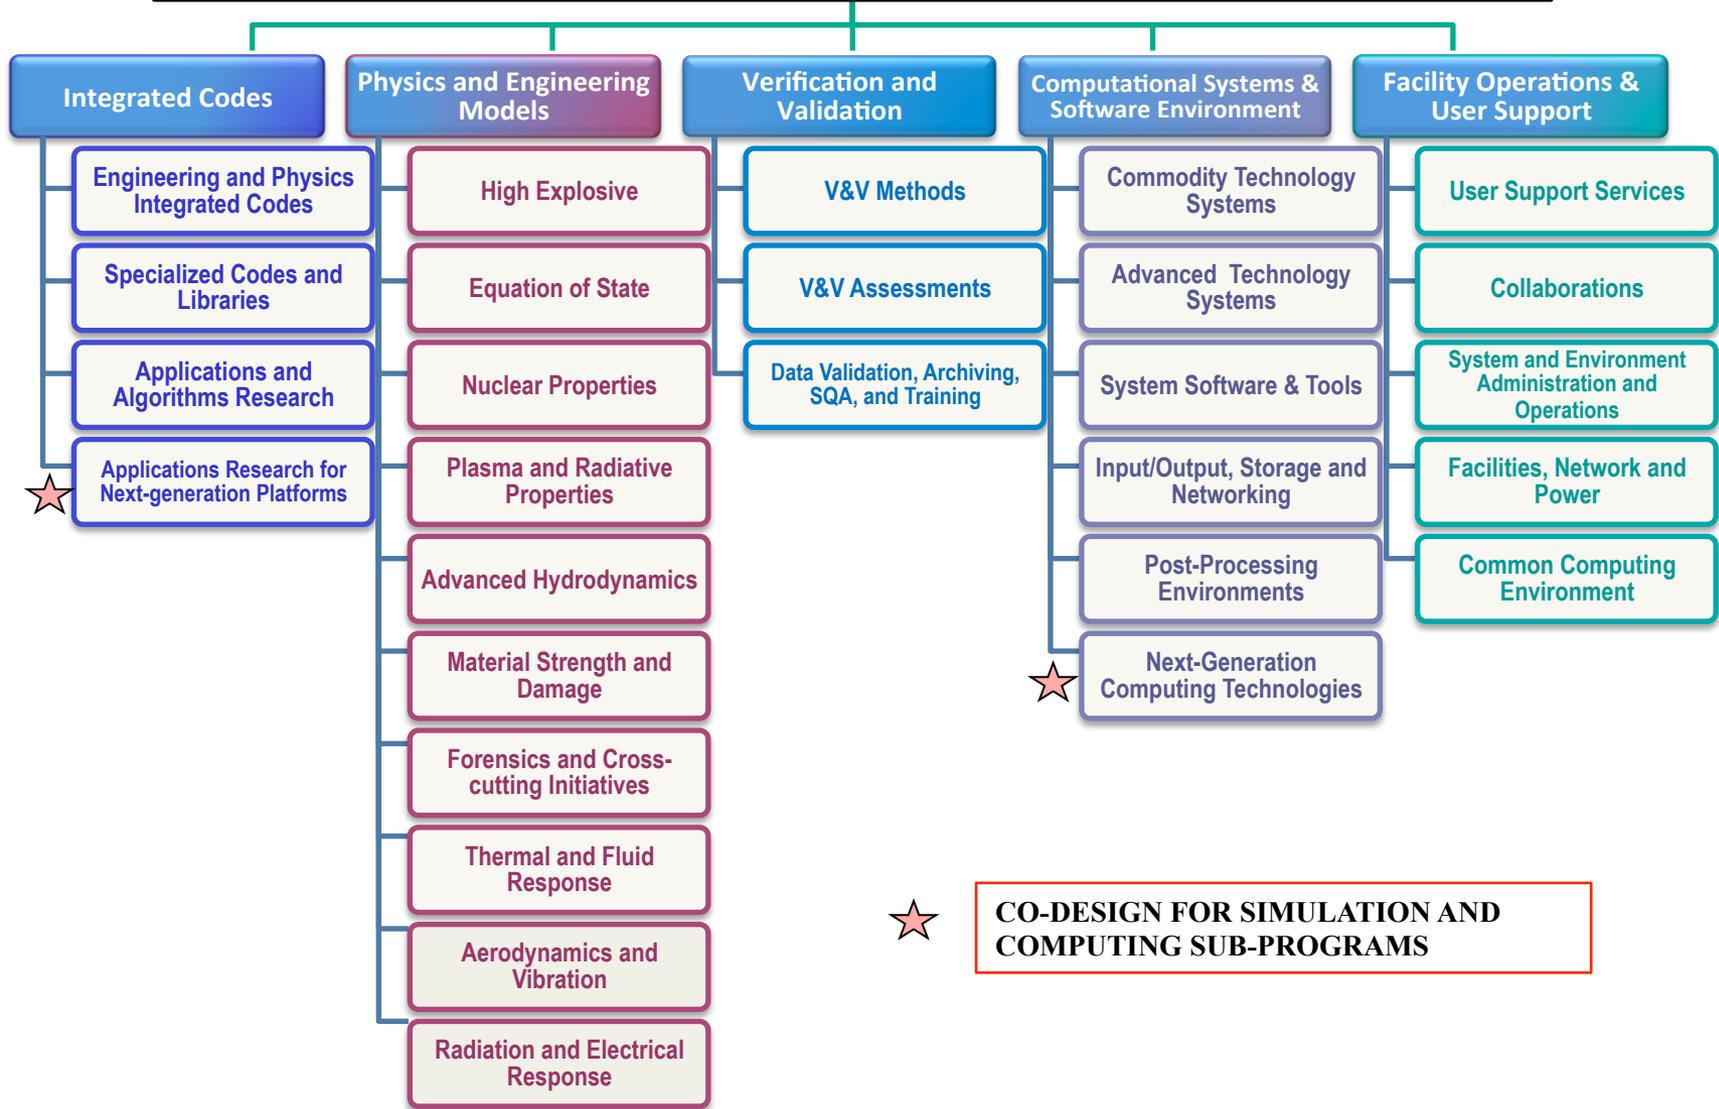


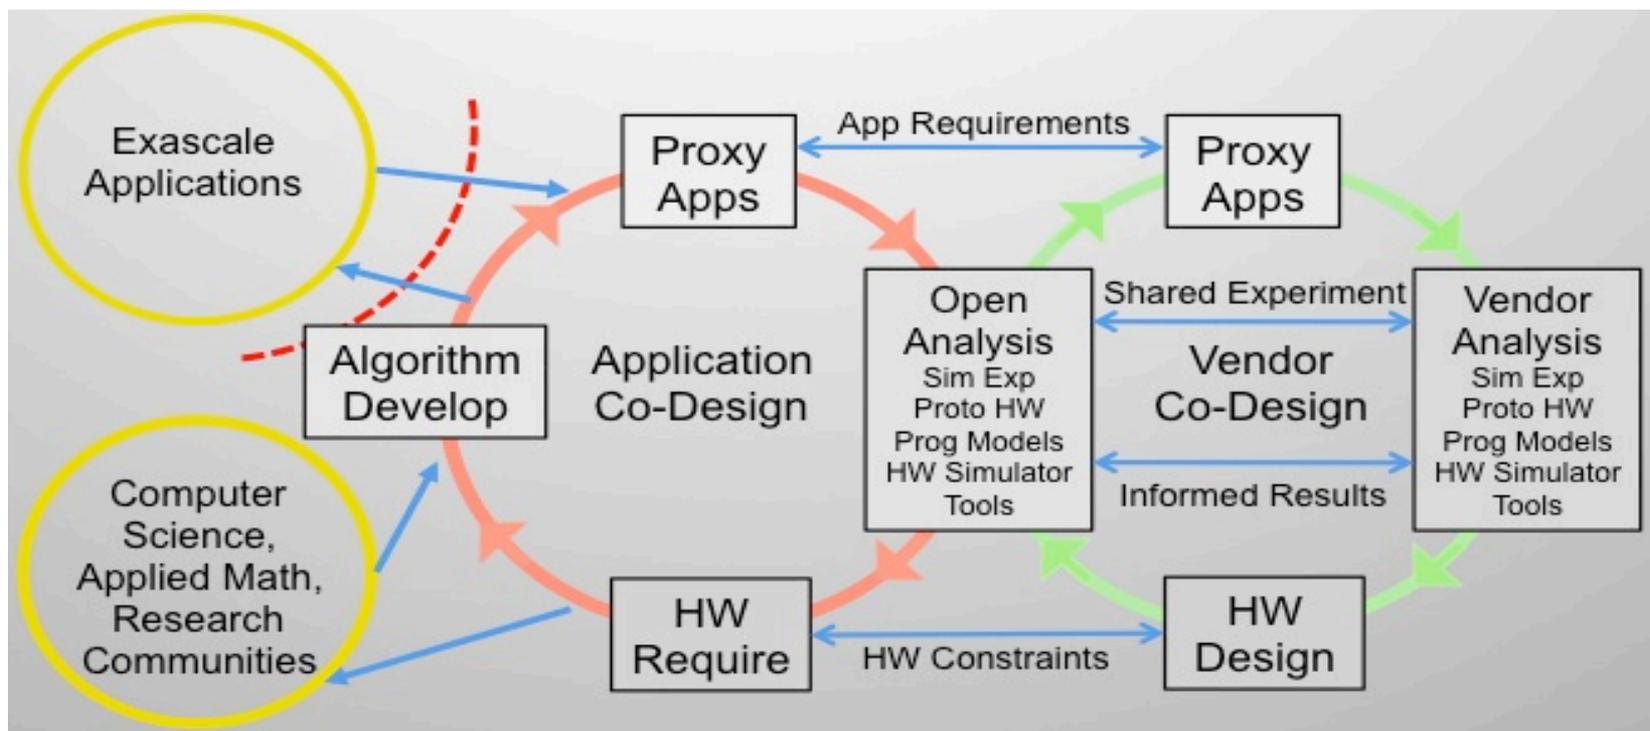
# Preparing for Next Computing Paradigm,

**Sriram Swaminarayan**  
**For Thuc Hoang**

## ADVANCED SIMULATION AND COMPUTING



- To prepare and execute an efficient transition of the ASC application code base and system software tools to the next paradigm of extreme-parallelism computing
  - collaborate closely with other US HPC programs, computing vendors, academia, and international organizations.



# Our co-design activities will depend on how broadly we define the term

Co-design is a tool for application developers to inform and influence vendor design of future architectures and system software. And conversely, provide app developers with early insight into computing trends so they can adapt early.



