

Erich Strohmaier, Berkeley Lab & TOP500

Wu Feng, Virginia Tech & Green500

with

Natalie Bates, EE HPC Working Group

Tahir Cader, HP & The Green Grid

Craig Steffen, NCSA

And others on the Compute System Metrics Team

# SETTING TRENDS FOR ENERGY-EFFICIENT SUPERCOMPUTING

ISC BoF; June 2011; Hamburg, Germany

# WHY WE ARE HERE

---

- × Context

- + Power consumption and facility costs of HPC are increasing.

- × “Can only improve what you can measure”

- × What is needed?

- + Converge on a common basis for:

- × METHODOLOGIES

- × WORKLOADS

- × METRICS

- for energy-efficient supercomputing, so we can make progress towards solutions.

# AGREEMENT IN PRINCIPAL

---

- ✘ Collaboration between Top500, Green500, Green Grid and EE HPC WG
- ✘ Evaluate and improve methodology, metrics, and drive towards convergence on workloads
- ✘ Report progress at SC and ISC

# COMPLIMENTARY METRICS

## × “Data Center” Measurements

### + Power Usage Effectiveness (PUE)

×  $PUE = \text{Total Facility Energy} / \text{IT Energy}$

### + Energy Reuse Effectiveness (ERE)

×  $ERE = (\text{Total E.} - \text{Reused E.}) / \text{IT Energy}$

AND

Today's focus

## × “Compute System” Measurements

### + Workload/Productivity Metrics

### + Useful Work / Energy Consumed

# PROPOSED METRIC GRANULARITY

---

- ✘ Measure behavior of key system components including compute, memory, interconnect fabric, storage and external I/O
  - + Workloads and Metrics might address several components at the same time
- ✘ Phased implementation planned

# PROPOSED WORKLOADS

---

- ✘ *Leverage well-established benchmarks*
- ✘ *Must exercise the HPC system to the fullest capability possible*
  
- ✘ Use High Performance LINPACK (HPL) for exercising (mostly) compute sub-system
- ✘ Use RandomAccess (Giga Updates Per second or GUPs) for exercising memory sub-system (?)
- ✘ *Need to identify workloads for exercising other sub-systems*
  - + e.g.: *SNIA has a Green Storage Initiative*

# UNIFY AND IMPROVE METHODOLOGY

- ✘ HPL and RandomAccess measurement methodologies are well established
- ✘ Green500 & TOP500 power-measurement methodology
  - + Similar, but not identical methodologies
- ✘ Issues/concerns with power-measurement methodology
  - + What to include in the measurement (e.g., integrated cooling)
  - + Node, rack or system level measurements
  - + Variation in start/stop times as well as sampling rates
  - + Need to increase vendor and/or supercomputing center power-measurement reporting
    - ✘ November 2010 Green500 List – 54% of the list were submitted /measured numbers, the other 46% were derived by Green500

# UNIFY AND IMPROVE METHODOLOGY

- ✘ Current power measurement methodology is very flexible, but compromises consistency between submissions
- ✘ Proposal is to keep flexibility, but keep track of rules used and quality of power measurement
- ✘ Tiers of power measurement quality
  - ✘ Sampling rate; more measurements means higher quality
  - ✘ Completeness of what is being measured; more of the system translates to higher quality
  - ✘ Common rules for start/stop times

## PROGRESS SINCE SC BOF

---

- ✘ Focus on improving power measurement methodology
  - + Survey key sites to clarify issues and concerns
  - + Proposal for tiers of power measurement quality
- ✘ Working paper drafted and under review
- ✘ Team broadened to include representatives from major sites and additional industry players
- ✘ MOU between The Green Grid, the Top500, the Green500 and the EE HPC WG

## ISSUES TO RESOLVE

---

- ✘ *Identify workloads for exercising other sub-systems; e.g., storage, I/O*
- ✘ *Still need to decide upon exact metric*
  - + *Classes of systems (e.g., Top50, Little500)*
  - + *Multiple metrics or a single index*
- ✘ *Implement tiers of power measurement quality*
- ✘ *Facilitate increased vendor and/or SC center power measurement reporting*
- ✘ *Include idle as well as fully exercised measurements*

# WHY WE ARE HERE

---

- ✘ To gather community momentum and support
- ✘ To solicit your feedback
- ✘ To ask you to participate by submitting power measurements for Top500 and Green500 Lists

# Energy Efficient HPC Working Group

## EE HPC WG

- Driving energy conservation measures and energy efficient design in high performance computing.
- Demonstrate leadership in energy efficiency as in computing performance.
- Forum for sharing of information (peer-to-peer exchange) and collective action.

<http://eehpcwg.lbl.gov>



Sponsored by:  
U.S. Department of Energy  
Energy Efficiency and Renewable Energy  
Sustainability Performance Office

**EXTRA SLIDES**

# 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](http://www.verneglobal.com/pdf/WP_FourMetrics.pdf)

# 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*
- ✘ *Idle as well as full stress*

# 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

# THE GREEN GRID BOARD MEMBERS

