

# BUILDING THE IESP ROADMAP

Bob Lucas, Pete Beckman, Franck Cappello, Jack Dongarra, Satoshi Matsuoka, Bernd Mohr, Terry Moore, Thomas Sterling, Kathy Yelick

# Purpose



- The IESP software roadmap is a planning instrument designed to enable the international HPC community to improve, coordinate and leverage their collective investments and development efforts.
- After we determine what needs to be accomplished, our task will be to construct the organizational structures suitable to accomplish the work

# Key Trends

# Requirements on X-Stack

- Increasing Concurrency
- Reliability Challenging
- Power dominating designs
- Heterogeneity in a node
- I/O and Memory: ratios and breakthroughs

- Programming models, applications, and tools must address concurrency
- Software and tools must manage power directly
- Software must be resilient
- Software must address change to heterogeneous nodes
- Software must be optimized for new Memory ratios and need to solve parallel I/O bottleneck

# Roadmap Components (1/2)

System SW	OS	Pete	Beckman
		Barney	MacCabe
System SW	Runtime	Jesus	Labarta
		Rajeev	Thakur
System SW	I/O	Alok	Choudhary
		Yutaka	Ishikawa
System SW	External Env	Giovanni	Aloisio
		?	?
System SW	System Mgmt	Bill	Kramer
		Bob	Wisniewski
Devel Env.	Prog. Models	Barbara	Chapman
		Mitsuhisa	Sato
Devel Env.	Frameworks	Mike	Heroux
		Robert	Harrison
Devel Env.	Compiler	Barbara	Chapman
		Mitsuhisa	Sato
Devel Env.	Libraries	Jack	Dongarra
		Anne	Trefethen
Devel Env.	Debugging	Wolfgang	Nagel
		David	Skinner

# Roadmap Components (2/2)

Applications	App Pioneers	Bill	Tang
		Richard	Kenway
Applications	Algorithms	Bill	Gropp
		Fred	Streitz
Applications	Data & Vis	John	Taylor
		Rick	Stevens
Cross-Cuts	Resilience	Franck	Cappello
		Sudip	Dosanjh
Cross-Cuts	Power	Satoshi	Matsuoka
		John	Shalf
Cross-Cuts	Performance	Bernd	Mohr
		Jeffrey	Vetter
Cross-Cuts	Programmability	Thomas	Sterling
		Hiroshi	Nakashima

## Priority Research Direction (use one slide for each)

### Scientific and computational challenges

**Brief overview of the underlying scientific and computational challenges**

### Summary of research direction

**What will you do to address the challenges?**

### Potential scientific impact

**What new scientific discoveries will result?**

**What new methods and techniques will be developed?**

### Potential impact on SCIENCE DOMAIN

**How will this impact key open issues in SCIENCE DOMAIN?**

**What's the timescale in which that impact may be felt?**

# CFD Simulation for mechanical and vibratory behaviour of the fuel assemblies inside a nuclear core vessel

Expert name/affiliation - email: Yvan Fournier/EDF – yvan.fournier@edf.fr

## Scientific and computational challenges

Computations with smaller and smaller scales in larger and larger geometries for a better understanding of physical phenomena

⇒ A better optimisation of the production (margin benefits)

**2007:** 3D RANS, 5x5 rods, 100 millions cells, 2 M cpu.hours (4000 cores during 3 weeks)

**2015:** 3D LES Full vessel (17x17x196 rods) unsteady approach, >50 billion cells, 1000000 cores during few weeks

## Software issues – short term (2009/2011)

Mesh generation, visualization

Scalability, load balancing

Solvers (multi-grid, better&simpler pre-conditioner, ...)

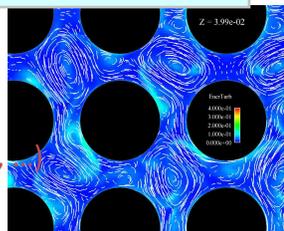
Mixing programming models (ex. MPI/OpenMP)

Stability and robustness of the software stack (MPI, ..)

API of scientific libraries (ex. BLAS!)

Standardisation of compiler optimisation level pragmas

Computing environment standardization (batch system, MPIExec,



## Software issues – long term 2015/2020

New numerical methods (stochastic, SPH, FV)

Scalability of linear solvers, hybrid solvers Code optimisation: wall of the collective communications, load balancing

Adaptive methods (may benefit all of computation/visualisation/meshing)

Data redistribution, IO (if flat MPI-IO model OK, good, otherwise require new “standard” data models)

Fault tolerance

Machine independent code optimisation & performance

## Impact of last machine change (x10 Gflops -> 100 Tflops)

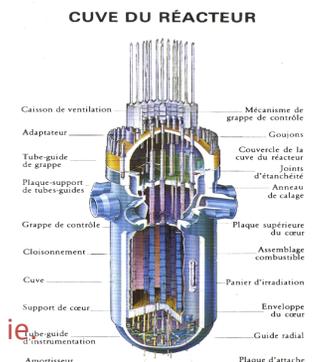
Pre/post adaptation

Reinforcement of the HPC expertise

Few extra “simple” programming rules

No rewriting, same solvers, same programming model, same software architecture thanks to technological evolution anticipation

Expected impact (100 Tflops -> Xpflops): ie  
2015 software issues



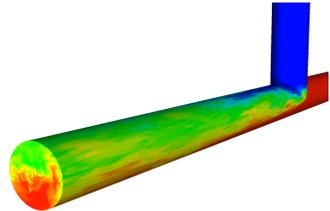
# Computational Challenges and Needs for Academic and Industrial



Consecutive thermal fatigue event

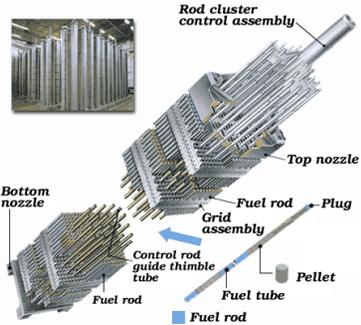
Computations enable to better understand the wall thermal loading in an injection.

Knowing the root causes of the event ⇒ define a new design to avoid this problem.



Computation with an L.E.S. approach for turbulent modelling

Refined mesh near the wall.



Part of a fuel assembly  
3 grid assemblies

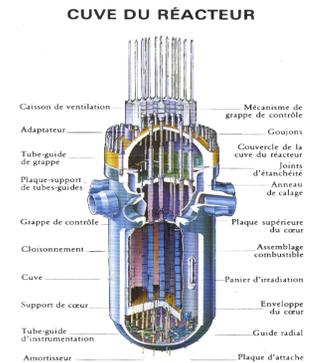
9 fuel assemblies

No experimental approach up to now

Will enable the study of side effects implied by the flow around neighbour fuel assemblies.

Better understanding of vibrational phenomena and wear-out of the rods.