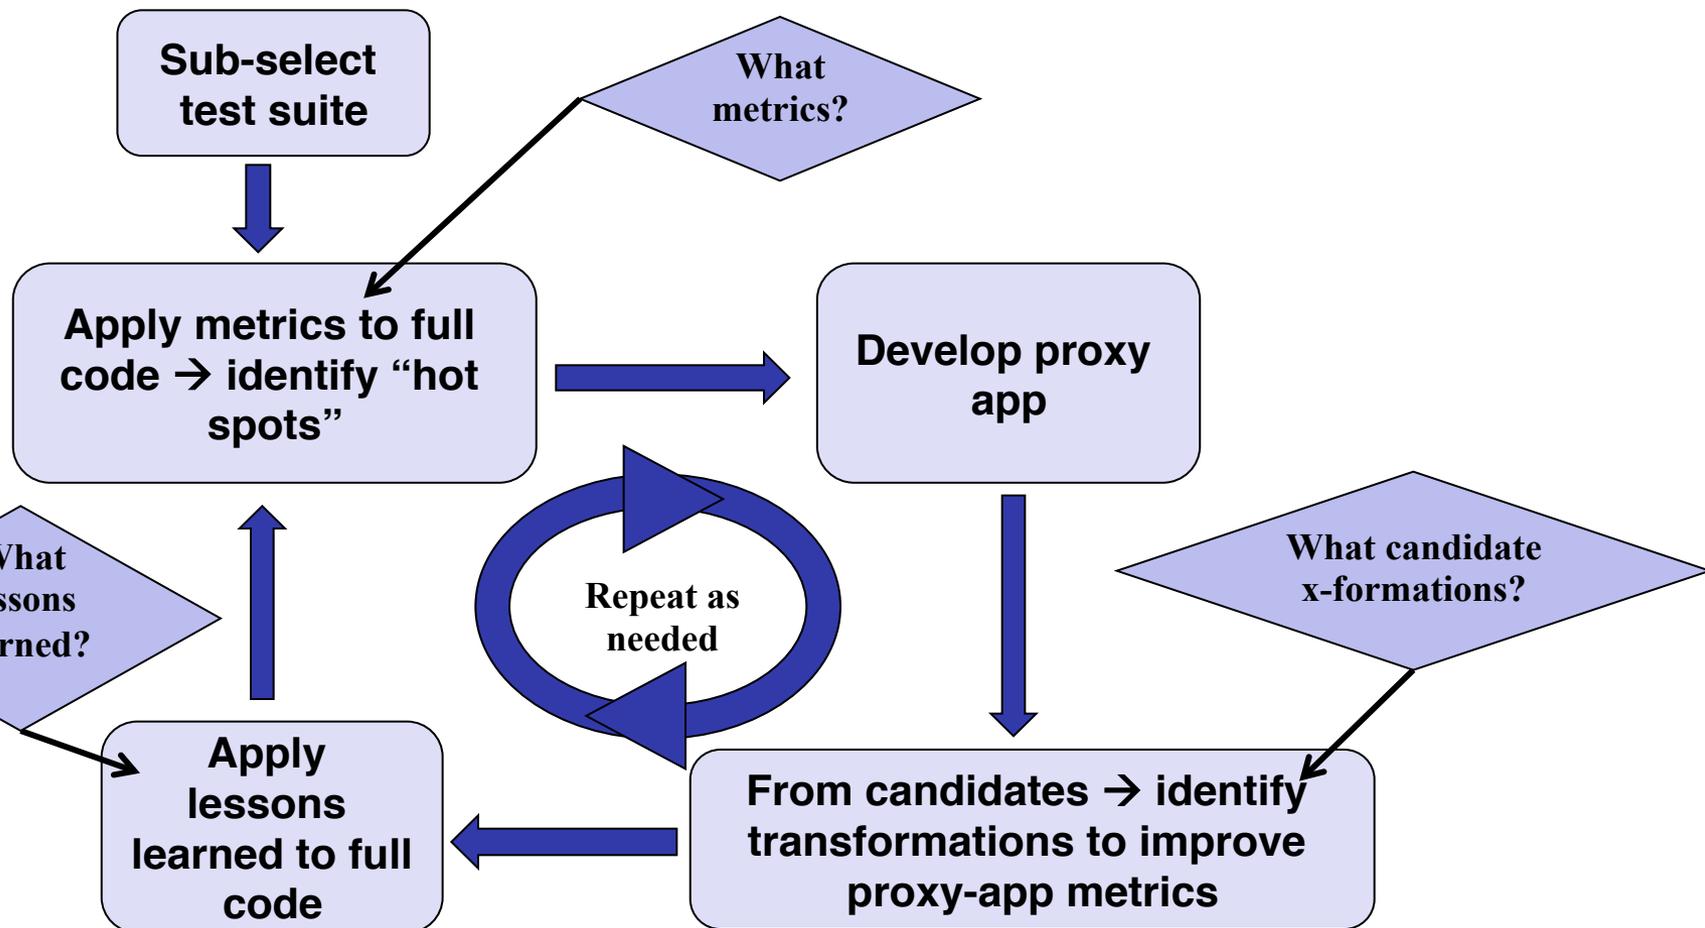
Co-design encapsulates a broad set of collaborative efforts, not limited to vendor engagements. Academia, the OSS community, language developers, and “big data” architects all provide a hand in helping us achieve exascale. Simply viewing the world through the eyes of another is the basis for a co-design relationship.

