

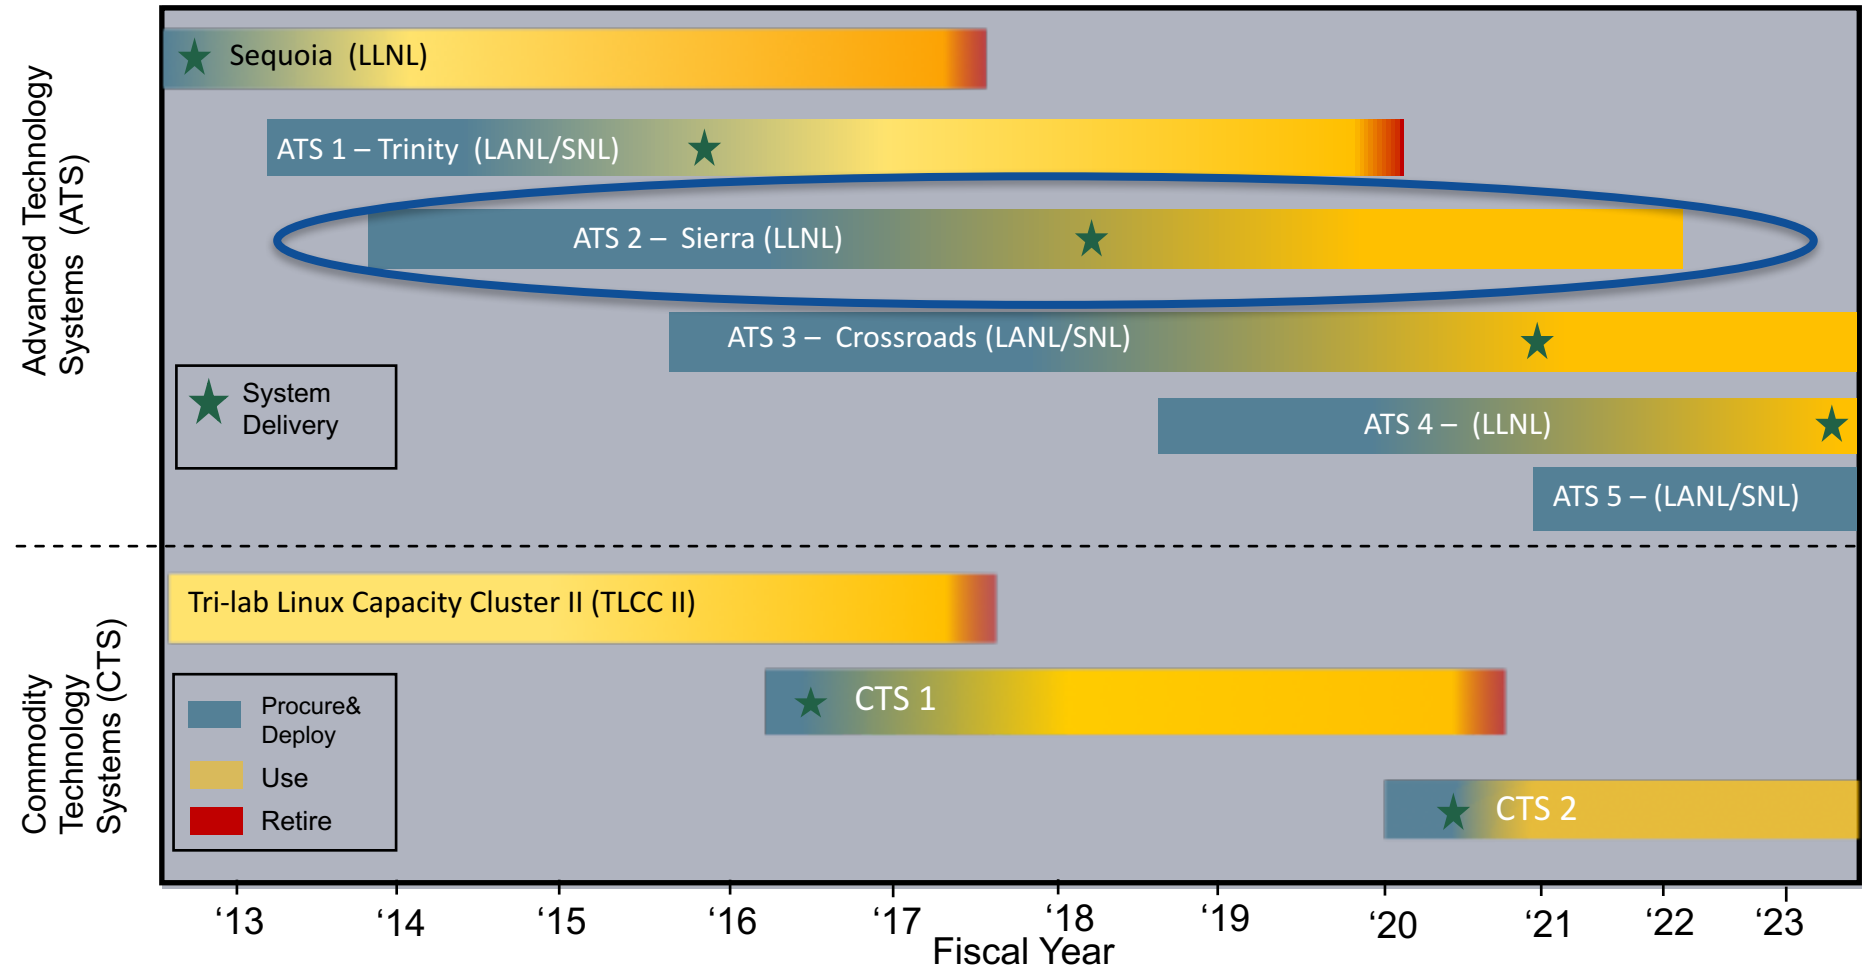
Sierra: The LLNL IBM CORAL System

Bronis R. de Supinski
Chief Technology Officer
Livermore Computing

September 11, 2017



Sierra will be the next ASC ATS platform



Sequoia and Sierra are the current and next-generation Advanced Technology Systems at LLNL

Sierra is part of CORAL, the Collaboration of Oak Ridge, Argonne and Livermore



Modeled on successful LLNL/ANL/IBM Blue Gene partnership (Sequoia/Mira)

LLNL's IBM Blue Gene Systems



BG/L



BG/P Dawn



BG/Q Sequoia

Long-term contractual partnership with 2 vendors

2 awardees for 3 platform acquisition contracts

2 nonrecurring eng. contracts

RFP

