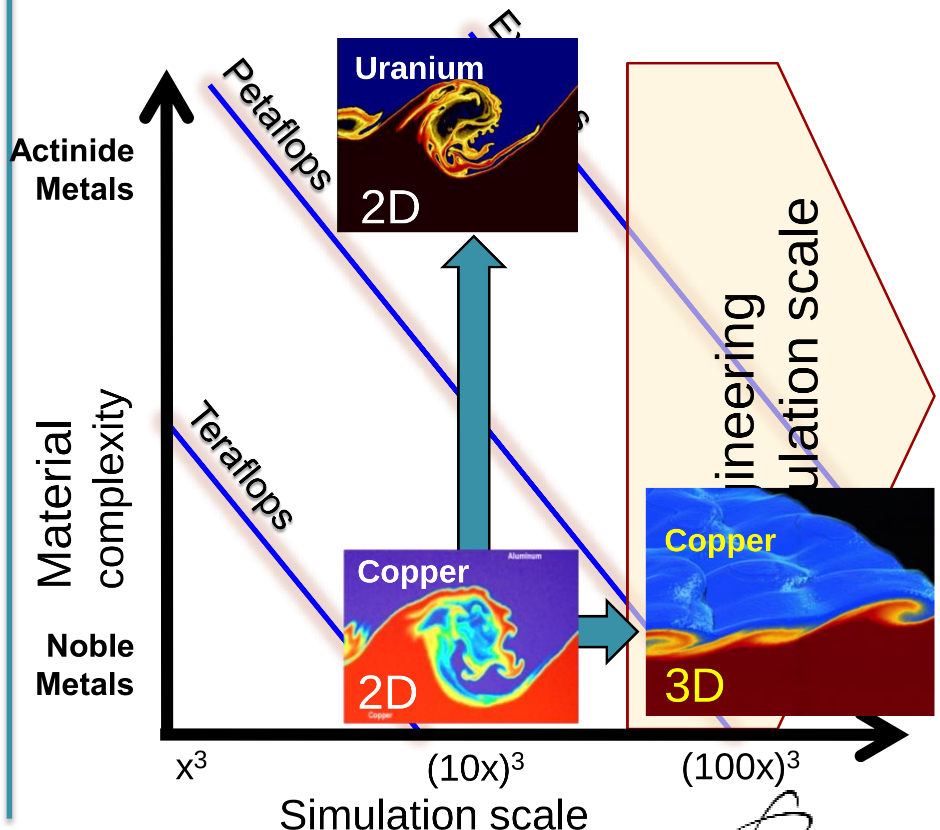

ASCI was successful, but (much) more work remains

"When we at ASCI first estimated what we would need by now in high-performance computing, we underestimated. In my view, we must continue to advance the power and resolution of our computers to do our mission; the ongoing weapon life-extension programs and our annual assessment of the deterrent depend on it."

Charles McMillan, Director,
Los Alamos National Laboratory
National Security Science, April 2013

The number of weapons designers with test experience continues to shrink

While our confidence in the stockpile remains high, the approach to underwrite that confidence has changed

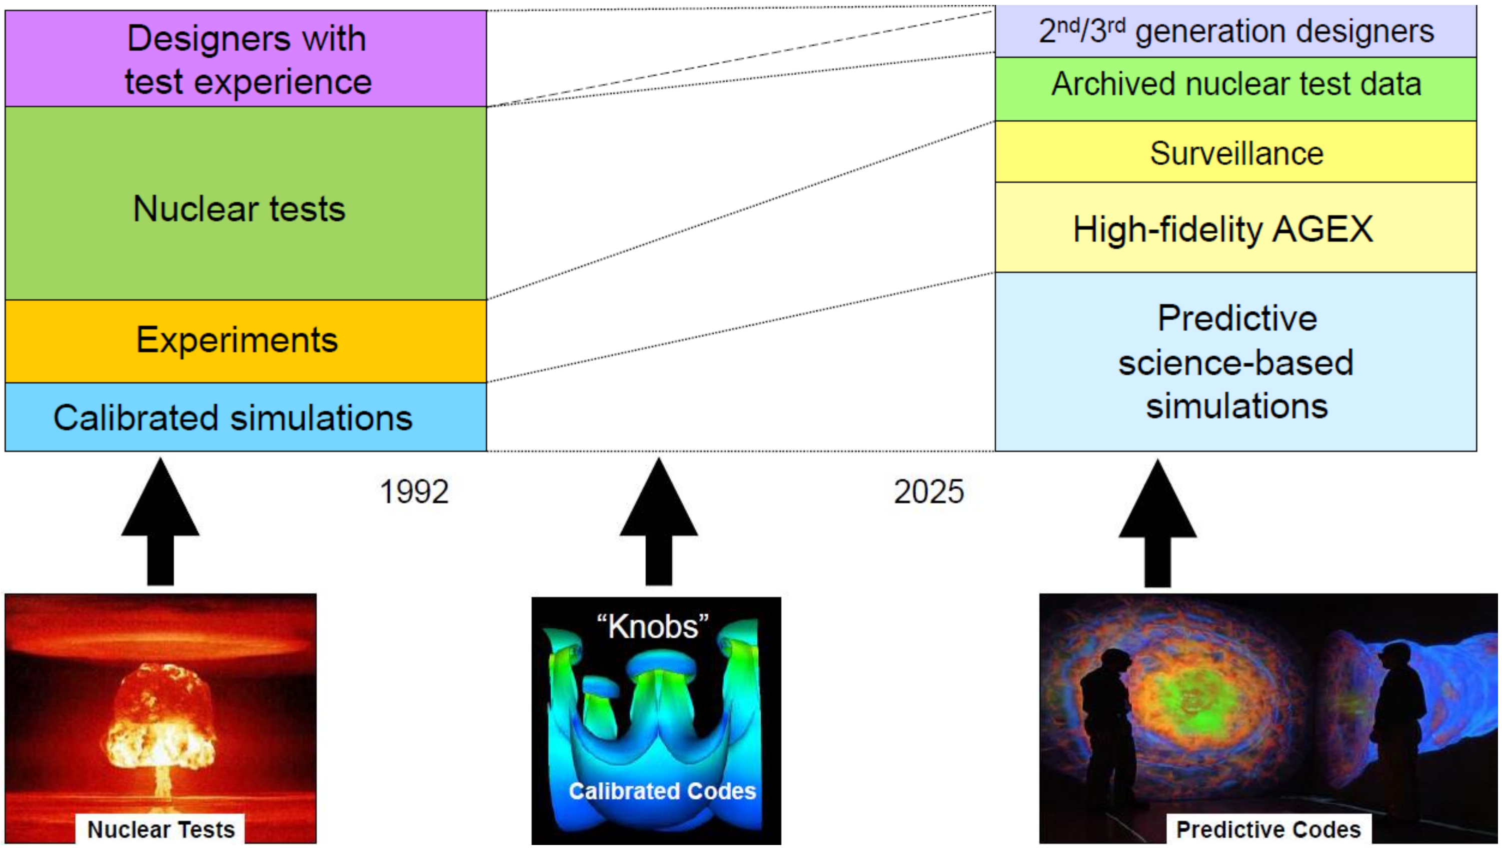

Code validation through experiments and test history remains essential

There were reasons why ASCI was successful and these should be kept in mind as we plan future initiatives….

