

Energy Efficiency Metrics

Natalie Bates, Energy Efficient HPC Working Group

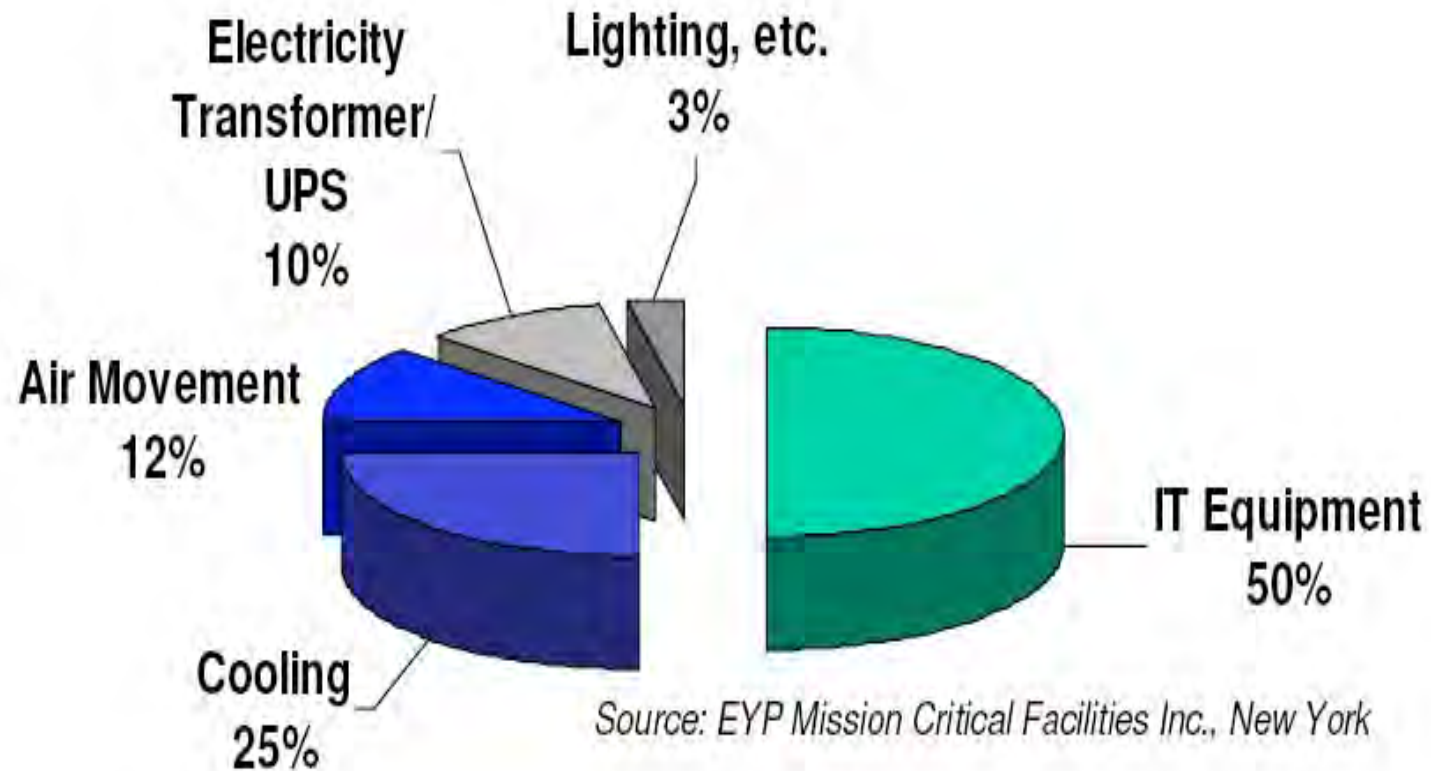
With help from

Michael Patterson, Intel and The Green Grid

Tahir Cader, HP and The Green Grid

Wu Feng, Virginia Tech and the Green500
and Erich Strohmaier, LBNL and the TOP500

Where Data Center Power Goes



Source: EYP Mission Critical Facilities Inc., New York

Complimentary Metrics

- Power Usage Effectiveness (PUE)
 - Data Center measurement
 - $PUE = \text{Total Facility Power} / \text{IT Equipment Power}$

AND

- Workload/Productivity Metrics
 - Compute System measurement
 - Useful work / Energy Consumed

Data Center Efficiency

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

- Expressed as Power Utilization Effectiveness (PUE)
 - Perfect PUE = 1.0
- Typical data center today = ~ 2.0
- Best Practices ~1.5
- World Class ~1.2-1.3

Global Harmonization of Metrics

- The Green Grid's Power Usage Effectiveness (PUE) metric has achieved significant worldwide adoption
 - $PUE = \text{Total Facility Power} / \text{IT Power}$
- In February of 2010, The US DOE hosted a meeting targeted at the Global Harmonization of Metrics
 - Attendees included The Green Grid, US DOE, US EPA, EU Code of Conduct, Japan's METI, and Japan's GIPC
 - At this meeting, the group agreed to “harmonize” on the PUE
- As a next step, the group agreed on the need for immediate work on a data center productivity metric

Agreement in principal

- Collaboration between Top500, Green500, Green Grid and EE HPC WG
- Improve methodology, metrics, instrumentation and testing in order to 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
- Report progress at SC and ISC

