LA-UR-22-22134

Approved for public release; distribution is unlimited.

Title:	Oppenheimer Science and Energy Leadership Program (OSELP) Computing Platform Strategy, Priorities, and Roadmap
Author(s):	Lujan, James Westley
Intended for:	Report
Issued:	2022-03-08

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Los Alamos NATIONAL LABORATORY

Oppenheimer Science and Energy Leadership Program (OSELP)

Computing Platform Strategy, Priorities, and Roadmap

Speaker
Jim Luján/High Performance Computing (ASC)

March 7, 2022

IS&T capabilities span LANL organizational structure

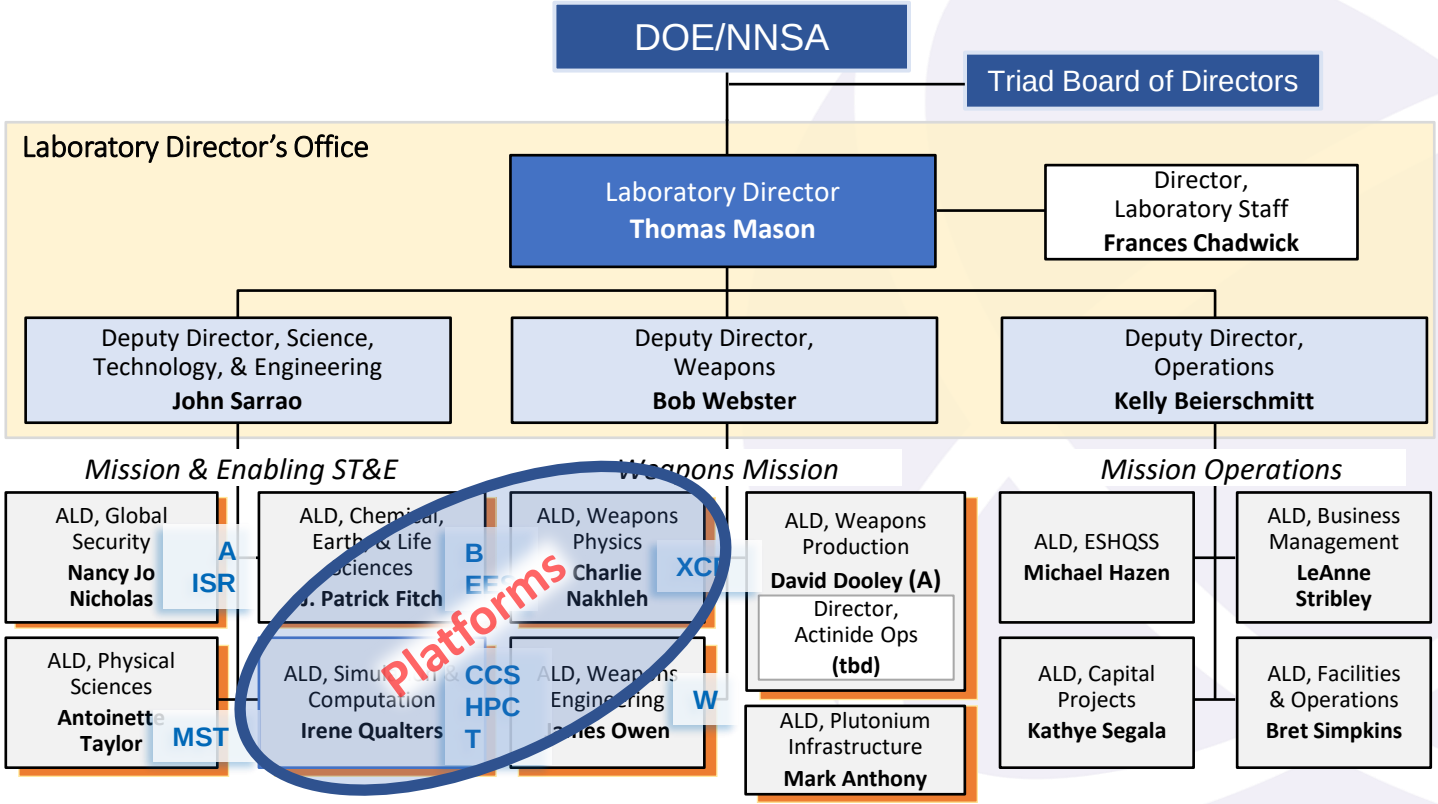

Divisions with >25% efforts in IS&T

ASC strategy is focused on three principal goals

- **Provide capabilities enabling a resilient deterrent and responsive complex.**
  - Rapid manufacturing and phase 1-5 cycles
  - Military characteristics for needs of a changing security landscape
  - “Headlights” for long-range national decisions

- **Provide capabilities needed to confidently assess and qualify the performance of nuclear weapons.**
  - Aging, new environments, new designs, new production

- **Build computing technologies for solving national security problems that are intractable today.**