- **Have clear programmatic requirements**
  - 3D validated codes to replace underground tests (UGTs)
- **Set a realistic but aggressive goals**
  - Initial goal: 100 TF 3D full-system calculation
- **Integrate planning to achieve the goals**
  - Manage all the elements needed for success and integrate these through planning
- **Enjoy freedom to exploit technological innovations and increased understanding**
  - Commodity clusters, advanced architecture systems, uncertainty quantification, and “predictive” simulation, advanced facilities

ASCI End Point: 100-TF IBM Purple

Courtesy: M. McCoy, LLNL

As we enter the next era of stewardship, challenges are evolving

Predicting with confidence an older & small stockpile

- **High accuracy in individual weapons calculations:**
  - Advanced physical models
  - 3-D to resolve features and phenomena
- **Uncertainty quantification:**
  - Ensembles of simulations to explore impacts of small variations

Expanding threats (& role for data)

- **Hostile Environments**
- **Foreign Assessments**
- **Technological Surprise**

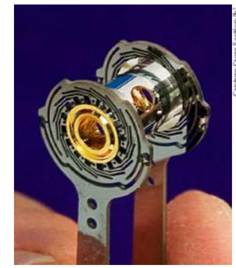

High Energy Density (HED) Experiments

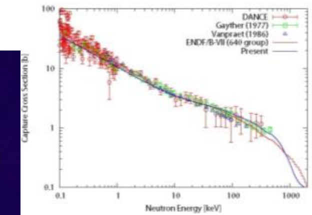

Nuclear Measurements

Dynamic Material Properties

Computing

Confidence

Computing enables resolution & fidelity at scale AND flexibility & agility

Los Alamos NATIONAL LABORATORY

High performance computing continues to be essential for Stewardship

ASC Strategic Objectives

- Robust Tools
  - Models, Codes, Techniques
- Prediction through Simulation
  - Verified and Validated Codes
- Balanced Operational Infrastructure
  - Platforms and Infrastructure

Partnering with SC and academic community helps enable success

Historic trends indicate that memory requirements increase dramatically with physics fidelity

- **Prediction fidelity over the last two decades is dominated by:**
  - Increased resolution
  - Increased dimensionality
  - Subgrid model complexity
- **Successful experimental validation of models is**
  - Universal at higher resolution and full dimensionality
  - Tuned when we rely on locally valid sub-grid models

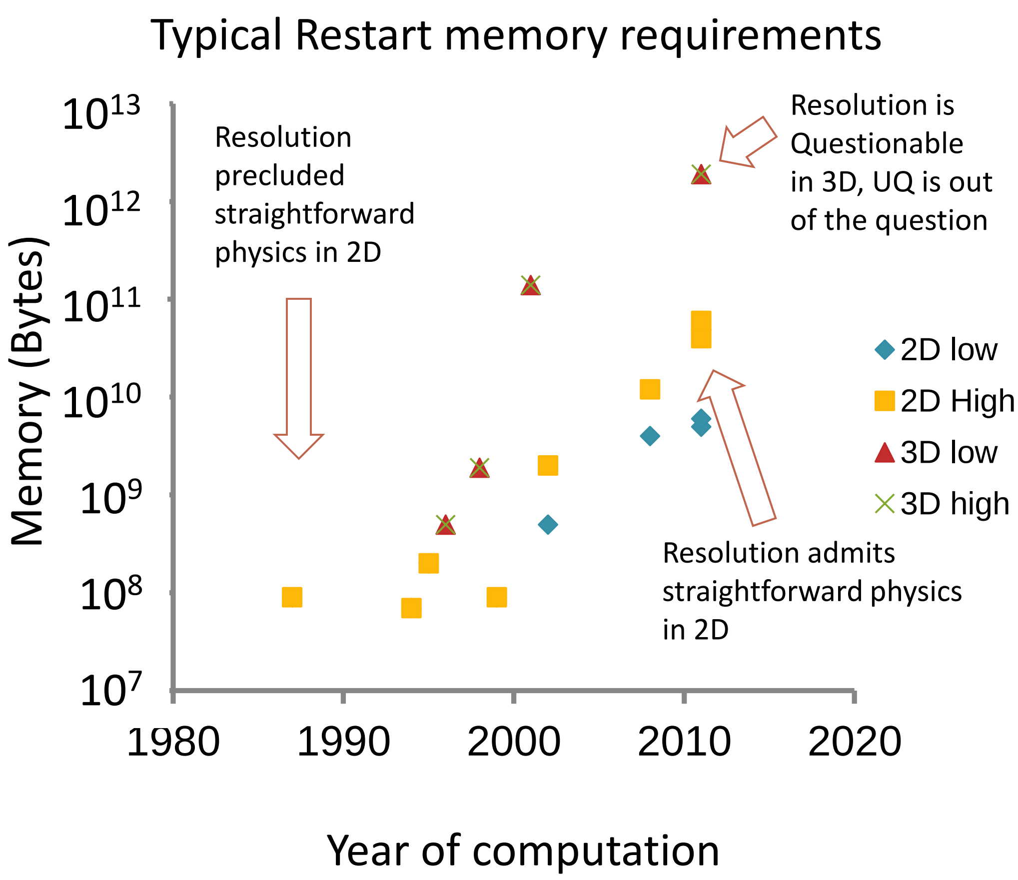

Increased resolution in simulation is changing our picture of how weapons work

The Tri-Lab has an integrated computing roadmap to meet its mission needs (& it's a busy time)