NRE contract

ORNL Summit contract (2017 delivery)

LLNL Sierra contract (2017 delivery)

NRE contract

ANL Aurora contract (2018 delivery)

CORAL is the next major phase in the U.S. Department of Energy's scientific computing roadmap and path to exascale computing

The Sierra system that will replace Sequoia features a GPU-accelerated architecture



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


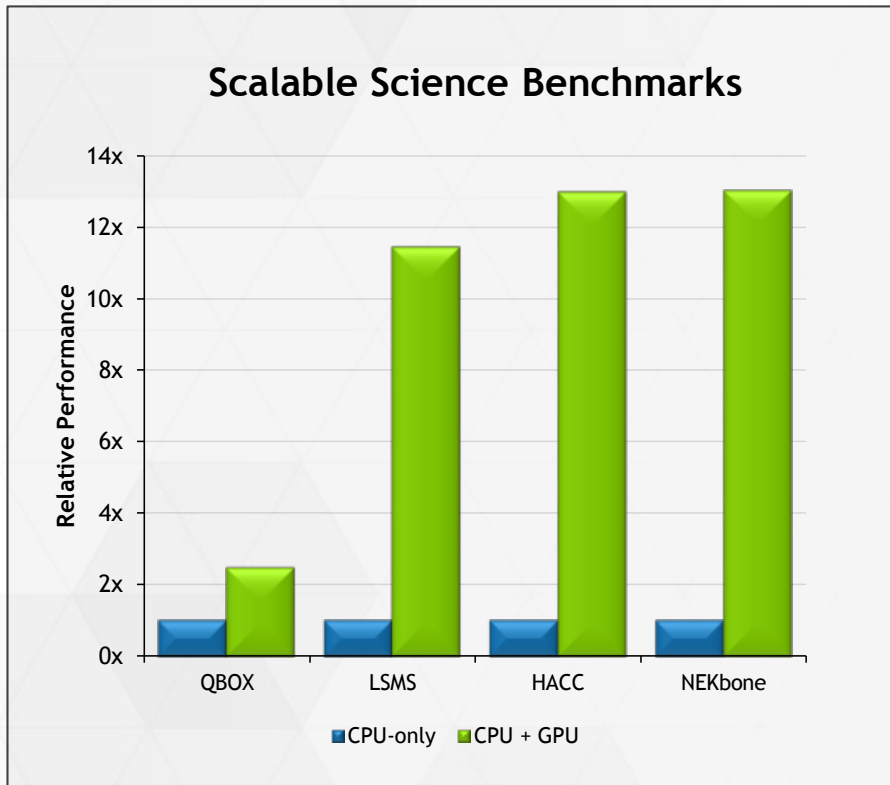
- Single Plane EDR InfiniBand
- 2 to 1 Tapered Fat Tree



GPFS File System

- 154 PB usable storage
- 1.54 TB/s R/W bandwidth

Outstanding benchmark analysis by IBM and NVIDIA demonstrates the system's usability