  - The focus of this goal is on high-performance computing platforms capable of simulating very complex physical systems in human learning times.

Trends in computing, sometimes misaligned with our interests, provide a basic challenge

- Demise of miniaturization led to pursuit of performance through the development of specialized processors.
- Commodity drivers for those specializations tend to run counter to our needs.

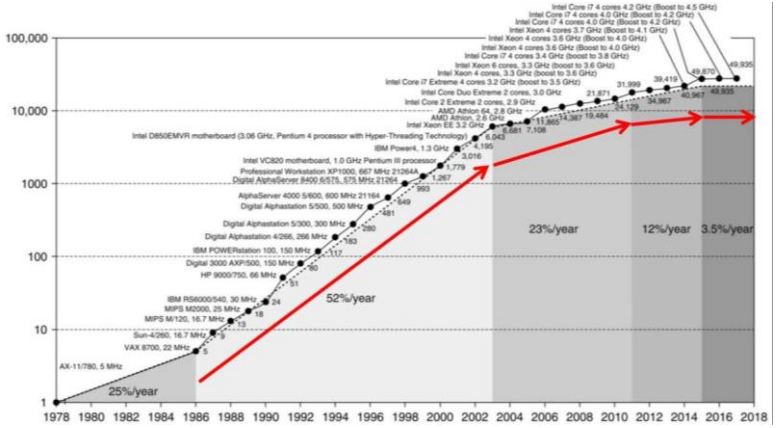

In the post-Moore era, we expect more special-purpose devices, because they will not have comparably performing general-purpose processors right around the corner to compete with. -- Leiserson et al

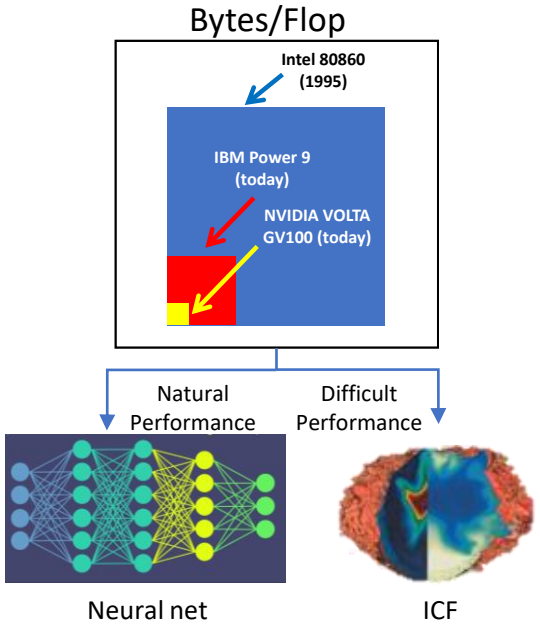

At the same time that technology is shifting, the nation is asking us to make a broader contribution

- **Changes on the world stage are motivating new questions**
  - Rebalance away from predictive capability for performance
  - The scope of the questions we are being asked has grown
    - We need a return to our roots in computational methods development
    - Multiphysics coupling and machine learning are both opportune fields