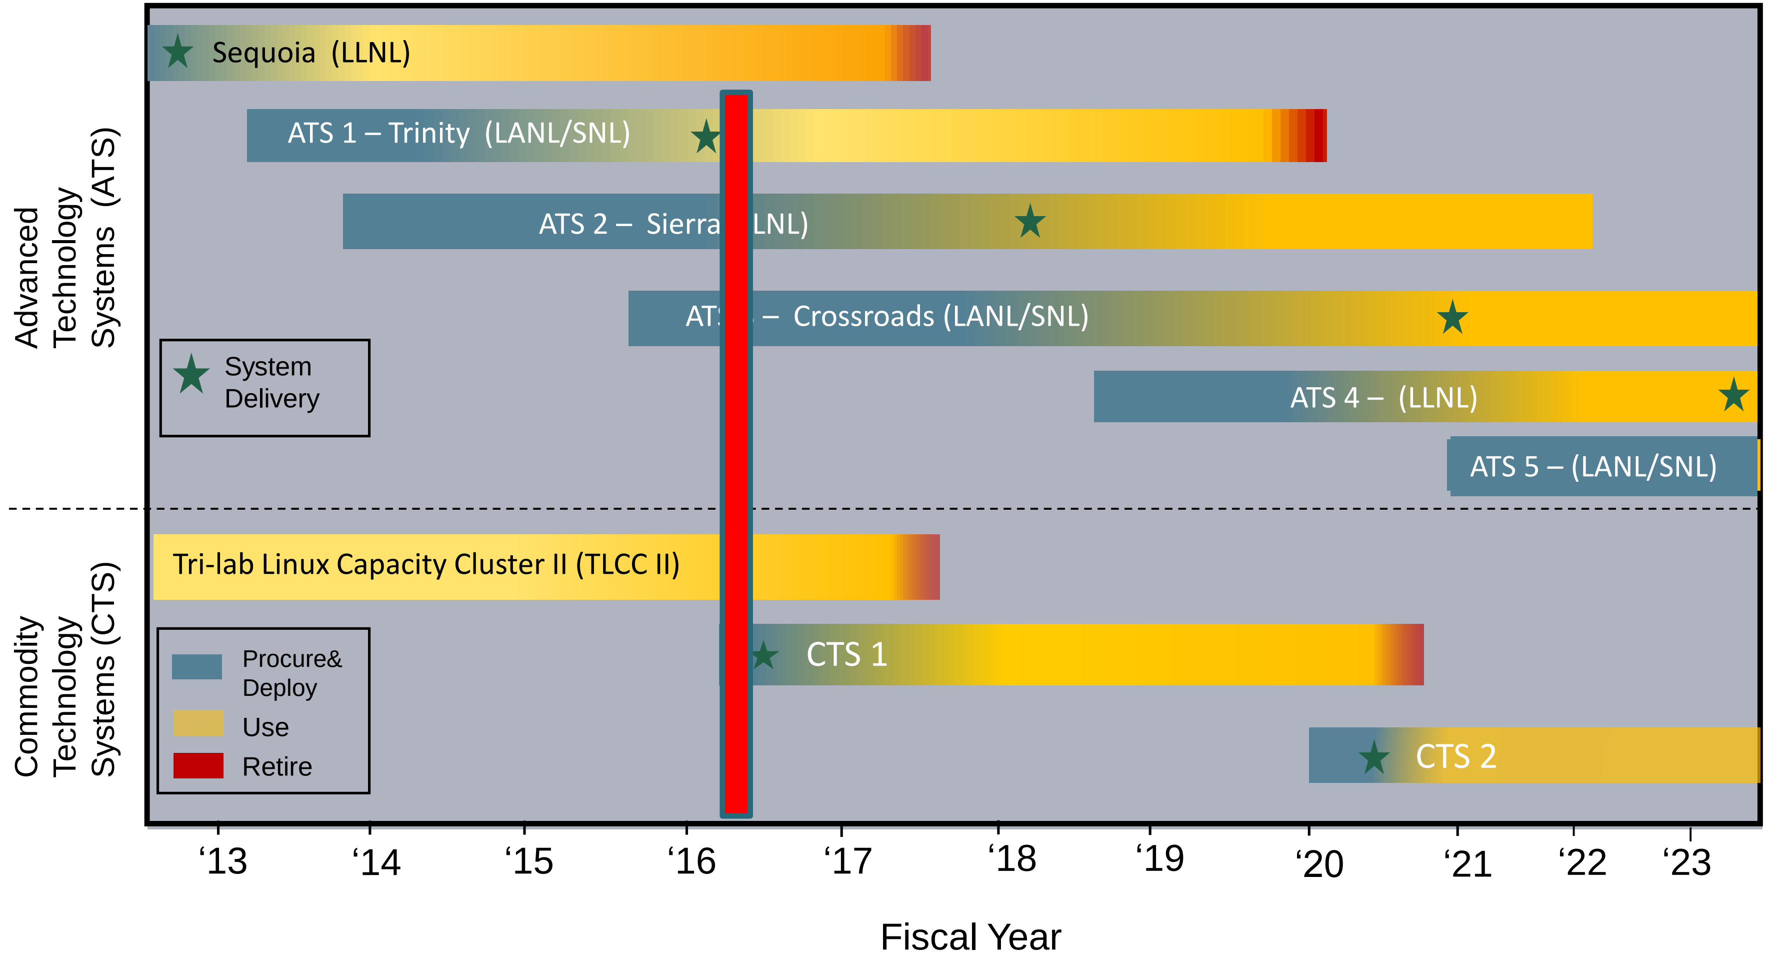

Los Alamos NATIONAL LABORATORY

Crossroads (ATS-3) will continue the technology disruption on the path to exascale

- **Specified by memory needed to satisfy the mission needs: 4 PB to 10 PB**
  - NOT flops
- **Technology disruptions expected:**
  - Processors will be many-core and heterogeneous
    - Light and heavy cores on the same silicon
  - Memory DIMMs will be replaced with stacked (3D) memory
    - Stacked memory is fast but of low capacity
  - Spinning disks will mainly be replaced by multiple layers of flash drives
- **Acquisition is ASC-ASCR partnership**

CROSS ROADS

Mission need is
- 3D
- High geometric fidelity
- High numerical fidelity
- Medium physics fidelity

APEX
LANL • LBNL • SNL

Los Alamos NATIONAL LABORATORY

Co-design of relevant applications (e.g., mesoscale materials) remains an important frontier

