

HPC Energy Efficiency: Why It Matters to LLNL

3rd Annual Workshop on Energy Efficient
High Performance Computing @ SC12

November 11, 2012

Robin Goldstone
HPC Advanced Technologies Office

 Lawrence Livermore
National Laboratory

LLNL-PRES-600793

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



Predictive simulation at LLNL is focusing on sustainable nuclear security, global and energy security and American competitiveness

- **Nuclear mission is to support deterrence**

- Stockpile aging away from test base

- **Broader national security role**


- Nuclear counterterrorism
- Nonproliferation
- Network and cyber security
- Space situational awareness
- Energy and Climate

- **American competitiveness**

- Critical for the U.S. - can't compete on wages
- Collapse the design and prototyping process
- Enable rapid investigation of many possibilities
- Enables U.S. industry to be at the frontier of discovery



Scientific discovery through predictive simulation: Wind tunnel tests at NASA Ames validate simulation predictions...



LLNL HPC Facilities Overview

- 28.3 PF of peak computational capability across multiple compute enclaves
- 50MW available power across ten data centers and ~100,000 ft²
- LEED Gold (B453) and LEED Silver (B451) certifications

System	Program	Manufacture/ Model	OS	Inter- connect	Nodes	Cores	Memory (GB)	Peak TFLOP/s
Unclassified Network (OCF)								26,238
Sierra	M&IC	Dell	Linux	IB QDR	1,944	23,328	46,656	261.3
Cab (TLCC2)	Appro	Appro	Linux	IB QDR	1,296	20,736	41,472	431.3
Hera (TLCC)	ASC+M&IC	Appro	Linux	IB DDR	864	13,824	27,648	127.2
Hyperion	ASC	Dell	Linux	IB QDR	1,152	9,216	11,520	90.3
Ansel	M&IC	Dell	Linux	IB QDR	324	3,888	7,776	43.5
RZZeus	ASC	Appro	Linux	IB DDR	288	2,304	6,912	22.1
Aztec	M&IC	Dell	Linux	N/A	96	1,152	4,608	12.9
Edge	M&IC	Appro	Linux	IB QDR	216	2,592	20,736	29.0
Sequoia	ASC	IBM BG/Q	Linux	5D Torus	98,304	1,572,864	1,572,864	20132.7
Vulcan	HPC-IC	IBM BG/Q	Linux	5D Torus	24,576	393,216	393,216	5033.2
Herd	M&IC	Appro	Linux	IB DDR	6	96	1,088	0.9
RZMerl	ASC	Appro	Linux	IB QDR	162	16	5,184	53.9
OCF Totals	12				129,228	2,043,232	2,139,680	26,238
Classified Network (SCF)								2,061
Dawn	ASC	IBM BG/P	Linux	IBM	36,864	147,456	147,456	501.4
Muir	ASC	Dell	Linux	IB QDR	1,296	15,552	31,104	174.2
Juno (TLCC)	ASC	Appro	Linux	IB DDR	1,152	18,432	36,864	162.2
Graph	ASC	Appro	Linux	IB DDR	576	13,824	72,960	110.6
Coastal	Other	Dell	Linux	IB DDR	1,152	9,216	27,648	88.5
Eos (TLCC)	ASC	Appro	Linux	IB DDR	288	4,608	9,216	40.6
Zin (TLCC2)	ASC	Appro	Linux	IB QDR	2916	46656	93,312	970
Inca	ASC	Dell	Linux	N/A	96	1,152	4,928	12.9
SCF Totals	8				44,340	256,896	423,488	2,061
Combined Total	20				173,568	2,300,128	2,563,168	28,299



Sequoia Highlights

- Specifications
 - 20 petaFLOP/s target - 16.32 PF achieved
 - Memory 1.5 PB, 4 PB/s bandwidth
 - 98K nodes, 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
- IBM Innovation: 3rd generation BlueGene/Q
- LLNL Innovation: liquid cooling and electrical distribution
- #1 on June 2012 Top500, Graph500 *and* Green500



Sequoia challenges:
Facilities
Hardware
Software
Applications scalability



Operational efficiency requires many considerations and ongoing planning

- Facilities personnel must be 1st class members of your HPC systems procurement and operations team – facilities cannot be an afterthought
 - The expertise of the LLNL facilities team was critical to the success of Sequoia
 - Facilities innovation can compliment platform innovations: a win-win situation!
- Not just FLOPs/watt but also FLOPs/ft²
 - LLNL used Purple, BG/L and Sequoia RFPs to drive impressive energy efficiency improvements

Platform	MW	SF	TF	TF/MW	SF/TF
Purple	4.80	8000	100	21	80.0
BG/L	2.50	4000	600	240	6.7
Dawn	2.00	2000	500	250	4.0
Sequoia	9.60	4000	20000	2083	0.2
K	12.70	18000	11280	888	1.6

- Focus on reduction in losses (electrical, mechanical)
 - It's not just about what you use – but also what you lose!*
- PUE: lower is better but need to evaluate ROI
 - E.g. B453 free cooling project can improve PUE from 1.27 to 1.15 for \$5M. Is it worth it?
- Future planning must consider site utility distribution and even upstream providers

When we run Linpack on our big systems, WAPA notices!

