National Nuclear Security Administration
DEFENSEPROGRAMS

Applying HPC and AI/ML Capabilities for Stockpile Stewardship Mission

Thuc Hoang
Director, Office of Advanced Simulation and Computing and Institutional Research and Development

October 26, 2021

NATIONAL NUCLEAR SECURITY ADMINISTRATION OFFICE OF DEFENSE PROGRAMS

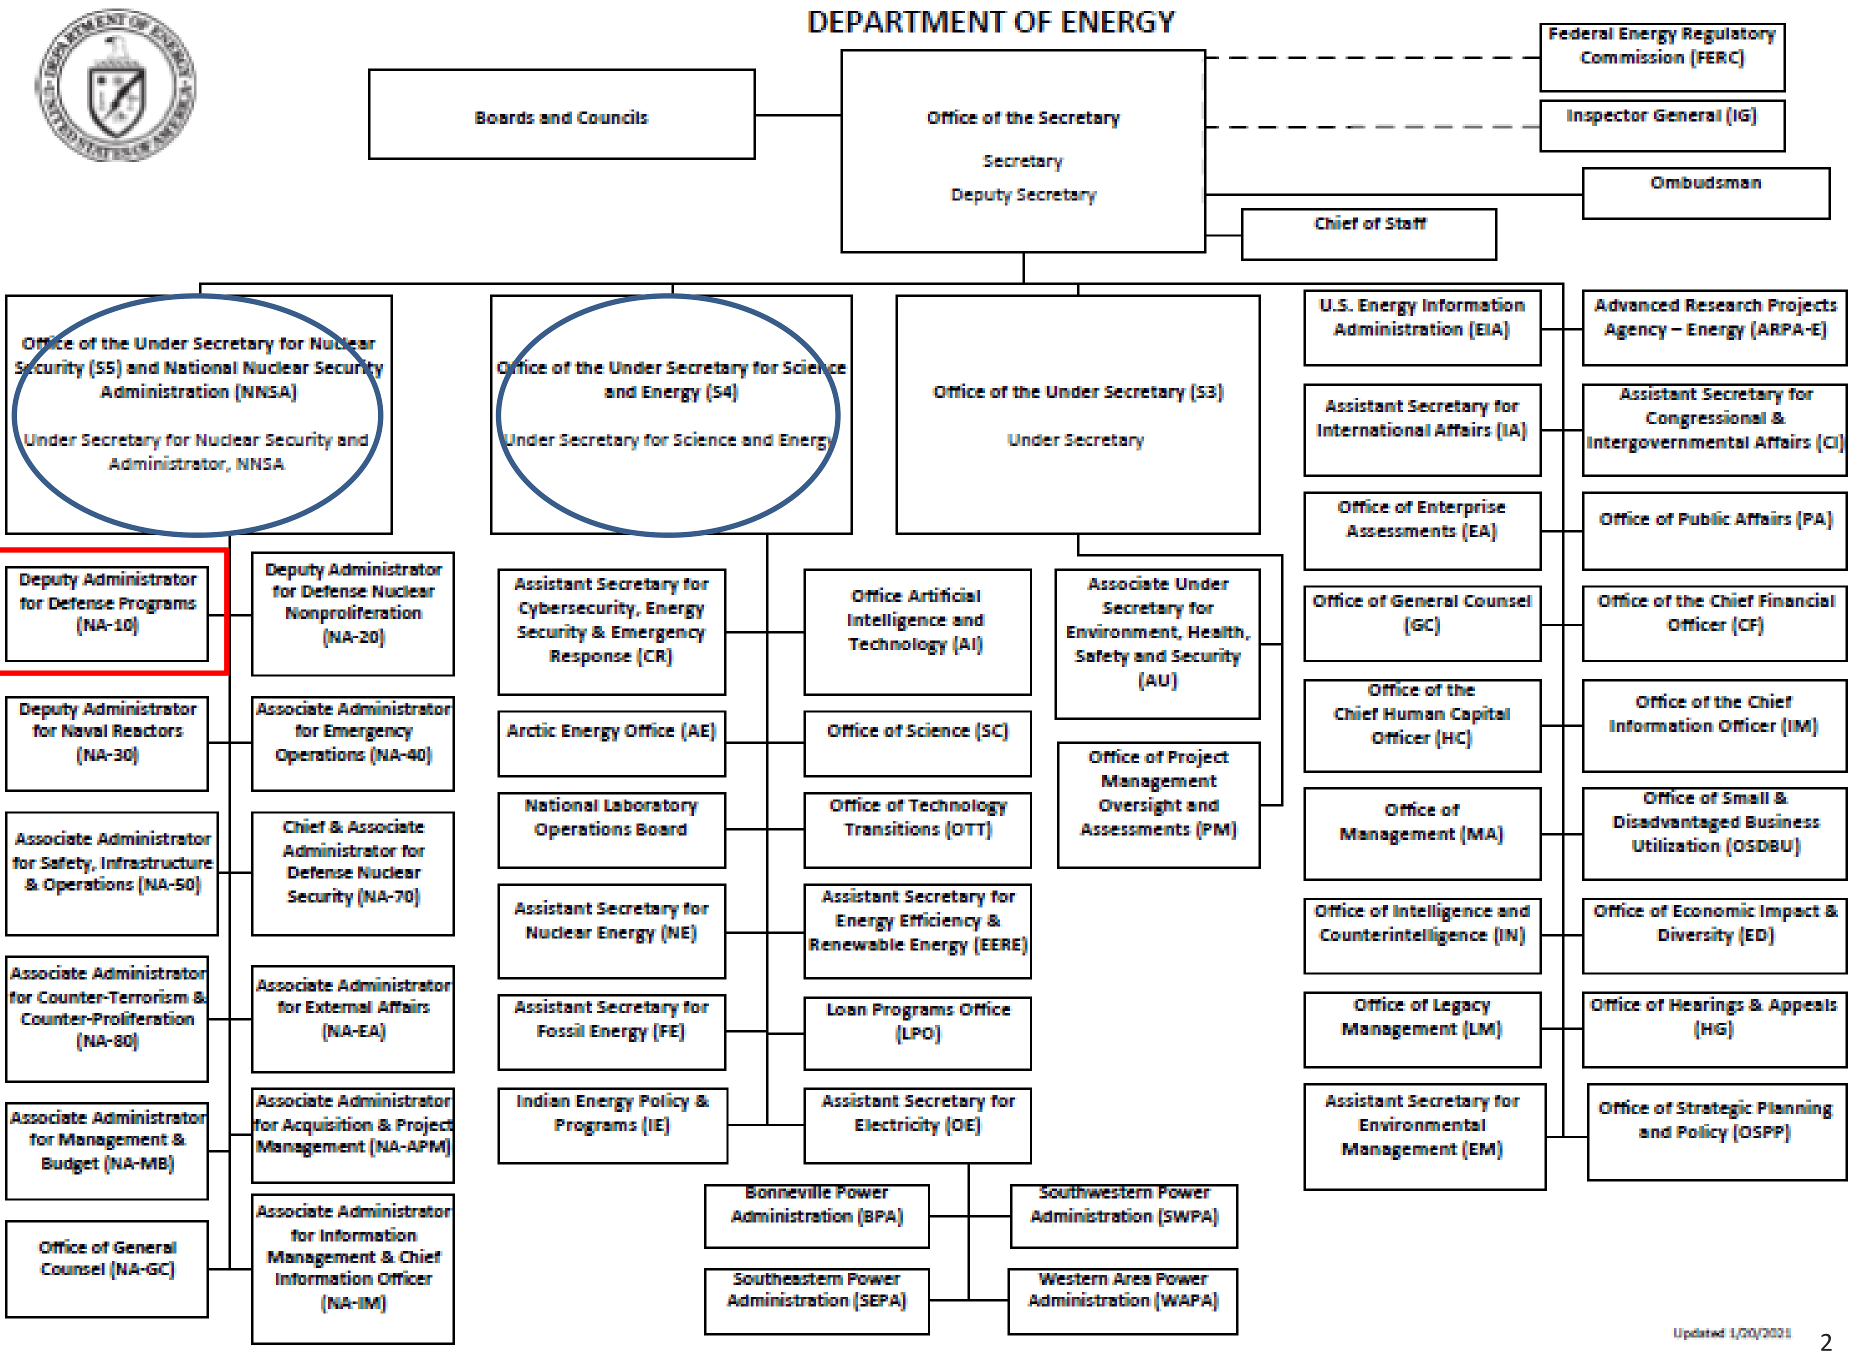
DEPARTMENT OF ENERGY
Boards and Councils
Office of the Secretary
Secretary
Deputy Secretary
Federal Energy Regulatory Commission (FERC)
Inspector General (IG)
Ombudsman
Chief of Staff
Office of the Under Secretary for Nuclear Security (S5) and National Nuclear Security Administration (NNSA)
Under Secretary for Nuclear Security and Administrator, NNSA
Office of the Under Secretary for Science and Energy (S4)
Under Secretary for Science and Energy
Office of the Under Secretary (S3)
Under Secretary
U.S. Energy Information Administration (EIA)
Advanced Research Projects Agency – Energy (ARPA-E)
Assistant Secretary for International Affairs (IA)
Assistant Secretary for Congressional & Intergovernmental Affairs (CI)
Office of Enterprise Assessments (EA)
Office of Public Affairs (PA)
Office of General Counsel (GC)
Office of the Chief Financial Officer (CF)
Office of the Chief Human Capital Officer (HC)
Office of the Chief Information Officer (IM)
Office of Management (MA)
Office of Small & Disadvantaged Business Utilization (OSDBU)
Office of Intelligence and Counterintelligence (IN)
Office of Economic Impact & Diversity (ED)
Office of Legacy Management (LM)
Office of Hearings & Appeals (HG)
Assistant Secretary for Environmental Management (EM)
Office of Strategic Planning and Policy (OSPP)
Deputy Administrator for Defense Programs (NA-10)
Deputy Administrator for Defense Nuclear Nonproliferation (NA-20)
Deputy Administrator for Naval Reactors (NA-30)
Associate Administrator for Emergency Operations (NA-40)
Associate Administrator for Safety, Infrastructure & Operations (NA-50)
Chief & Associate Administrator for Defense Nuclear Security (NA-70)
Associate Administrator for Counter-Terrorism & Counter-Proliferation (NA-80)
Associate Administrator for External Affairs (NA-EA)
Associate Administrator for Management & Budget (NA-MB)
Associate Administrator for Acquisition & Project Management (NA-APM)
Office of General Counsel (NA-GC)
Associate Administrator for Information Management & Chief Information Officer (NA-IM)
Assistant Secretary for Cybersecurity, Energy Security & Emergency Response (CR)
Office Artificial Intelligence and Technology (AI)
Arctic Energy Office (AE)
Office of Science (SC)
National Laboratory Operations Board
Office of Technology Transitions (OTT)
Assistant Secretary for Nuclear Energy (NE)
Assistant Secretary for Energy Efficiency & Renewable Energy (EERE)
Assistant Secretary for Fossil Energy (FE)
Loan Programs Office (LPO)
Indian Energy Policy & Programs (IE)
Assistant Secretary for Electricity (OE)
Associate Under Secretary for Environment, Health, Safety and Security (AU)
Office of Project Management Oversight and Assessments (PM)
Bonneville Power Administration (BPA)
Southwestern Power Administration (SWPA)
Southeastern Power Administration (SEPA)
Western Area Power Administration (WAPA)
Updated 1/20/2021

