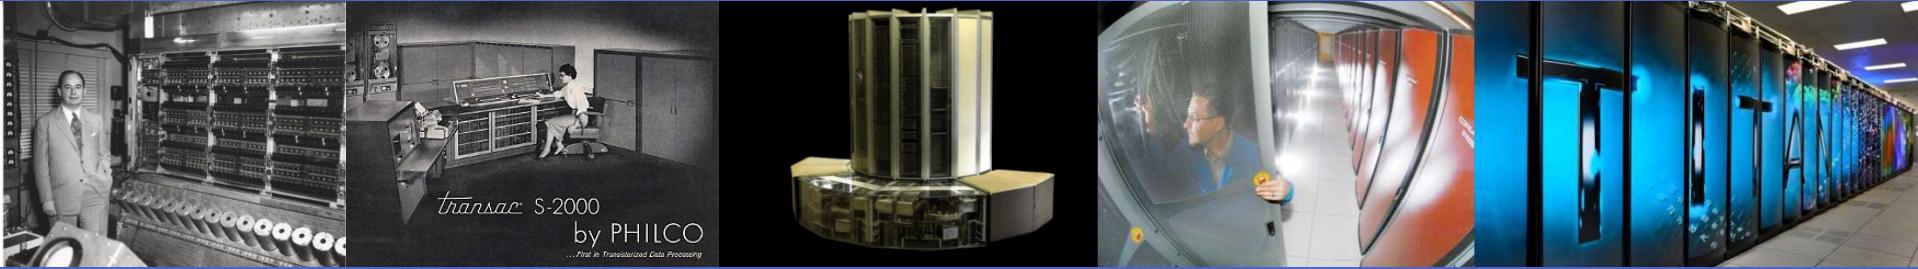


*Exceptional service in the national interest*



## Historical Impact of Government Investment in High Performance Computing

IEEE International Conference on Rebooting Computing  
San Diego, CA  
October 17, 2016

Rob Leland  
Vice President, Science & Technology  
Chief Technology Officer  
Sandia National Laboratories



Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Outline

- **Some context**
- **My thesis**
- **Historical review of major transitions in computing**
- **Laying out the pattern**
- **Looking forward**

# Some context

- **Why might this be important?**

- Many indicators suggest we are in need of a major advance in computing
- If we better understand how such advances have come about in the past, we are more likely to generate similar advances in the future

- **What are we really talking about?**

- Thesis: government investment has played a key role in creating transitions in HPC
- High Performance Computing (HPC) means ...
  - *Computing systems operating at or near the maximum performance achievable in a given technological era.*
  - Simulation, analytics, cryptanalysis, streaming, ...
  - Hardware, software, algorithms, applications ... the full stack

- **How can we make use of this thesis?**

- What approach/mindset is suggested by history?

# My thesis (in more detail) regarding evolution of HPC

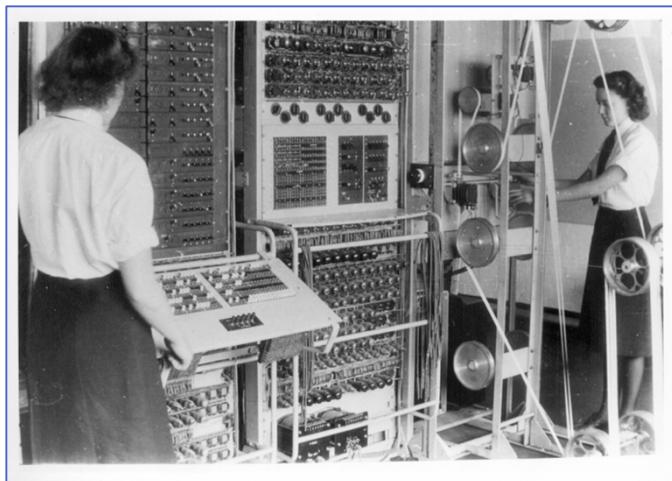


- **Periodic shifts in architecture and programming models**
  - Response to national security context in which ...
  - Prevailing capability seen as insufficient for future defense or intelligence needs
- **Overlaying shifts with government investments shows interesting pattern**
  - Government funded pilot/prototype typically *preceded* shift by 5-7 years
- **Suggests government investment has played a key role in catalyzing shifts**
  - Resonates for many in field
  - Particularly if they lived first hand through an example
- **Acknowledge**
  - Based on subjective assessment of historical pattern
  - Correlation NOT= causation
  - Other contributing factors
  - Relative contribution often not clear
  - Some major government investments did not lead to a shift
  - Viewed from predominantly U.S. government perspective with a DOE bias

