Trinity: Next-Generation Supercomputer for the ASC Program

Doug Doerfler
Sandia National Laboratories
Scalable Computer Architectures Department
Trinity Architecture Lead

SAND 2014-2739C
Unlimited Release

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Topics Covered

- ASC Computing Strategy
- Partnerships
- Joint Procurement Process
- Mission Drivers
- Workload Analysis
- Status

Status

- Formal Design Review completed April 2013
- Independent Project Review (Lehman) completed May 2013
- Trinity/NERSC8 RFP released August 2013
- Technical Evaluation of the proposals completed September 2013
- Initial negotiations for both systems completed November 2013
- NNSA Independent Cost Review completed Jan 2014
- NERSC8 is in the final approval stages
  - Should be announced "soon"
- Trinity went back to the proposing vendors for a Best and Final Offer (BAFO)
  - Target delivery date of September 2015 is unchanged

ASC computing strategy

- Approach: Two classes of systems
  - Advanced Technology: First of a kind systems that identify and foster technical capabilities and features that are beneficial to ASC applications
  - Commodity Technology: Robust, cost-effective systems to meet the day-to-day simulation workload needs of the program
- Investment Principles
  - Maintain continuity of production
  - Ensure that the needs of the current and future stockpile are met
  - Balance investments in system cost-performance types with computational requirements
  - Partner with industry to introduce new high-end technology constrained by life-cycle costs
  - Acquire right-sized platforms to meet the mission needs

Sandia National Laboratories

Advanced Technology Systems

- Leadership-class platforms
- Pursue promising new technology paths
- These systems are to meet unique mission needs *and* to help prepare the program for future system designs
- Includes Non-Recurring Engineering (NRE) funding to enable delivery of leading-edge platforms
- Trinity will be deployed by ACES (New Mexico Alliance for Computing at Extreme Scale, i.e. Los Alamos & Sandia)