Defense Programs is Structured to Support the United States Nuclear Deterrent

Maintain the Current Stockpile | **Deliver LEPs** | **Prepare for the Future**

Life Extension Programs (LEPs) prevent operational gaps while enhancing safety, security, and use control.

Production Modernization provides strategic investments to modernize infrastructure and manufacturing capabilities.

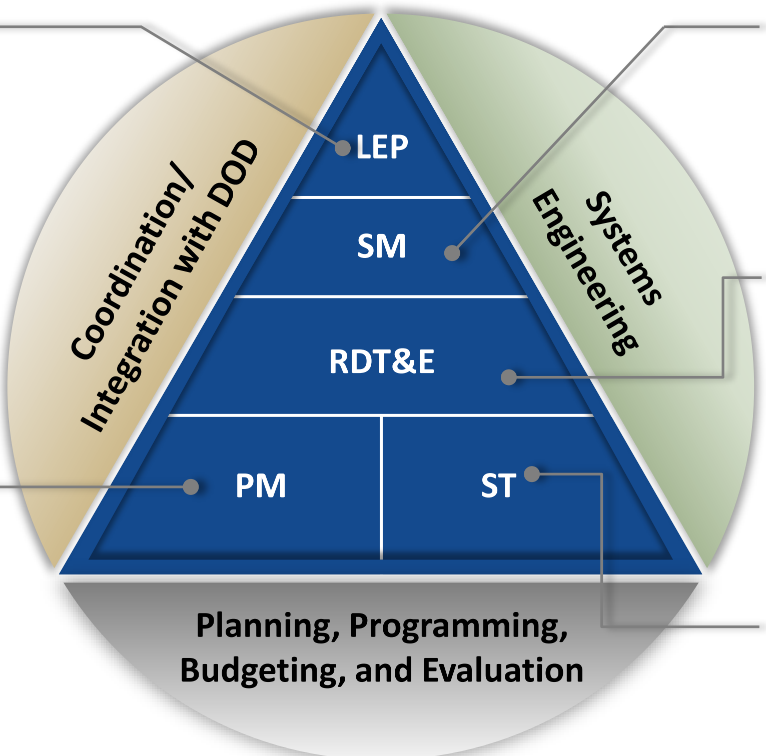

Stockpile Management sustains the Nation's nuclear weapon stockpile.

Research, Development, Test, and Evaluation provides tools and capabilities for stockpile assessment and certification, including the development of predictive capabilities.

