SC14 Energy Efficiency High Performance Computing Working Group (EEHPCWG) 5th Annual Workshop

The Drive for Energy Efficiency in HPC

November 17, 2014

Dona Crawford, Associate Director Computation

LLNL-PRES-664281
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

EEHPCWG's Vision for Energy Efficiency in HPC

- Encourage the HPC community to proactively develop energy efficient solutions
- Reduce operating costs through increased energy efficiency
- Serve as a forum for sharing of information (peer-to-peer exchange)
- Drive action through standards and improved processes

EEHPCWG is well aligned with LLNL's goal to improve efficiencies. LLNL participates in a multitiude of areas across the working group.

LLNL Strives to Improve System Usage Effectiveness

- Go beyond Power Usage Effectiveness (PUE) of the facility
  - Seek to achieve Total Power Usage Effectiveness through new metrics, iTUE and TUE
  - LLNL participated in Sequoia TUE case study
  - Results indicate that iTUE measuring capabilities need to be specified in future HPC procurements

ITUE (Sequoia) = 1.24
PUE (Facility) = 1.27
TUE (Sequoia) = ITUE x PUE

TUE = 1.24 × 1.27 = 1.57

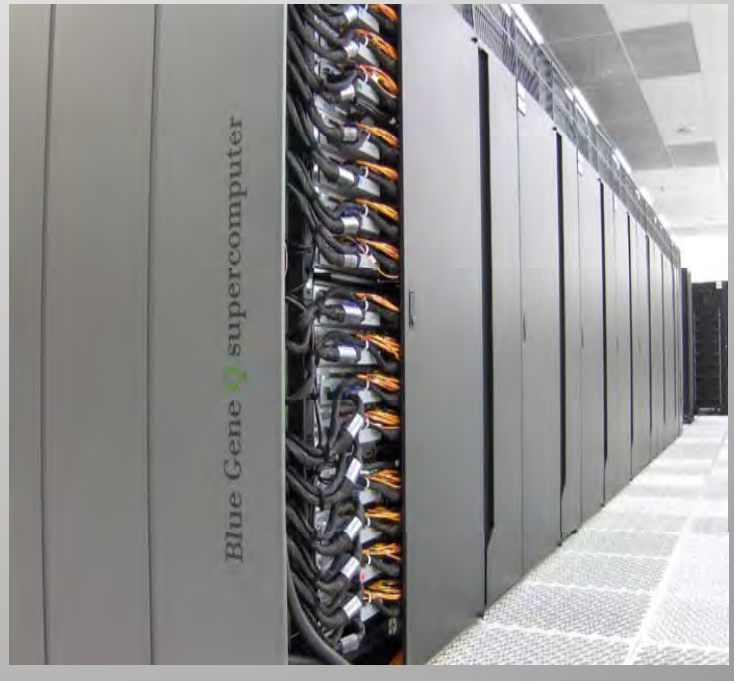

LLNL's Results with L2 vs. L3 Power Methodology

- While attempting to perform L3 power methodology measurements for Linpack, it was discovered that 10% of the meters were misconfigured (*) and not reporting adequately for Sequoia
- All L3 requirements otherwise met but the measurement was downgraded to L2

Aspect	L3 Requirement	Sequoia
Aspect 1a: granularity of power measurements	Continuously integrated total energy	Met*
Aspect 1b: timespan of power measurements	A time series of equally spaced integrated total energy values	Met*
Aspect 1c: reported measurements	Core phase avg power, >10 energy measurements within core phase, whole app avg power, idle power	Met
Aspect 2: machine fraction	Whole machine	Met
Aspect 3: subsystems included	All participating subsystems	Met
Aspect 4: power measurement point	Upstream of power conversion	Met

LLNL and EEHPCWG Demand Response Team

- EEHPCWG's Demand Response Team was started due to LLNL's utility, Western Area Power Administration (WAPA), approaching LLNL during Sequoia's commissioning
  - WAPA noticed spikes in power from 9MW to 180kW
  - LLNL implemented a power management monitoring system
  - New tools allow for further grid analysis
  - Continue to simulate data analytics on HPC loads to understand applications for future systems
  - Ultimately move from static to dynamic power management for improved grid analysis through extensive active case studies