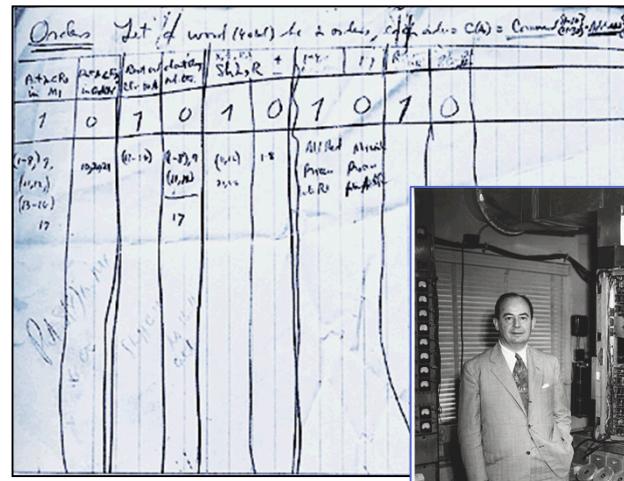
# Early electronic computing era

- **Electronic computing developed to meet military needs in WWII**

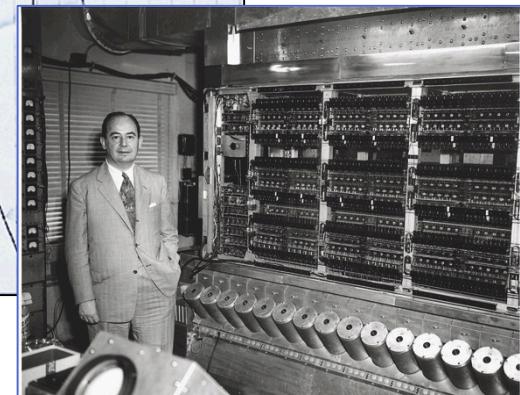
- Colossus, Bletchley Park, 1943 ... code breaking, dedicated function
- ENIAC, U. Pennsylvania, 1945 ... ballistics tables, plugboard programming
- Von Neumann
  - Initiated 1946 at Institute for Advanced Studies (IAS)
  - Vacuum tubes, oscilloscopes, assembly language ... many operational challenges
  - But momentous – flexible stored program, reliability architecture, hydrogen bomb
  - IAS, Princeton, 1951
  - MANIAC, LANL, 1952
  - ORACLE, ORNL, 1953



## *Women's Royal Naval Service operating Colossus during World War II*



*First line of code written for  
the von Neumann Digital  
Computing Project*



## *von Neumann with the IAS machine*

# What lessons to draw?

## ■ National urgency

- World War II needs
- Experimental feedback – Feynman's approach not scalable
- Cold War drivers

## ■ Technological opportunity

- Mature technology base (vacuum tubes)
- Innovation – tube memory, physical layout, reliability architecture
- Engineering expertise

## ■ Conceptual advance

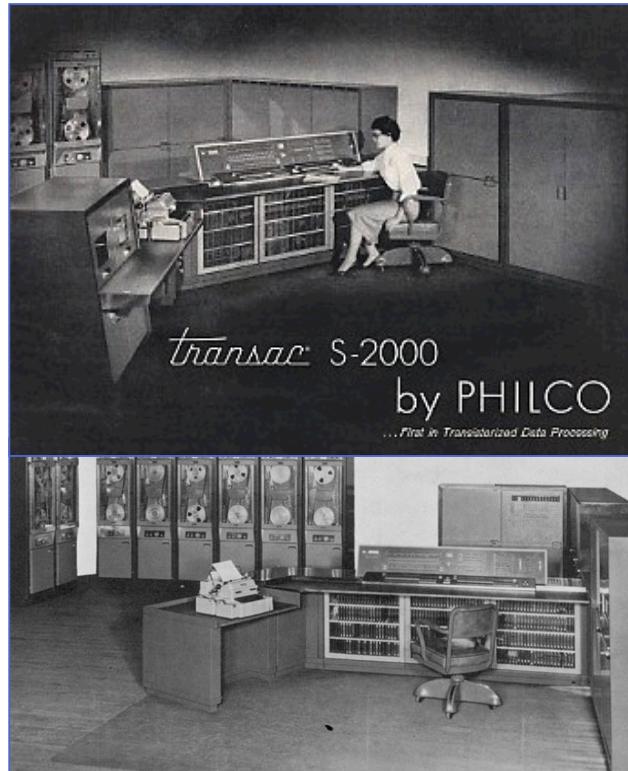
- Previous theoretical advance (Turing)
- Stored (flexible) programming