The whole vessel reactor



**Computations with smaller and smaller scales in larger and larger geometries**  
 ⇒ a better understanding of physical phenomena ⇒ a more effective help for decision making  
 ⇒ A better optimisation of the production (margin benefits)

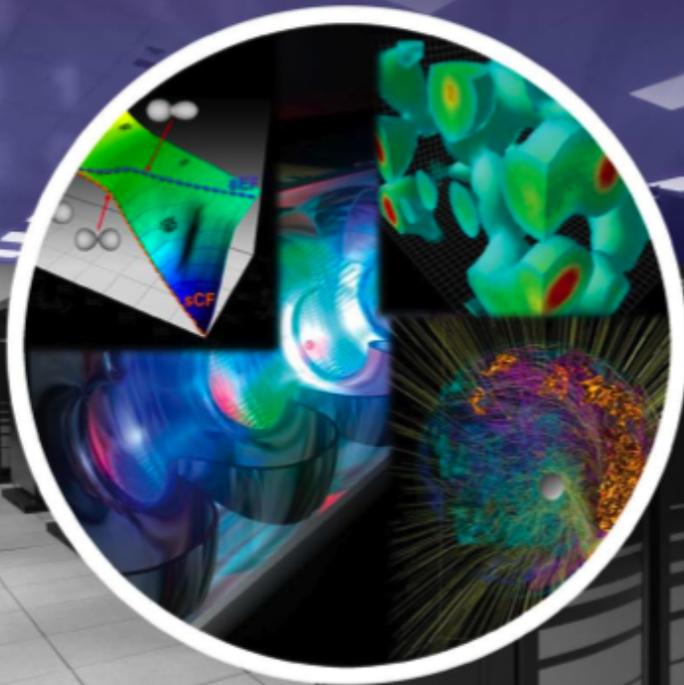


Fujitsu VPP 5000	Cluster, IBM Power5	IBM Blue Gene/L « Frontier »	600 Tflops during 1 month	10 Pflops during 1 month
1 of 4 vector processors	400 processors	20 Tflops during 1 month		
2 month length computation	9 days			
# 1 Gb of storage	# 15 Gb of storage	# 200 Gb of storage	# 1 Tb of storage	# 10 Tb of storage
2 Gb of memory	25 Gb of memory	250 Gb of memory	2,5 Tb of memory	25 Tb of memory
Power of the computer	Pre-processing not parallelized	Pre-processing not parallelized	... ibid. ...	... ibid. ...
		Mesh generation	... ibid. ...	... ibid. ...
		IESP/Application Subgroup	Scalability / Solver	... ibid. ...
				Visualisation

# Scientific Grand Challenges

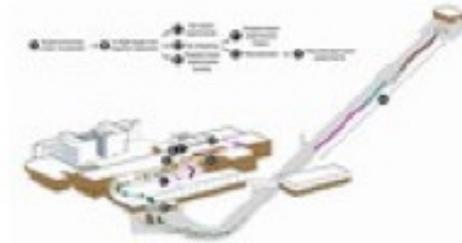
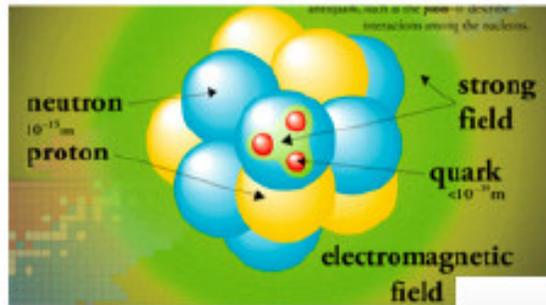
FOREFRONT QUESTIONS IN NUCLEAR SCIENCE AND  
THE ROLE OF HIGH PERFORMANCE COMPUTING

January 26-28, 2009 • Washington, D.C.



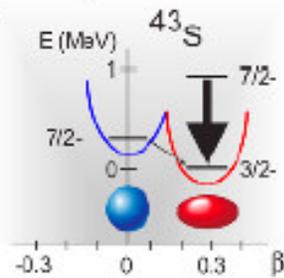
*Sponsored by the Office of Nuclear Physics and the Office of Advanced Scientific Computing Research*

# Understanding Nuclear Structure



$^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$   
 $^{132}\text{Sn}$  structure

*Ab initio* structure  
in light nuclei



$^{78}\text{Ni}$  structure

$^8\text{Be}(\alpha, \gamma)^{12}\text{C}$

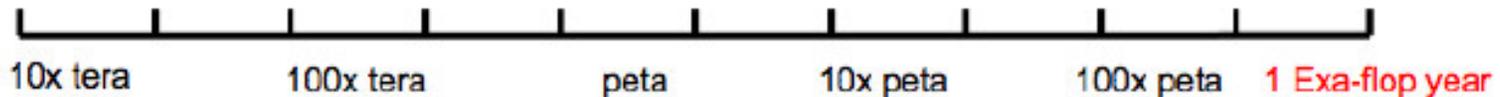
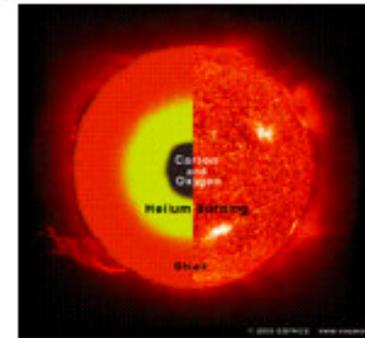


Image courtesy of James P. Vary (Iowa State University).