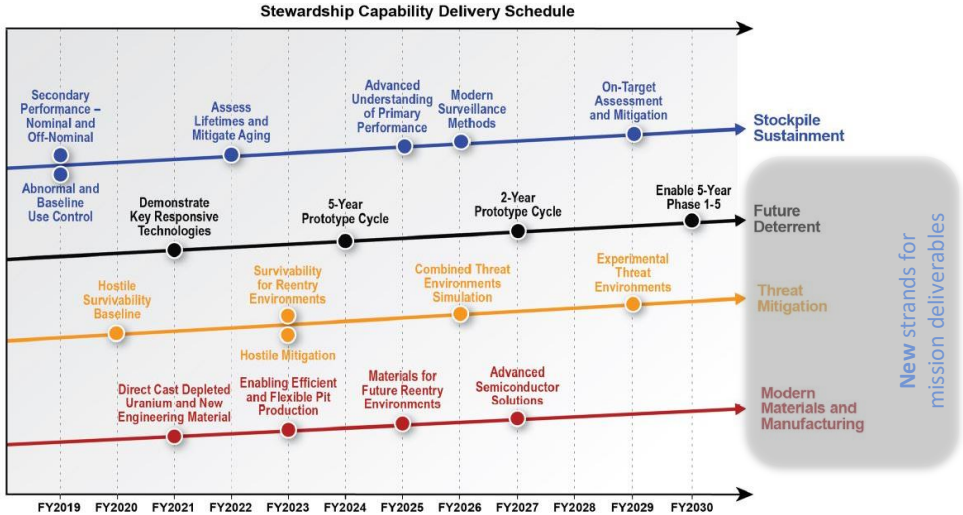

Enduring Mission Skews LANL Computing

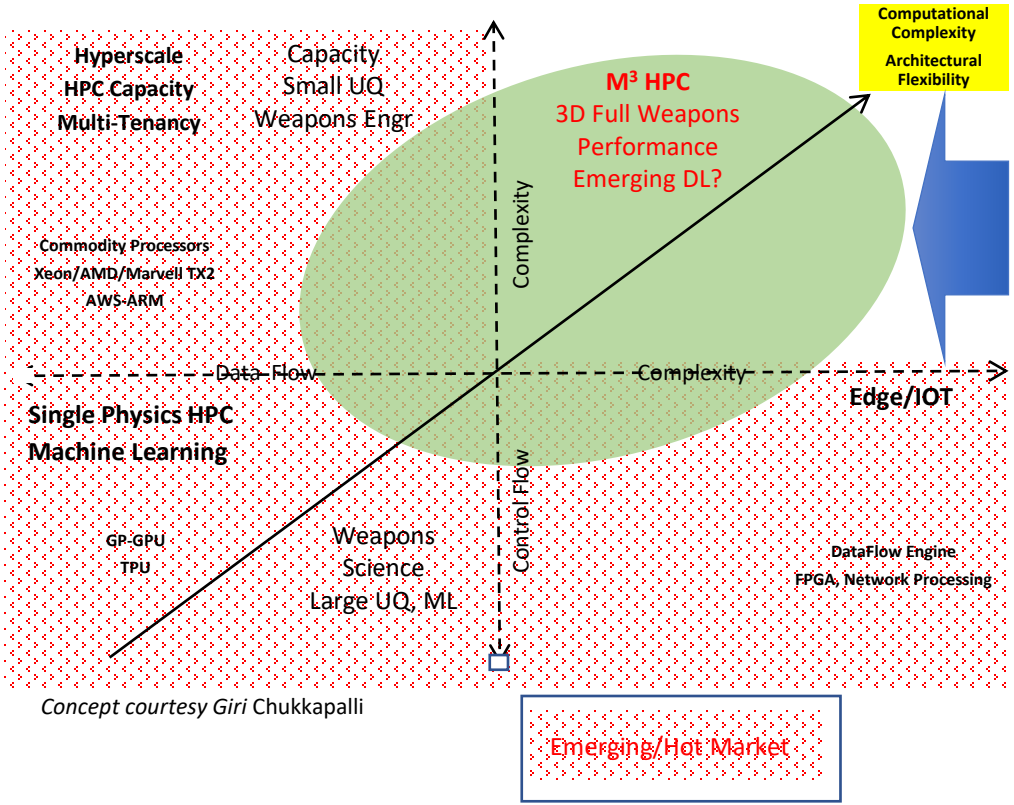

- Current Weapons Perf:
  - PB working set
  - Multi-physics, resolutions, link-scales
  - 3D Unstructured/sparse
  - ½ machine 6 month duration

Enable human time scale exploration of most complex Multi-Physics/Resolution/Link-Scale problems

