Run Rules for The Green500

Power Measurement of Supercomputers

Version 0.9

Below is version 0.9 of the specification for measuring the power and energy consumption of qualified supercomputers (or high-end computers). A system must meet all of the identified criteria if it is to be considered as meeting the requirements for entry in *The Green500 List* and *The Little Green500$^{(beta)}$ List*.

1. Definitions

Below is a description of supercomputers and other terms relevant to *The Green500 List*.

1.1 Top500 List (or Top500): The list of the fastest supercomputers in the world, relative to the high-performance LINPACK (HPL) benchmark. The Top500 is maintained by the University of Mannheim, the University of Tennessee, and NERSC/LBNL.

1.2 High-Performance Linpack (HPL): HPL is a portable implementation of the LINPACK benchmark for distributed memory systems. The benchmark can be downloaded from http://www.netlib.org/benchmark/hpl/.

1.3 Supercomputer:

i. In general, a supercomputer is defined as a computer that is at the forefront of current processing capability, particularly with respect to the speed of calculation.

ii. In the context of *The Green500 List*, a supercomputer is a computing system that is fast enough to appear of the latest Top500 List.

iii. In the context of *The Little Green500$^{(beta)}$ List*, a ("little") supercomputer is a computing system that achieves performance on the HPL benchmark at a high-enough level to have secured entry into the oldest Top500 list released within 19 months of the release date of *The Little Green500$^{(beta)}$ List* that is being submitted to. (For additional details, see the "Qualifying Supercomputers" section.)

1.4 FLOPS: Floating-point operations per second.

1.5 R_{Peak}: The theoretical *peak* performance (measured in FLOPS) for a supercomputer.

1.6 R_{max}: The *max*imum performance (measured in FLOPS) achieved by a supercomputer while running HPL benchmark.

1.7 On Mode/Active Power: The supercomputer is connected to a power source and is capable of executing user programs (e.g., HPL benchmark).

1.8 Sleep Mode/Low Power: The supercomputer is connected to a power source and powered on, but currently in a reduced power state after receiving instructions from systems software or hardware. Systems that have been idle for a prolonged period of time may transition to this mode to conserve power and energy. The system returns to On Mode (or Active Power) with full operational capability upon sensing or receiving a request from the system or end user.

1.9 Off Mode/Power Standby: The supercomputer is connected to a power source with limited or no activity and is waiting to be switched to On Mode by a direct signal from the user/system. This mode may or may not be present in current systems. Hibernating systems that still require small amounts of power characterize this mode of operation.

1.10 Hard Off Mode: A condition where the supercomputer is still plugged in but has been disconnected from an external power source. This mode is engaged via a "hard off switch." While in this mode, the supercomputer will not draw any electricity and will usually measure zero watts when metered.

1.11 Disconnected: The supercomputer has been unplugged and is therefore disconnected from all external power sources.

1.12 Under Load: The supercomputer is "under load" when it is operating on user-defined parallel tasks. In this document, the defined parallel task is the HPL benchmark used by the Top500.

1.13 Peak Power: The theoretical peak power consumed (measured in kilowatts or kW) by a supercomputer.

1.14 R_{max} Power: The average power consumption (measured in kilowatts or kW) of a supercomputer while achieving R_{max}.

1.15 Subcomponent: A subcomponent is that part of a supercomputer which can be measured in isolation for power consumption. For example, this may be a single chassis, a rack, or any physical enclosure that is larger than the fundamental unit of a chassis. The important criterion here is the ability to measure power in isolation for a period of time while the subcomponent is engaged collectively with the entire supercomputer to perform a task. Single subcomponent power consumption under load will be used to estimate system-wide power consumption. The entire system may be considered as a "subcomponent" under this definition. *The subcomponent that is measured must consume at least 1 kilowatt of power.*

2. Qualifying Supercomputers

All reported results must follow the test methodology described in Section 3. In addition to those requirements each list has its own rules, which are enumerated below.