# Microscopic Nuclear Fission Solutions

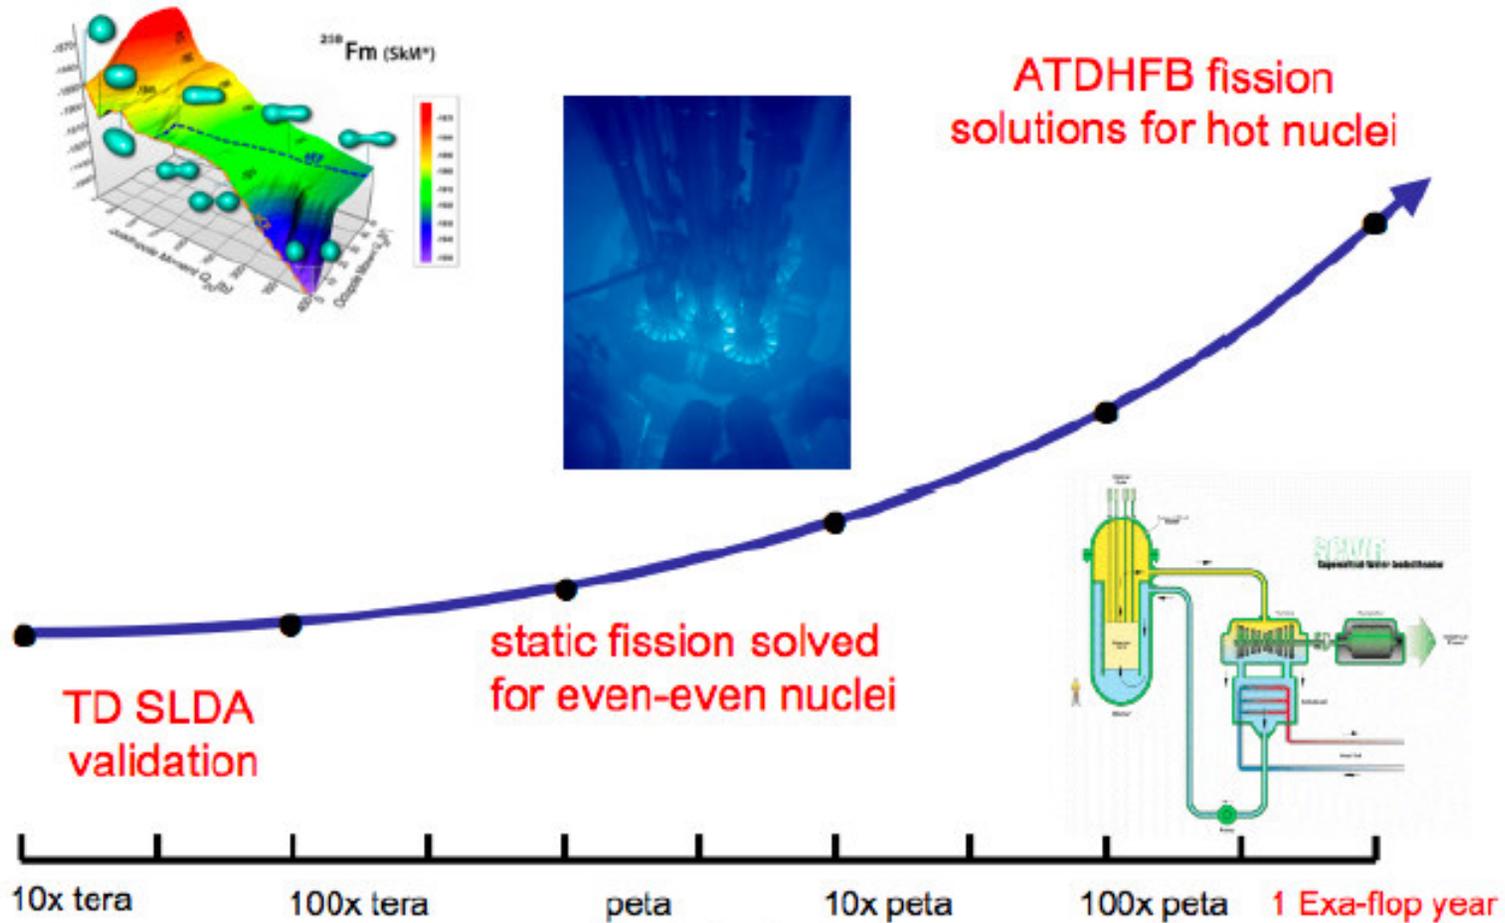


Image courtesy of James P. Vary (Iowa State University).