Proposed metric granularity

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

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*

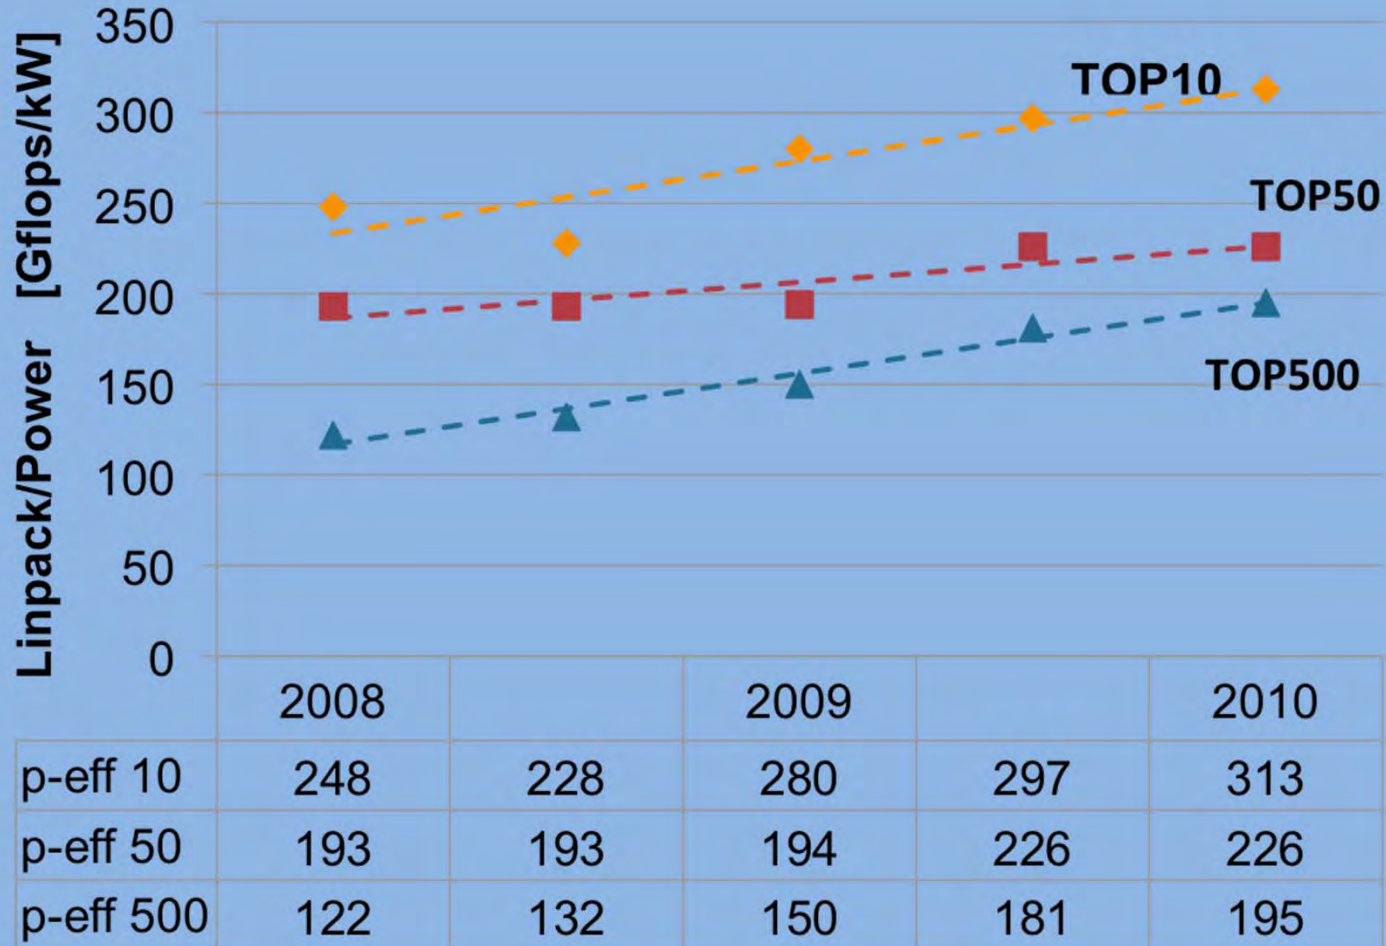
Proposed workloads

- *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*

Unify and continuously improve methodology

- HPL and RandomAccess measurement methodologies 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 SC center power measurement reporting*
 - *June 2010 Green500 List has 225 systems with derived power numbers*

Power Efficiency



June 2010: Green500 Top 10

- Accelerator-based supercomputers now occupy the top eight slots
- Three IBM (QPACE) machines tie for first place

Green 500	Top 500	
#1	#131	Forschungszentrum Juelich (FZJ)
#1	#132	Universitaet Regensburg
#1	#133	Universitaet Wuppertal
#4	#2	National Supercomputing Centre in Shenzhen
#5	#35	IBM Poughkeepsie Benchmarking Center
#5	#88	DOE/NNSA/LANL
#7	#3	DOE/NNSA/LANL
#8	#19	Chinese Academy of Sciences
#9	#331	Mississippi State University
#10	#381	Banking (M)

Recommendations

- If you don't already do it, start measuring and tracking your Data Center PUE
- When purchasing new equipment, ask vendors for workload/productivity energy efficiency performance
 - Use HPL and RandomAccess
 - And/Or, some other workload specific to your user community