Projections included code changes that showed tractable annotation-based approach (i.e., OpenMP) will be competitive

Sierra system architecture details have recently been finalized with Go decision



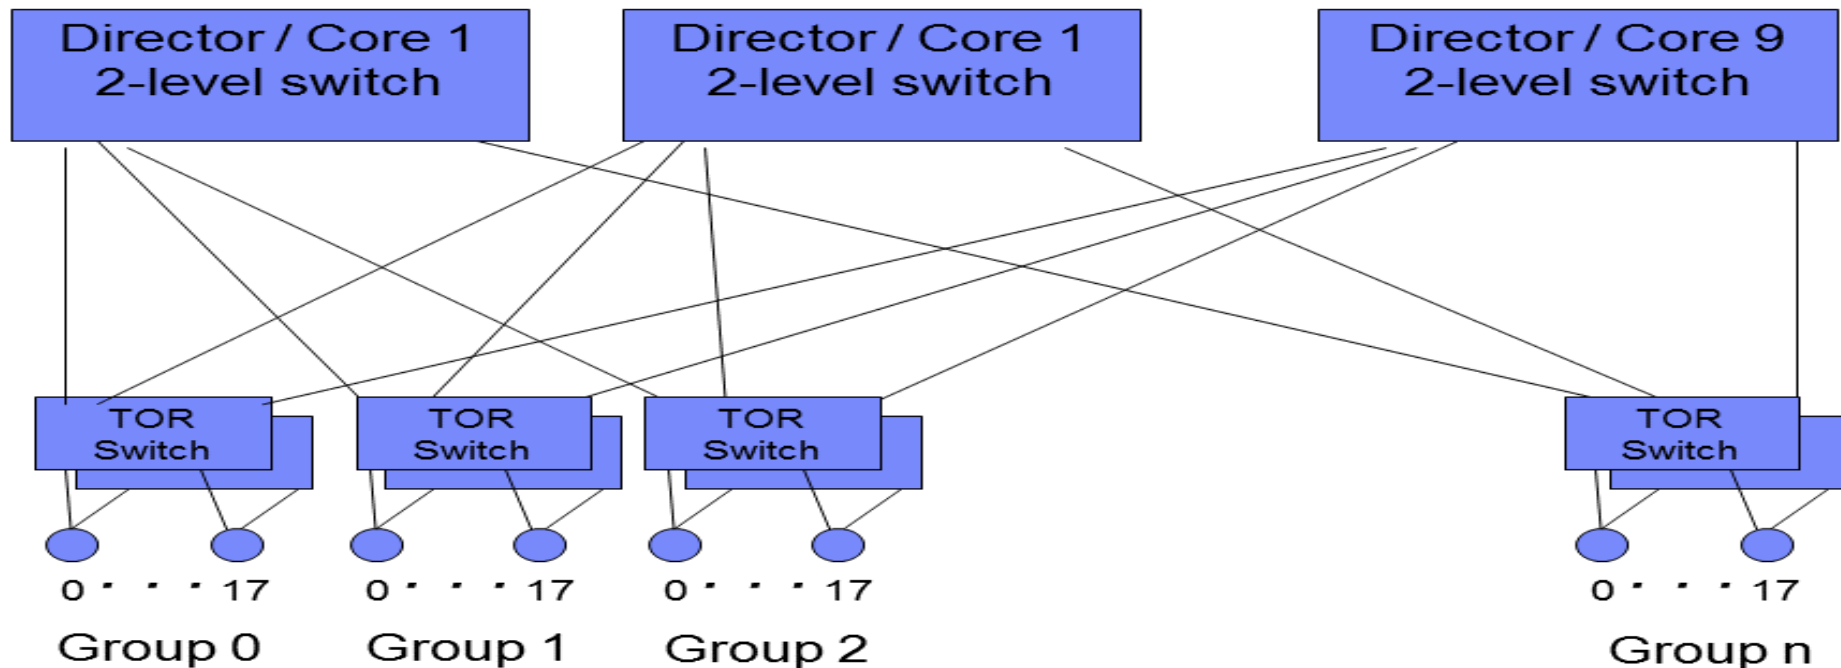
	Sierra	uSierra
Nodes	4,320	684
POWER9 processors per node	2	2
GV100 (Volta) GPUs per node	4	4
Node Peak (TFLOP/s)	29.1	29.1
System Peak (PFLOP/s)	125	19.9
Node Memory (GiB)	320	320
System Memory (PiB)	1.29	0.209
Interconnect	2x IB EDR	2x IB EDR
Off-Node Aggregate b/w (GB/s)	45.5	45.5
Compute racks	240	38
Network and Infrastructure racks	13	4
Storage Racks	24	4

