

Sequoia Integration

EEHPC Working Group Webinar
September 20, 2012

Anna Maria Bailey Livermore Computing Facility Manager



LLNL-PRES-550311

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Sequoia Overview

- Sequoia Machine
- Sequoia Integration Challenges
 - Facility Integration
 - Innovative Solutions Needed
 - Construction and Commissioning Challenges
 - Planned vs. Reality
 - Grove File System Integration
 - Systems Integration



Ninety-six Sequoia racks powered up
as of April 26, 2012

Sequoia Machine

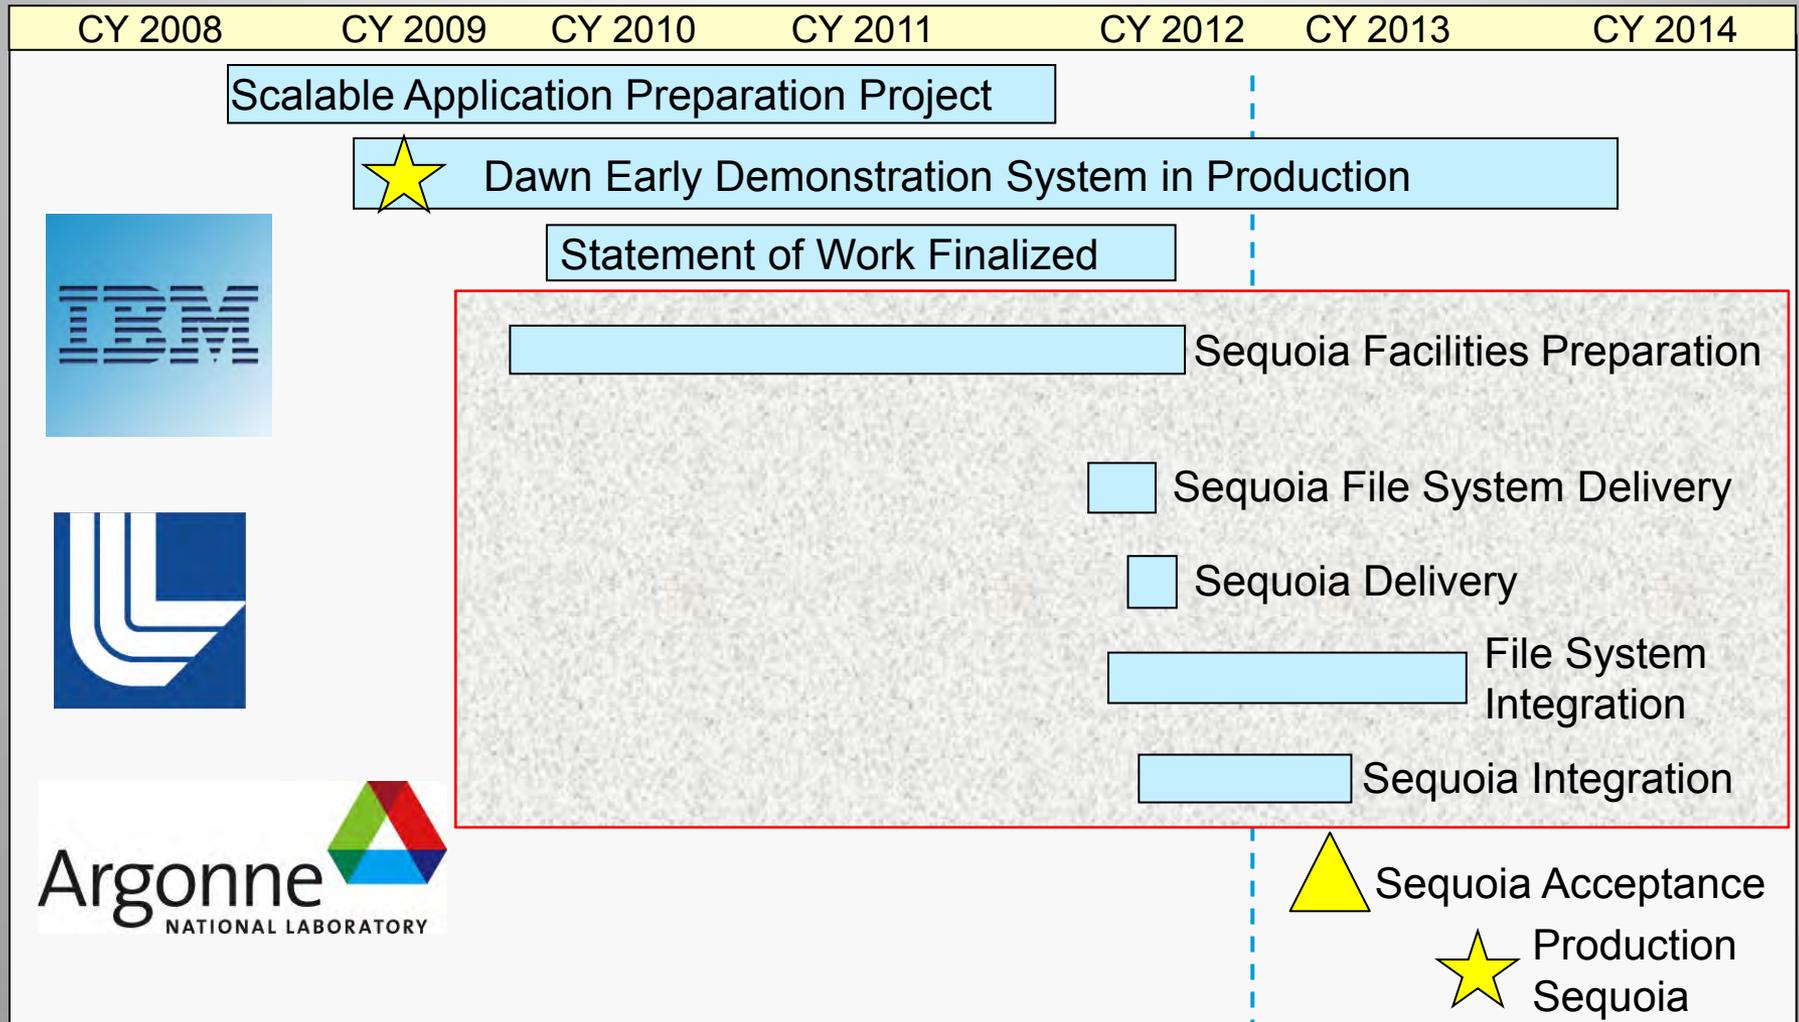
- Sequoia statistics
 - 20 petaFLOP/s target - Achieved 16.32 PF in June, 2012
 - Memory 1.5 PB, 4 PB/s bandwidth
 - 1.5M cores
 - 3 PB/s link bandwidth
 - 60 TB/s bi-section bandwidth
 - 0.5–1.0 TB/s Lustre bandwidth
 - 50 PB disk
- Power – 9.6 MW in 4,000 ft²
- Cooling - 91% liquid cooled and 9% air cooled
- Third generation IBM BlueGene
- Challenges
 - Hardware, software and applications scalability



