

# Tools and Scalable Application Preparation Project

Computation ERC

May 6–8, 2009



**W. Scott Futral**

Development Environment Group Leader

**S&T Principal Directorate - Computation Directorate**

Lawrence Livermore National Laboratory

# The Scalable Application Preparation (SAP) Project supports code teams' preparations for Sequoia



## **Mission:**

Pro-actively assist applications code developers in preparing for Sequoia

## **Goal:**

Code teams have the information and support they need to succeed in using Sequoia

## **Methods:**

Explore, Communicate (with teams and vendor), Evaluate, Document, Train, Support



# The Scalable Application Preparation (SAP) Project promotes activities to ensure success on Sequoia



- Ongoing ANL/IBM/LLNL BlueGene collaboration
  - Team provides expertise in compilers, debuggers, performance tools
  - Providing access to IBM experts, including an on-site IBM applications analyst
  - Leverage ExaCT, ISCR Applications Testbed
- 
- Collaborating closely with application teams
  - Engaging third-party vendors, university research partners, and the open source community
  - Addressing support issues via LC Hotline, User Training, and Documentation



# Simulators are a key component in our plans to prepare for new Sequoia features



- Transactional memory and speculative execution offer potential performance benefits
  - Early access is useful for exploration and education
- IBM's Software Transactional Memory Compiler (stmxl) is now available on LC uP and Purple systems
- Remote access to a hardware device simulator for the Sequoia processor is scheduled for Spring 2010
- Allows execution of user binaries and early access to runtime software (MPI & OpenMP)
- IBM on-site analyst is the simulator POC and will assist code teams with access and usage



# SAP activities are under way and will continue through early years of Sequoia production lifetime



**Alan Gara, IBM Fellow and Chief Architect for Blue Gene, presents Sequoia overview**

- 'From Here to Sequoia, a Code Developer Perspective' presented on 3/23/09
- 90 participants from 6 Labs registered
- Introduction to Sequoia and motivator for code developers
- 8 talks covered system architecture to Dawn account requests



## SAP coming events include training on tools

- TAU Workshop (May 26 & 27 on-site)
  - Focus is on using the TAU tool for performance experiments on multi-core architectures, including Dawn BG/P.
- IBM HPCToolkit Presentation (May 15 & 16 on-site)
  - Presentation and research visit by I-hsin Chung of IBM
- Initial Dawn Science Runs
  - These commenced as soon as acceptance completed
- Code team discussions continue
  - These have been initiated and will benefit from the Sequoia Seminar presentation just completed on 3/23



# LC and Computations stand ready to address key challenges Sequoia poses for applications teams



- BG/L experience informs Dawn/Sequoia scalability
- OpenMP & Posix threads experience on Linux/AIX
- Integrated codes currently run at Purple capability scale
- Dawn early experience for code development
  - SMP parallelism
  - Python
  - Larger memory per core than BG/L
  - ISCR Applications Testbed
  - Production use of UQ codes expected



**Sequoia will be a Tri-Lab ASC resource**

S&T Principal Directorate - Computation Directorate

# Through Computations' tools development work and SAP Project, Livermore is prepared and equipped to 'scale' Sequoia