# Neutron Star – Extreme Nuclei

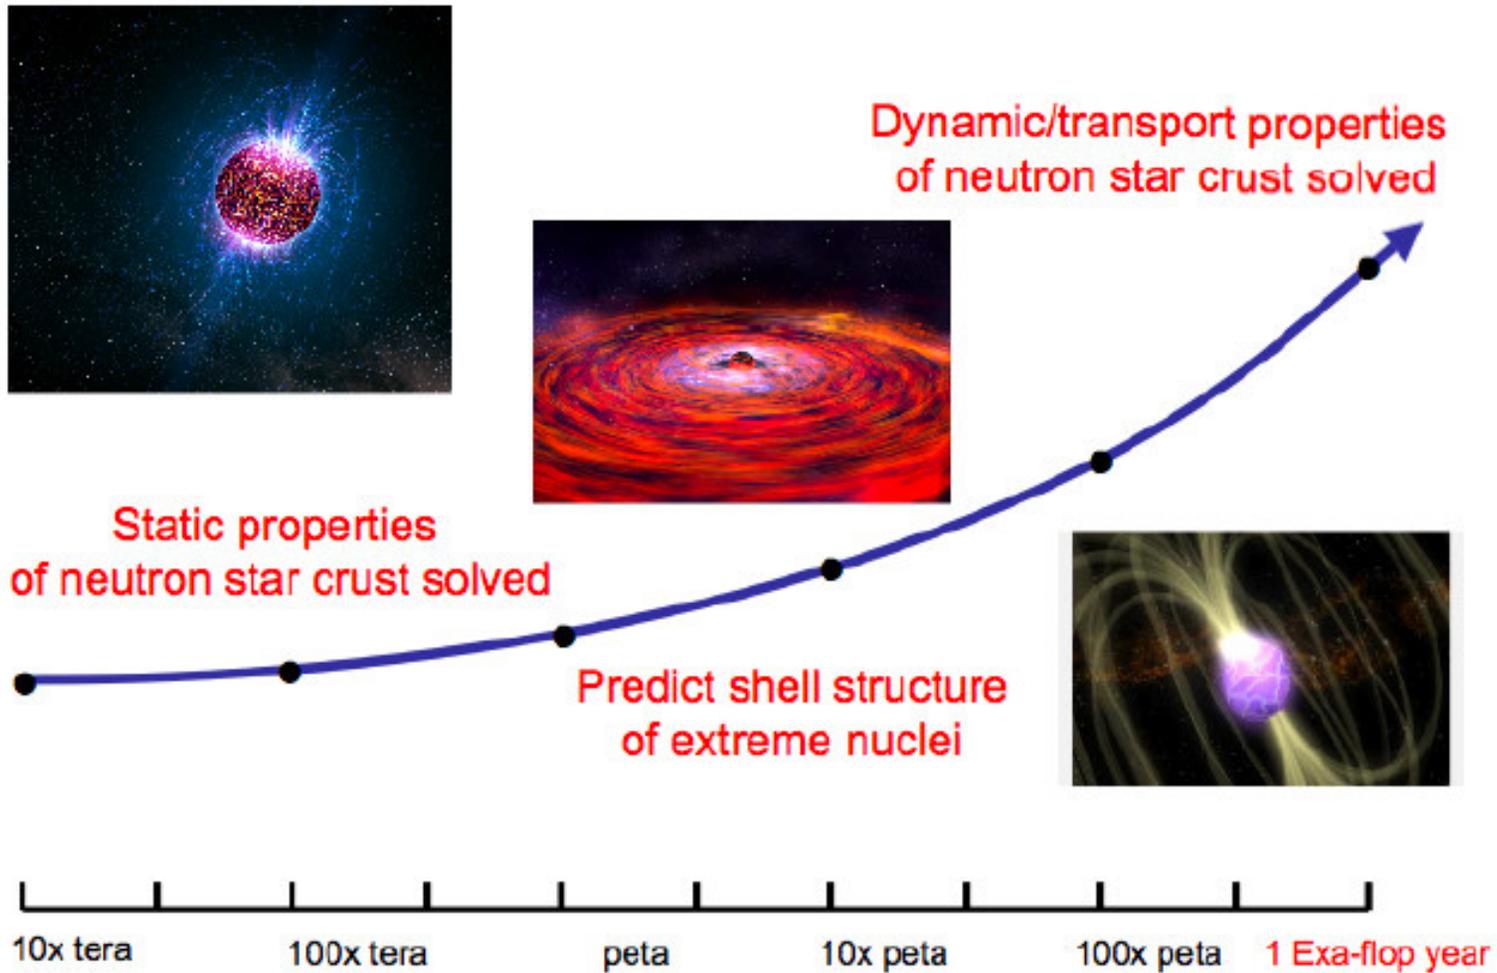


Image courtesy of James P. Vary (Iowa State University).

# Quantum Chromodynamics -Cold

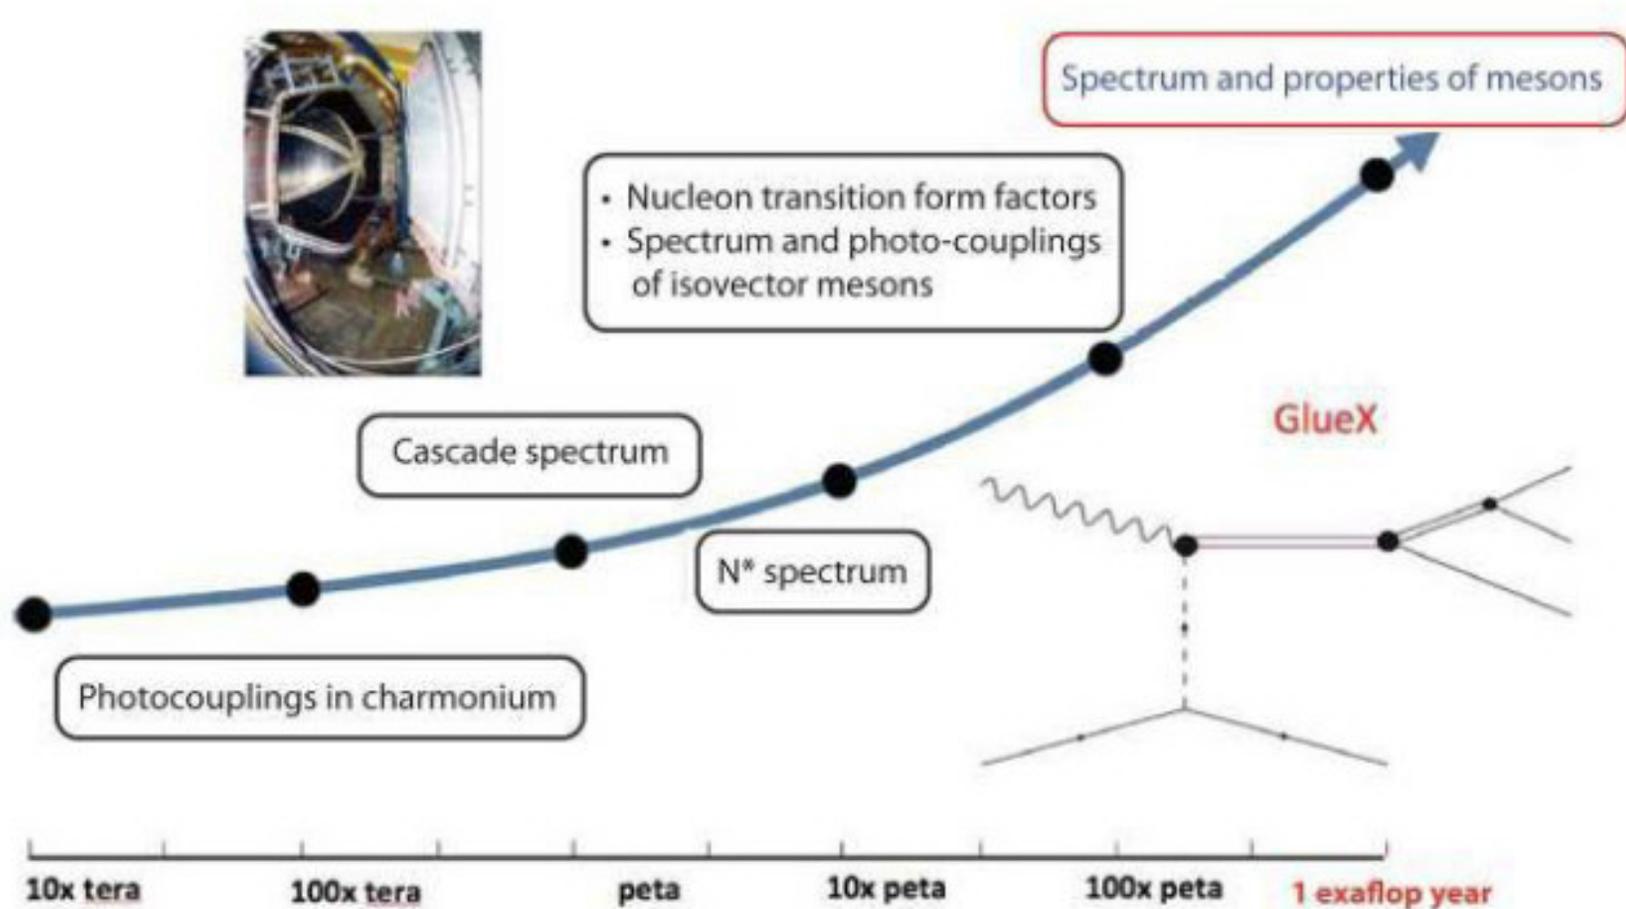


Image courtesy of Thomas Luu (LLNL) and David Richards (Jlab)