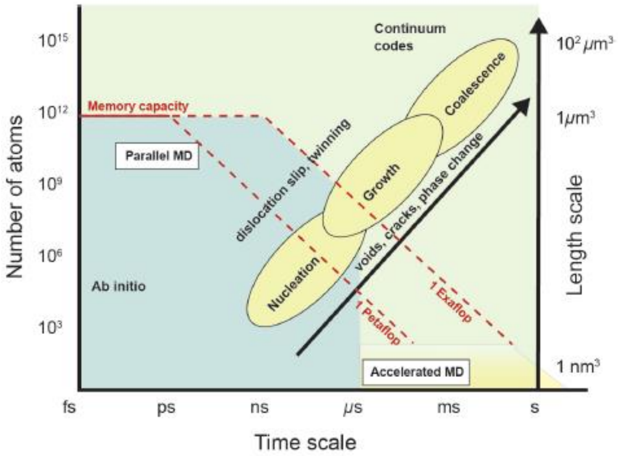

Mesoscale materials phenomena need extreme-scale computing

Variable-resolution models are synergistic with multi-probe, in-situ, transient measurements

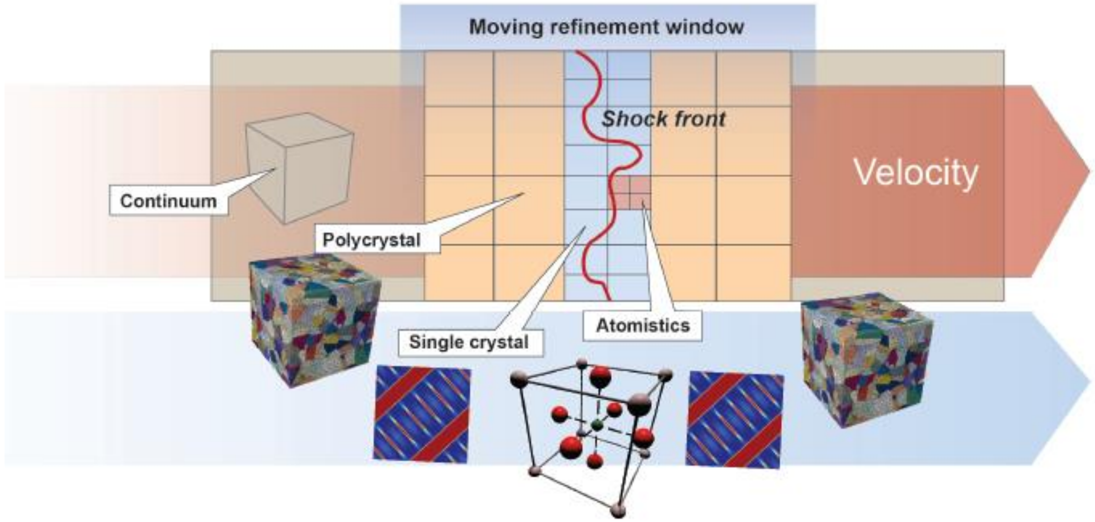

The development of validated models will reduce uncertainty in integrated codes and provide predictive descriptions of newly manufactured materials & components

MaRIE will provide critical data to inform and validate advanced modeling and simulation to accelerate qualification of advanced manufacturing – move from “process-” to “product-based”

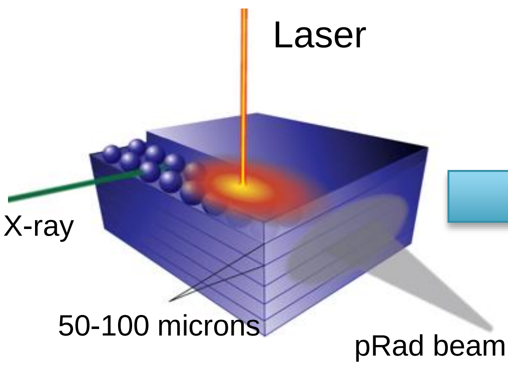

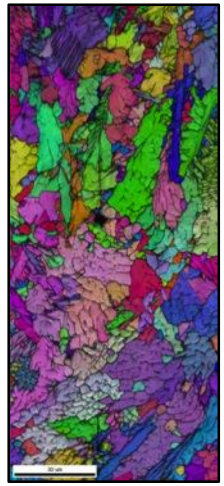

Additively Manufactured 316L Stainless Steel

Damage under shock loading unable to coalesce leading to a tougher, more shock resistant material

Prototype Gas Bottle

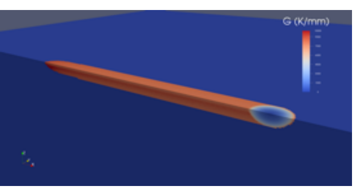

Process Modeling

Microstructure Modeling

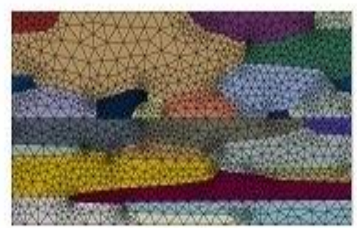

Properties Modeling

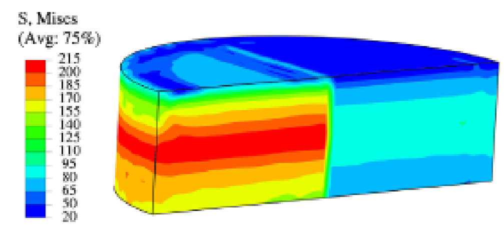

Performance Modeling

Los Alamos' stewardship strategy relies on high performance computing and co-design

Exascale will allow a predictive capability with simulations of multiple coupled physical process

- **Exascale capability is needed to perform:**
  - Detailed highly resolved 3D boosted nuclear performance simulations
    - Validate understanding against experimental data
    - Discriminate physical mechanism dominance, to assure emergent behavior is correctly predicted.
  - Detailed highly resolved 3D nuclear safety simulations
  - Lower-resolution 3D performance simulations with uncertainty quantification (UQ)
    - UQ achieved through in situ capabilities or ensembles
    - Sub-grid models capture unresolved physics
  - Molecular dynamics simulations of adequate resolution

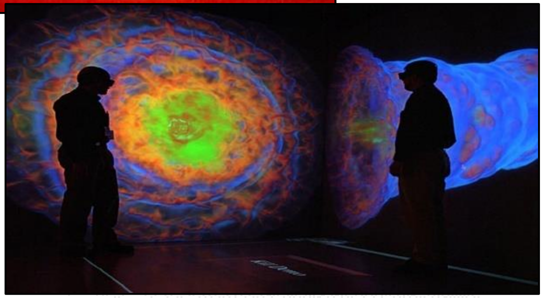

Maintaining stockpile reliability AND safety requires highly resolved 3D simulation to evaluate that balance.