ASC Platform Timeline

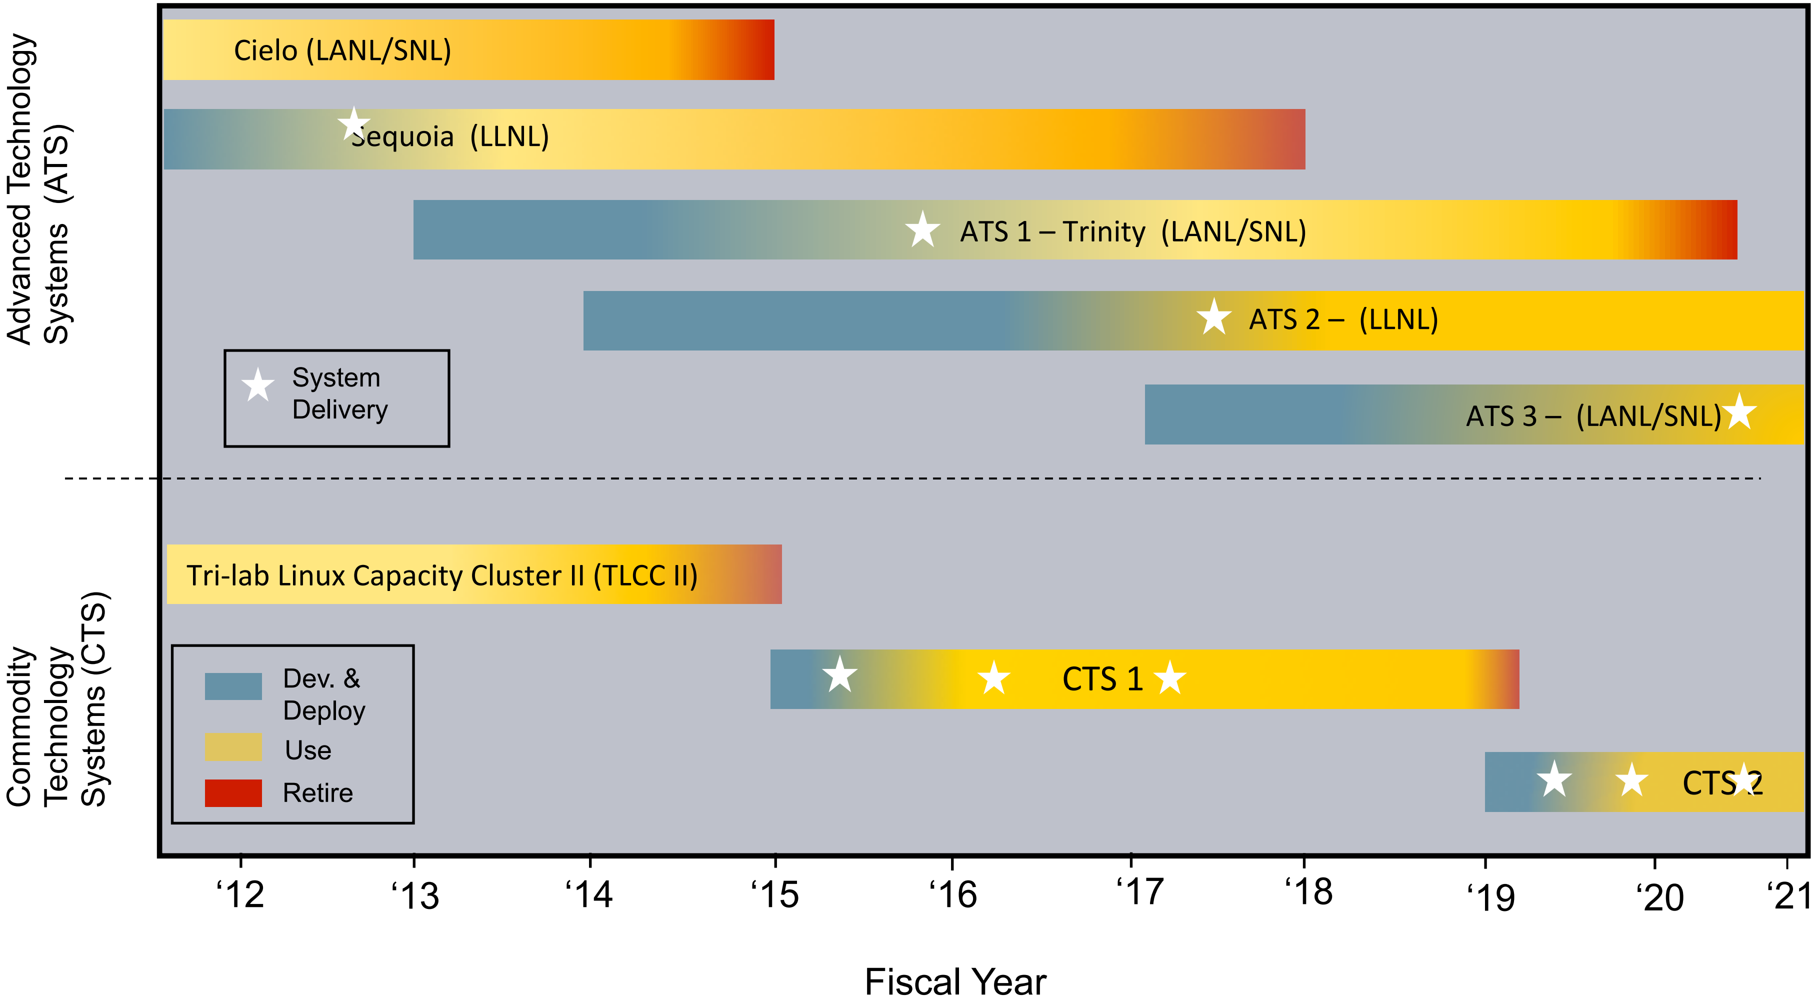

The ACES partnership and supercomputing is vital at both Los Alamos and Sandia

- SNL/LANL MOU signed March 2008 to integrate and leverage capabilities
- Commitment to the shared development and use of HPC to meet NW mission needs
- Continued world leadership in HPC more broadly for both institutions also a key goal
- Major efforts are executed by project teams chartered by and accountable to the ACES co-directors
  - Cielo delivered ca. 2011
  - Trinity delivery in 2015 is now our dominant focus
- Both Laboratories are fully committed to delivering a successful platform as it's essential to the Laboratories