2.1. Green500 List: In order to qualify for *The Green500 List*, a supercomputer must achieve performance at least as high as the 500th-ranked system in the current version of the Top500 List. The power reported with submissions to *The Green500 List* must be the power for achieving the *performance reported to The Green500 List*, but need not match the corresponding R_{max} and power reports to the Top500 List. It is required that all the cores in the supercomputer be used for executing the HPL benchmark for a Green500 submission, unless fewer cores were used to report to the Top500 List, in which case that number is acceptable. It is also encouraged to submit both the R_{max} power (see section 1.N) and the power at which the entire supercomputer operates at highest energy efficiency in terms of performance per watt. In short, if you wish to apply dynamic voltage and frequency scaling (DVFS) or some other methodology, which affects performance without dropping off the bottom of the Top500 List, and still using the entire machine, that is allowed.

2.2. Little Green500$^{(beta)}$: To be eligible for *The Little Green500$^{(beta)}$ List*, a supercomputer must be at least as fast as the 500th-ranked supercomputer on the oldest Top500 List released within 19 months of the release date of *The Little Green500$^{(beta)}$ List* that you are submitting for, as noted in the diagram below. The power reported for *The Little Green500$^{(beta)}$ List* can be for any subcomponent (refer to Section 1.O.) of the supercomputer. Further, techniques such as DVFS can be used for reporting power and performance values to *The Little Green500$^{(beta)}$ List*. However, the subcomponent must match the full performance requirements.

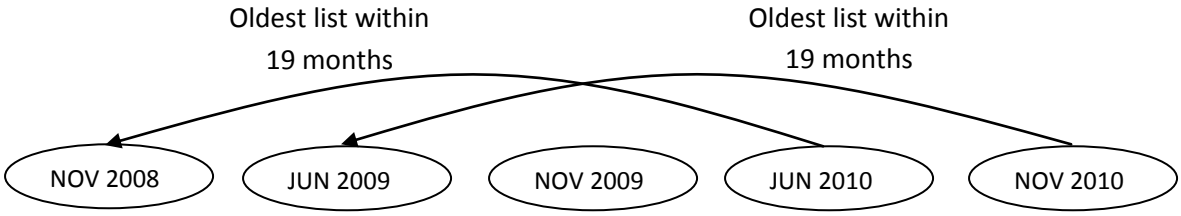

Example Little Green500 List references:
Jun 2010 machines must meet the bottom of the Nov 2008 Top500, and Nov 2010 must meet the bottom of the Jun 2009 Top500

3. Test Methodology

Supercomputer Testing Set-up, Methodology, and Documentation: We utilize, where possible, existing and widely accepted industry practices for measuring supercomputer performance and power use under normal or typical operating conditions. The testing and measurement methods below are based upon the US Environmental Protection Agency's specifications for ENERGY STAR compliance adapted for the Green500, where necessary. Submitters are required to perform tests and self-certify supercomputers that qualify for the Green500 based on Section 2 of this document.

The power requirements of the equipment shall be measured from the electrical outlet or power source to the supercomputer under test. The average true power consumption of any single node shall be measured during the On Mode/Active Power state of the system. When performing measurements to self-certify a system for *The Green500 List*, the supercomputer under test must adhere to the test methodology described in Section 3.2.

To ensure consistent means for measuring the power consumption of supercomputers, the following protocols must be followed:

- **Supercomputer Testing Set-Up and Conditions:** Outlined below in Section 3.1 are the ambient test conditions and measurement protocols that must be followed when measuring the power of the supercomputer.
- **Supercomputer Testing Methodology:** The actual steps for measuring On Mode/Active Power are provided in Section 3.2.
- **Supercomputer Testing Documentation**: Documentation requirements for submission of data to the Green500 are detailed in Section 3.3. This protocol ensures that outside factors do not adversely affect the test results and that the test results can be consistently reproduced. Manufacturers may elect to use in-house or independent laboratory facilities to provide test results.

3.1 Supercomputer Testing Set-Up and Conditions

A. **Supply Voltage and Frequency:** The supply voltage and frequency for all measurements should be +/- 1% of the rated AC voltage and frequency. If the rated AC voltage or

frequency is unclear from the specifications, it is recommended that voltage and frequency is selected based on the country in which the supercomputer is installed.

