

Energy Efficiency Metrics and Benchmarks at Exascale: The Application Perspective


Natalie Bates,
Energy Efficient HPC Working Group

With special help from

Jim Laros, Chung-Hsing Hsu, Thomas Ilsche, Daniel
Hackenberg, Steve Poole, Parks Fields, Barry Rountree

Energy Efficient HPC Working Group

- **Mobilizing the HPC community to accelerate EE HPC**
 - Sharing of best practices and taking collective action.
 - Only membership criteria is ‘interest’ in driving EE HPC.



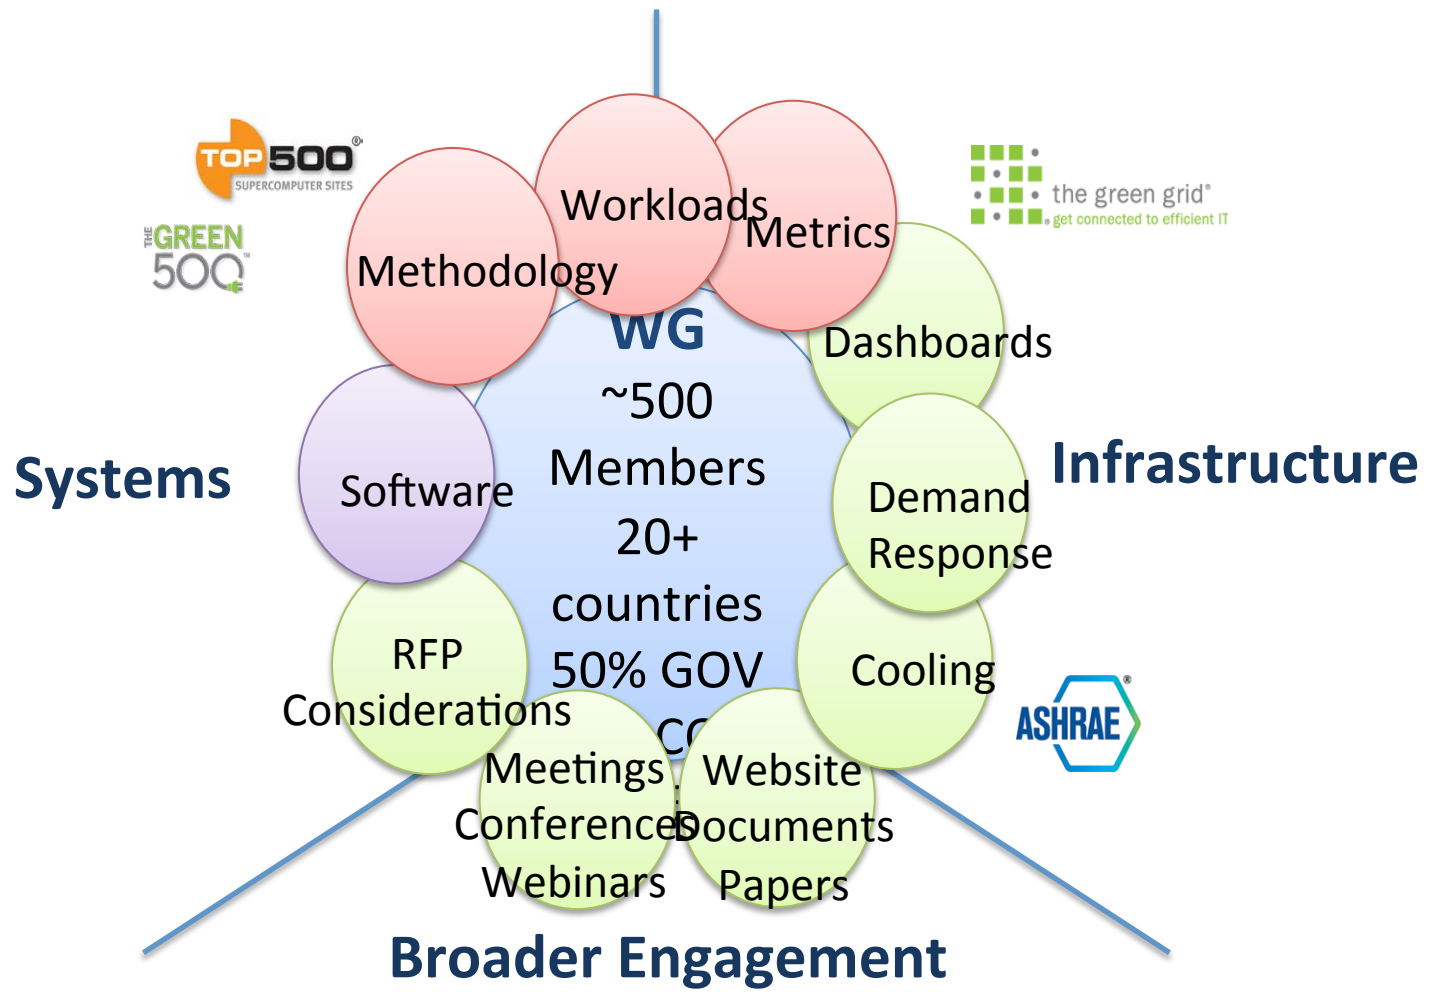
<http://eehpcwg.lbl.gov>

natalie.jean.bates@gmail.com

Energy Efficient HPC Working Group

Team Activities:

<http://eehpcwg.lbl.gov/calendar/minutes>



Procurement considerations document


□ 2013

- Measurement locations and frequencies
- Management and control
- Benchmarks
- General objectives
- Air and liquid cooling

□ 2014

- Timestamping and clocks
 - Temperature measurement
 - More... power distribution, characterization
- <http://eehpcwg.lbl.gov/documents/procurement-considerations-presentations>

Measurement: Reported Values and Internal Sampling



Measurement Requirement Example:

	MANDATORY	IMPORTANT	ENHANCING
System	≥ 10 per sec	≥ 100 per sec	$\geq 1,000$ per sec
Node	≥ 100 per sec	$\geq 1,000$ per sec	$\geq 10,000$ per sec
Component	$\geq 1,000$ per sec	$\geq 10,000$ per sec	$\geq 1,000,000$ per sec

Measurement Frequency Requirements (voltage and current)

	MANDATORY	IMPORTANT	ENHANCING
System	≥ 1 per sec	≥ 10 per sec	≥ 100 per sec
Node	≥ 10 per sec	≥ 100 per sec	$\geq 1,000$ per sec
Component	≥ 100 per sec	$\geq 1,000$ per sec	$\geq 10,000$ per sec