## ■ Government response

- Scientific leadership
- Government investment

# Mainframe era

- **Government investments for military and intelligence applications led to civilian “mainframe” market:**
  - Data processing
  - Weather, climate modeling
  - Scientific R&D
- **IBM701, 1953, close copy of IAS machine**
  - Vacuum tube-based system
- **Solid state computing**
  - Philco's transistorized computer, 1955, NSA funded
  - Transac S-2000, 1957, first commercial version
- **NSA and AEC (DOE predecessor) are primary sponsors of HPC mainframe development throughout 1960's**
- **CDC 6600, 1966, federal sponsorship**
  - Seymour Cray design
  - Fast logic, RISC, functional parallelism
  - Order of magnitude performance improvement
  - Inspired famous TJ Watson “janitor” memo



A large-scale, high speed, all transistorized, electronic data processing system.

**STORAGE:**

Memory Capacity — a basic core memory unit of 4,096 words expandable to 65,536 words in steps of 4,096 words.

Magnetic Drum capacity is 32,768 words (262-144 alpha-numeric characters). Up to 256 drums allowed.

**INPUT/OUTPUT:**

Magnetic Tape, Punched Paper Tape (5, 6, 7 or 8-channel), Punched Cards (80 column), High Speed Printer (Off-line).

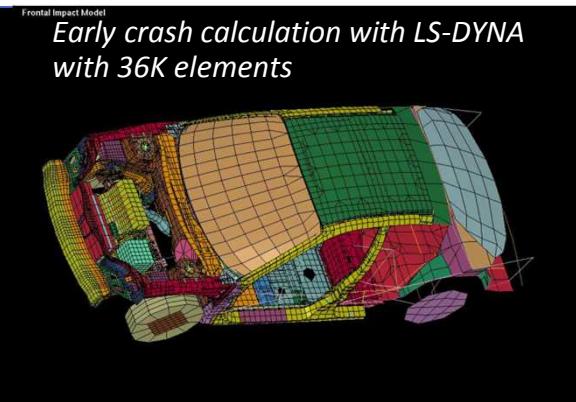
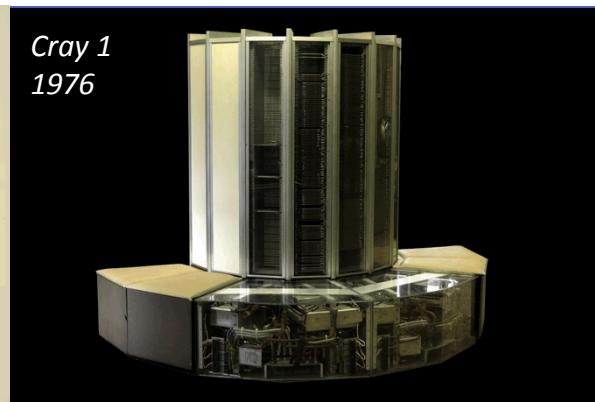
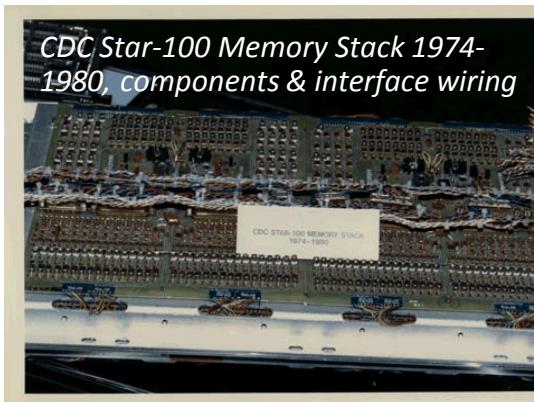
**PRICES:**

Monthly rental — \$30,000.  
Purchase price — \$1,450,000.