# Quantum Chromodynamics - Hot

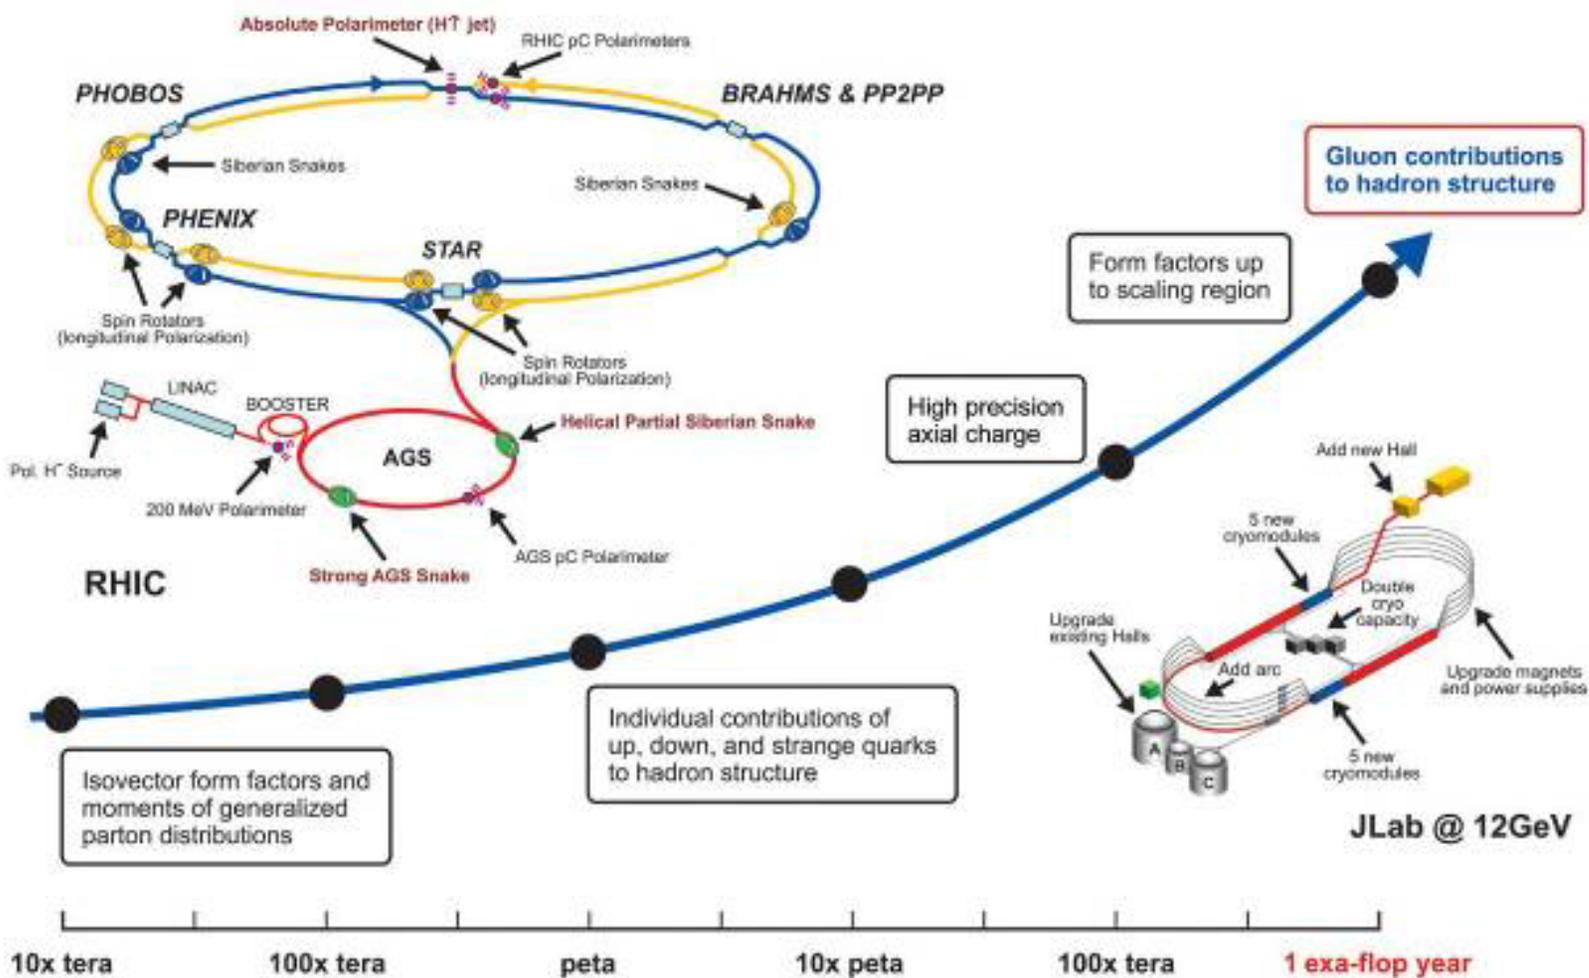