Sequoia statistics:

20 PF/s target
Memory 1.5 PB,
1.5M cores, 6M threads
50 PB disk, 512GB/s

Sequoia – Long Road to Production with Early Planning



Sequoia Demands Detailed Facilities Integration

- Mechanical requirements:
 - 91% liquid cooled and 9% air cooled
 - New tertiary process loop needed to meet requirements
 - Air-cooling requirements: 1700 CFM/rack= 163,000 CFM total
- Electrical requirements
 - 100 kW/rack= 9.6MW
 - Completed 15MW upgrade in 2011
 - New innovative electrical distribution required to minimize under-floor congestion
- Physical requirements
 - Mechanical distribution, power distribution, cable trays, fire protection
 - Space: 96 racks in 4,000 ft²
 - Weight: 4,500 lbs/racks=210 tons total
 - Approximately equivalent to 30 adult elephants



Sequoia - Liquid Cooling Facility Innovation

- 91% Liquid Cooled and 9% Air Cooled
 - Liquid cooling inlet requirements - 64F to 74F
 - New tertiary loop needed required
 - Water pressure = 25 psi
 - Water temperature rise = 20F/rack
 - GPM/rack = 25 to 30 gpm
 - Stainless steel or copper specified
 - Polypropylene sustainable piping
 - Total project savings - \$2M
 - Maintain B-453 LEED Gold status
 - Provides efficient flow
 - Reduced heat gain and loss
 - Minimized environmental chemical impacts – Ensure ISO 14001 compliance



Sequoia - Electrical Distribution Innovation

- Sequoia electrical – 100 kW/rack = 9.6MW
 - (4) 480V, 3 Phase, 60A line cords/rack
 - Custom innovative electrical distribution required to minimize under floor congestion and reduce electrical losses
 - Reduced electrical distribution by 75%
 - Total project savings of \$1M
- Sequoia file system (Grove) requires custom power solutions
 - File system racks demand multiple power connections for operation
 - Places demands on electrical distribution and increases system losses
 - Innovative solution to reduce distribution by 66%



Sequoia – Planned vs. Reality

Facility Construction and Commissioning Challenges

- Future rack floor supports are scheduled to ship in two days
 - IBM notifies LLNL that the weight of the racks has increased by 25%—the stands have to be redesigned
 - The rack floor stands need to be installed first prior to electrical and mechanical systems
- Contractor rearranges the schedule and works out of sequence until design modifications are implemented



Sequoia – Planned vs. Reality

Facility Construction and Commissioning Challenges

- Sequoia water piping system complete at end of November
 - Flushing and filling to begin in December.
- LLNL receives notification that the primary Hetch Hetchy water supply will not be available beginning December 4th due to system construction for minimum of 4 weeks.
- Secondary system has unacceptable water quality properties.
- LLNL has a source of demineralized water used heavily in laboratory facilities
 - This facility has only one 1” line available
 - Filling and flushing process was slow.