ACES (NNSA) is partnering with LBL/NERSC (Office of Science) on the procurement

- Strengthen the alliance between NNSA and SC on road to exascale
- Show vendors a more united path on road to exascale
- Shared technical expertise between labs
- *Should* gain cost benefit
- Saves vendors money/time responding to a single RFP, single set of technical requirements
- Outside perspective reduces risk -- avoids tunnel vision by one lab
- More leverage with vendors by sharing information between labs
- Benefits in production, shared bug reports, quarterly meetings
- Less likely to be a one-off system with multiple sites participating

Why is the NNSA/ASC and SC/ASCR collaborating?

- The April 2011 MOU between SC & NNSA for coordinating Exascale activities was the impetus for ASC and ASCR to work together on the proposed Exascale Computing Initiative (ECI).
- While ECI is yet to be realized, ASC & ASCR program directors made strategic decisions to co-fund and collaborate on:
  - Technology R&D Investments: FastForward and DesignForward
  - System Acquisitions: Trinity/NERSC-8 and CORAL
  - Great leveraging opportunities to share precious resources (budget & technical expertise) to achieve each program's mission goals, while working out some cultural/bureaucratic differences.
- The Trinity/NERSC-8 collaboration will proceed with joint RFP and selection, separate system awards, attendance at other system's project reviews and collective problem solving.

This is a collaboration of two separate projects resulting in two distinct systems