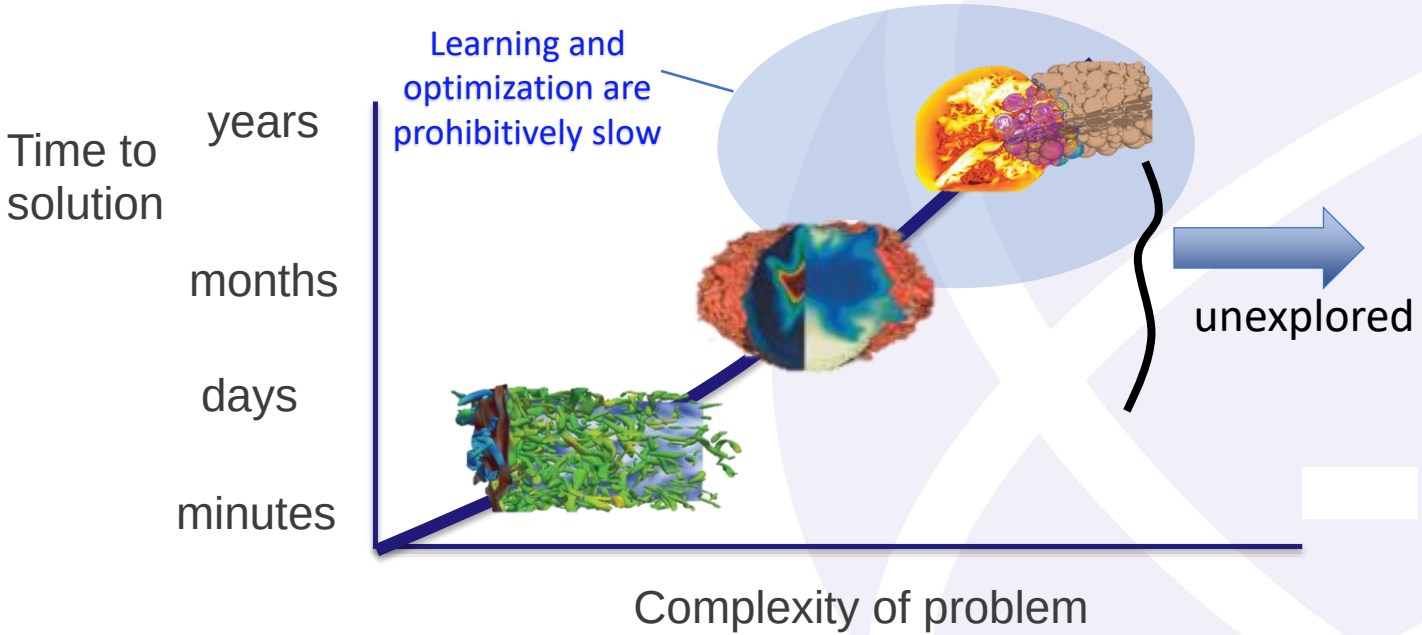

Runs that currently take 1 PB DRAM and 10k nodes / ½ million cores for 6 months need to run in 6 days
Need to move from 1% efficient to 30% efficient in 6 years

Our strategy: make the best of what industry offers, and shape tomorrow's technologies