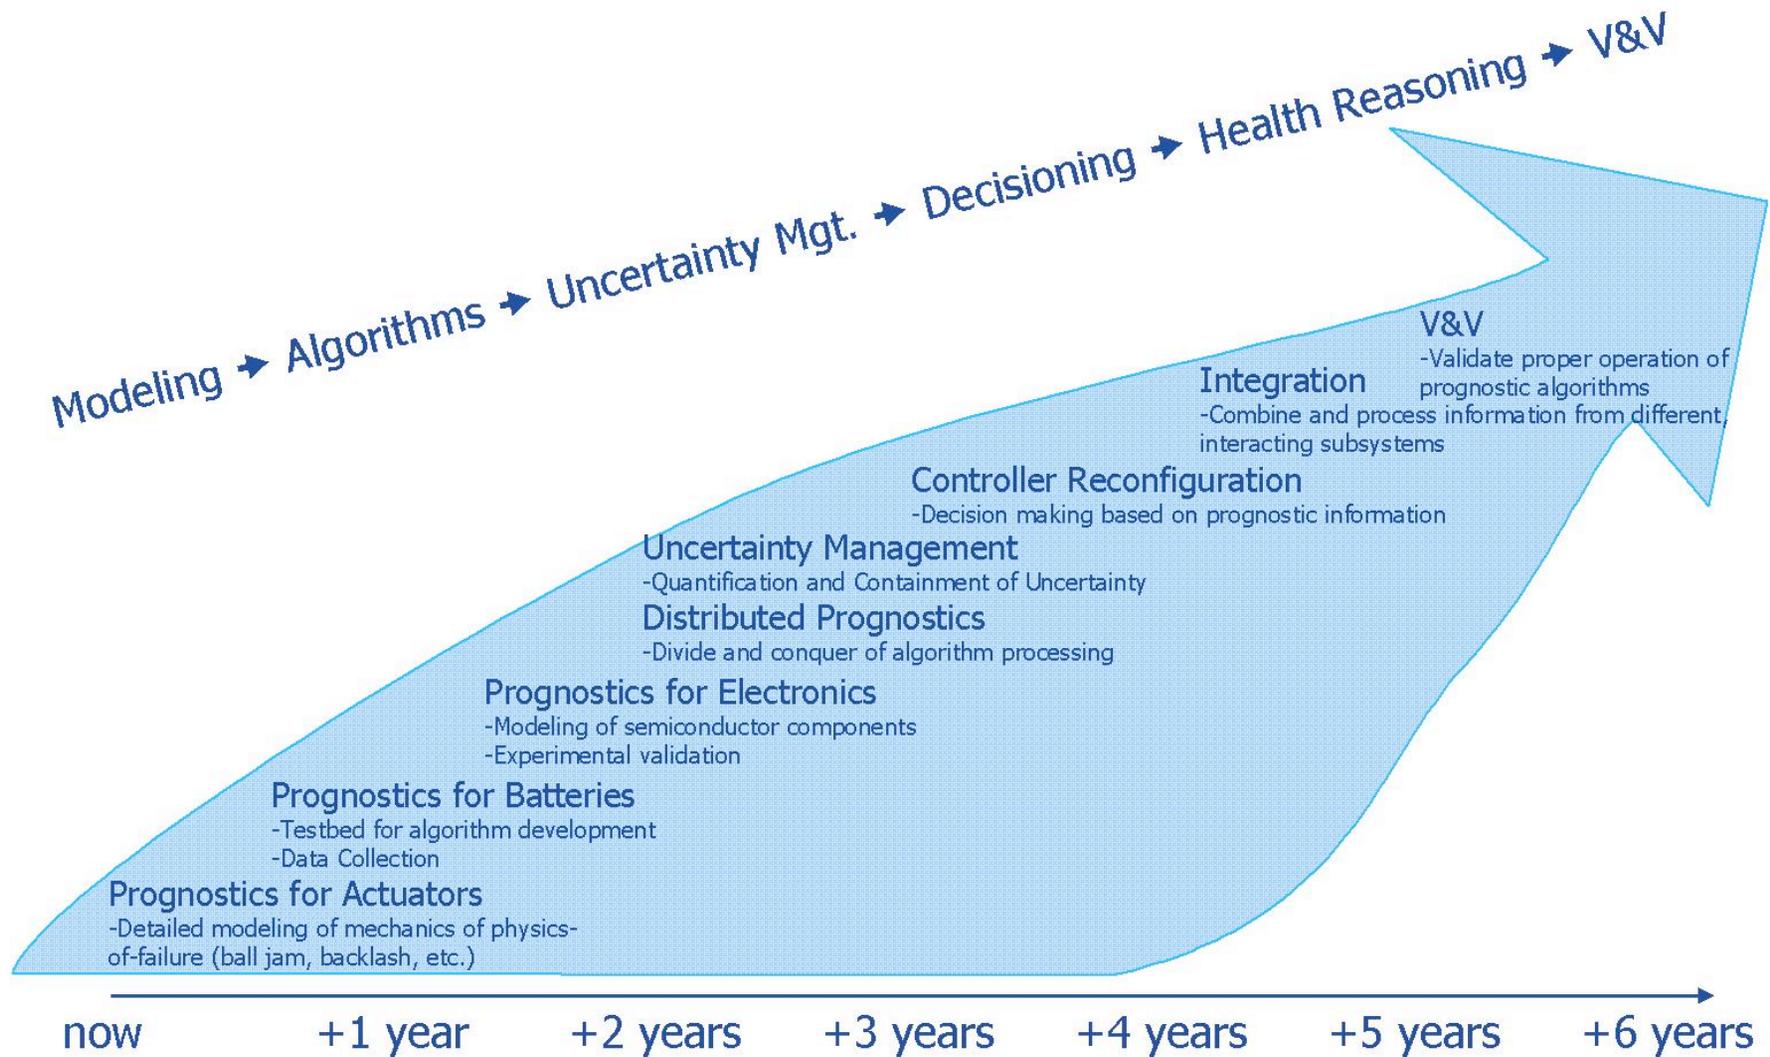
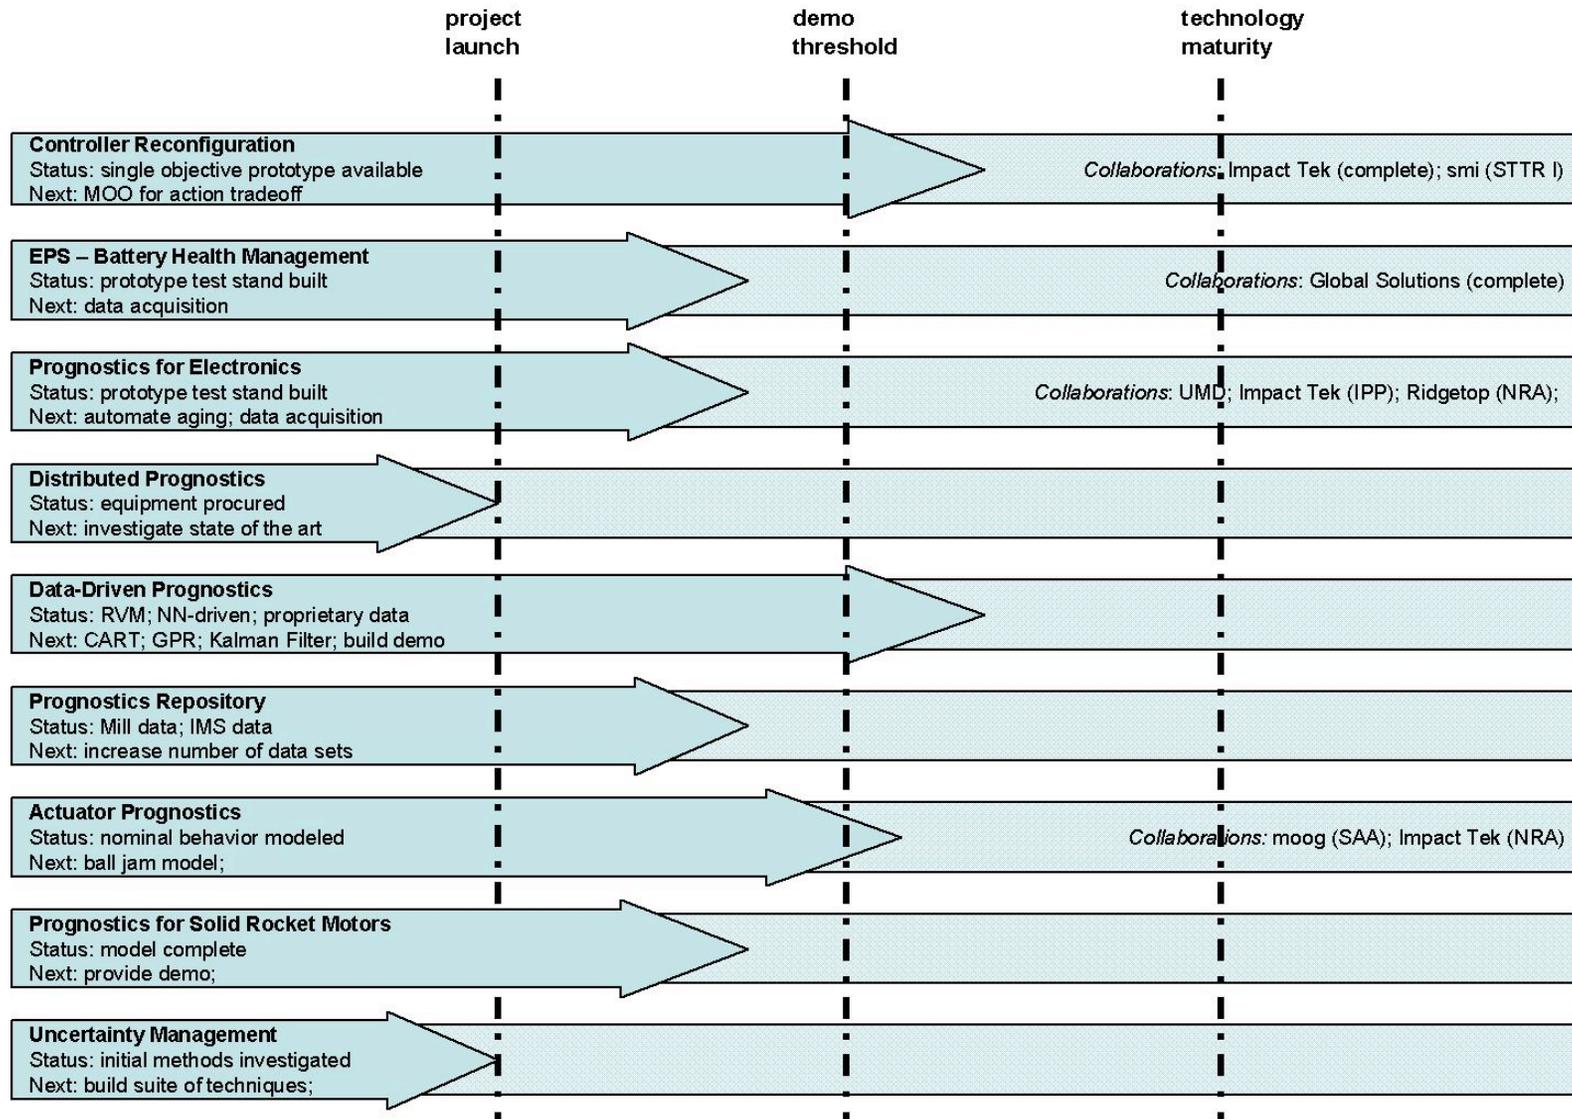