These are working numbers; the final configuration will only be set once the system is fully installed

LLNL and ASC platform chose to use a tapered fat-tree for Sierra's network



- Full bandwidth from dual ported Mellanox EDR HCAs to TOR switches
- Half bandwidth from TOR switches to director switches
- An economic trade-off that provides approximately 5% more nodes



This decision, counter to prevailing wisdom for system design, benefits Sierra's planned UQ workload: 5% more UQ simulations at performance loss of < 1%



Sierra architectural decisions reflect its planned UQ workload

- Sierra is contrasted with ORNL's Summit system
 - Summit will feature 3 Voltas per Power9 (i.e., 6 GPUs per node)
 - Summit will provide a full bandwidth fat-tree
 - Summit will include 2X Sierra's main memory per node
- Sierra's workload will focus on uncertainty quantification
 - Multiphysics ensemble calculations that stress throughput
 - Will fit each physics package into 64 GiB memory (or less)
 - Aggregate memory footprint under reduced
 - Relatively low network demand, placed to minimize contention
- Sierra architectural decisions support this workload
 - Traded network and memory for compute nodes



These tradeoffs improve Sierra's effectiveness by about 5%




Early Access systems provide critical pre-Sierra generation resources



Early Access Compute Node

- “Minsky” HPC Server
- 2xPOWER8+ processors
- 4xNVIDIA GP100; NVLINK
- 256GB SMP (32x8GB DDR4)



Early Access Compute System



- 18 Compute Nodes
- EDR IB switch fabric
- Ethernet mgmt
- Air-cooled

Early Access to CORAL Software Technology

Based on IBM’s current technical programming offerings with enhancements for new function and scalability. Includes:

- Initial messaging stack implementation
- Initial Compute node kernel
- Compilers:
 - o Open Source LLVM
 - o PGI compiler
 - o XL compiler

Initial Early Access GPFS

- GPFS Management servers
- 2 POWER8 Storage controllers
- 1 GL-2; 2x58 6TB SAS drives or
- 1 GL-4; 4x58 6TB SAS drives
- EDR InfiniBand switch fabric

- Enhancement to GL-6, 6x58 6TB SAS drives in progress



Three LLNL Early Access systems support ASC and institutional preparations for Sierra



Shark
597.6TF



Manta
597.6 TF

Compute

Compute

Infrastructure
Storage

Compute

Compute

Infrastructure
Storage

Compute nodes on Ray have 1.6TB NVMe drives →



Compute

Compute


Compute

Infrastructure


Storage

Ray:
896.4 TF

CORAL NRE drove Spectrum Scale advances that will vastly improve Sierra file system



NRE + Hardware
Advancements



Original Plan of Record

- 120PB delivered capacity
- Conservative drive capacity estimates
- 1TB/s write, 1.2TB/s read
- Concern about metadata targets

Modified (Cost Neutral)

- 154PB delivered capacity
- Substantially increased drive capacities
- 1.54 TB/s write/read
- Enhanced Spectrum Scale metadata perf. (many optimizations already released)