➢ **Challenge:** Computing technology trends are working against us

Make compute tech our friend

We and Livermore have advanced a strategy that explores two paths enabled by commercial trends

Accelerated Architectures
Enabled by Industry-Driven Architectural Trends
Codesigned to Application Demands

Lawrence Livermore National Laboratory

Complex Simulation Architectures
Leveraging Component-Level
Customization Tailored for Application

Los Alamos NATIONAL LABORATORY

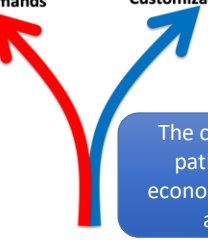

The challenge with this path is in finding an economical way to shape architectures.

NVIDIA partnership as a path to success

Live with it (efficiently?)

Successful efforts to rewrite & refactor existing capabilities for current and near-term platforms

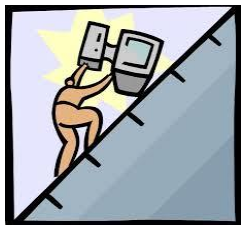

Each project ports its code(s), for each new architecture

Abstraction layers
Automate the port?

ATDM investment suggests efficiencies

➢ **Challenge:** Changes on the world stage are rapidly motivating new questions

Codesign new compute tech:
understand when it has enabled a break through

Innovate new approaches:
get back to our roots…
look beyond existing capabilities

Targeted efforts:
help to resolve critical needs

Our plan for advances in platform technology build on a recent industrial capability for specialization

The world is adopting customization as the path to performance in the post-Moore era

- ARM is at the center of a broad industrial capability for tailoring

Processor IP

2.5D integration

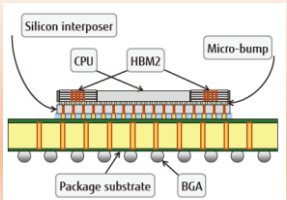

The Fab economy

GlobalShuttle MPW Schedule:

Tech. Node	Jan	Feb	Mar	Apr	May
12nm			18 MAR 1423		
22nm	13 JAN 2232		02 MAR 2233		11 MAY 2234