B. **Total Harmonic Distortion (Voltage):** < 2% THD

C. **Ambient Temperature:** 20 degrees C +/- 5 degrees C

D. **Relative Humidity:** 30-80%

E. **Line Impedance:** < 0.25 ohm

F. **Power Measurement Test Conditions:** The HPL runs should comply with all the run rules of the Top500 List. The power measurement procedures should meet all the conditions described in section 2.1 and 2.2. The R_{max} power measurement (if reported) should be taken simultaneously with the HPL execution for submission to the Top500 List.

G. **Power Measurement Protocols:** Supercomputer power shall be measured in watts during execution of the HPL benchmark, following the rules for entry into the Top500 List. The system must be in On Mode/Active Power for at least 15 minutes prior to launching HPL to warm-up the system. HPL should be run for a minimum of 10% of the overall execution time of the benchmark prior to any power measurements. A true RMS power meter with a crest factor of at least five shall be used to measure the power consumption of each randomly chosen subcomponent meeting the definition in Section 1.O. Measurements shall be taken for a period of no less than 20% of the overall execution time of the benchmark and instantaneous power must be recorded with a sampling rate of at least one measurement per second during the same period. This experiment must be repeated at least one more time on the same subcomponent. If the average power of the two readings on the same node **varies by 1%**, the experiment should be repeated on a different subcomponent. This is required to ensure results are reliable and repeatable. Variances within this threshold require no further experiments.

3.2. Supercomputer Testing Methodology

Test Method: Below are the steps for measuring the true power requirements of the test unit in On Mode/Active Power. Submitters are required to test their supercomputers using analog or digital interface monitors with sampling rates of at least one measurement per second.

A. Connect the test unit (e.g., single rack) to the electrical outlet or power source and test equipment. For subcomponents with an external power supply, the external power supply (as opposed to a reference power supply) must be used in the test.
B. Power-on all test equipment and either verify that the wall outlet power is within specifications or adjust the AC power source output as described in Section 3.1.A. (e.g., 115V ~1%, 60Hz ~1%).
C. Check for normal operation of the test unit (e.g. single rack).
D. Allow the unit under test to reach operating temperature (at least 15 minutes).
E. Launch HPL benchmark on the full system.
F. Set the power meter current range. The full-scale value selected multiplied by the crest factor rating (lpeak/lrms) of the meter must be greater than the peak power reading from the oscilloscope (or digital multimeter).
G. Allow the readings on the power meter to stabilize (approximately 10% of the overall execution time) and then begin recording the true power reading in watts from the power

meter. This will require an additional PC or laptop to interface with the power meter and record measurement data electronically.

H. Record power measurements for at least 20% of the overall execution time to calculate FLOPS/watt.
I. Calculate total FLOPS/watt of the supercomputer by multiplying subcomponent's FLOPS/watt by the number of "subcomponent" as defined in Section 1.O.
J. Record the test conditions and test data.

3.3. Supercomputer Testing Documentation

- **Submission of Supercomputer Data to Green500:** Submissions can be directly made at http://www.green500.org/. Power measurement data should be obtained using the aforementioned protocols.

- **Effective Date**: A supercomputer will be listed in *The Green500 List* and *The Little Green500$^{(beta)}$ List* as long as the conditions in section 2 are met.

- **Future Specification Revisions**: Green500 reserves the right to change the specification should technological and/or market changes affect its usefulness to consumers, industry, or the environment. We will continually assess the market in terms of energy efficiency and new technologies. The Green500 will attempt to recognize the most energy efficient models of supercomputers in the marketplace and recognize those manufacturers who have made efforts to further improve energy efficiency.

- **Recommended Power Measurement Equipment**: It is recommended that power measurement equipment from the list of accepted power measurement devices of Standard Performance Evaluation Corporation (SPEC) be used. This list of accepted devices can be viewed at http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Device_List.html.

To provide constructive feedback, please send e-mail to "feedback –at– green500 –dot– org"