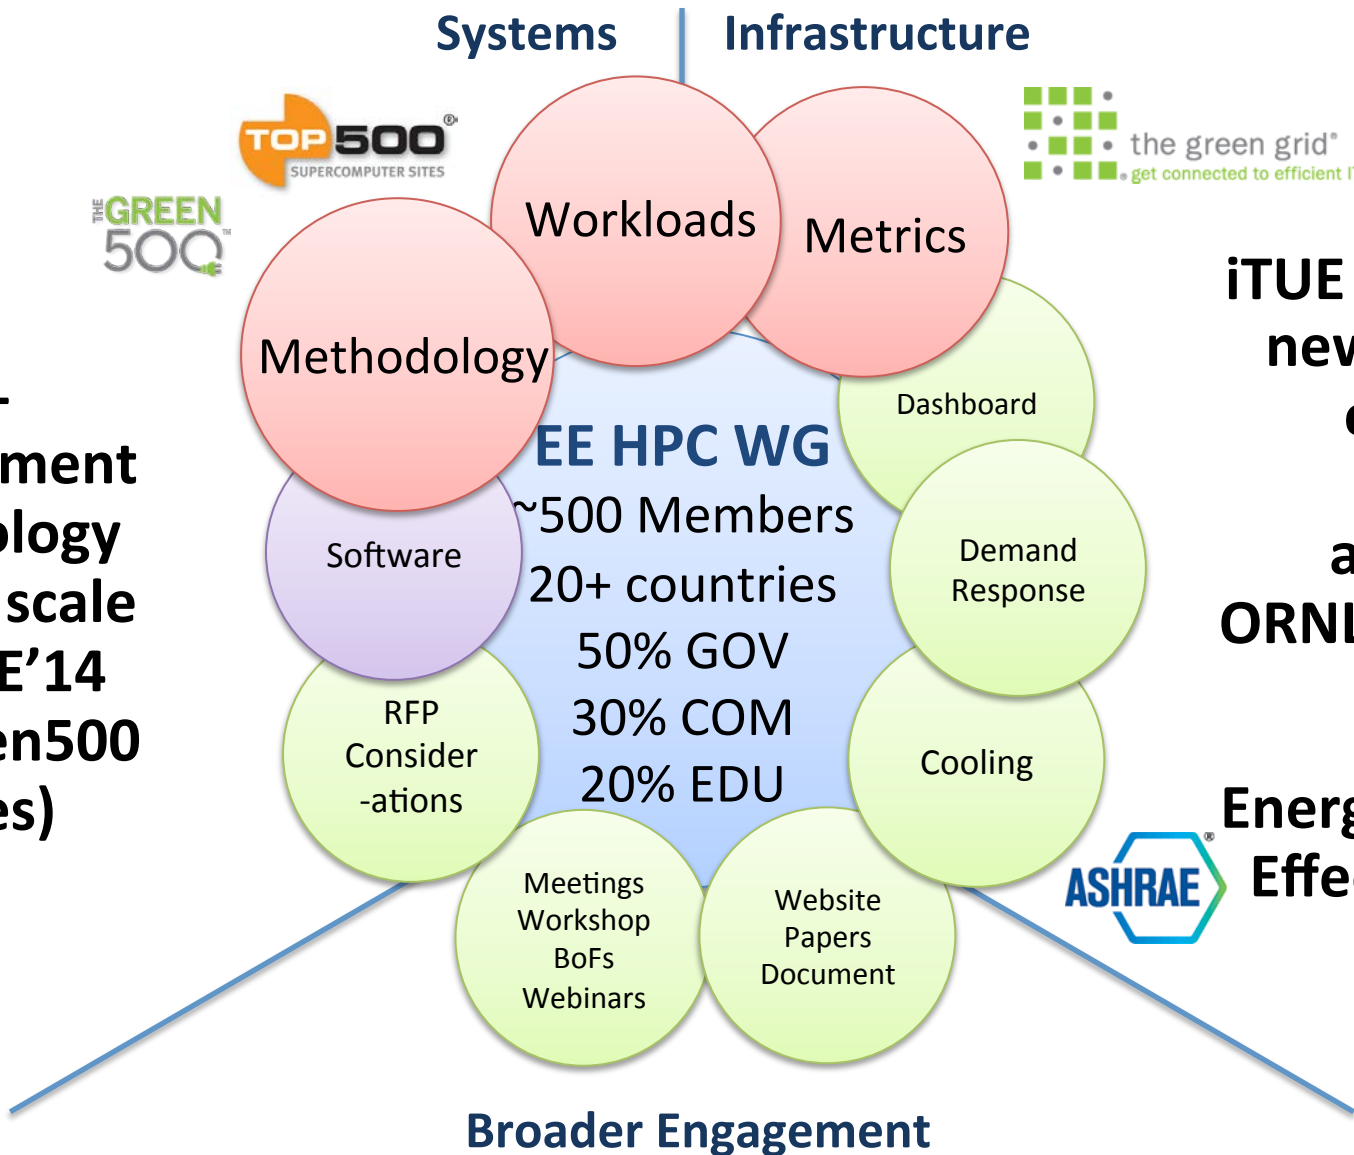
Reported Frequency Requirements (discrete power)

Common Power API

- ❑ Led by United States Department of Energy's Sandia National Laboratory with community participation
- ❑ Includes use cases models as well as specification
- ❑ Standard, comprehensive API for platform power measurement and control
- ❑ Multiple layers of abstraction for multiple types of users
- ❑ For more information, contact Suzanne M Kelly smkelly@sandia.gov

Energy Efficient HPC Working Group

A power-measurement methodology for large scale HPC (ICPE'14 and Green500 Resources)