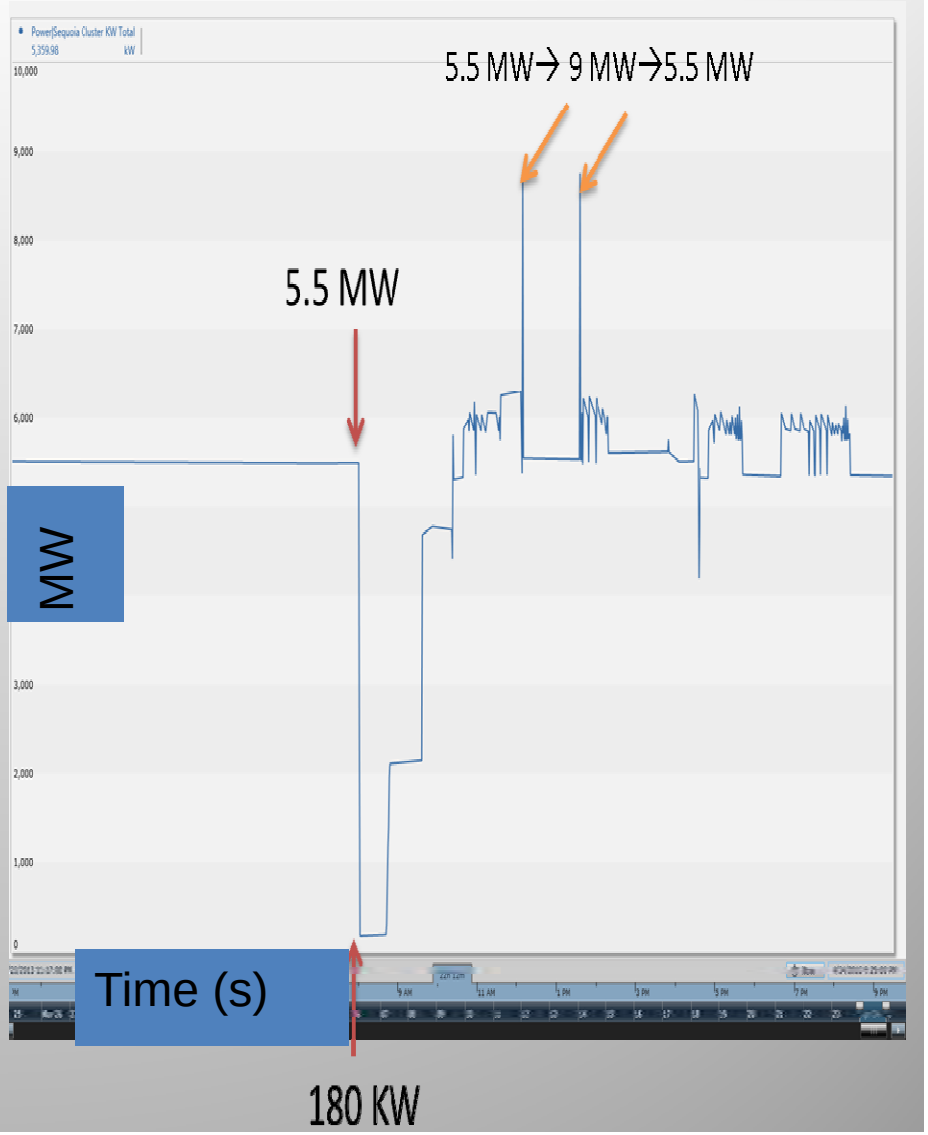

HPC and the Grid Future Focus

- Continue to work with WAPA to implement live feeds from HPC datacenters at LLNL to minimize unwanted large block load swings
  - WAPA performs all LLNL grid scheduling – Long and short term forecasting – Load and Grid Integration is the future

Environmental Monitoring System

Load Integration to Respond to Demand Changes

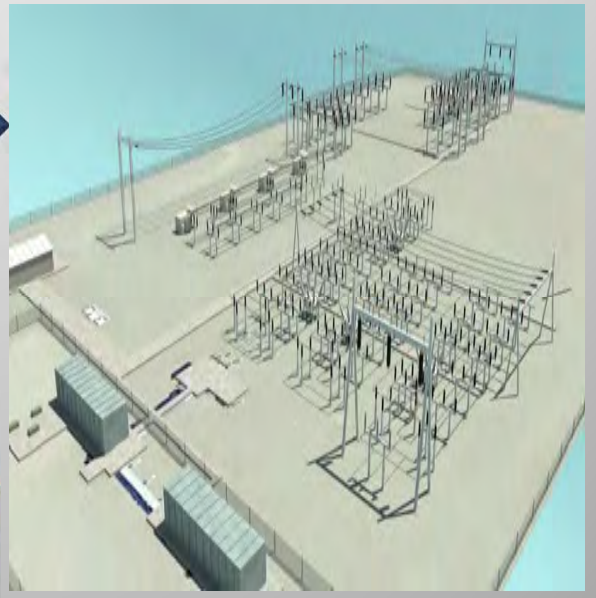

Power Management System

Grid Integration to Implement Required Energy Transactions

Why EEHPCWG?

- 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 for Systems
  - Power Usage Effectiveness (PUE) to Total Usage Effectiveness (TUE)
  - Demand Response Intiatives
  - Data Center Energy Management Dashboard

EEHPCWG is driving HPC energy efficiencies.

Lawrence Livermore
National Laboratory