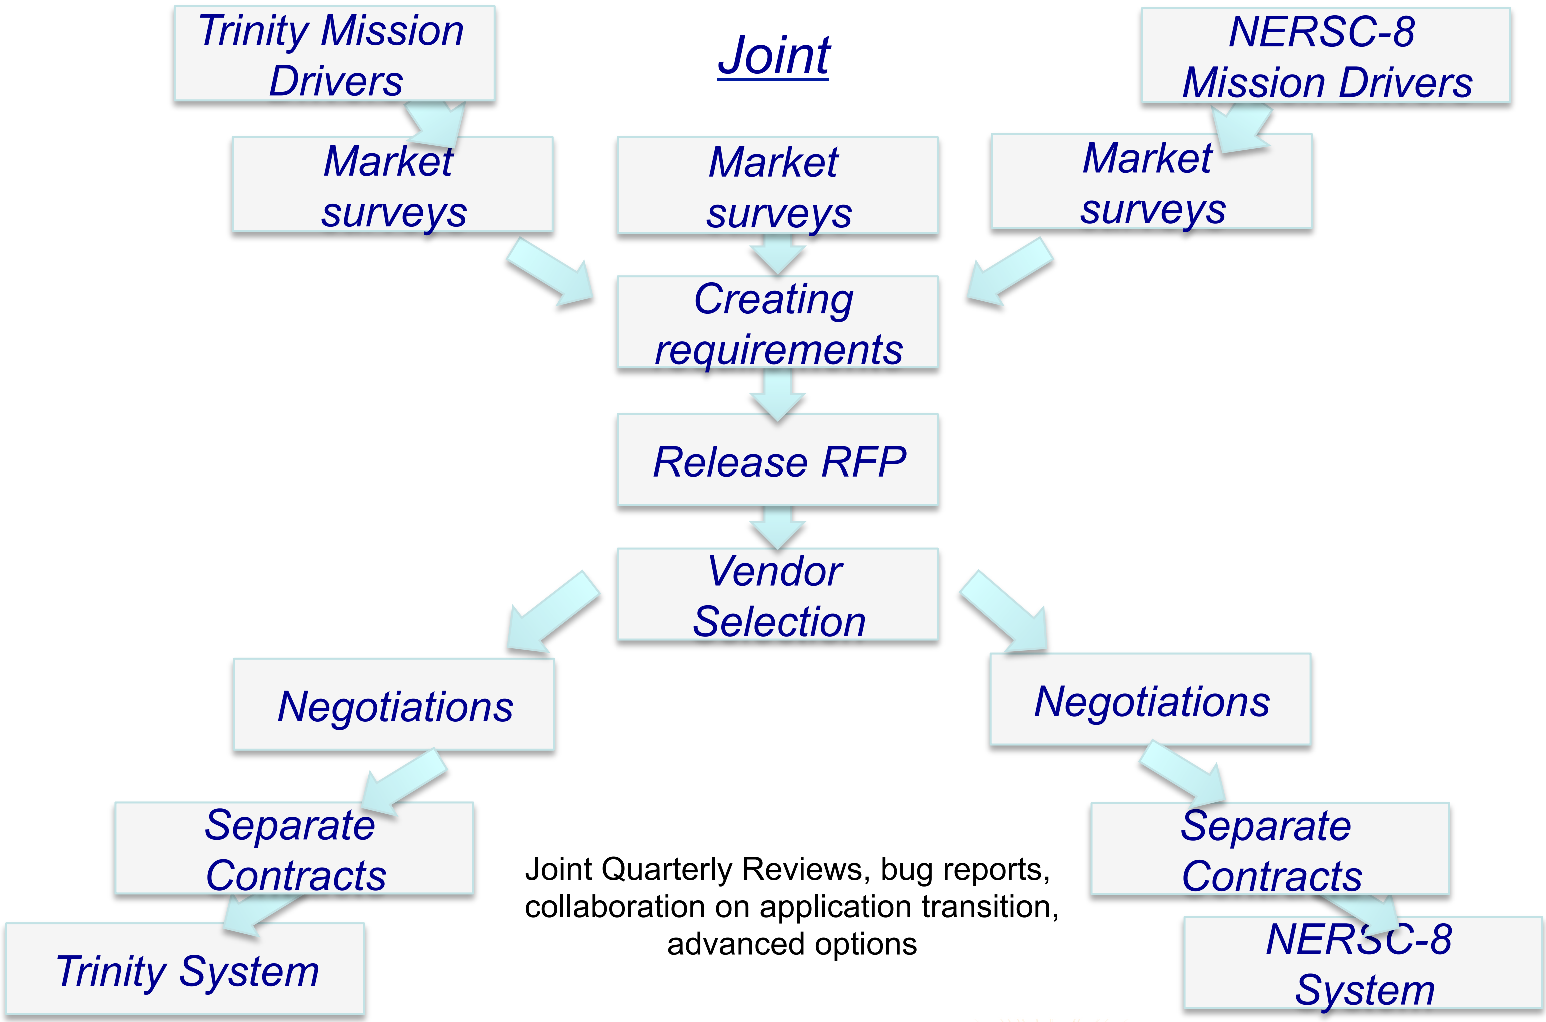

Target System Configurations

	Trinity	NERSC-8
Memory Capacity	2 PB to 4 PB	1 PB to 2 PB
Capability Improvement	8 to 10x over Cielo	8 to 10x over Hopper
Sustained System Performance (SSP)	20 to 60x over Hopper	10x to 30x over Hopper
JMTTI	> 24 hours	> 35 hours
File System BW metric – time to dump 80% RAM	20 mins	30 mins
File System disk capacity	> 30x main memory	> 20x main memory
Power	< 12+3 MW	< 6 MW
Off-platform I/O	> 140 GB/s	> 180 GB/s

Trinity Mission Need

- Mission Need developed with tri-Lab input
  - Meetings with stakeholders at all 3 Laboratories
  - Meetings/Telecons with tri-Labs and NNSA ASC Program
  - ACES technical staff
  - Integrated mission need
- Mission Requirements are primarily driving memory capacity
  - 2-4 PB of aggregate main memory
- Mission Need document (Trinity CD0) includes requirements, drivers, budget information, schedules, constraints, and limitations
- Trinity CD0 approved and signed in December 2012 by Donald Cook, Deputy Administrator for Defense Programs, NNSA

Workload Requirements indicate a critical need for significant increase in computing resources for ASC Program

TRINITY

- 19 workloads from LANL, LLNL & SNL
- 6 mission drivers
- Not all requirements identified
- Totals by year:

Mapping ASC Mission Goals to Cielo/Trinity Jobs	
Mission	**Computations**
Stockpile Stewardship	Directed Stockpile Work
	ASC Validation and Verification
Nonnuclear Work	Office of Secure Transport
Weapons Science	Science Campaigns, predictive capability framework
	High Energy Density Physics
Global Security	Global Security

Year	2016	2017	2018	2019	2020
Compute (Cielo Years)	27.3	49.6	74.4	79.3	79.3

Cielo Workload Analysis: Job Size Breakdown for CCC-3

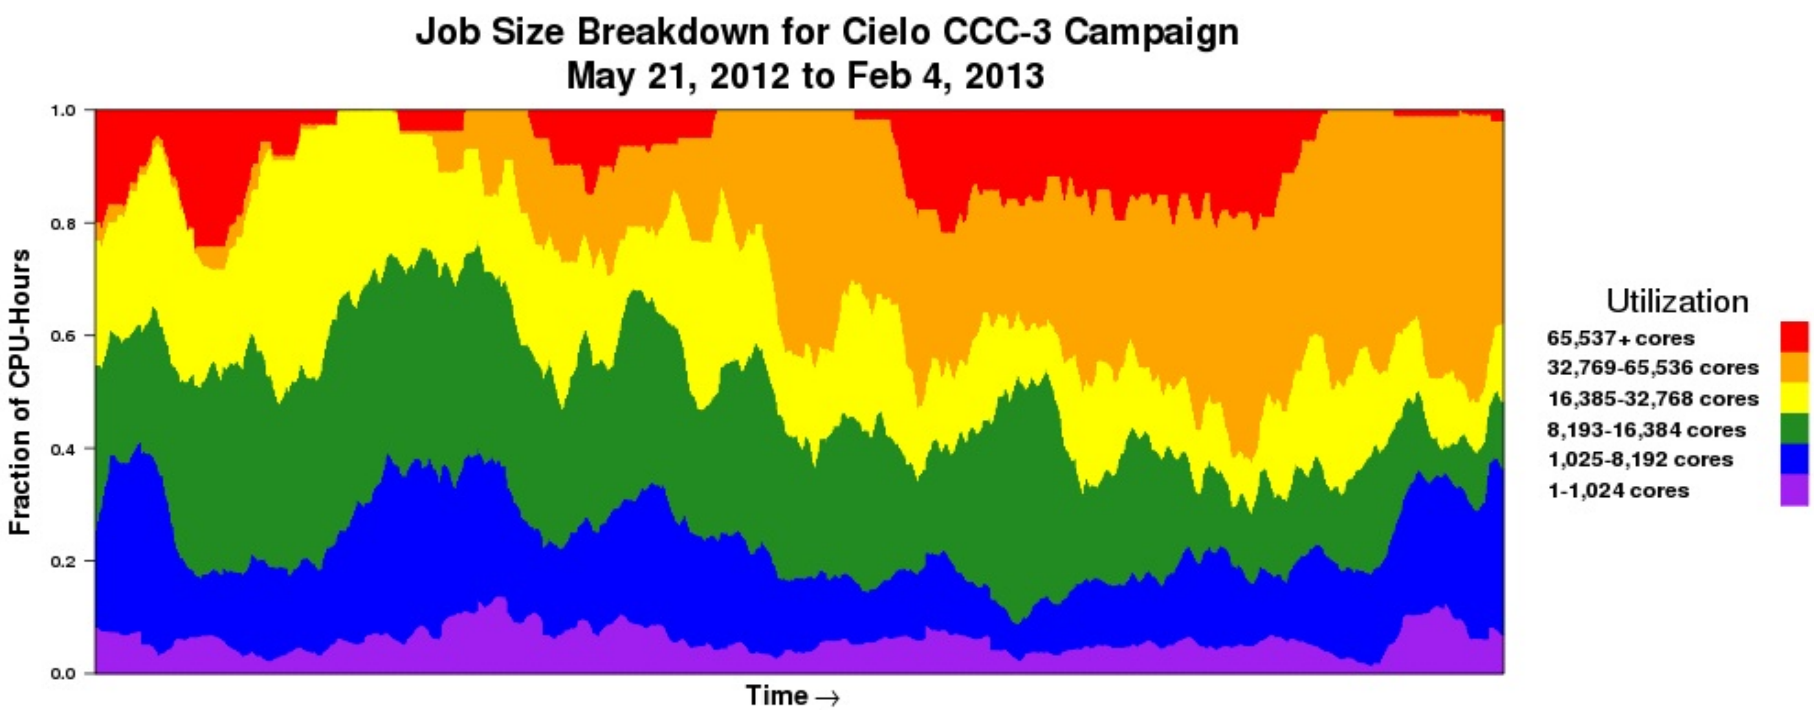

Trinity is Sized for Large 3D Jobs

Table 1. ASC 3D Simulation Memory Requirements over Time

Timeframe/ Fidelity	Physics fidelity	Geometric fidelity	Numerical fidelity	Restart file memory size
1980	Not applicable Could not run the 3D simulations due to memory size limitations	Not applicable	Not applicable	Not applicable
1990	Low	Low Not all items included	Low Cell size was large	$3x10^{9}$ or 3 Gbytes
2000	Low	Medium Not all items included	Medium	$3x10^{11}$ or 300 Gbytes
2010	More physics added	Medium	Medium	$8x10^{13}$ or 80 Tbytes
2015 Projected	Increased physics	Higher	Higher	$7.5x10^{14}$ or 750 Tbytes

Point Design:
Current LANL 3D problem runs on ¼ of Cielo today with about 80 TBytes of memory

Projected memory size on Trinity of this increased resolution problem is about 750 TBytes of memory on Trinity

Trinity target memory size is 2 PBytes (minimum) to 4 PBytes to run several (2-4) jobs of this size (1/4 to ½ of system)