Image courtesy of BNL, Thomas Luu (LLNL) and David Richards (JLab)

# NASA Example



# NASA Example



CMOS Node		130nm	90nm	65nm	45nm	32nm	
Scaling Technology	Lithography	KrF + HT	ArF + HT → ArF + PSM → ArF + Immersion			EUV or ArF+HRI*	
	Translation	Isolation	STI(HDP-CVD)				STI (another gap fill)
		Gate Electrode	poly Si			Metal	
		Insulator	SiON			SiON/High-k → High-k	
		Junction	RTA → CoSi <sub>2</sub> Spike Anneal		Spike Anneal → NiSi Laser Anneal		Raised S/D NiSi /Laser Anneal
	Interconnection	Barrier /Metal	PVD-Ta/Cu			CVD-Ta(N)/Cu	CVD or Direct plating
		Inter-Layer Dielectric	FSG(k=3.6)	SiOC(k=3.0 → 2.6:option)		low-k (k=2.6-2.4)	low-k (k=2.4-2.2)
mobility ↑ SD-resistance ↓		(100) P-channel TR		Strain Si Channel → StrainSi On Insulator(SSOI) → (110)-face-TR		Ge channel? Metal-SD?	
		:Booster Technology		*HRI:High Refractive Index Materials:Lens, Fluid and Resist			

# Technology Roadmap

Category	Applications	2007				2008				2009				2010			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
HV BiCMOS (40V)	LCD gate AC/DC DC/DC controller			0.5um Stack Via													
BiCD (40V)	DC/DC converter WLED driver			1um/1P2M (LTBCDII) BV>40V, Rds=80mohm-mm2@Vgs=5V				In Product verification									
				1.0um/2P2M (LBCD040III) BV>50V, Rds=95mohm-mm2@Vg=5V; w/Schottky				In shuttle verification									
				0.5um/2P2M (LBCD040IV) BV>50V, Rds=60mohm-mm2@Vg=5V; w/Schottky													
BiCD (600V)	AC/DC			1.0um/1P2M (LBCD600I) BV>700V, Rds=0.2ohm-cm2@Vgs=15V; w/Schottky													
				0.5um/2P2M (LBCD600II) BV>700V, Rds=0.2ohm-cm2@Vgs=15V; w/Schottky													
ESD+Filter	ESD protection							ESD+Filter ESD>8KV									

o Focus on the HV market, from 40V to 600V Power Up Your IC



# Fundamentals of Technology Roadmapping

Marie L. Garcia  
Olin H. Bray

Strategic Business Development Department  
Sandia National Laboratories  
P.O. Box 5800  
Albuquerque, NM 87185-1378

E-mail:	<a href="mailto:mgarci@sandia.gov">mgarci@sandia.gov</a>	<a href="mailto:ohbray@sandia.gov">ohbray@sandia.gov</a>
Phone:	(505) 843-4191	(505) 843-4205
FAX:	(505) 843-4223	(505) 843-4223

## Abstract

Technology planning is important for many reasons. Globally, companies are facing many competitive problems. Technology roadmapping, a form of technology planning, can help deal with this increasingly competitive environment. While it has been used by some companies and industries, the focus has always been on the technology roadmap as a product, not on

# Four Goals for IESP



- **Strategy for determining requirements**

- clarity in scope is the issue



- **Comprehensive software roadmap**

- goals, challenges, barriers and options

- **Resource estimate and schedule**

- scale and risk relative to hardware and applications

- **A governance and project coordination model**

- Is the community ready for a project of this scale, complexity and importance?
- Can we be trusted to pull this off?