*In FY16 \$257K  
\$12.4M*

# Vector era

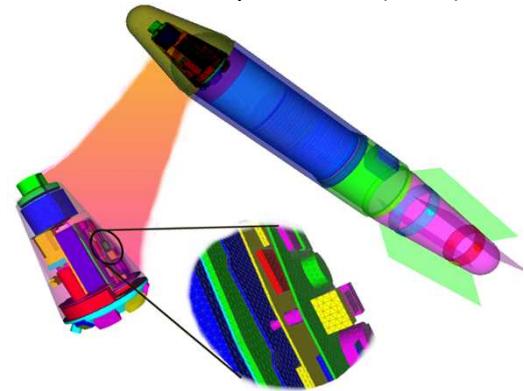
- In the 1970s, National Lab needs outpaced performance growth of mainframes, driving development of vector-based (data parallel) systems
- **CDC STAR-100, LLNL**
  - Million word memory
  - Legacy codes a poor fit
  - Lab overhauls codes to take advantage of vector machine architecture
- **Seymour Cray leaves CDC to form Cray Research**
  - Cray 1 builds on STAR-100, adds optimized scalar processing unit, increased memory size
- **Critical partnerships between Labs and computing suppliers**
  - New machines and software for vector systems
  - Aids expansion of supercomputing into new areas (oil and gas, aerospace, automotive, ...)



# Distributed memory MPP era

- **Attack of the killer micros, 1980's**
  - Personal computer market drives key components to commodities
  - Used as building blocks in Massively Parallel Processing (MPP) architectures
- **High Performance Computing and Communications Initiative, 1991**
  - Multi-agency federal program (DoD, DOE, NASA, NSF, ...)
  - Early MPP machines, algorithms, tools, applications, education
  - Fiber optical network, web browser, ...
- **DOE/NNSA Accelerated Strategic Computing Initiative, 1995**
  - Accelerated shift from vector systems to MPP systems
  - ASCI Red, 1996, Intel/SNL – first terascale system
  - Blue Gene, 2005, IBM/LLNL – advanced power efficiency – led to commercial product (IBM BG family)
  - Red Storm, 2005, Cray/SNL – high bandwidth interconnect – led to commercial product (Cray XT family)
  - Linux cluster commoditization of MPP technology – low cost capacity systems -- multiple commercial products