Stockpile Stewardship at the Exascale:
Understanding boost

Nuclear Performance

- Exascale, for the first time, will allow:
  - 3D boost simulations with multiple coupled physical processes at unprecedented resolution
  - Detailed highly resolved 3D nuclear safety simulations
  - UQ performed in 3D at lower resolution with sub-grid models to capture unresolved physics

Boost is just one example where detailed high resolution 3D simulations are required to improve predictive capability of nuclear weapon performance

Weapons Science

Exascale will allow resolution of important length scales with appropriate fidelity possibly *in situ* with performance simulation

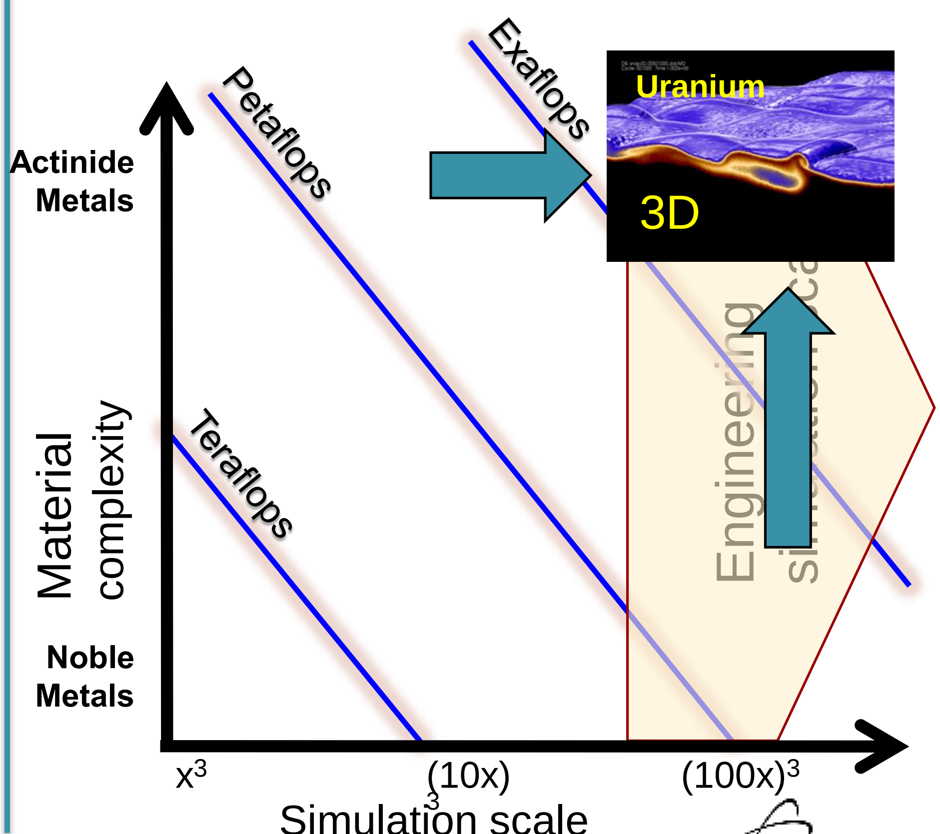

ASC got an early start on Exascale through ATDM: Structure of pre-exascale program