Sequoia – Planned vs. Reality

Facility Construction and Commissioning Challenges

- Filter socks trap sediment as process closed-loop water flows through pumps
- Initial socks replaced hourly are clean and pass
- Final socks running overnight fail as particles are encountered



Sequoia – Planned vs. Reality

Facility Construction and Commissioning Challenges

- Source of particles is unclear - red small shavings appeared
 - Anticipated small blue/green shavings from pipe
- Question: “What is the source of the red flakes? Answer: “The paint on the pumps”
 - Impeller was over-sprayed at the factory
- Isolate, drain, inspect, scrub and clean pumps and then refill and back to the socks.

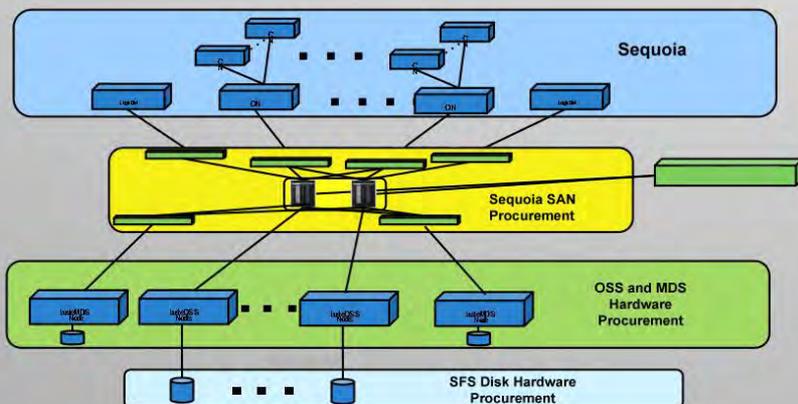


Sequoia File System – Grove is Massive Risk Mitigation – Proven Hardware

- Requirements
 - 50PB file system
 - 500GB/s minimum, 1TB/s stretch goal
 - QDR InfiniBand SAN connection to Sequoia—768 IB links into core
 - Must integrate with existing Ethernet infrastructure