Japan and Europe have shown great potential for HPC

- US Exascale effort bet on a single path with its $2B investment

European Processor Initiative

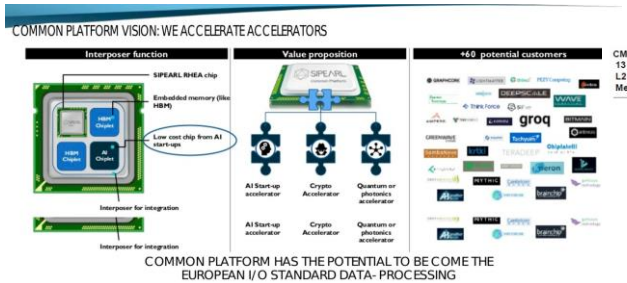

Japanese (MEXT) – Fugaku

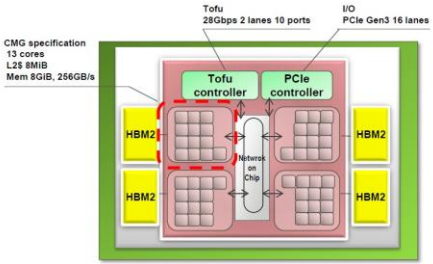

NNSA ASC Computing Strategy – 2022

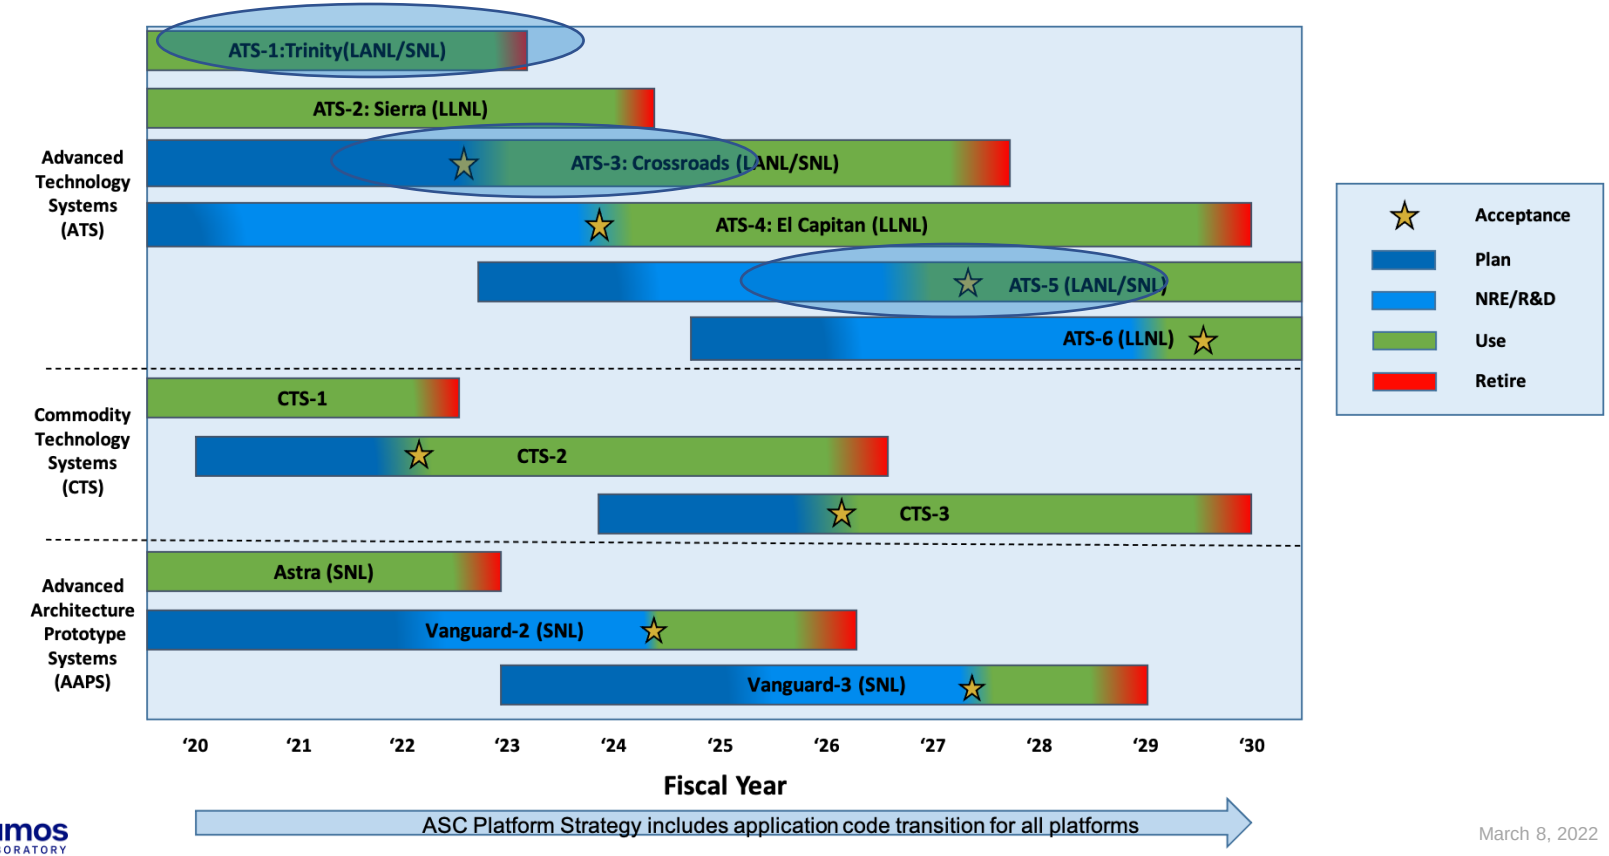

Crossroads will Enable an Increase in Predictive Capability for ASC while Extending State of the Art