*SALINAS weapon system simulation;  
ASCI Red, 3000 processors (2002)*



*Cray Red Storm, SNL*



*IBM Blue Gene, LLNL*

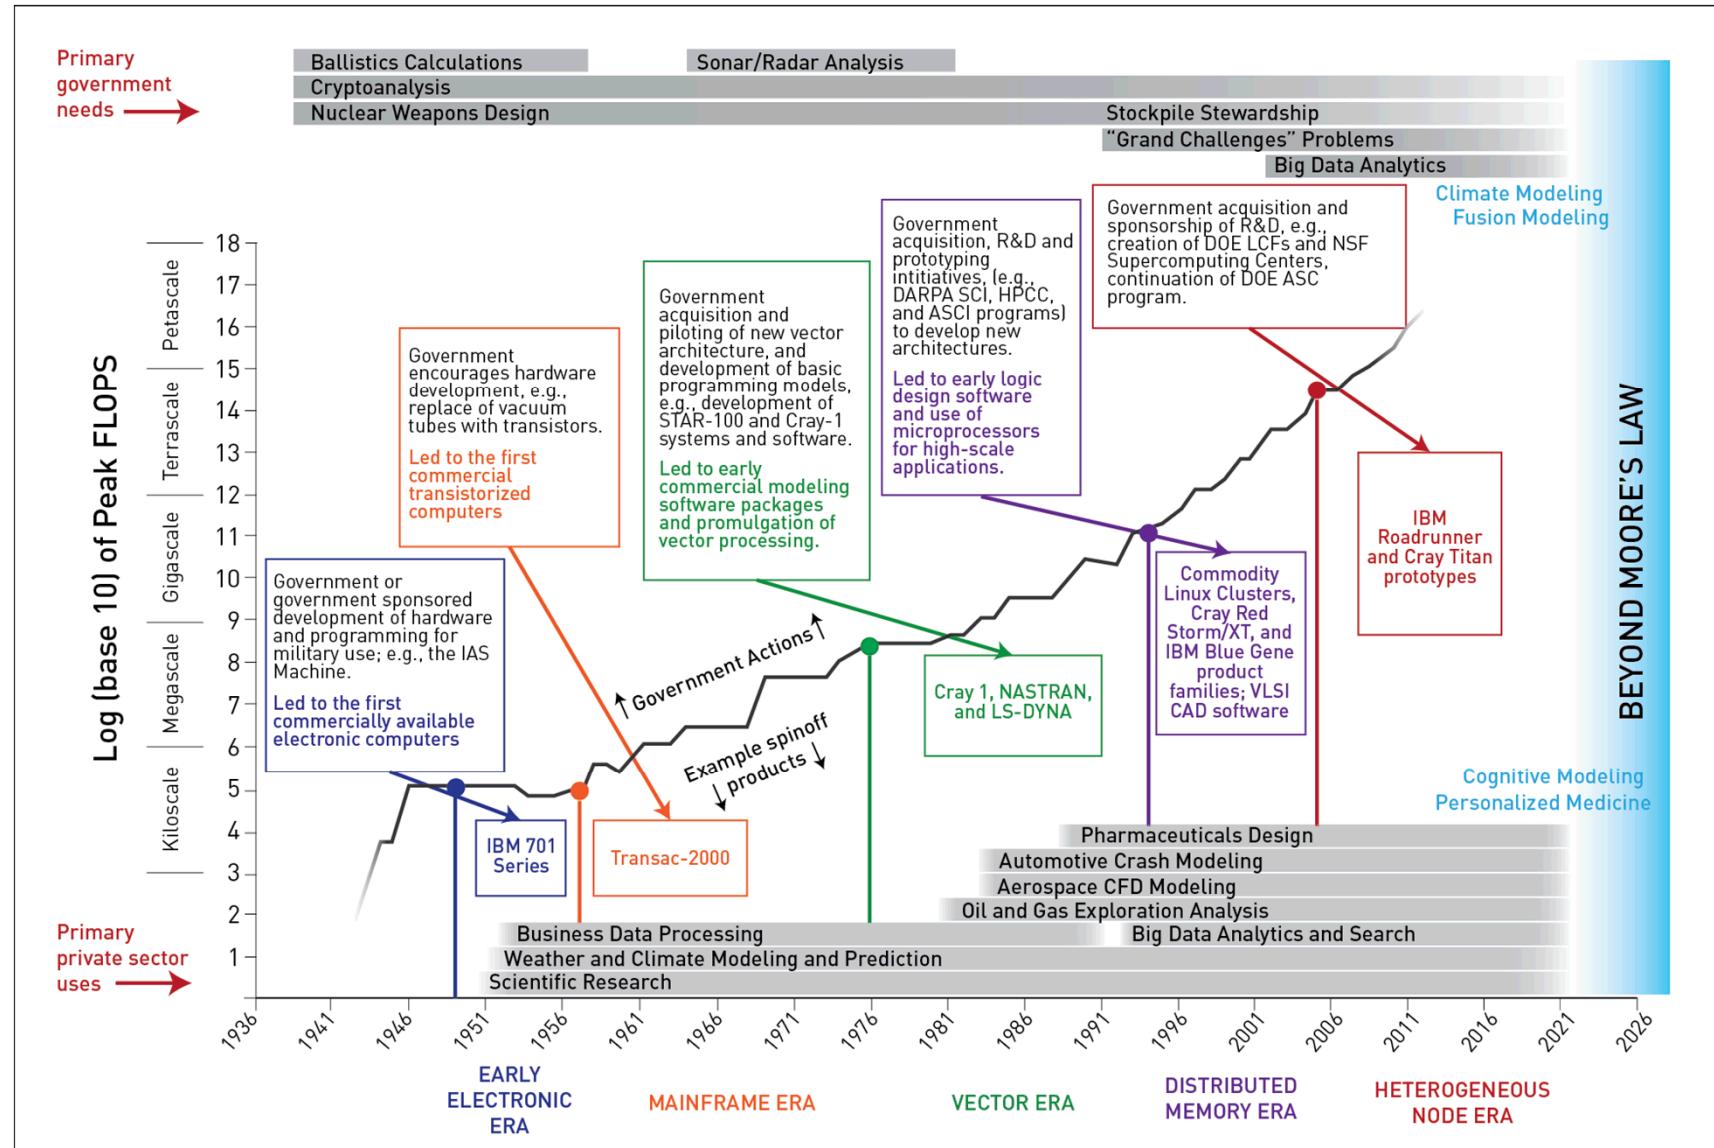


# Heterogeneous node/many core era

- **Shift from scaling up to scaling out**
  - Add processing cores to a single chip to extend gains from existing technology
  - Add accelerator chips to node to leverage commodity technology from other markets
  - Software development focus on associated parallel processing challenges
- **DOE/NNSA ASCI is formalized as the ASC Program, 2004**
  - Roadrunner, 2008, IBM/LANL – hybrid CPU/accelerator architecture
  - Commodity Technology Systems – 7 petaflop capacity -- “Scalable Unit” Linux clusters
  - Advanced Technology Systems – goal of 40+ petaflop capability
    - E.g. Trinity, 2015, Cray/LANL/SNL, 11 petaflops, “burst buffer” storage, advanced power management
- **DOE/SC ASC Research Program (ASCR) to supply HPC**
  - National Leadership Computing Facilities at ORNL/ANL, 2004
    - Titan, 2012, Cray/ORNL – hybrid CPU/GPU processor architecture
    - Mira, 2012, IBM/ANL – third generation Blue Gene