ASC Program is in RDT&E

Secure Transportation provides safe and secure shipment of nuclear weapons, weapons components, and special nuclear material.

Advanced Simulation and Computing (ASC)

ASC is predictive science through simulation: the people, state-of-the-art computational platforms, and simulation tools used in the annual certification of nuclear weapons stockpile:

- **Prediction Through Simulation**. Deliver verified and validated physics and engineering codes to enable simulations and risk-informed decisions of nuclear weapons performance, safety, and reliability.
- **Robust Tools**. Develop robust models, codes, and computational techniques to support stockpile needs such as Significant Finding Investigations, Life Extension Programs, annual assessments, as well as evolving future requirements.
- **Balanced Operational Infrastructure**. Implement a balanced computing strategy of platform acquisition and operational infrastructure to meet Directed Stockpile Work and Stockpile Stewardship Program needs for production and advanced simulation capabilities.

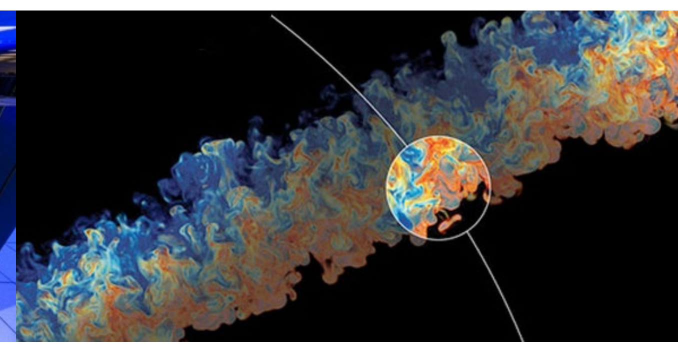

Confidence in Stockpile Assessment Rests on Our People and Tools

Designers with test experience

Nuclear tests

Experiments

Empirical simulations

1992

NNSA Exascale System

2023

2025

2nd/3rd generation designers

Archived nuclear test data

Surveillance

High-fidelity AGEX

Predictive science-based simulations

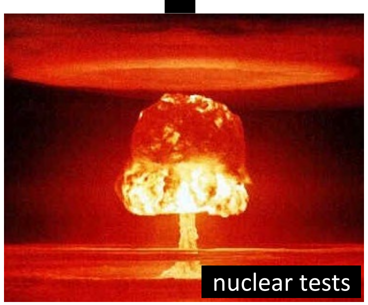

Calibrated codes evolve into predictive capabilities

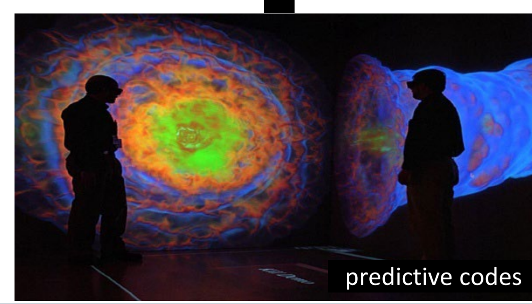

Our confidence, once based on nuclear tests, now must be founded on improved physics understanding incorporated into predictive simulations validated against historic nuclear test data and focused experiments

ASC Program
National Work Breakdown Structure (nWBS)

ADVANCED SIMULATION AND COMPUTING (ASC)

- **Integrated Codes**
  - Engineering & Physics Integrated Codes
  - Specialized Codes and Libraries
  - Applications and Algorithms Research
  - Applications Research for Next-generation Platforms
- **Physics & Engineering Models**
  - Materials Response
  - Transport, Plasmas, Atomic, Nuclear
  - Engineering Science
  - Integrated Modeling & Application
- **Verification & Validation**
  - V&V Methods
  - V&V Assessments
  - Data Validation, Archiving, SQA, and Training
- **Advanced Technology Development & Mitigation**
  - Next-Gen Code Development & Application
  - Next-Gen Architecture & SW Development
  - Inter-Agency CoDesign
- **Computational Systems & Software Environment**
  - Commodity Technology Systems
  - Advanced Technology Systems
  - System Software & Tools
  - Input/Output, Storage & Networking
  - Post-Processing Environments
  - Next-Gen Computing Technologies
  - Beyond Moore's Law
- **Facility Operations & User Support**
  - System and Environment Administration and Operations
  - Common Computing Environment
  - Collaborations
  - Special Purpose Facilities, Systems, Operations, & Support

ASC Mod/Sim Capability is Unique Within National Security and Science Communities

- The suite of NNSA codes are the embodiment of all **our nuclear weapons knowledge** – captured in the data, models, phenomenology and expert staff.
- Simulation codes are tuned for **specific architectures to run in a reasonable time** (e.g., less than a year for large problems).
- Nuclear weapons codes are very large and complex and support very high consequence decisions:
  - Long life-time projects with >1 million lines of code, originally 15+ years of development by large teams and over $1B investment
  - Transforming millions of lines of code to mitigate impact can take 3-5 years
- Tightly-coupled, multi-physics codes contain a set of distinguishing characteristics:
  - Mechanics, particle transport, nuclear reactions, thermodynamics, electromagnetics, multi-scales, etc.
  - Verification and validation, together with quantification of uncertainties, is essential & multiplying needs by ~1000X

Laboratory	Million Lines of Code*
LANL	16.5
LLNL	15.7
SNL	18.2
Trilab Totals	**50.4**

* not including math libraries
* not including new codes

Typical single code's line count:
- Multi-physics: 1.3 million
- Engineering: 3.5 million
- Safety: 1.0 million

High-fidelity simulations are required to either establish or negate the need for return to underground testing

Increasing code complexity, coupled with architecture uncertainty, is driving NNSA Labs to a new, modular approach

- Codes today are >5x larger than pre-ASCI (1995) codes
- Modular infrastructure can be reused by other codes
- Central Data Store allows individual components to be modified/rewritten independently

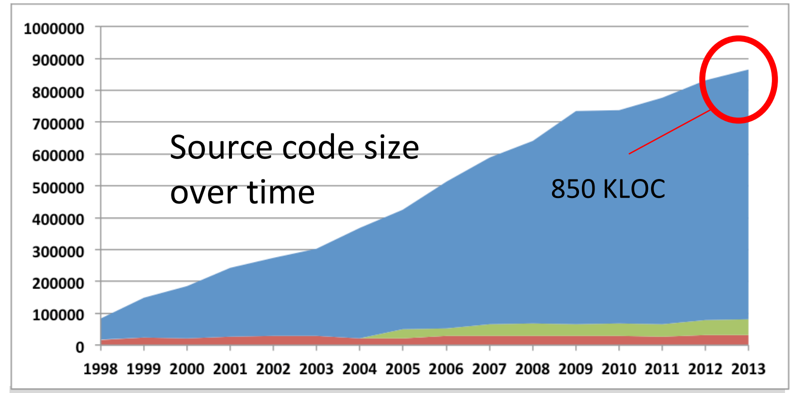

Average of 60K lines of code (KLOC) per year of additional code to support a representative integrated code