- Provide a significant capability improvement over current platforms (>> Sequoia, > Trinity) in key physics areas
  - Improvement is a function of performance (total time to solution), increased geometries and physics capabilities
- Increased capabilities drive computational improvements
  - Higher fidelity models -> increased aggregate memory capacity
  - While sustaining time to solution -> increases in computational capabilities, memory bandwidth & scaling characteristics
- Improve performance of current and future programming models, overall science workflow efficiency, and manageability (reliability, resilience, energy, etc.)
- Extend procurement methodology yet again beyond Trinity
  - No FLOPS -> app performance
  - No IO speeds -> 90% efficient

Intel Sapphire Rapids (SPR) **w/HBM**
All NVMe based file system

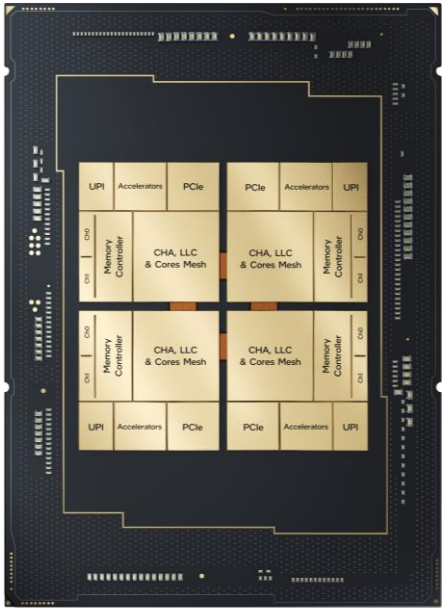

A new institutional investment helps position the US on a path for long-term gains in computing technologies

- Based on new line of NVIDIA technologies
  - $80M, delivery in early FY23
  - First new family in 2 decades
- Broad spectrum of advanced technologies for laboratory applications
  - Multiphysics, Machine Learning, Analytics…
  - Platform for hardware studies
- Since it is an institutional system, we can dedicate time to partnerships
  - Universities, DOD and industry

We and NVIDIA have assigned leaders for a codesign partnership

- Libraries and Compilers
- Deep Learning
- Storage
- Networking
- Light House Applications
- Sparsity, Simulations
- Characteristics + Profiling
- Architectures
- Testbeds

A partnership with NVIDIA answers the hardest question: How do we influence much larger industry efforts?