# Goals for IESP



## □ **Develop a comprehensive community software roadmap for Exascale systems**

- Identify those software capabilities that will be needed for fully functional exascale systems, what are the barriers and how can we overcome them
- Determine which elements will occur naturally and which elements need R+D investment
- Determine those components that have solid starting points and which that need *ab initio* efforts
- Determine which components are suitable for an open community development model

# Goals for IESP

---

- **Develop an estimate of the resources required and timeline needed to develop the required software**

- Need to put the software element of exascale in appropriate budget and schedule context
- Need to understand the risks (technical, schedule and organizational)

- ◆ □ Need to distinguish between the applications software efforts and the systems software

- ◆ □ The software timeline should be aligned with that of the hardware (and precede it where possible)

# Deliverables:



- Note: Slice-n-dice: Need two perspectives
  - ▣ Application needs (science requirements)
  - ▣ System hardware / software capabilities
  - ▣ Key regional interests
  
- For this workshop:
  - ▣ Quad Chart(s) (two to four)
  - ▣ PPT slide with major capabilities
  - ▣ PPT outline of roadmap text
  
- In two weeks:
  - ▣ 2 pages of roadmap text

## **Priority Research Direction (use one slide for each)**

### **Key challenges**

**Brief overview of the barriers and gaps**

### **Summary of research direction**

**What will you do to address the challenges?**

### **Potential impact on software component**

**What capabilities will result?**

**What new methods and components will be developed?**

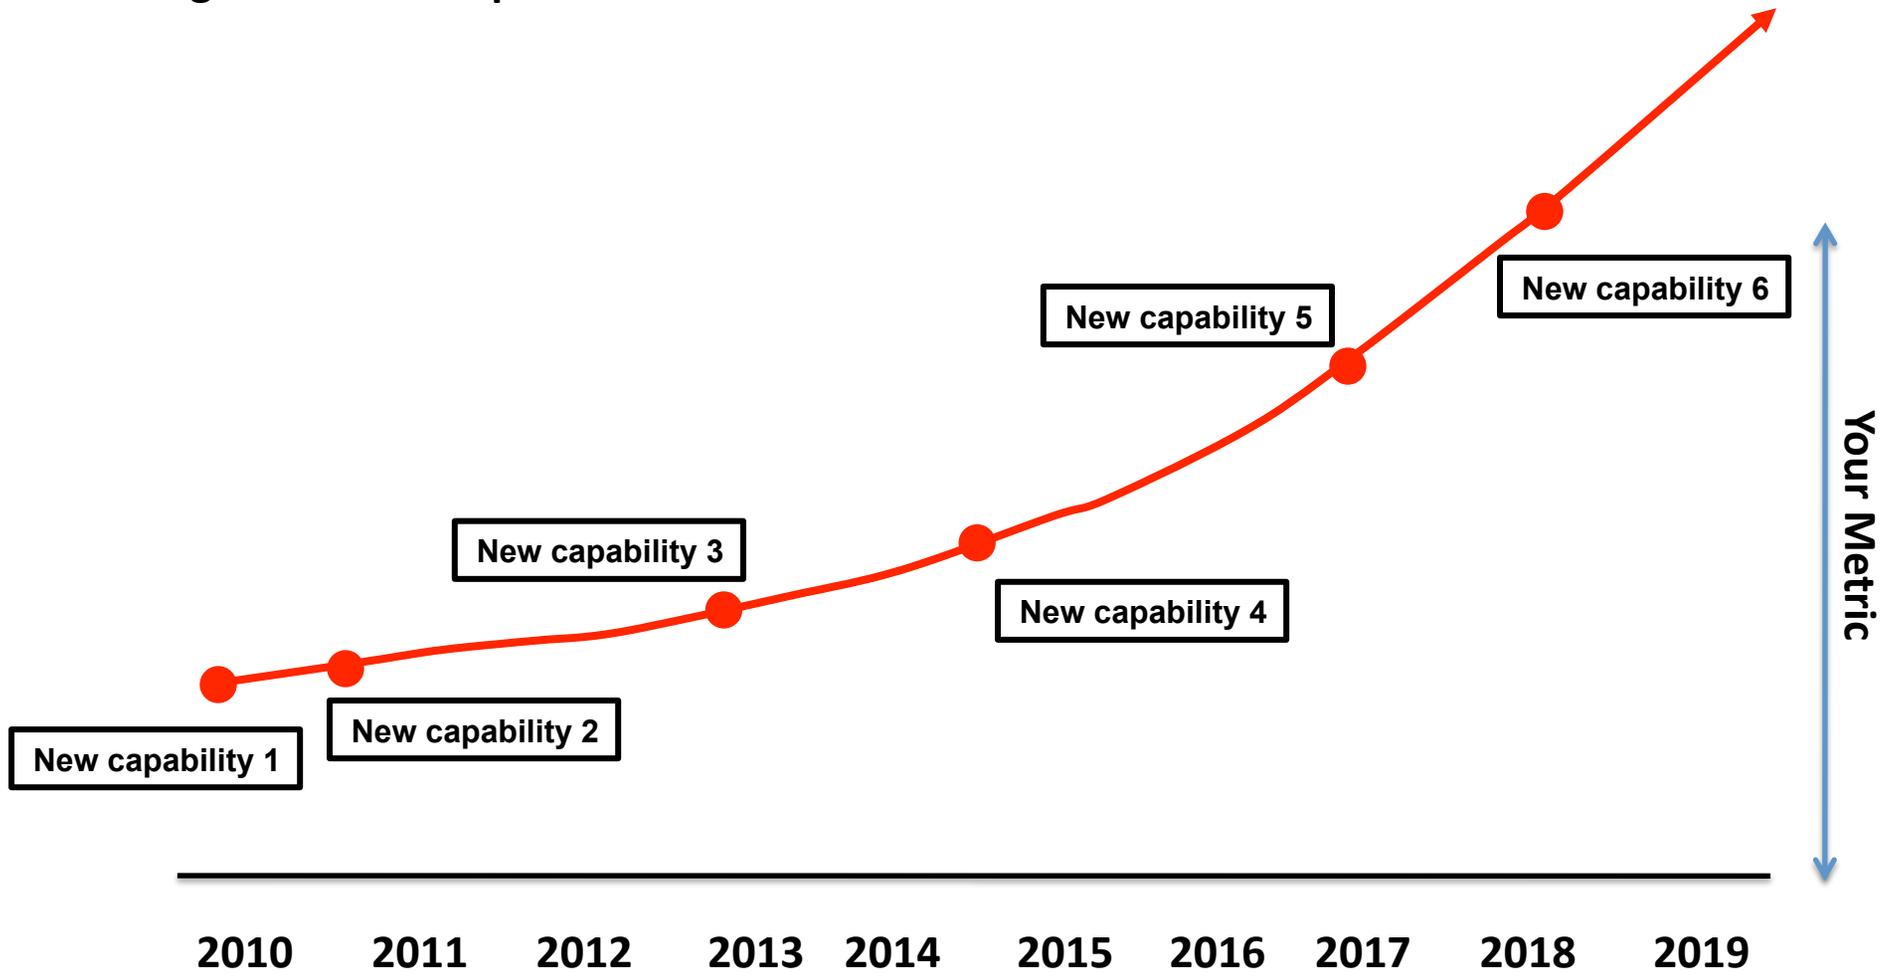
### **Potential impact on usability, capability, and breadth of community**

**How will this impact the range of applications that may benefit from exascale systems?**

**What's the timescale in which that impact may be felt?**

# 4.x <component>

<single short description of the area>



## 4.x <component>

- Technology drivers
- Alternative R&D strategies
- Recommended research agenda
- Cross-cutting considerations