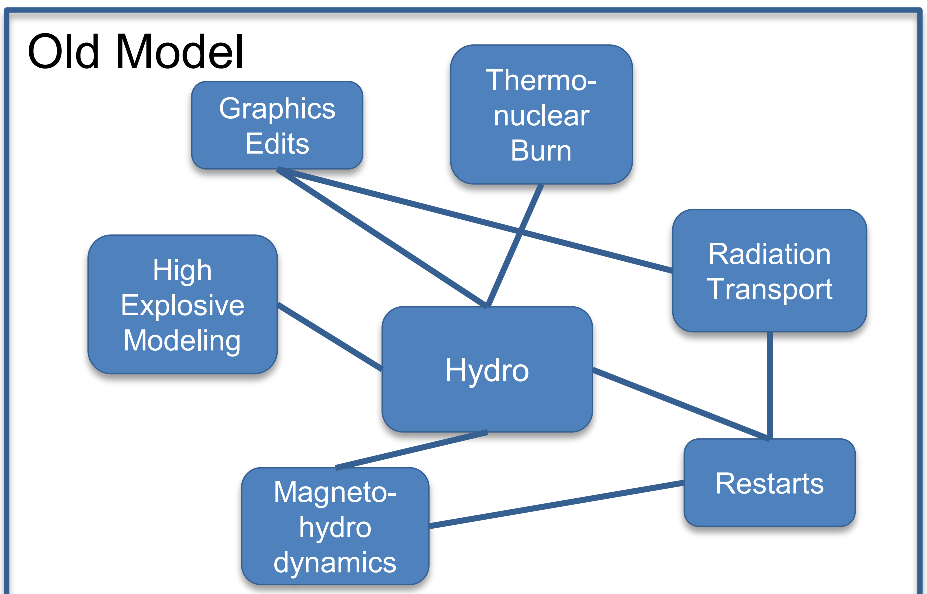

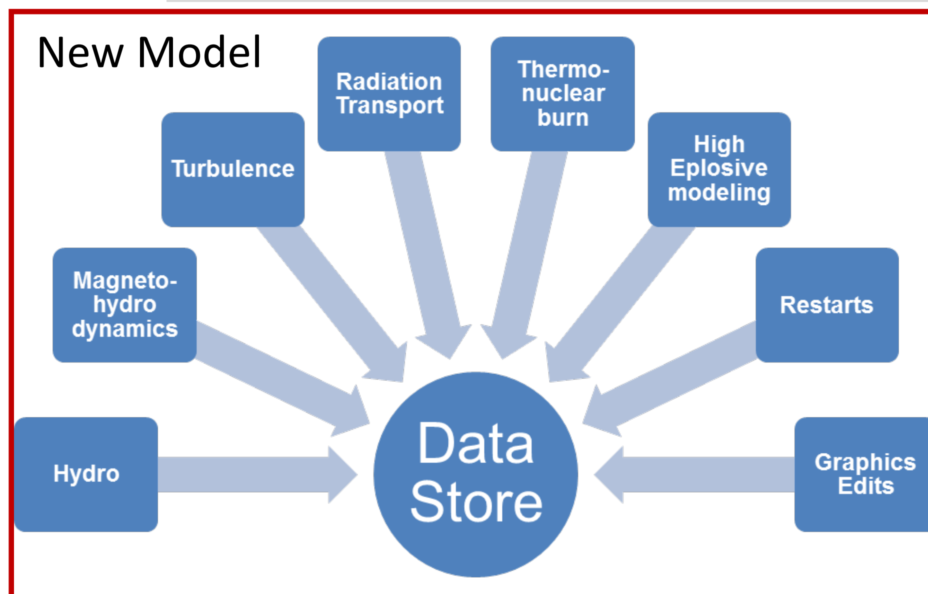

Current ASC Computing/Architecture Environment

Sierra @ LLNL

Compute Node
2 IBM POWER9 CPUs
4 NVIDIA Volta GPUs
NVMe-compatible PCIe 1.6 TB SSD
256 GiB DDR4
16 GiB Globally addressable HBM2 associated with each GPU
Coherent Shared Memory

Compute Rack
Standard 19"
Warm water cooling

Compute System
4320 nodes
1.29 PB Memory
240 Compute Racks
125 PFLOPS
~12 MW

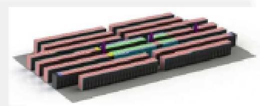

Components

IBM POWER9
- Gen2 NVLink

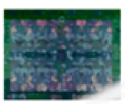

NVIDIA Volta
- 7 TFlop/s
- HBM2
- Gen2 NVLink

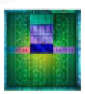

Mellanox Interconnect
Single Plane EDR InfiniBand
2 to 1 Tapered Fat Tree

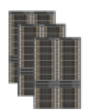

GPFS File System
154 PB usable storage
1.54 TB/s R/W bandwidth

Trinity @ LANL

Trinity Cray XC40 Specifications	
Intel Xeon E5-2698v3 "Haswell"	Intel Xeon Phi 7250 "Knights Landing"
9436 nodes	9984 nodes
Dual socket, 16 cores/socket, 2.3 GHz	1 socket, 68 cores, 1.4 GHz, > 3 Tflops/KNL
128 GB DDR4	96 GB DDR4 + 16GB HBM
1.15 PB on-node memory	1.12 PB on-node memory
Aries "Dragonfly" interconnect	

Trinity HSW
32 cores + HT [64]
AVX 256-bit [4 doubles]

+

Trinity KNL
68 cores + 4 HW [272]
AVX-512 512-bit [8 doubles]

Cray Sonexion Storage System
78 PB Usable, ~1.6 TB/s

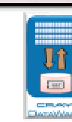

Cray DataWarp
576 Burst Buffer Nodes
3.7 PB, ~3.3 TB/s

Astra @ SNL

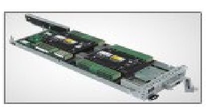

HPE Apollo 70 Cavium TX2 Node

HPE Apollo 70 Chassis: 4 nodes

HPE Apollo 70 Rack

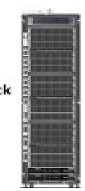

18 chassis/rack

72 nodes/rack

3 IB switches/rack (one 36-port switch per 6 chassis)

36 compute racks (9 scalable units, each 4 racks)

2592 compute nodes (5184 TX2 processors)

Tri-lab CTS-1 Deployments

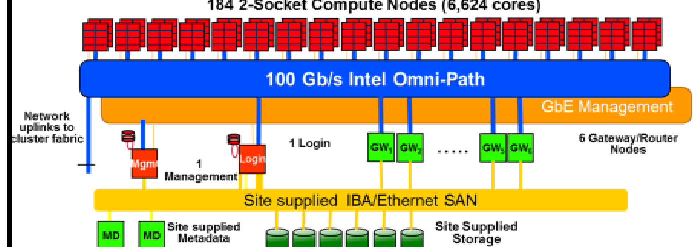

ASC & Institutional Computing:

LLNL- 52 SUs
LANL- 37.5 SUs
SNL- 34.5 SUs

Scalable Unit (SU) Parameters: 192 nodes; 232 TF/s; 24.5 TB DRAM
- 184 dual socket compute nodes; 128 GB DDR4-2400 SDRAM
- HPC 100 Gb/s network fabric; scalable to 4K+ end points
- Compute and Gateway nodes remote boot from Management nodes
- Site supplied parallel storage and SAN
- Options for on node PCIe accelerators and NVRAM
- Base for building multi-SU clusters

ASC Platform Timeline

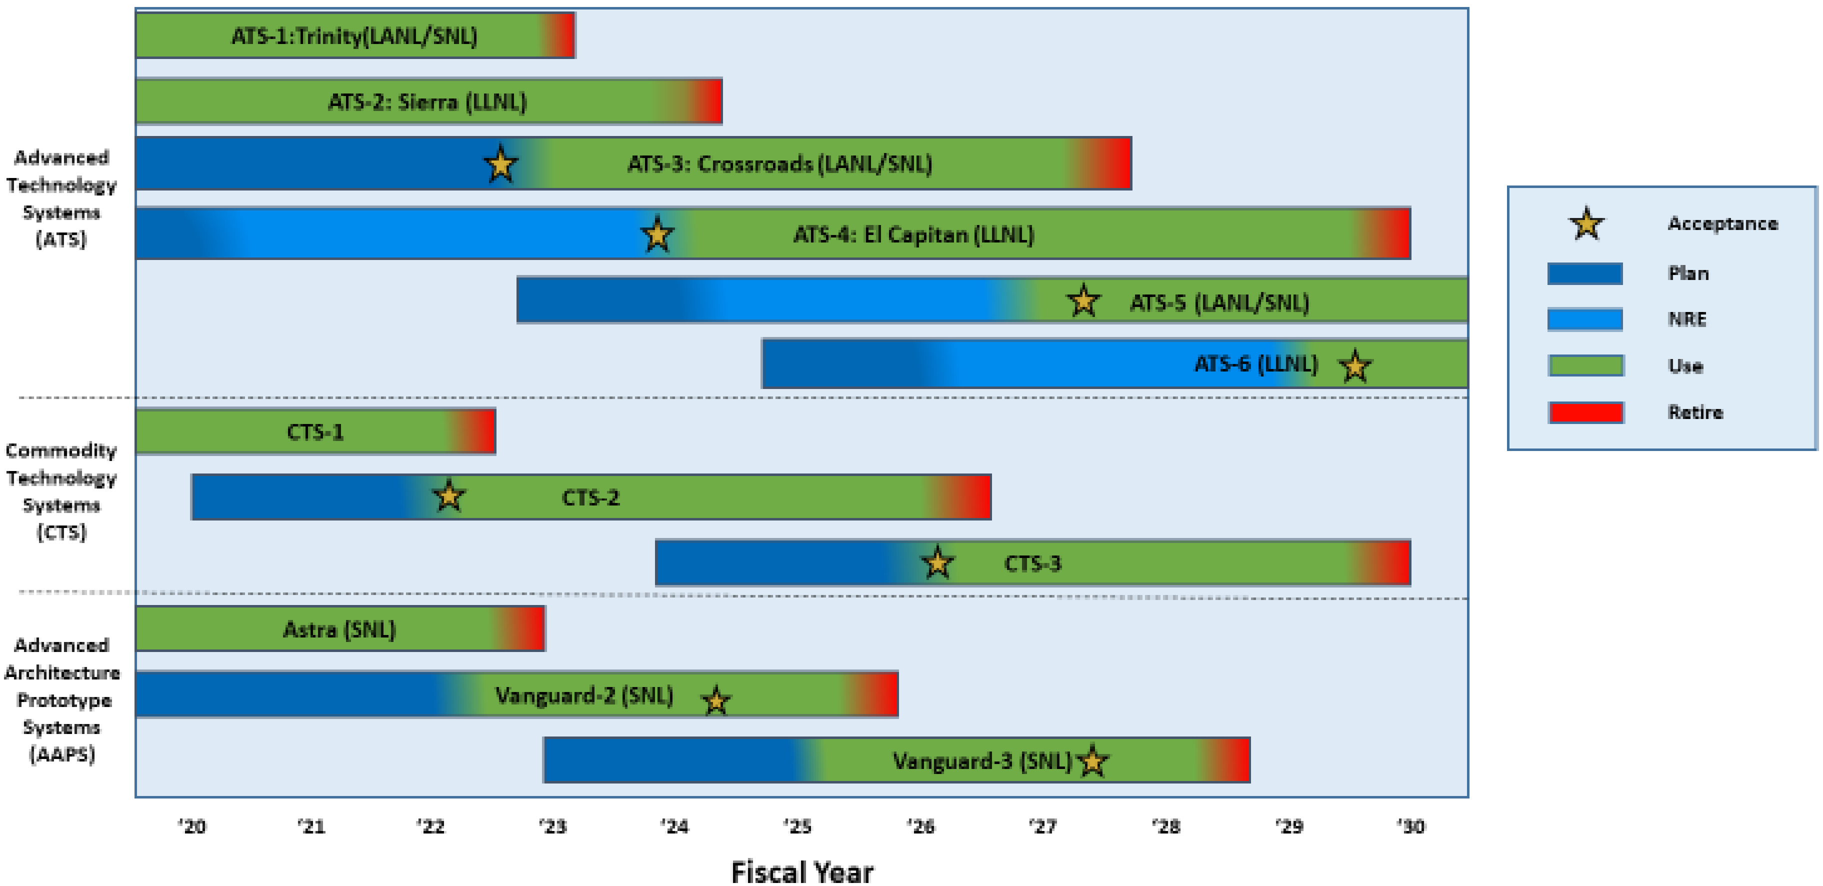

February 2021

We Must Keep Up with Evolving Architectures to Meet Mission Needs

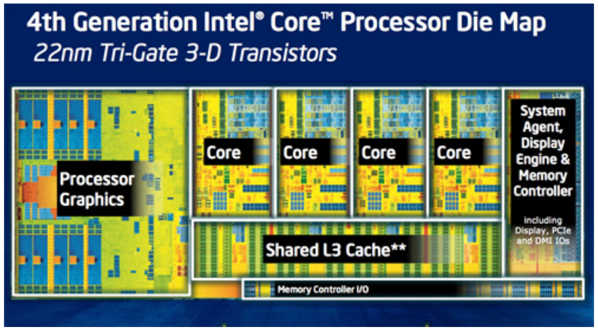

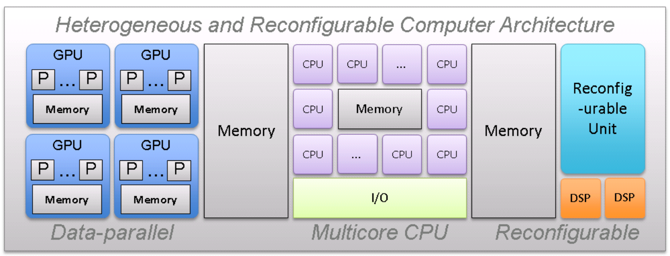

Homogeneous Many-Core & Heterogeneous Architectures

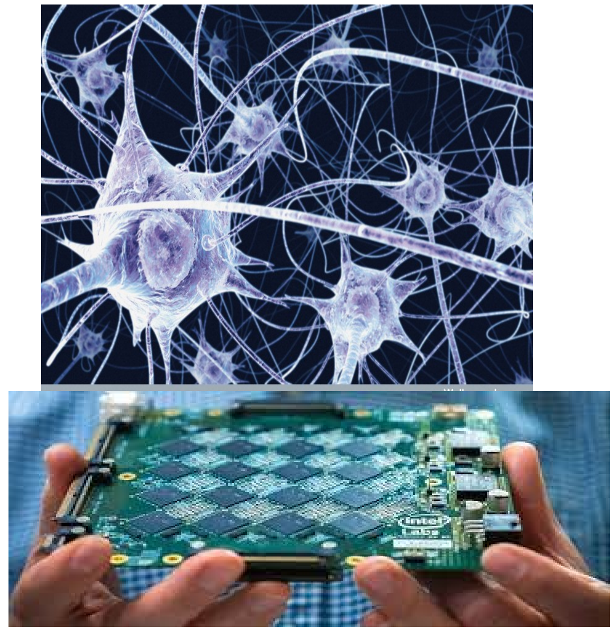

Neuromorphic systems

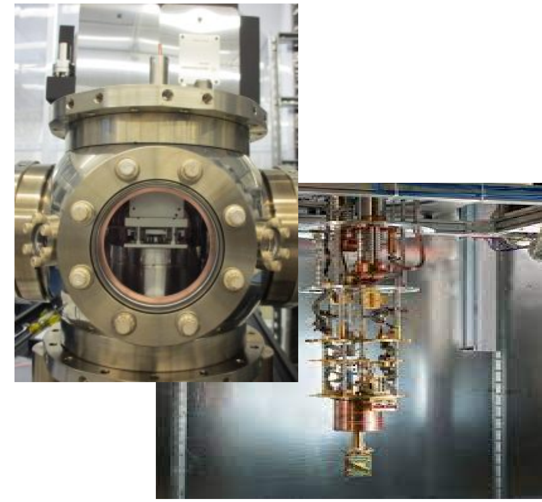

Quantum Computing

New memory architectures

The computing landscape will continue to evolve; we must capitalize on future technologies

SNL Project on Credibility for Scientific Machine Learning: Data Verification & Model Qualification

Problem

- High-consequence ND workflows require rigorous V&V and UQ
- Need evidence-based credibility for scientific machine learning when used in ND workflows.

Technical Approach

- Rigorously assess ML models using Predictive Capability Maturity Model (PCMM) workflow
  - Statistical properties of input features when training a ML model
  - Examine associated sources of noise, e.g., measurement noise.
- Exemplar: Stronglink Aging Classification
  - Predict health of devices using waveforms collected using non-destructive tests.
  - Identify key input features in time or Fourier domains, assess the quality of training datasets, and decompose uncertainty into its aleatoric and epistemic components.

Deployment:

- A standalone implementation and integration with Dakota.