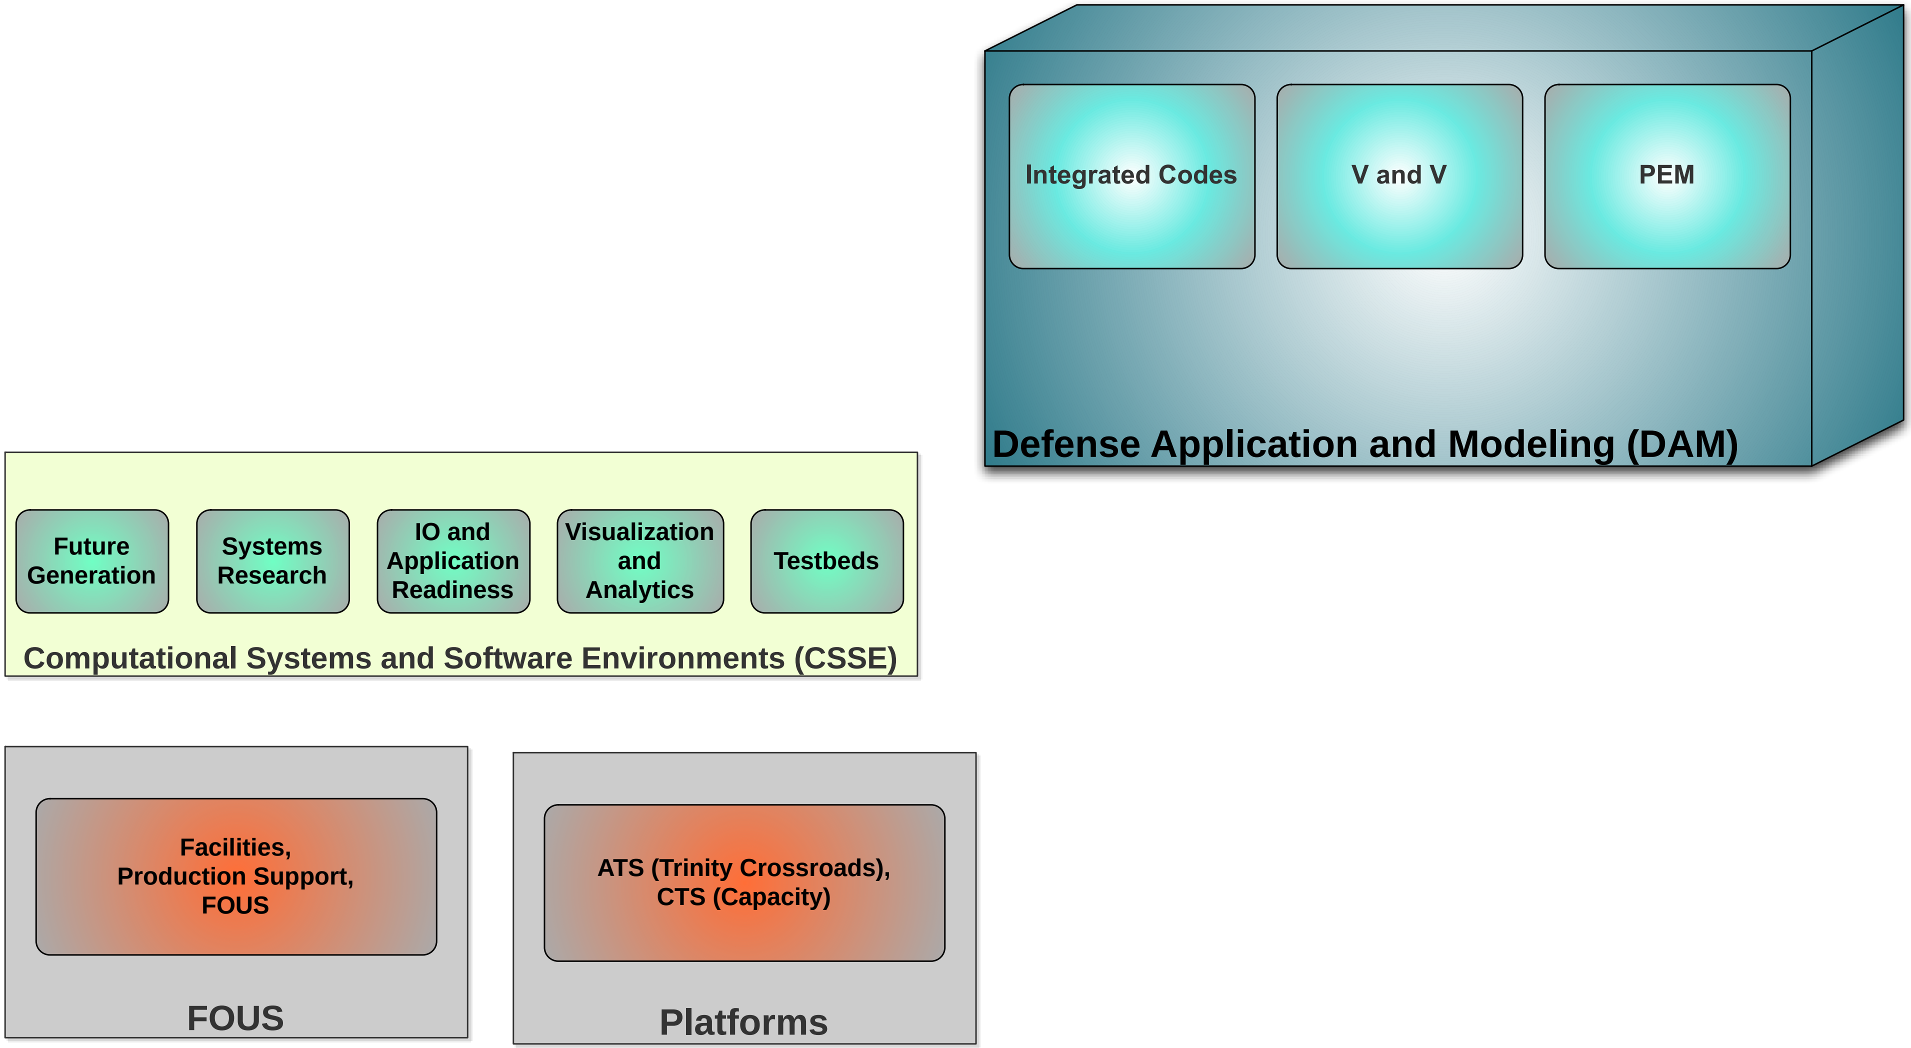

ASC got an early start on Exascale through ATDM: Responding to mission challenges & new technology

Advanced Technology Development and Mitigation (ATDM)

- Next Generation Code Development & Applications (CDA)
- Next Generation Architecture and Software Development (ASD)
- Future High Performance Computing Technologies

Defense Application and Modeling (DAM)

- Integrated Codes
- V and V
- PEM

Computational Systems and Software Environments (CSSE)

- Future Generation
- Systems Research
- IO and Application Readiness
- Visualization and Analytics
- Testbeds

FOUS

- Facilities, Production Support, FOUS

Platforms

- ATS (Trinity Crossroads), CTS (Capacity)

ASC got an early start on Exascale through ATDM: Strong complementarity with ASCR research