Close partnerships lead to ecosystem improvements

NVIDIA Volta GPUs (GV100) provide the bulk of Sierra's compute capability



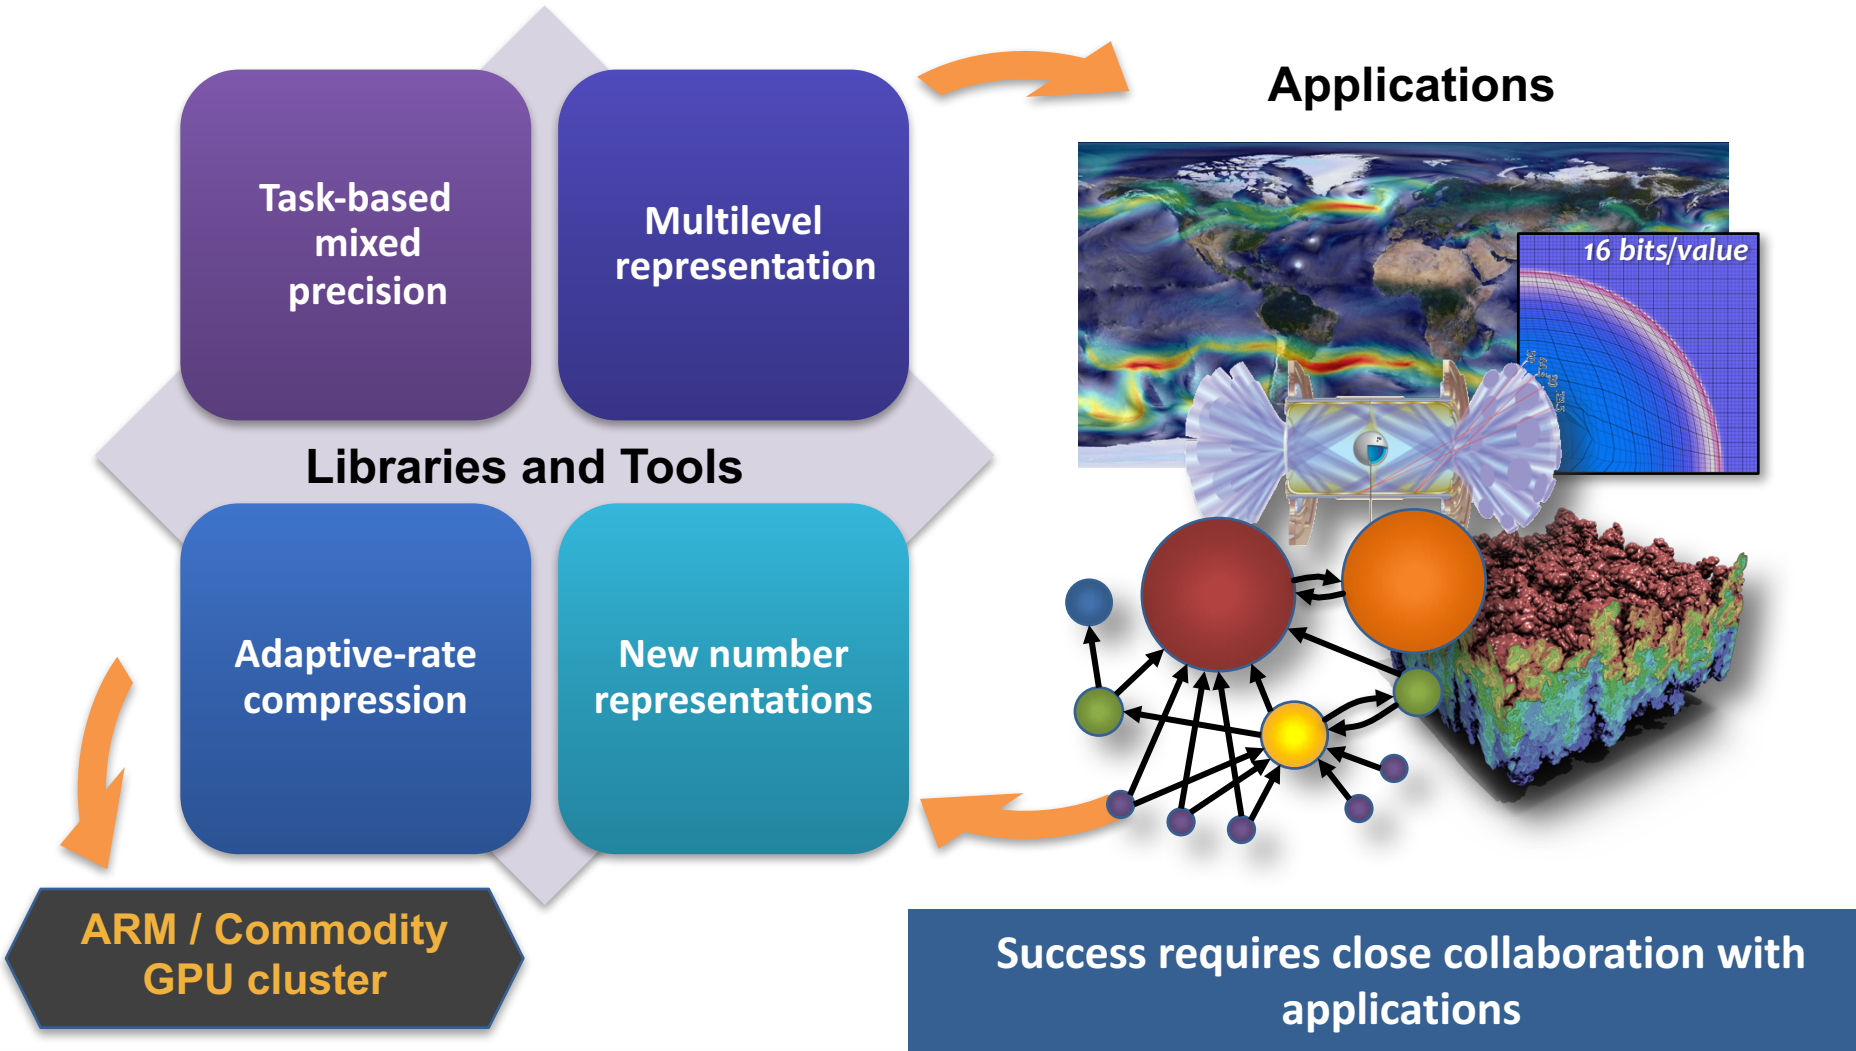
SMs	80
FP64 Units (per SM)	32
FP32 Units (per SM)	64
Tensor Cores (per SM)	8
Register File (per SM)	256KiB
L1/Shared Memory (per SM)	128KiB
Double Precision Peak (TFlop/s)	7 (7.5)
Single Precision Peak (TFlop/s)	14 (15)
Tensor Op Peak (TOp/s)	120
HBM2 Bandwidth (GB/s)	898
NVLINK BW to CPU/Other GPU	75 (60)



To realize Sierra's full potential, we must exploit the tensor operations. The commoditization of machine learning will make this an enduring challenge.