Characteristics of PCMM Elements

Representation and Geometric Fidelity
What features are neglected because of simplifications or stylizations?

Physics and Material Model Fidelity
How fundamental are the physics and material models and what is the level of model calibration?

Code Verification
Are algorithm deficiencies, software errors, and poor SQE practices corrupting the simulation results?

Solution Verification
Are numerical solution errors and human procedural errors corrupting the simulation results?

Model Validation
How carefully is the accuracy of the simulation and experimental results assessed at various tiers in a validation hierarchy?

Uncertainty Quantification and Sensitivity Analysis
How thoroughly are uncertainties and sensitivities characterized and propagated?

DATASHEETS

Verification of training datasets: documenting composition, collection process, preprocessing, cleaning, labeling, uses, distribution, and maintenance.

EXEMPLAR

Stronglink Component Aging Classification

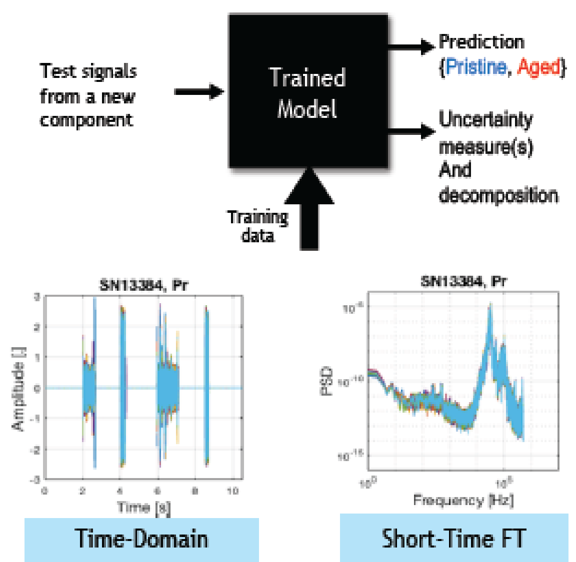

LANL High Explosive Calibration Workflow

- **Goal**: Create a ML-based workflow for agile HE model calibration and development.
- **Method**: Both experimental and simulation data are used in training and validation. Multiple ML optimization efforts and database work to transform HE analysis.
- **Key result**: Workflow for HE calibrations is working based on CalData database and ML tool for calibration based on FUNCLE and Magpie.

HE data and modeling workflow
PEM-HE point of view

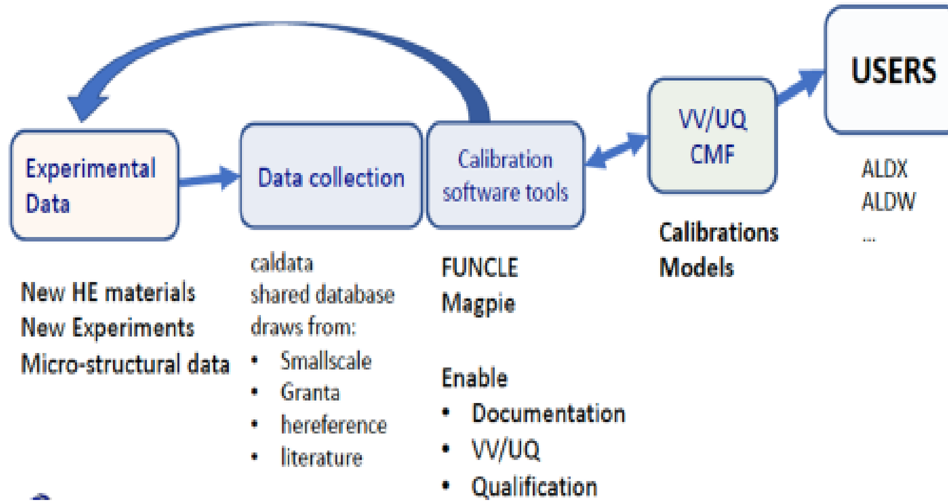

Calibration workflow starts with experimental data. This work develops CalData database which houses curated data. Using these data, EOS and rate law calibrations are made with Magpie and FUNCLE. Long term vision for this tool set is to enable VV/UQ and qualification work.