Computational Workload by Driver

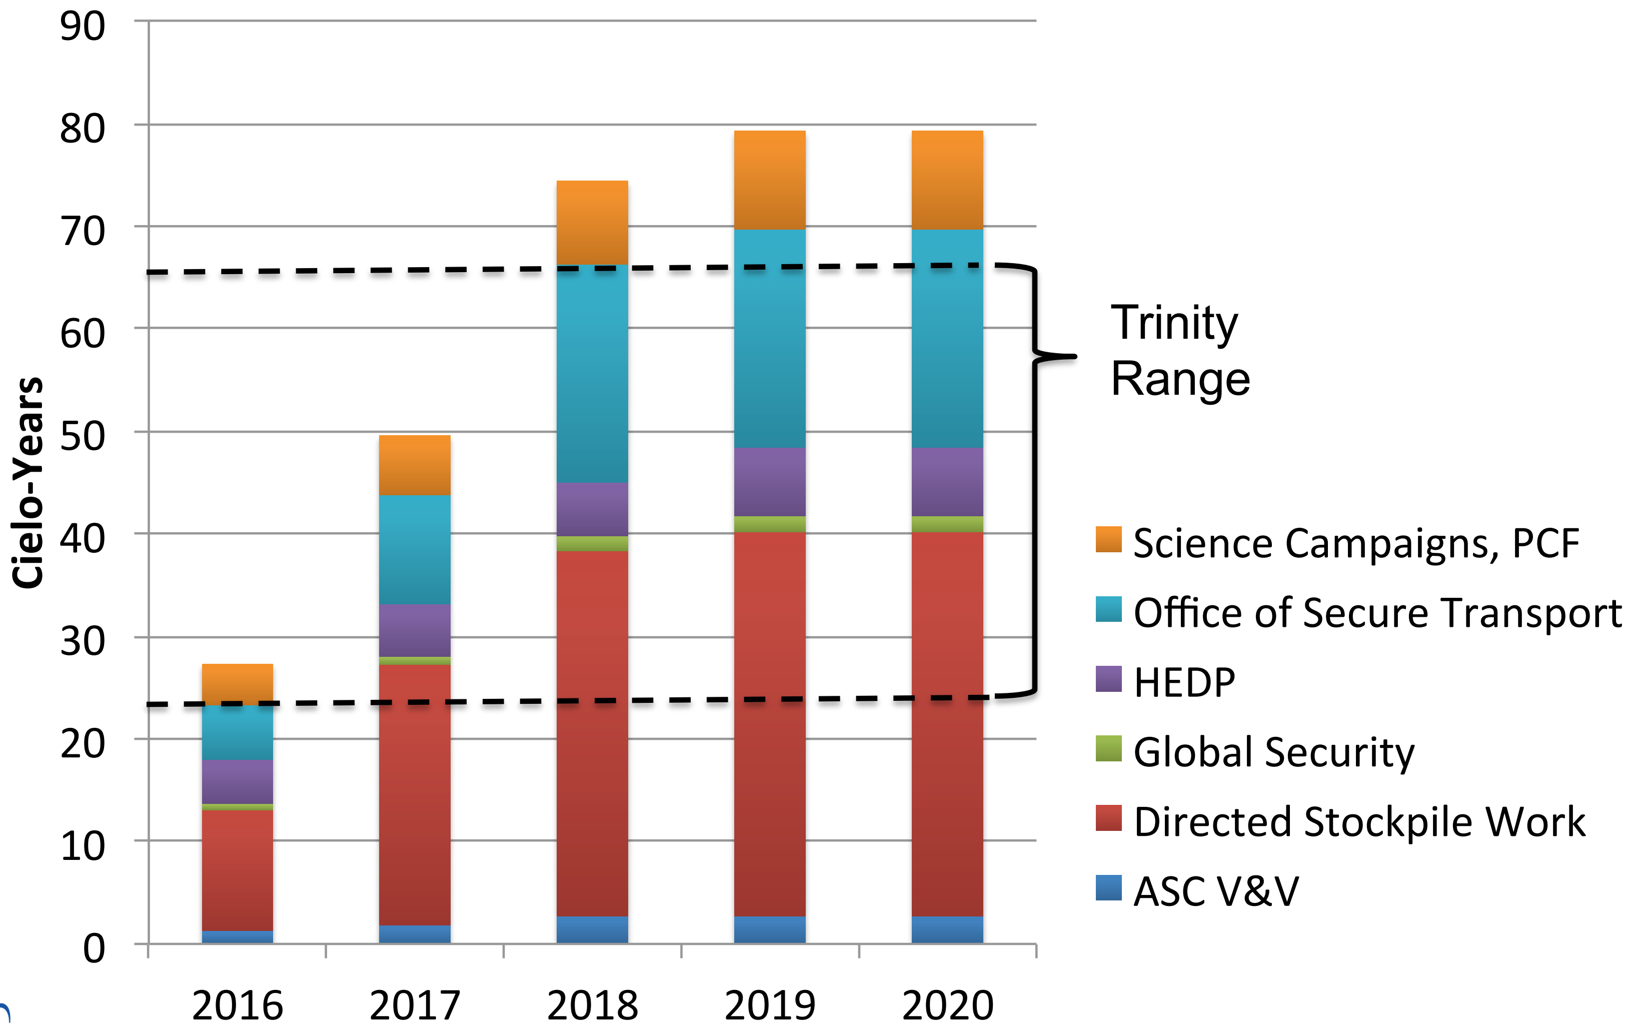

Memory Requirements are Increasing

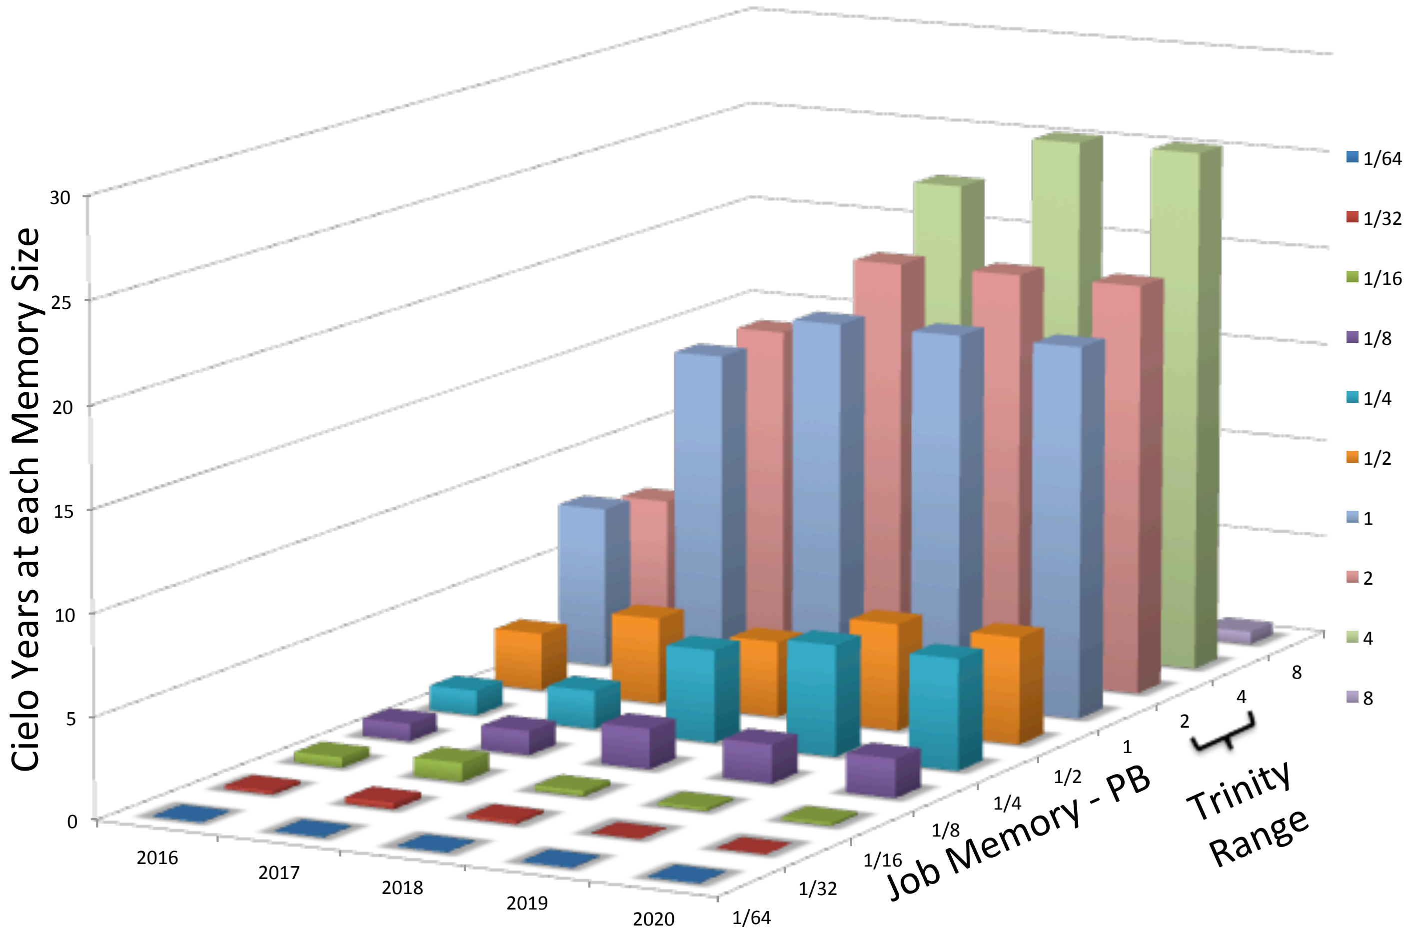

Status

- Formal Design Review completed April 2013
- Independent Project Review (Lehman) completed May 2013
- Trinity/NERSC8 RFP released August 2013
- Technical Evaluation of the proposals completed September 2013
- Initial negotiations for both systems completed November 2013
- NNSA Independent Cost Review completed Jan 2014
- NERSC8 is in the final approval stages
  - Should be announced “soon”
- Trinity went back to the proposing vendors for a Best and Final Offer (BAFO)
  - Target delivery date of September 2015 is unchanged

Special thanks to the following for contribution of slides used in this talk

- Thuc Hoang – ASC HQ
- Manuel Vigil – LANL
- Jim Lujan – LANL
- Katie Antypas – LBL
- Nick Wright – LBL
- and a cast of others …