**DOE labs have been doing co-design for decades. What's different this time are the time scales (pre-procurement), and the willingness of the app developers to entertain major changes to their process (assuming sufficient lead time!)**

# Proxy Applications Represent an Initial Thrust at the three ASC Labs

- Applications of interest to NNSA are not open and sharable with the community
  - Proxy apps are a requirement for us to work with the community
  - Multi-physics proxy apps *may* fall under export control
- Proxy apps are a useful learning tool internal to the labs
  - Integrated codes have become so large, explorations of new concepts are difficult
  - Offer our own developers a platform for exploration of data layouts, algorithm changes, programming models, etc...
- Ensuring lessons learned in proxy apps are apropos to our internal codes will represent a major effort
  - Do the proxy apps truly represent the full apps? E.g. Instruction mix, data layouts, degree of parallelism, etc...

# Proxy app explorations must tie back to the full applications they are representing



**When does this approach “converge”?**

- **FastForward: June 2012-2014**
  - **AMD:** heterogeneous processor & 2-level memory
  - **IBM:** advanced memory concept
  - **Intel:** core energy efficiency & processing-near-memory
  - **Intel/WhamCloud:** storage reliability, I/O API, burst buffer mgmt
  - **Nvidia:** memory hierarchy, processor/packaging/programming
  
- **DesignFoward-Interconnect: October 2013-2015**
  - **AMD:** interconnect architectures and associated execution models
  - **Cray:** open network protocol standards
  - **IBM:** energy-efficient interconnect architectures and messaging models
  - **Intel:** interconnect architectures and implementation
  - **Nvidia:** interconnect architectures for massively threaded processors

# Tri-lab Advanced System Technology Testbeds - A Co-Design Tool

- **Not for production computing cycles - but can be provided to Test Pilot users**
- **Both hardware and software are intended to be highly dynamic**
- **Closer to prototypes and technology development drivers**
- **Multiple nodes available but more important to explore a diverse set of architectural alternatives, than push large scale**
- **Available for PSAAP access, via SARAPE and coordinated by the CRT**

Hostname	Compton	Curie	Shannon	Teller	Volta
CPU	Dual Socket Intel E5-2670 (Sandy Bridge) 2.6 GHz 8-core	Interlagos	Sandy Bridge	AMD A10-5800K (Piledriver) 3.8GHz Quad-core	Intel Xeon E5-2695 V2 dual socket (Ivy Bridge) total 24 cores 2.4 GHz
Accelerator	Pre-Production Intel Xeon Phi Knights Corner (KNC) 2 per node	Nvidia Kepler K20X	Nvidia Kepler K20X 2 per node	Radeon Northern Islands (on die integration)	None
GPU cores	57 1.1GHz cores	2688 732 MHz cores	2688 732 MHz cores	384 800MHz cores	N/A
Nodes	42	52	32	104	56
Interconnect	Mellanox QDR IB	Gemini	Mellanox QDR IB	QLogic QDR IB	Aires
Other	80GB SSD per node	Full featured RAS sys <b>On SRN</b>	Full PCI Gen 3 NVIDIA GPU Direct	Integrated CPU/GPU+ 256GB SSD/node power monitoring capability	Full featured RAS system including power monitoring and control capabilities

- Coordination for one-on-one collaborations between Centers and the tri-labs should be managed through the Center-TST channel
- FastForward program: <https://asc.llnl.gov/fastforward/>
- Co-design & Proxy Apps documents: <http://proxyapps.lanl.gov/>
- CSSE, FastForward & DesignForward: [Thuc.Hoang@nnsa.doe.gov](mailto:Thuc.Hoang@nnsa.doe.gov)



- **Rob Neely, LLNL**
  - Design space spanned by proxy applications
  - A sampling of LLNL's co-design activities with proxy application
- **Richard Barrett, SNL**
  - SNL's award winning Mantevo proxy application suite
  - Validation methodology
  - Some examples to show co-design connection between app and mini-app
- **Sriram Swaminarayan, LANL**
  - LANL's co-design activities
  - A sampling of LANL's proxy applications