LLNL Inertial Confinement Fusion (ICF) Modeling

AI-driven hierarchical simulation for optimal design

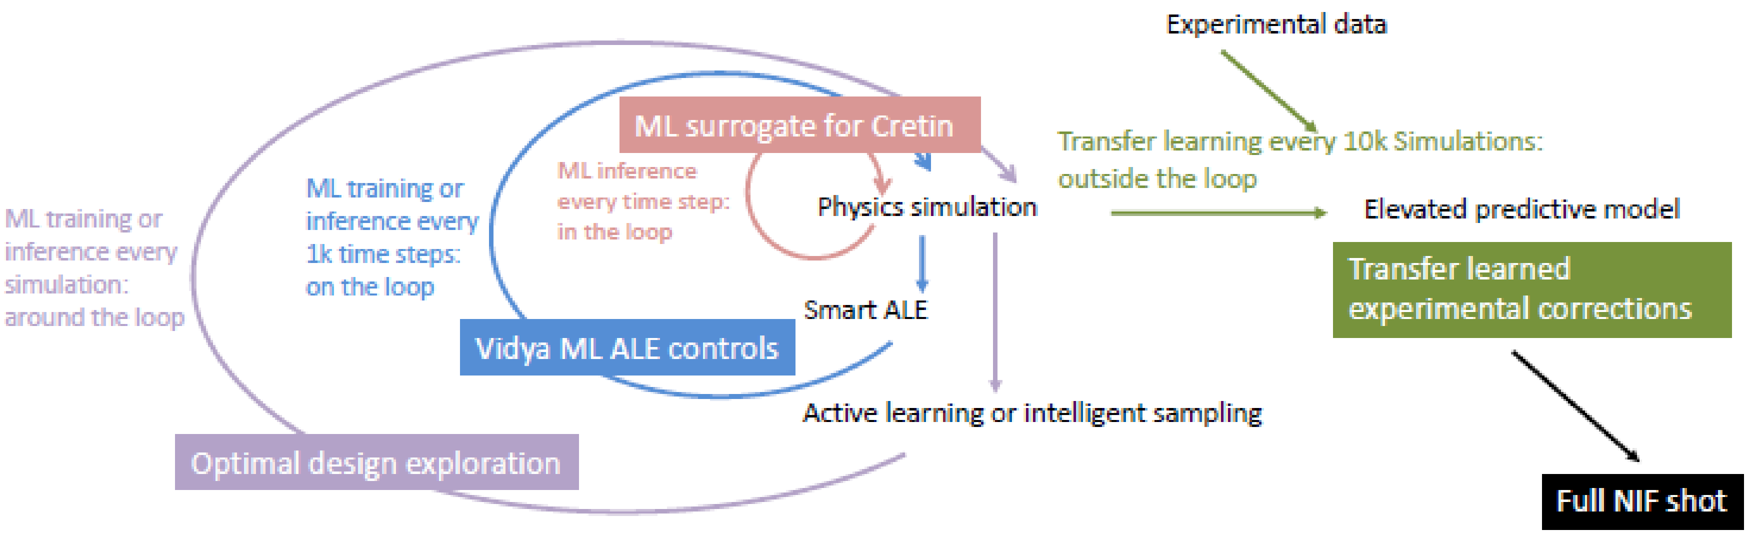

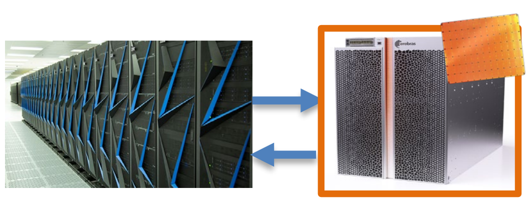

NNSA's Lassen Supercomputer

Cerebras System

- Radiation opacities near gold wall requires complex calculations
- These calculations can consume 90% of the runtime
- A Deep Neural Network, trained offline, has demonstrated 10X speedups relative to inline opacity calculations

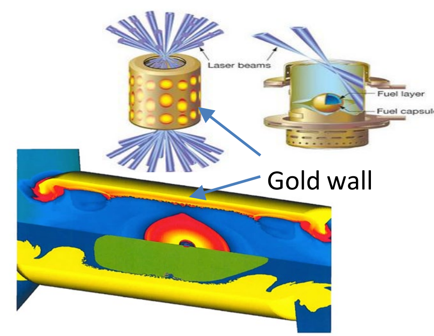

We have demonstrated another ~20X speedup with AI optimized hardware, with >100X possible