**iTUE and TUE,
new energy-
efficiency
metrics
applied at
ORNL's Jaguar
(ISC'13)**

**Energy Re-use
Effectiveness
(TGG)**

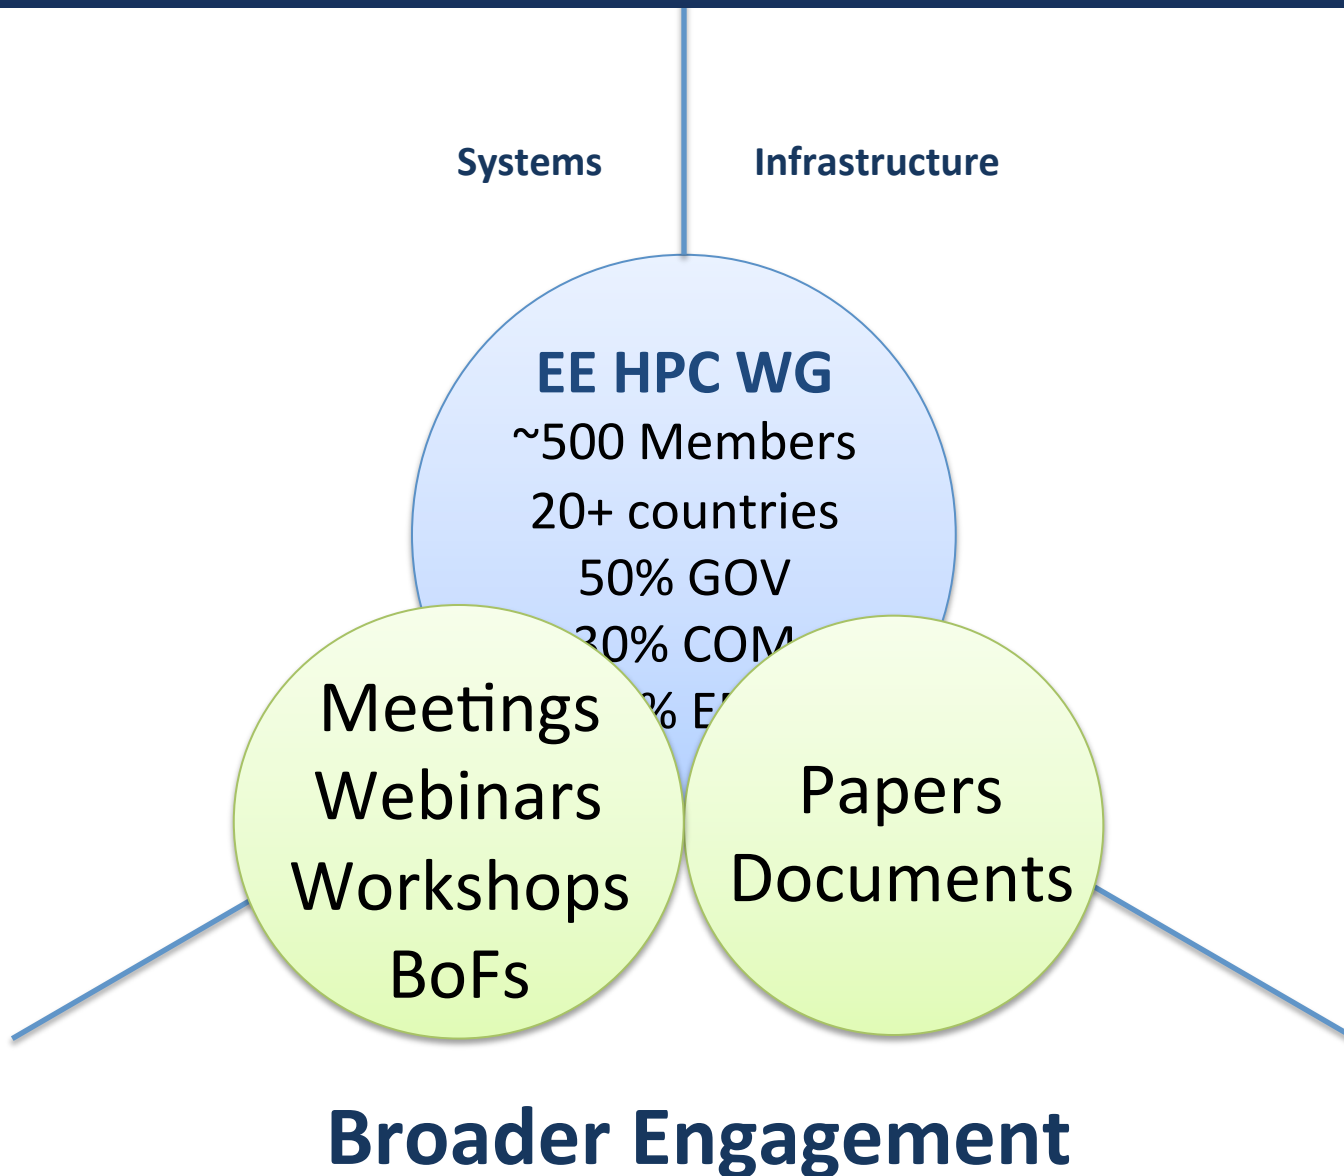
Energy Efficient HPC Working Group

<http://eehpcwg.lbl.gov>

natalie.jean.bates@gmail.com

Back-up

Energy Efficient HPC Working Group



Energy Efficient HPC Working Group

Capturing 'best practices' in system procurement that considers energy efficiency, especially on power and energy measurement and management capabilities. (Website)

Systems

Infrastructure

Capturing 'best practices' in cooling support for systems, especially on guidelines and commissioning; closely work with ASHRAE. (SC'10 and Website)