CROSSROADS

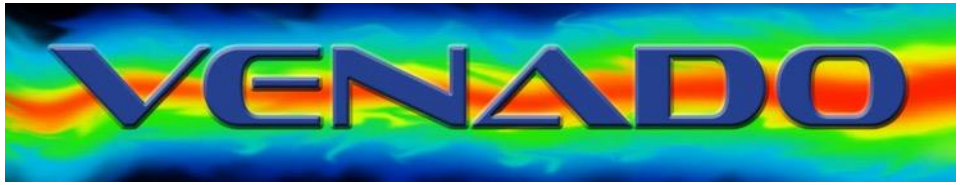
VENADO

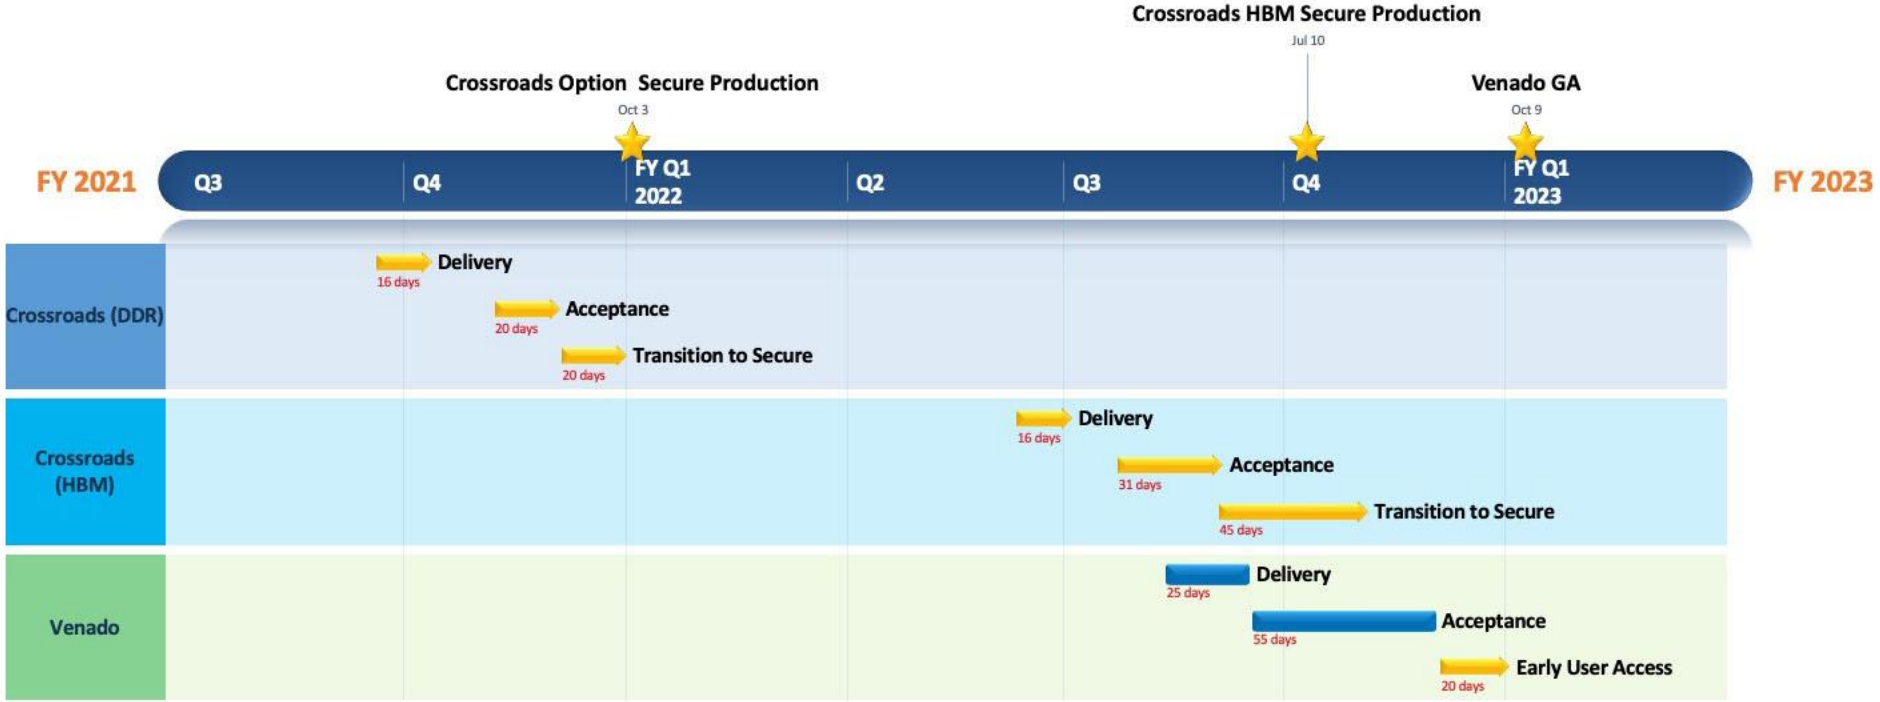
Crossroads HBM Secure Production
Jul 10
Crossroads Option Secure Production
Oct 3
Venado GA
Oct 9
FY 2021
Q3
Q4
FY Q1 2022
Q2
Q3
Q4
FY Q1 2023
FY 2023
Crossroads (DDR)
Delivery
16 days
Acceptance
20 days
Transition to Secure
20 days
Crossroads (HBM)
Delivery
16 days
Acceptance
31 days
Transition to Secure
45 days
Venado
Delivery
25 days
Acceptance
55 days
Early User Access
20 days