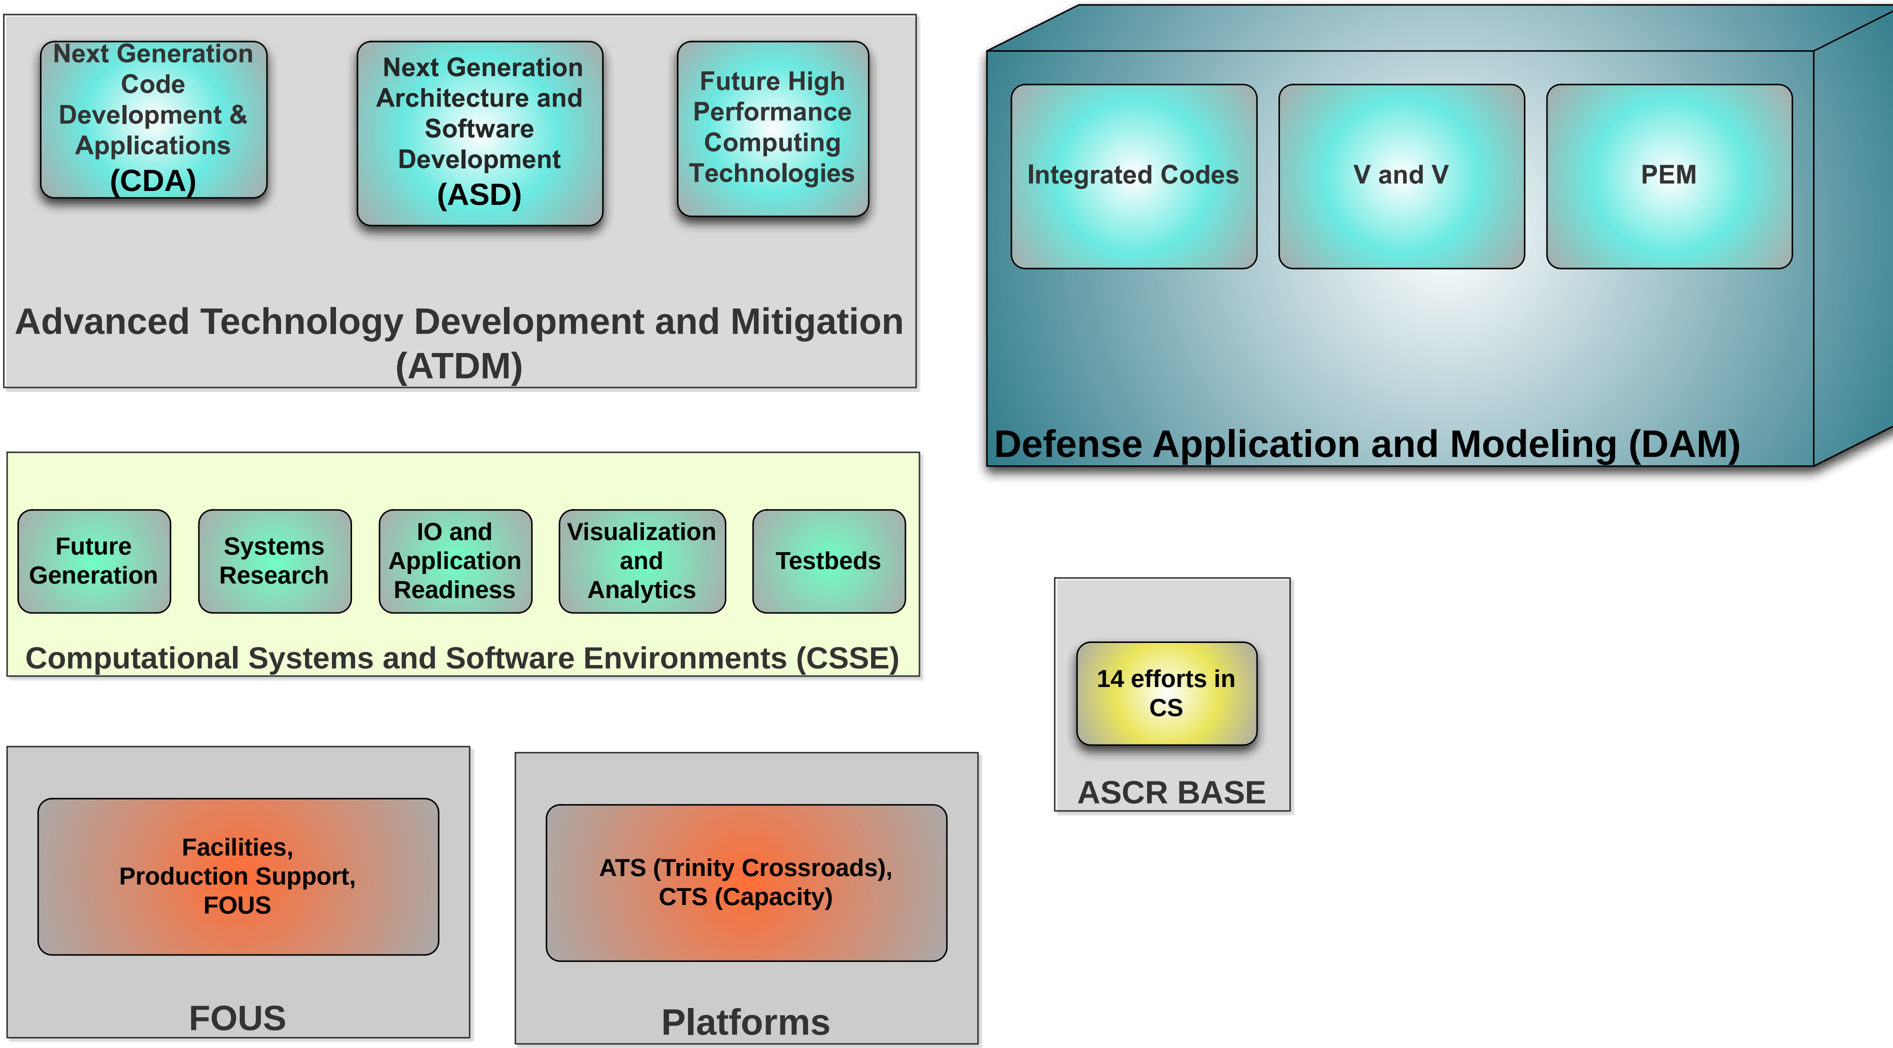

Integration is a blessing and a curse (and important for success)