Jeff Hittinger's LDRD team is exploring techniques to exploit capabilities in traditional simulations



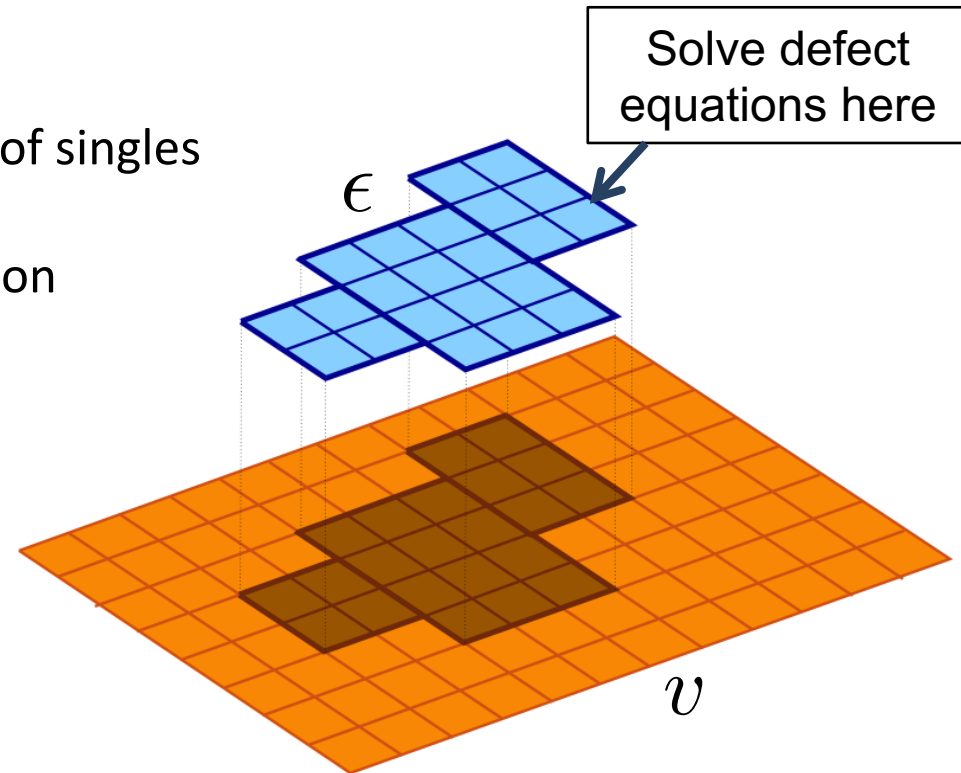
A promising direction is the potential for an AMR-like dynamic, local mixed precision



■ Dynamic Mixed Precision

- Hierarchical representation: sum of singles
- Block-based refinement
- Most calculations in single precision
- Key issues
 - Refinement criteria
 - Propagation of round-off error
 - Cost/benefit

■ Error transport techniques to understand error evolution:



$$\underset{\text{double}}{u} = \underset{\text{single}}{v} + \underset{\text{single}}{\epsilon^{(0)}} + \underset{\text{single}}{\epsilon^{(1)}} + \dots$$