# Historical investment patterns and eras of HPC



# Looking forward

- **National urgency**
  - Increasing international competition
  - Erosion of Moore's Law (and lead time to change course)
  - Rise of Big Data
  - Coming to the end of the MPP era
- **Technological opportunity**
  - New devices
    - TFETs, carbon nanotubes, spintronics, ...
  - New architectures and packaging
    - 3D, reconfigurable, superconducting, ...
- **Conceptual breakthrough**
  - "New" computational paradigms
    - Approximate, analogue, reversible, neuromorphic ...
  - New mathematical formulations
- **Government response**
  - National Strategic Computing Initiative (NSCI)
  - Executive Order issued July 2015
  - Ten agency joint vision, strategy, roadmap
  - Five strategic objectives ...



**Strategic drivers for change**

# NSCI Strategic Objectives

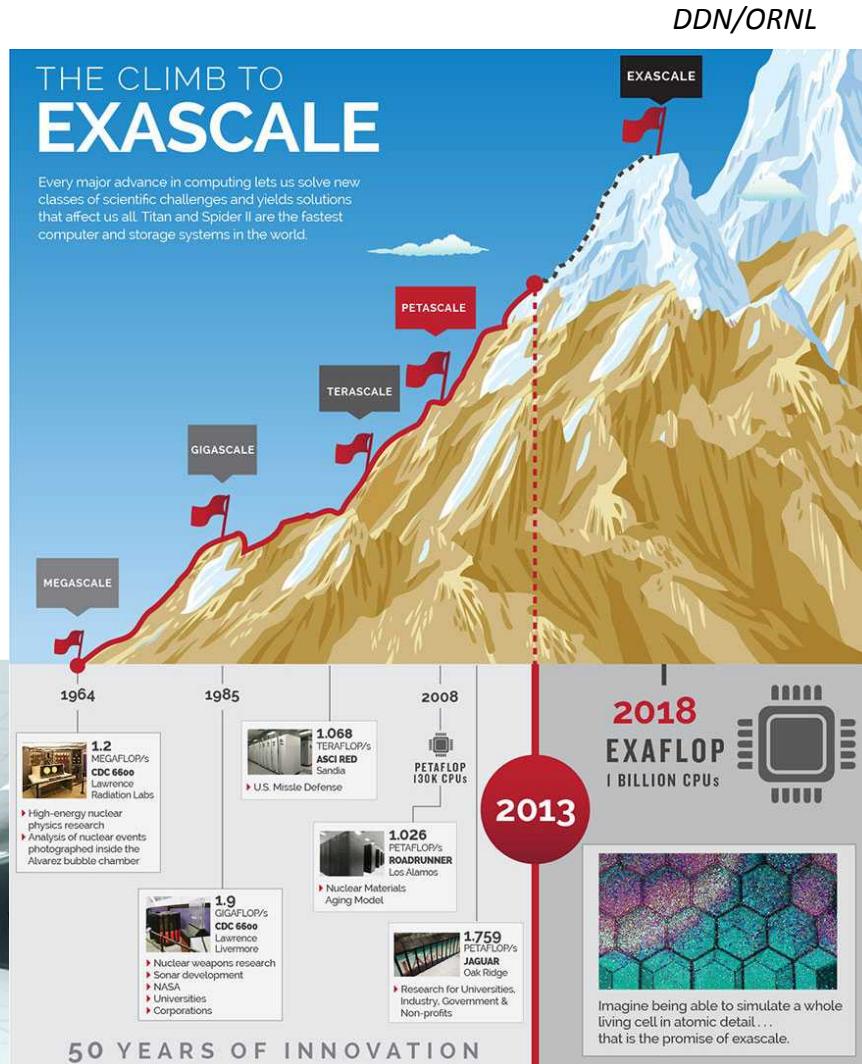


- (1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
- (2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
- (3) Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post- Moore's Law era").
- (4) Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
- (5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