Future Computing Paradigm

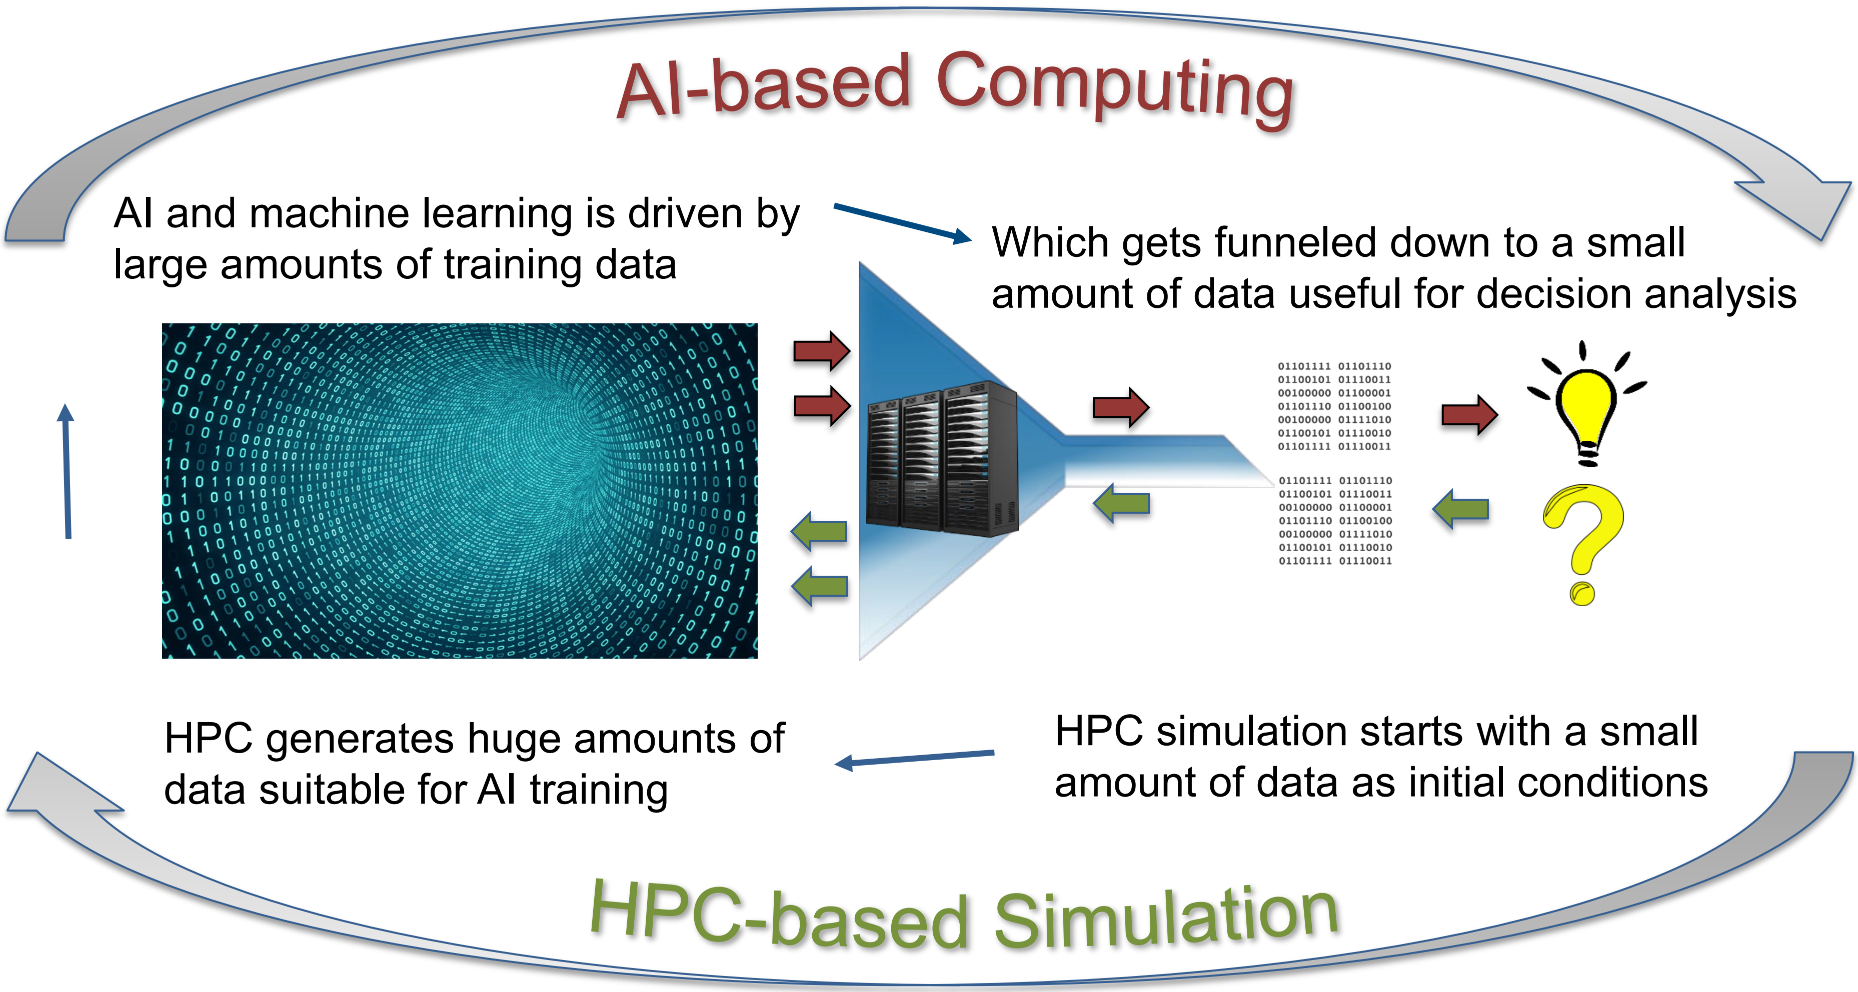

Courtesy: Jim Brase, LLNL

Innovate - Collaborate - Deliver

- AI-driven methods for designing, manufacturing and deploying products have the potential to revolutionize NNSA workflows:

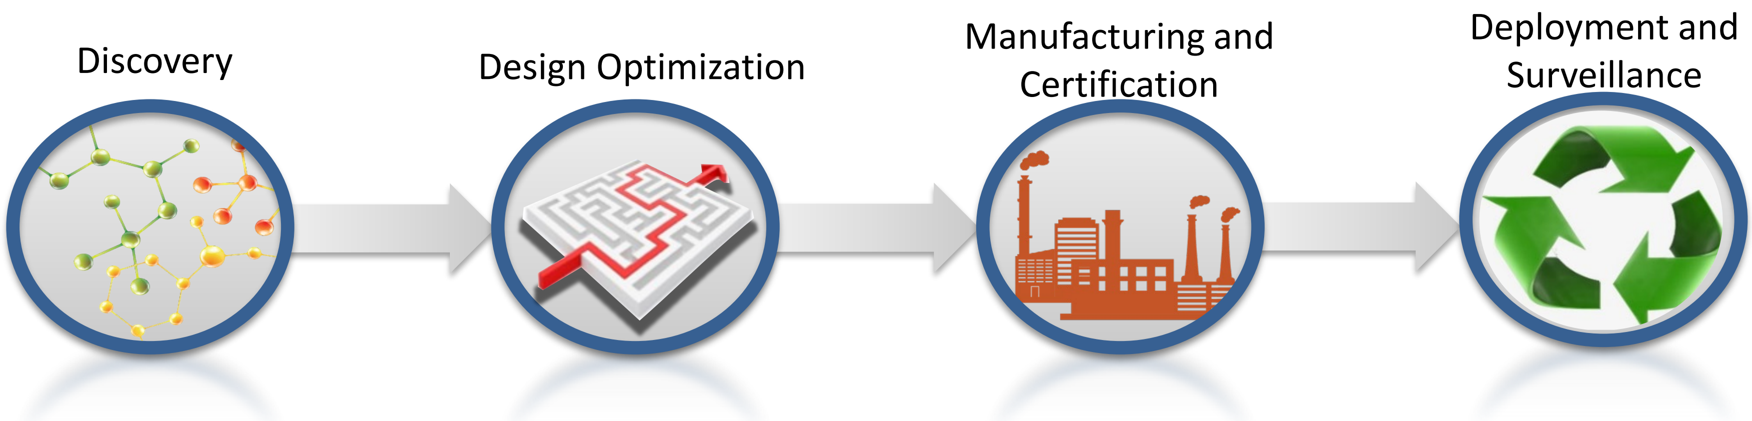

- High-performance computing (HPC) will continue to be the computing technology foundation for NNSA stockpile stewardship, plus other advanced architectures:
  - Surrogate models generated by powerful accelerator-based HPC systems provide faster models for many-query applications.
  - Current "exascale computing" technology allows ensemble calculations to be possible, with the latter being the nucleus of Uncertainty Quantification.

Full engagement and true collaborations as public-private partnerships between government, industry, and academia continue to be necessary

Thank you