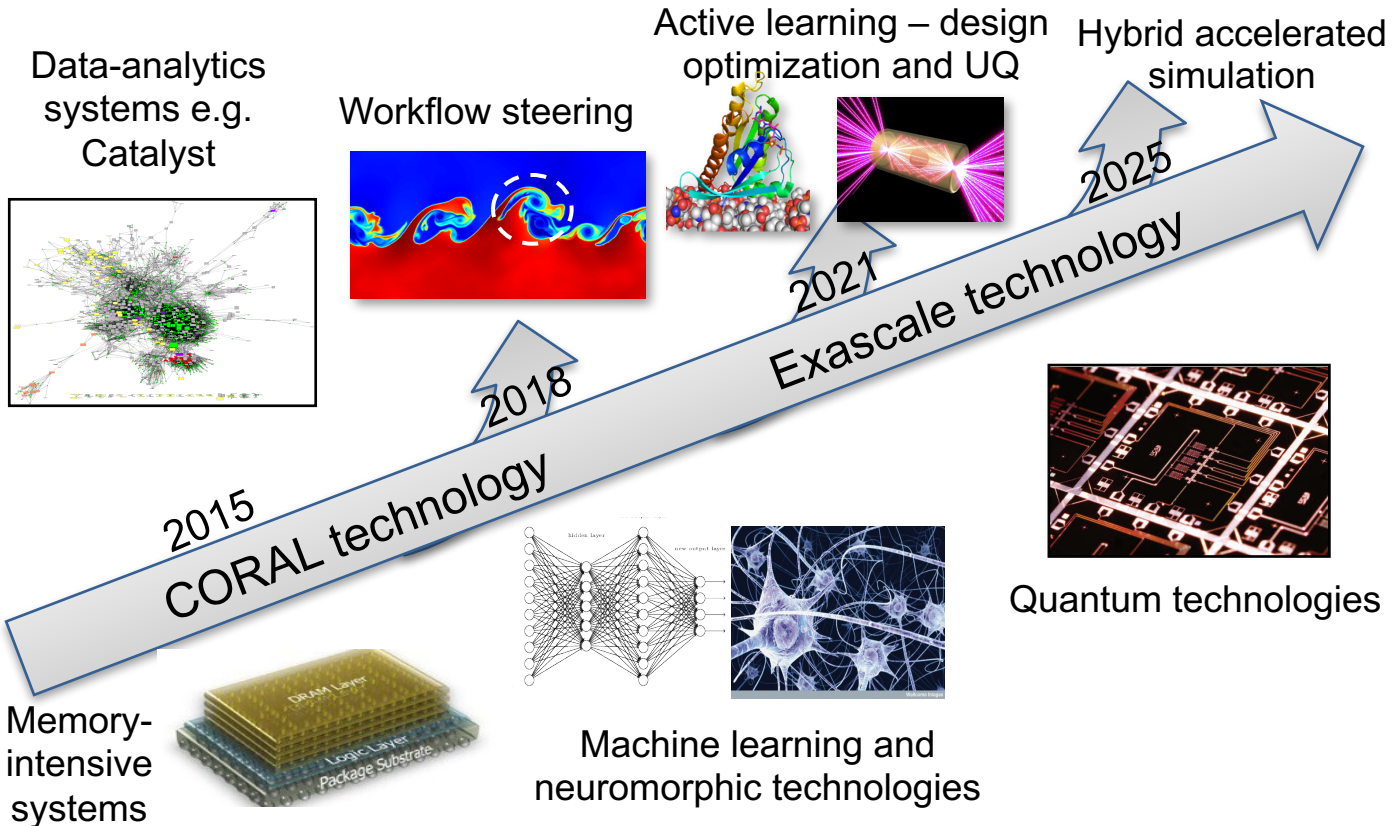
Advanced technology insertion will establish new directions for high-performance computing