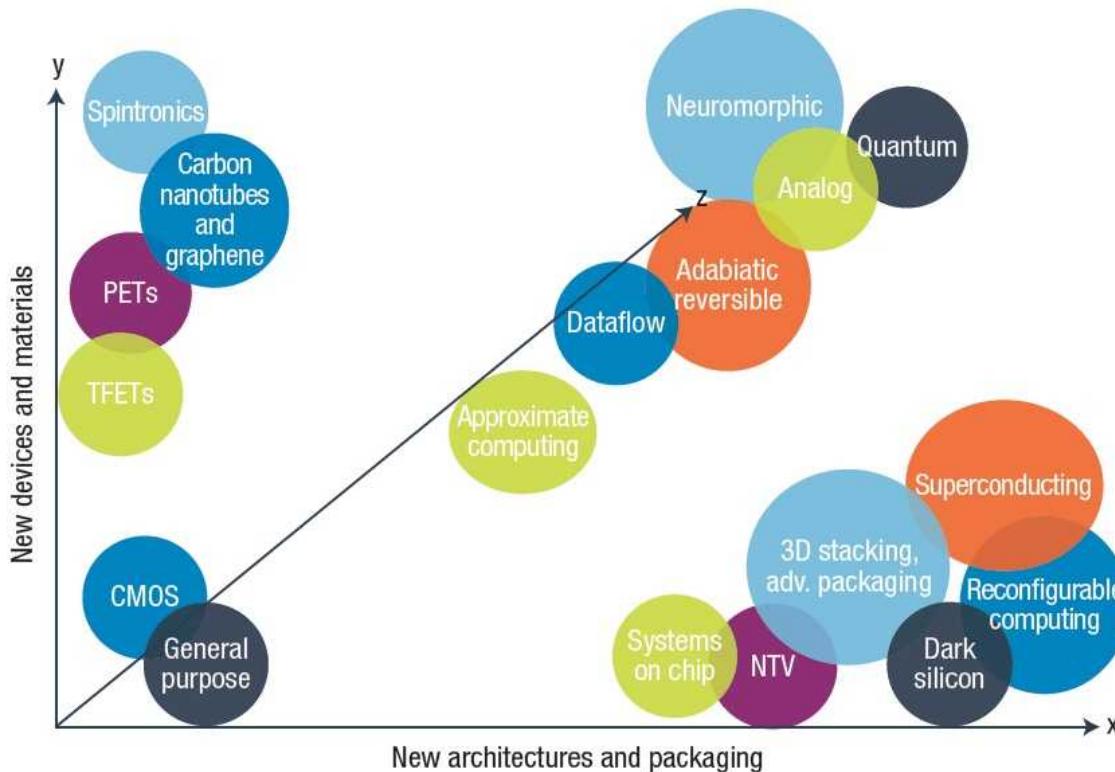
# Capable Exascale Computing

- **DOE Exascale Computing Project**
  - Full HPC stack
  - Partnership between DOE/SC and NNSA
  - Design criteria:
    - 50x improvement in application performance
    - 20-50 MW peak power
    - 1 week average MTBF
    - ...
- **Corresponds ~ end of CMOS roadmap**
  - Likely commercialization as petascale racks
  - What comes next?

ORNL parallel coordinate visualization



# Rebooting computing in the Beyond Moore's Law era



**FIGURE 1.** Technology scaling options along three dimensions. The graph's origin represents current general-purpose CMOS technology, from which scaling must continue. All the dimensions, which are not mutually exclusive, aim to squeeze out more computing performance. PETs: piezo-electric transistors; TFETs: tunneling field-effect transistors; NTV: near-threshold voltage.

Shalf and Leland, Computing Beyond Moore's Law, Computer 48:12, 2015.

# Conclusion

- **On the cusp of another transition in HPC**
- **Key attributes for success**
  - National urgency
  - Technological opportunity
  - Conceptual advance
  - Government response
- **All elements in place**
- **Oh brave new world!**
  - Recognize the opportunity
  - Do our part