

Wu Feng, Virginia Tech & Green500
Erich Strohmaier, Berkeley Lab & TOP500

with
Natalie Bates, EE HPC Working Group
Tahir Cader, HP & The Green Grid
John Shalf, Berkeley Lab & TOP500

SETTING TRENDS FOR ENERGY-EFFICIENT SUPERCOMPUTING

SC BoF; November 2010; New Orleans, LA, USA

Erich Strohmaier, Berkeley Lab & TOP500

Wu Feng, Virginia Tech & Green500

with

Natalie Bates, EE HPC Working Group

Tahir Cader, HP & The Green Grid

SETTING TRENDS FOR ENERGY-EFFICIENT SUPERCOMPUTING

SC BoF; November 2010; New Orleans, LA, USA

WHY WE ARE HERE

- × Context
 - + Power consumption of HPC and facility costs are increasing.
 - × “Can only improve what you can measure”
 - × What is needed?
 - + Converge on a common basis between different research and industry groups for:
 - × METRICS
 - × WORKLOADS
 - × METHODOLOGIES
- for energy-efficient supercomputing, so we can make progress towards solutions.

AGREEMENT IN PRINCIPAL

- ✘ Collaboration between Top500, Green500, Green Grid and EE HPC WG
- ✘ Improve methodology, metrics, instrumentation and testing
- ✘ Evaluate new technologies for HPC compute system energy efficiency
- ✘ Report progress at SC and ISC

COMPLIMENTARY METRICS

- ✘ Power Usage Effectiveness (PUE)
 - + “Data Center” Measurement
 - + $\text{PUE} = \text{Total Facility Power} / \text{IT Power}$

AND

Today's focus

- ✘ Workload/Productivity Metrics
 - + “Compute System” Measurement
 - + $\text{Useful Work} / \text{Energy Consumed}$

STAKEHOLDERS AND PURPOSE?

- ✘ Stakeholders

- + *HPC computer system designers and procurement decision makers*, including users, data center and facilities designers/managers

- ✘ Purpose

- + *Unite the community behind energy efficiency metric(s) for HPC systems* that form the basis for comparing and evaluating individual systems, product lines, architectures and vendors

WHAT IS A METRIC?

- ✘ A basis for comparison
- ✘ A reference point against which other things can be evaluated
- ✘ A measure

WHAT MAKES A METRIC EFFECTIVE ?

- ✘ Granular enough
 - + Individual components
 - + Analyzed in manageable chunks
 - + Assigned to specific parties for improvement
- ✘ Intuitive, obvious, and clear
- ✘ Scientifically accurate and used precisely
- ✘ Sufficiently flexible to respond to new technology developments
- ✘ Vendor-neutral
- ✘ Inexpensive and worthwhile

Stanley, J.R., Brill, K.G. and Koomey, J. "Four Metrics Define Data Center 'Greenness,'" Uptime Institute White Paper http://www.verneglobal.com/pdf/WP_FourMetrics.pdf

PROPOSED METRIC GRANULARITY

- ✘ Measure behavior of key system components including CPU, memory, storage and I/O
 - + Capture idle as well as fully-loaded utilization

WHAT IS A WORKLOAD?

- ✘ The application or benchmark software designed to exercise the HPC system to the fullest capability possible

AGREEMENT TO USE WORKLOAD METRICS

- ✘ Use Workload-based Metrics to Represent HPC Energy Efficiency
 - + Use workload-based for numerator and measured power during workload run for denominator
- ✘ Examples
 - + Green500 “FLOPS per Watt”
 - + SPEC FP / Measured Watt
 - + Green Grid “Productivity Proxies”
- ✘ *Still need to decide upon exact metric*
 - + *Classes of systems (e.g., Top50, Little500)*
 - + *Multiple metrics or a single index*

PROPOSED WORKLOADS

- ✘ *Leverage well-established benchmarks*
- ✘ Use High Performance LINPACK (HPL) for measuring compute energy efficiency
- ✘ Use RandomAccess (Giga Updates Per second or GUPs) for measuring memory energy efficiency
- ✘ *Need to identify workloads for measuring energy efficiency of storage and I/O*

WHAT IS A METHODOLOGY?

- ✘ The system of methods followed
- ✘ A way of doing something, especially a systematic way; implies an orderly logical arrangement (usually in steps)
- ✘ A measurement procedure

UNIFY AND IMPROVE METHODOLOGY

- ✘ HPL and RandomAccess measurement methodologies are well established
- ✘ Green500 and TOP500 power-measurement methodologies
 - + Similar, but not identical methodologies
 - + Yield very similar results
- ✘ Issues and concerns
 - + Define how to isolate integrated power/cooling systems from system power measurement
 - + Need to increase vendor and/or supercomputing center power-measurement reporting
 - ✘ June 2010 Green500 List
 - ★ 275 systems with measured #s, 225 systems with derived #s.

PROGRESS SINCE ISC BOF

- ✘ Stakeholders, purpose, and key definitions
- ✘ Focus only on compute system equipment in measurement (not facility or power/cooling)
- ✘ Workload categories are compute, memory, storage and I/O subsystems
 - + Compute: HPL
 - + Memory: RandomAccess
- ✘ Methodology for workloads are previously defined and documented
- ✘ Unify and continuously improve Green500 and TOP500 power measurement methodology

ISSUES TO RESOLVE

- ✘ *Identify workloads for measuring energy efficiency of storage and I/O*
- ✘ *Still need to decide upon exact metric*
 - + *Classes of systems (e.g., Top50, Little500)*
 - + *Multiple metrics or a single index*
- ✘ *Define how to isolate integrated power/cooling systems from system power measurement*
- ✘ *Need to increase vendor and/or SC center power measurement reporting*
 - + *June 2010 Green500 List has 225 systems with derived power numbers*

WHY WE ARE HERE

- × Context
 - + Power consumption of HPC and facility costs are increasing.
 - × “Can only improve what you can measure”
 - × What is needed?
 - + Converge on a common basis between different research and industry groups for:
 - × METRICS
 - × WORKLOADS
 - × METHODOLOGIES
- for energy-efficient supercomputing, so we can make progress towards solutions.

EXTRA SLIDES

THE GREEN GRID BOARD MEMBERS