LLNL 25-year HPC Sustainability Master Plan

- Goal: Identify and implement efficiencies in order to:
“Turn Megawatts into PetaFLOPs” -> ExaFLOPs, YottaFLOPs ...
- Initial efforts identified “low hanging fruit” with big payoff
 - Raise air supply temperature
 - Raise chilled water supply
 - Address air leakage issues
 - Rebalance the computer floor
- Sequoia liquid cooling and electrical distribution innovations yielded \$3M+ project cost savings
- Facility gap analysis and strategy - Evaluate structural, electrical and mechanical capabilities then decide:
Repurpose, Refurbish or Replace
- Future innovation: HPC scalable modular building concept





\$3.5M
Savings!

Sustainable HPC modular solutions: Scalable building concept (patent pending)




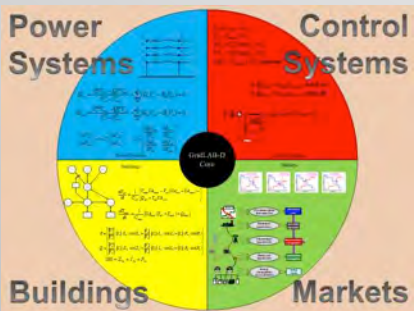
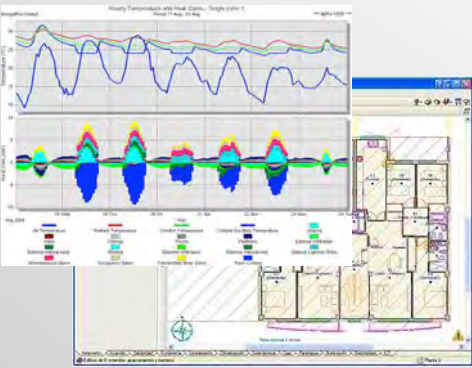
- Scalable building concept to meet changing HPC demands
- Design supports scalable HPC infrastructure
 - Scale power
 - Scale square footage and structure
 - Scale cooling solutions
 - Deploy facility power management system
- Minimize the use of cooling towers and chillers
- Utilize closed-loop cooling concepts with expanded liquid cooling temperature ranges
- Concept is more cost effective than containerized solutions based on detailed cost analysis

Sustainable HPC modular solutions: Design features



- Leverage TSF complex successes
 - Clear, unencumbered space
 - Scalable mechanical and electrical infrastructure
- LEED Gold Certification Goal
 - Advance liquid cooling options
 - Deploy free cooling
 - Cool roof
 - Scalable electrical distribution
 - Accommodate increased weights
 - Decrease space impediments
 - Target PUE: 1.17
- Initial build out
 - 3 to 6MW in 6,000 SF
- Total build out capability
 - 15MW in 24,000 SF

HPC opportunities for sustainable facility and utility planning



- Exploring projects in modeling/analytics to improve facility efficiency and reliability
 - Currently working with HPC Data Analytics and Computational Mathematics Groups
- EnergyPlus: model building energy heating and cooling to support LEED Gold certification
- GridLab-D: electricity power flow modeling to improve reliability, load shedding, etc.
- Overture: Understand fluid and thermal flows to improve building energy efficiency

How is EE HPC WG helping?

- WG brings together experts in both infrastructure *and* HPC computing systems
- Leadership role in providing analysis and recommendations in topics of importance to HPC community
 - Warm water cooling recommendations
 - Power measurement methodology
 - RFP guidelines
 - PUE -> TUE
 - Data Center Energy Management Dashboard



EE HPC WG is well aligned with LLNL strategy and priorities for energy efficiency.



Agenda

Start	Stop	Presenter	Topic
09:00	09:20	R. Goldstone	Opening remarks HPC Energy Efficiency: Why It Matters to LLNL
09:20	10:10	B. Tschudi and D. Martinez	Foundations for an Energy Efficient Data Center
10:10	10:30		Break
10:30	11:00	M. Patterson	Metrics Overview and Update
11:00	12:00	P. Kogge	Architectural Trends and Projections for Energy Efficiency
12:00	01:00		Lunch
01:00	01:45	J. Shalf	Analysis of Application Requirements and Impact on Energy Use
01:45	02:30	S. Matsuoka	Case Study: GPUs and TSUBAME
02:30	03:00	M. Ellsworth	New ASHRAE Thermal Guidelines for Air and Water Cooling
03:00	03:15		Break
03:15	04:15	H. Huber and G.I. Meijer	Case Study: LRZ Liquid Cooling, Energy Management, Contract Specialities
04:15	05:15	S. Hammond and N. Dube	Case Study: NREL and the Net-Zero Data Center
05:15	05:30	J. Rogers	Closing Remarks