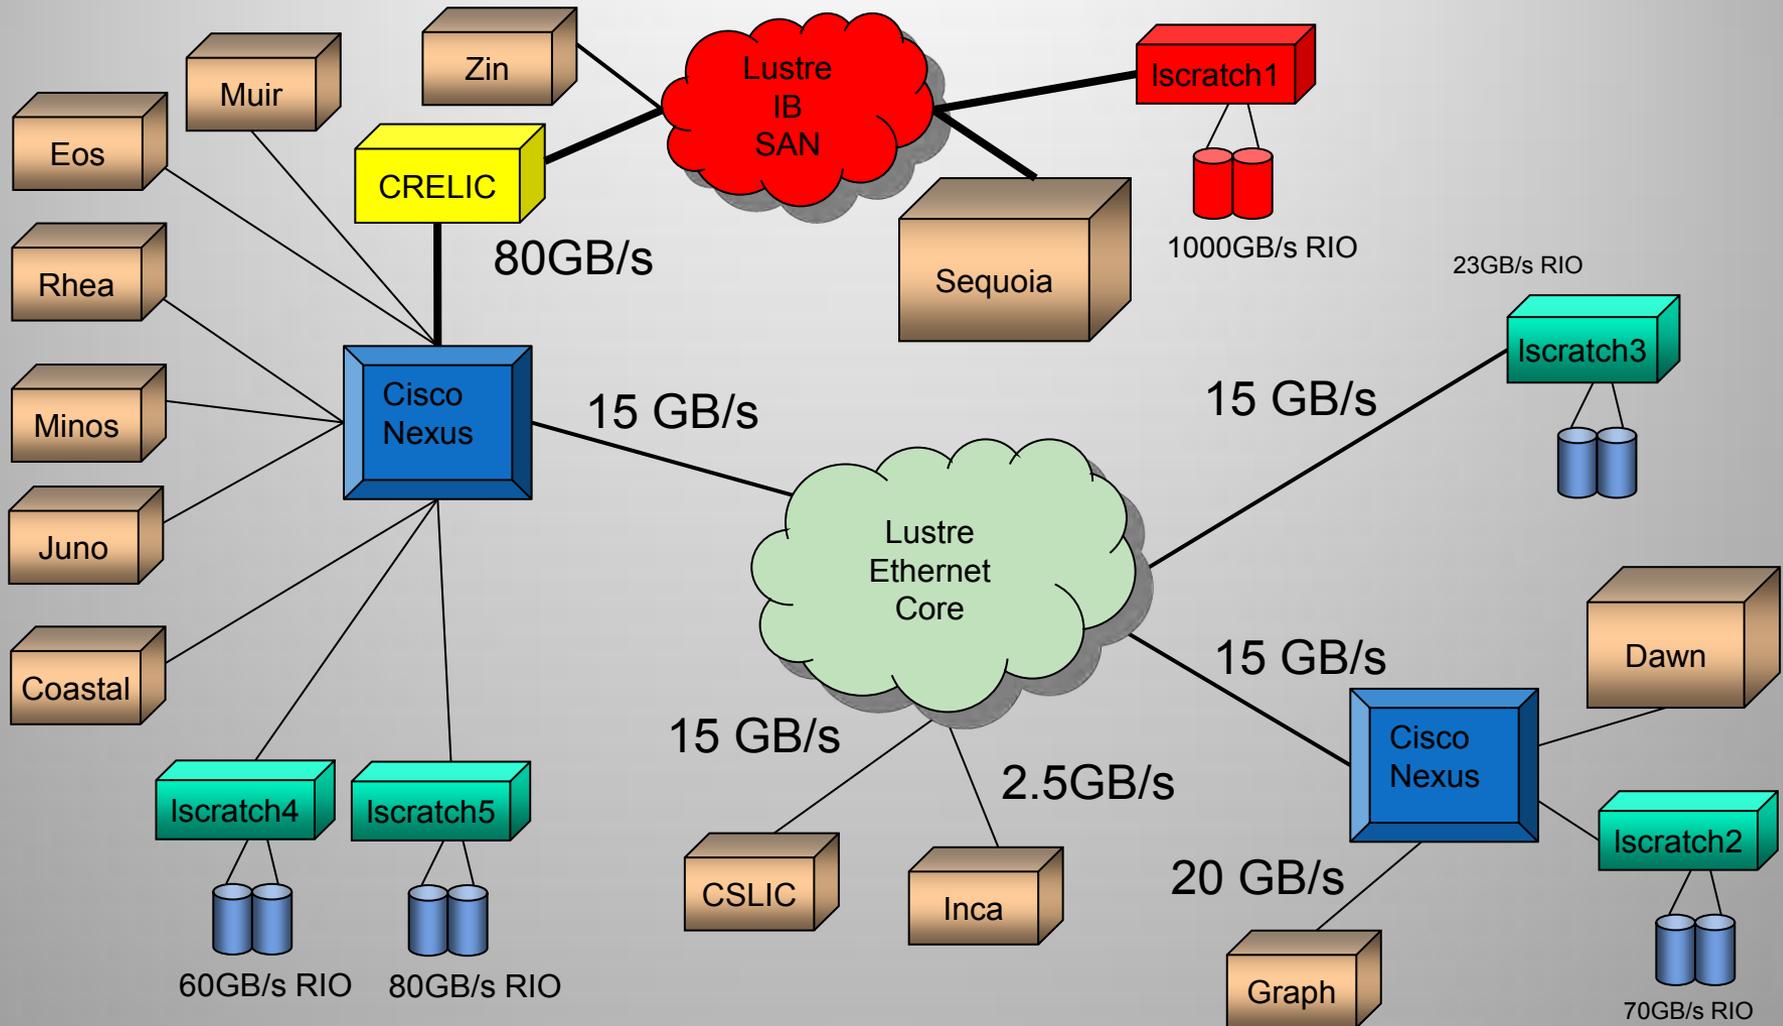


- 384 Netapp E5400s
 - 60 3TB SAS drives per bay, 4U
 - Dual RAID controllers



- 768 Appro GreenBlade OSS Nodes
 - Intel SandyBridge, TLCC2
 - Dual socket, 8 core @2.6GHz

Sequoia File System – Grove is Bridged to Existing Ethernet Infrastructure



Systems Integration – Sequoia, Grove and Facility Co-Located Construction and Activation Activities



Sequoia Planning, Innovation and Hard Work are the Key Elements of the Sequoia Journey

- It has been a long road and we are still integrating all of the systems
- Sequoia's success will be measured by the simulation and science achievements it enables

Questions

Anna Maria Bailey, PE
Lawrence Livermore National Laboratory
7000 East Ave PO Box 808 L-554
Livermore, CA 94550

Phone - (925)423-1288
Email - bailey31@llnl.gov