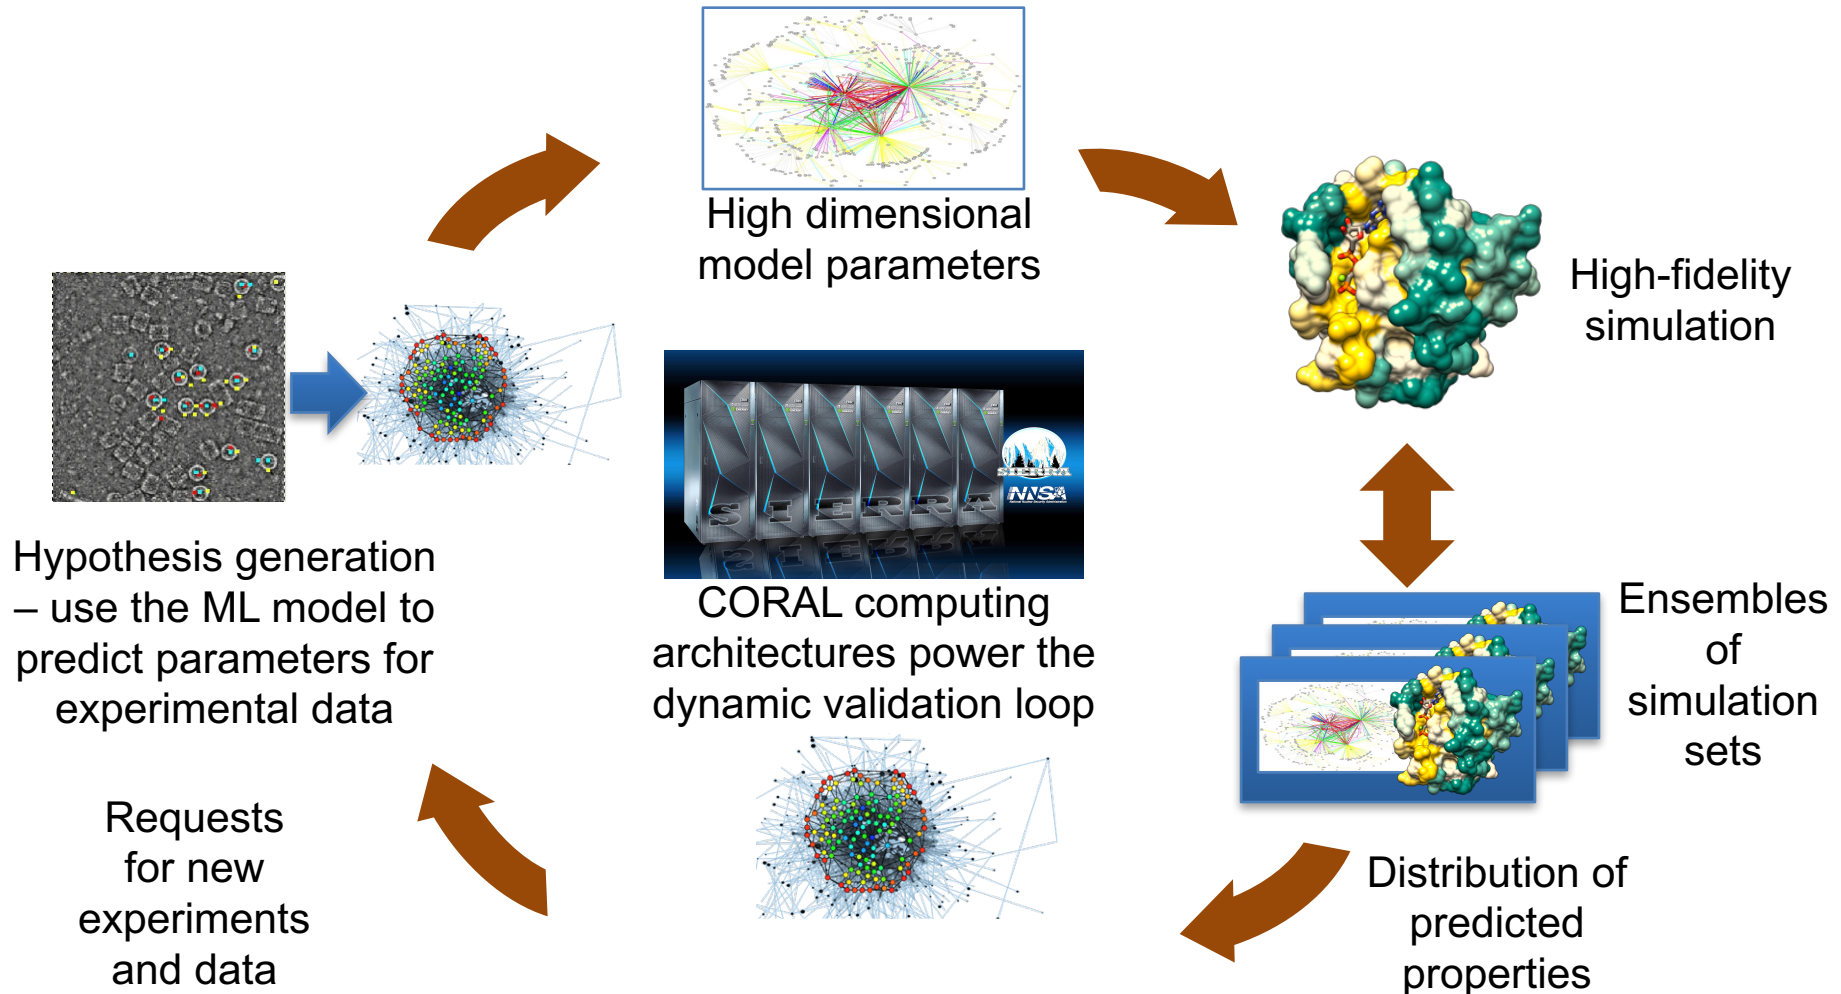
Analytics and simulation convergence



Cognitive simulation systems



Integrated machine learning and simulation could enable dynamic validation of complex models



Sierra and its EA systems are beginning an accelerator-based computing era at LLNL



- The advantages that led us to select Sierra generalize
 - Power efficiency
 - Network advantages of “fat nodes”
 - Balancing capabilities/costs implies complex memory hierarchies
- Planning a similar, unclassified, M&IC resource
 - Same architecture as Sierra
 - Up to 25% of Sierra’s capability
- Exploring possibilities for other GPU-based resources
 - Not necessarily NVIDIA-based
 - May support higher single precision performance

We have multiple projects planned to foster a healthy ecosystem