EE HPC WG

~500 Members
20+ countries
50% GOV
30% COM
20% EDU

Procurement Considerations

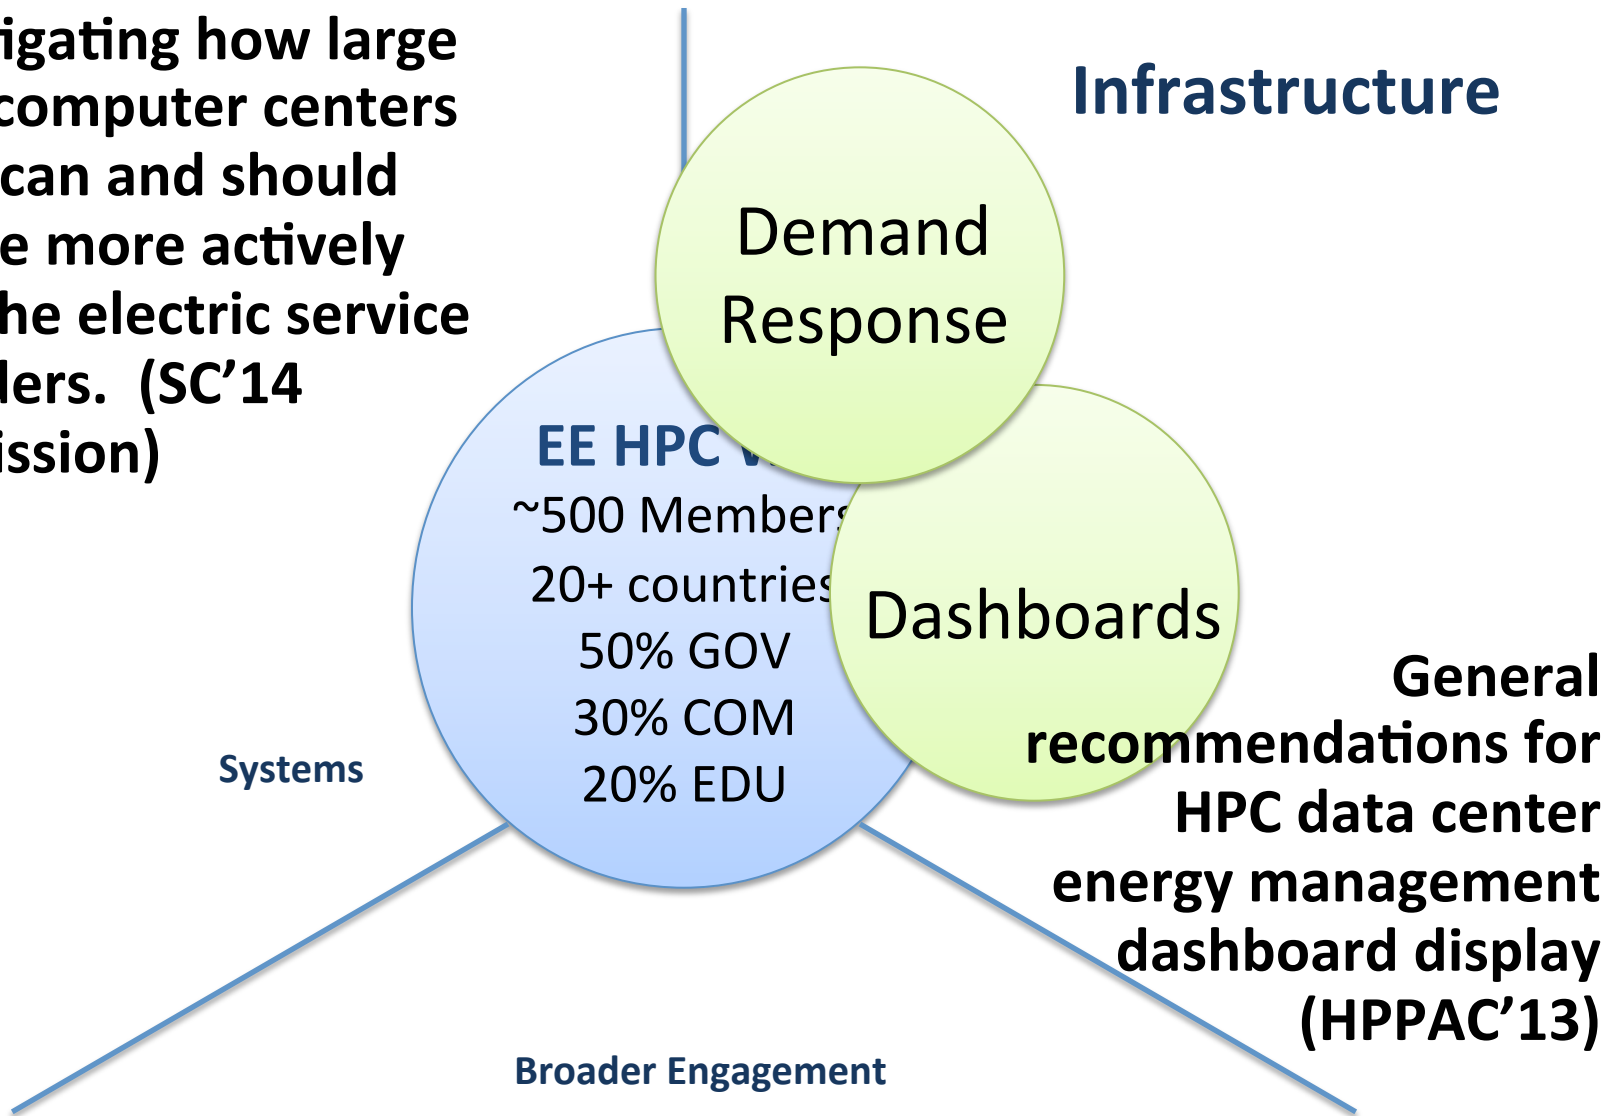
Cooling



Broader Engagement

Energy Efficient HPC Working Group

Investigating how large supercomputer centers have, can and should engage more actively with the electric service providers. (SC'14 Submission)



Energy Efficient HPC Working Group

Systems

Infrastructure

Software for
improving
energy
efficiency

Software

EE HPC WG

~500 Members

20+ countries

50% GOV

30% COM

20% EDU

Software for
other purposes
runs efficiently

Broader Engagement

Workloads and Benchmarks

- ❑ Cross-Team Concern:
 - ❑ Power Measurement Methodology, Procurement, TUE, SW Teams
- ❑ 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
- ❑ Exercise the HPC system to the fullest capability possible
- ❑ Use High Performance LINPACK (HPL) for exercising (mostly) compute sub-system

Next Steps

- Solicit and compile other potential workloads
 - RandomAccess (Giga Updates Per second or GUPs) for exercising memory sub-system
 - SystemBurn, ORNL
 - FIRESTARTER, University of Dresden