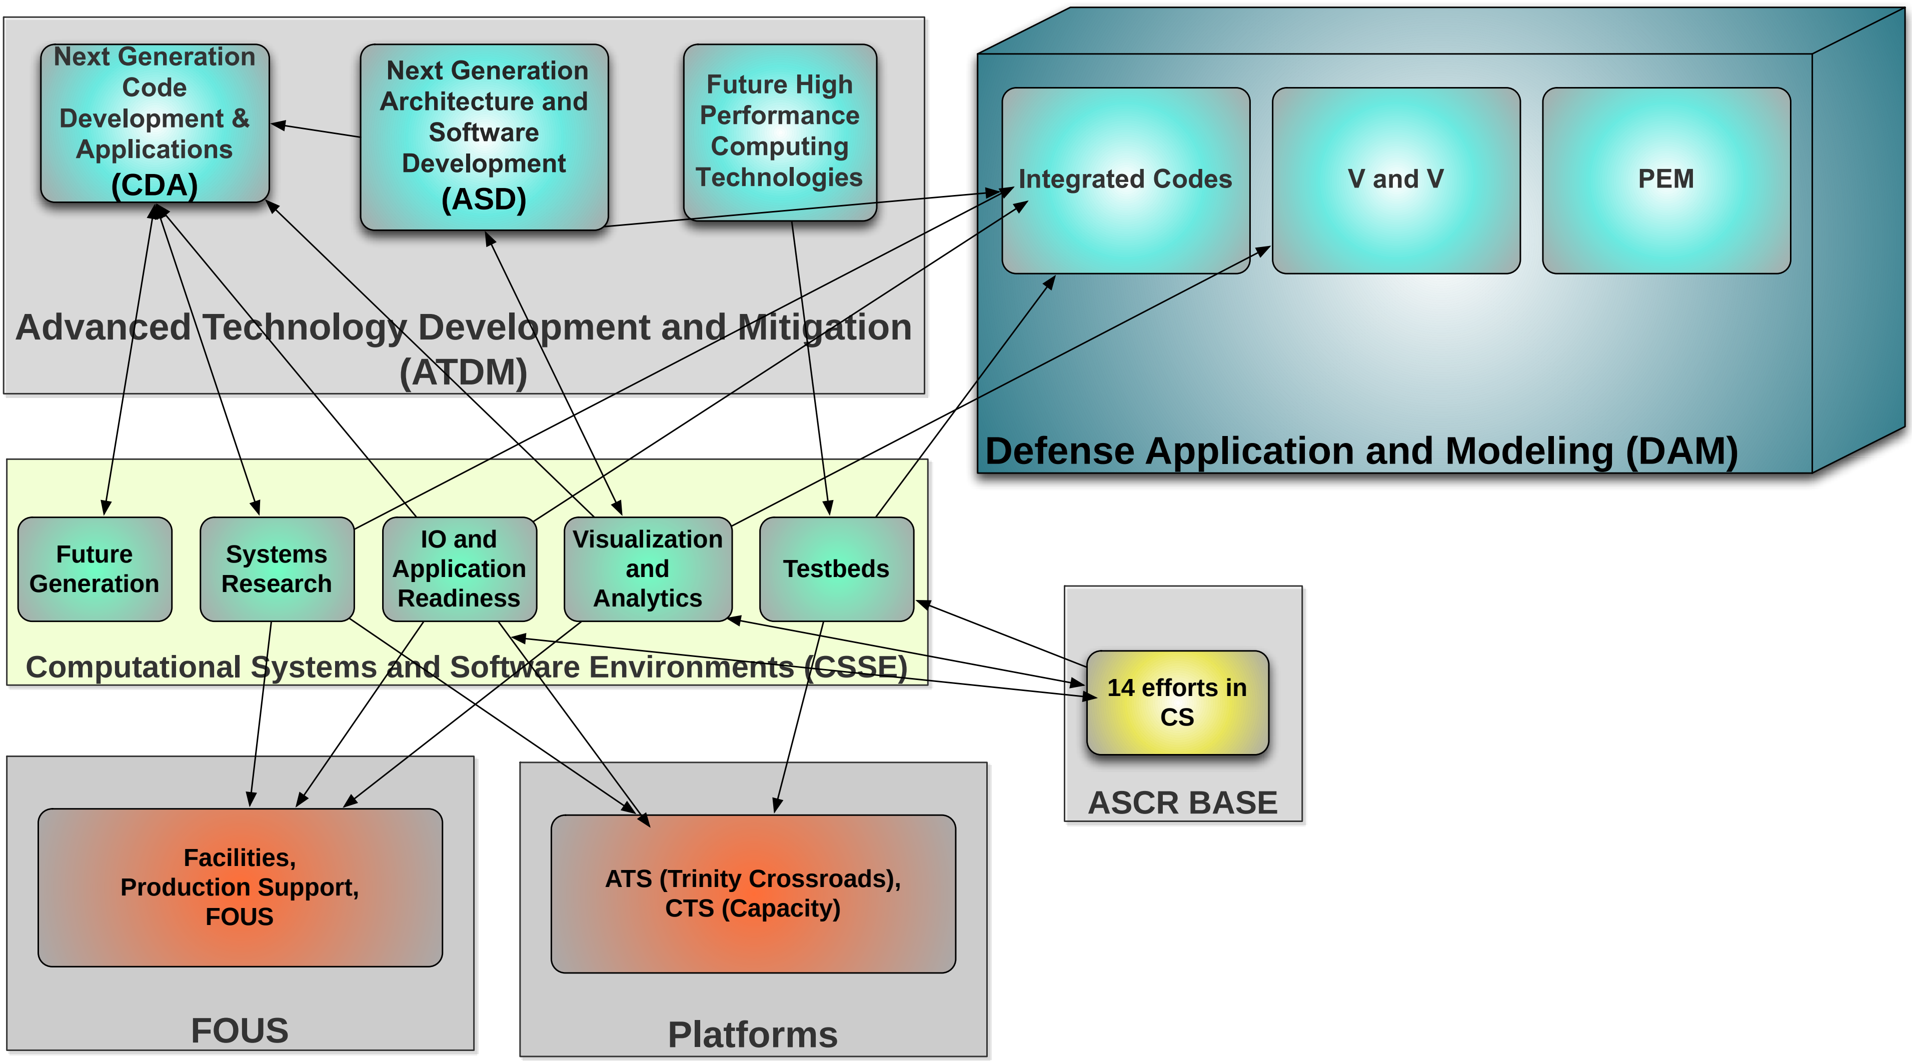

ECP is a key element of ensuring the future role of HPC in stewardship

- **From a national security perspective, exascale is important and urgent**
  - Resolution and Fidelity at scale is needed
  - Flexibility and Agility creates options and mitigates technological surprise
- **Exascale Computing Project is the right modality for success**
  - Lab lead project spanning SC and DP
- **A holistic approach focused on capable exascale is the right approach**
  - Focus on performance, not flops
  - Need to enable the ecosystem, not stunts
  - Computation and data analytics are converging & we should move fast
- **The Labs are committed and we're putting our best talent on the effort**

McMillan to ECP CD-1 Review, 9/16

The United States is retaining weapons in the stockpile beyond their original design life. How long will they be viable? What are their lifetimes?

As weapons age,

- components can deteriorate due to operational environments
  - Vibrations
  - Temperature cycling
  - Humidity
- materials can change as a result of intrinsic radiation and chemical reactions
  - Plastics can become brittle
  - Adhesives can weaken
  - Metals can corrode, metal coatings can deteriorate
- material physical properties can change
  - Loss of ductility, elasticity or strength

Weapon	Date Entered Service
B61-3	1979
B61-4	1979
B61-7	1985
B61-11	1997
B83	1983
W80-1	1982
W76	1978
W78	1979
W87	1986
W88	1989

Result can be formation of gaps, cracks, swelling and material displacement

How long can the current stockpile be sustained?

Weapons eventually need refurbishment, and undergo a life extension. Involves mix of reused, remanufacture, and replacement components

- Increasing number of materials in original weapon builds no longer commercially available
- Reuse of components can imply reemploying decades-old components
- Processes for remanufacturing components have evolved
- Experience base with particular replacement components may be limited
- Weapon performance must be certified under hostile nuclear threat environments
  - Blast, particle debris, x-rays, gamma rays, neutron dose & dose rate, EMP

B83 Bomb - total parts = 6,519

- When practical, we are to examine advancing weapon safety and use control

What are options that can be developed and certified without further nuclear testing?