

# The Green500 List

## Power Measurement of High-End Clusters

### Version 0.1

Wu-chun Feng and Kirk Cameron  
{feng, cameron}@cs.vt.edu

Last Updated November 12, 2006

Below is version 0.1 of the specification for measuring the power and energy consumption of qualified high-end clusters. A supercomputer must meet all of the identified criteria if it is to be considered as meeting the requirements for entry in the Green500.

1. **Definitions:** Below is a description of supercomputers and other terms relevant to the Green500 list.
  - A. Top500 list (or Top500): The list of most powerful computer systems maintained by the University of Mannheim, the University of Tennessee, and NERSC/LBNL.
  - B. Supercomputer: In the context of this document, a supercomputer is a computer system that achieves sustained performance on the LINPACK benchmark at a high-enough level to warrant entry in the current Top500 list.
  - C. FLOPS: Floating point operations per second
  - D. Peak FLOPS (or Peak): The theoretical peak floating point operations per second for a supercomputer.
  - E. Rmax FLOPS (or Rmax Peak): The sustained speed of the LINPACK benchmark for a supercomputer listed in the current Top 500 list.
  - F. On Mode/Active Power: The supercomputer is connected to a power source and is capable of executing user programs (e.g. LINPACK).
  - G. 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. Idle systems characterize this mode of operation. The system returns to On Mode with full operational capability upon sensing a request from the user/system.
  - H. 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.
  - I. 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 0 watts when metered.
  - J. Disconnect: The supercomputer has been unplugged and is therefore disconnected from all external power sources.
  - K. Peak Power: The theoretical peak power consumed by a supercomputer, measured in Kilowatts.
  - L. Rmax Power: The sustained (average) power consumption of the LINPACK benchmark for a supercomputer listed in the current Top 500 list, measured in Kilowatts.
  - M. 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 LINPACK benchmark used to rank the Top500.
  - N. Single node: A single node of the supercomputer is a subcomponent that can be measured in isolation for power consumption. For example, this may be a single 1U, 2U, 3U, 4U, or 5U unit in the case of a cluster or an entire cabinet in the case of a constellation. The important criterion here is the ability to measure power in isolation for a period of time

while the component is engaged collectively with the entire supercomputer to perform a task. Single node power consumption under load will be used to estimate system-wide power consumption. The entire system may be considered a “single node” under this definition provide power is measurable excluding other systems or components including HVAC.

2. **Qualifying supercomputers:** In order to qualify for the Green500, a supercomputer must meet the definition in Section 1.A and results must follow the test methodology described in Section 3.
3. **Test Methodology:**

***Supercomputer Testing Set-up, Methodology, and Documentation:*** We utilize, where possible, existing, 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 Agencies 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. Supercomputers unchanged from a previous list that still meet the qualifying criteria of Section 2, may remain qualified without the submission of new test data, assuming the specification remains unchanged.

The power requirements of the equipment shall be measured from the 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, the supercomputer under test must be in the same condition as when running the LINPACK benchmark for consideration in the Top500 list.

To ensure consistent means for measuring the power consumption of supercomputers, the following protocol must be followed, which has three main components:

*Supercomputer Testing Set-up and Conditions:* Outlined below in Sections A through H are the ambient test conditions and measurement protocols that must be respected when performance power measurements.

*Supercomputer Testing Methodology:* The actual steps for measuring On Mode/Active Power are provided in Section I, below.

*Supercomputer Testing Documentation:* Documentation requirements for submittal of qualified product data to the Green500 are detailed in Section J, below.

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 to provide test results.

### **Supercomputer Testing Set-up and Conditions**

General Criteria (**Stolen from EPA, need to check**)

Supply Voltage:

North America: 115(~1%) Volts AC, 60 Hz (~1%)

--**Something about 220V power, etc. for all of these**

Europe: 230 (~1%) Volts AC, 50 Hz (~1%)

Australia/New Zealand: 230 (~1%) Volts AC, 50 Hz (~1%)

Japan: 100 (~1%) Volts AC, 50 Hz (~1%)/60 Hz (~1%)

Total Harmonic Distortion (Voltage): < 2% THD

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

Relative Humidity: 30-80%

Line Impedance: <.25 ohm

*Power Measurement Test Conditions:* It is preferred that power measurements be taken simultaneously with LINPACK measurements used for submission to the Top500 list. If this is impossible, submitters must certify that conditions during measurement were identical to conditions for final LINPACK scoring and measured LINPACK values during testing must be provided along with an explanation as to why measurements were not taken during LINPACK qualifying tests.

*Power Measurement Protocols:* Supercomputer power shall be measured in watts during execution of the LINPACK benchmark as described for entry into the Top500 list. The system must be in On Mode/Active Power for about 15 minutes prior to launching LINPACK to warm up the system. LINPACK shall be run for a minimum of a 10-minute period prior to any power measurements being taken. A true RMS power meter with a **crest factor of at least five** shall be used to measure the power use of each randomly chosen single node meeting the definition in Section 2.N. Measurements shall be taken starting at any point after the 25-minute warm up for a period of no less than 10 minutes.

During the test measurement interval of at least 10 minutes, the following must be recorded by electronic means: maximum power use, minimum power use, average power use. **If the variance along the 10 minute test measurement interval exceeds XXX**, the experiment should be repeated two more times on different single nodes (as possible). This is required to ensure results are reliable and repeatable. Variances within this threshold require no further experiments.

### **Supercomputer Testing Methodology**

*Test Method:* Following are the test 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 1 measurement per second.

1. Connect the test sample (e.g. single node) to the outlet or power source and test equipment. For single nodes with an external power supply, the external power supply (as opposed to a reference power supply) must be used in the test.
2. Power on all test equipment and properly adjust power source voltage and frequency.
3. Check for normal operation of the test unit (e.g. single node) and leave all adjustments set to LINPACK Top500 score settings.
4. Ensure the unit is in On Mode/Active Power by using the ON/OFF switch on the test unit. Allow the unit under test to reach operating temperature (approximately 15 minutes).
5. Launch LINPACK benchmark on full system. Verify (and record) that Rmax Peak has been reached and is being sustained.
6. Either verify that the wall outlet power is within specifications or adjust the AC power source output as described in Section XXX (e.g. 115V ~1%, 60Hz ~1%).
7. Set the power meter current range. The full-scale value selected multiplied by the crest factor rating (I<sub>peak</sub>/I<sub>rms</sub>) of the meter must be greater than the peak current reading from the oscilloscope (or digital multimeter).
8. Allow the readings on the power meter to stabilize (10 minutes) 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.
9. Record power measurements and total Rmax Peak for 10 minutes to calculate FLOPS/Watt for single node.
10. Calculate total FLOPS/Watt by multiplying single node FLOPS/Watt times number of “single nodes” as defined in Section 2.N.
11. Record the test conditions and test data.

### **Product Testing Documentation**

*Submittal of Qualified Supercomputer Data to Green500:* Submittal procedures will be integrated in the Top500 reporting procedures. Please follow the Top500 guidelines for reporting LINPACK

measurement and certification. Power measurement data (obtained using the aforementioned protocol) will be requested as part of the Top500 procedures.

4. **Effective Date:** Qualification for the Green500 will follow all deadlines applicable to the Top500 list. Green500 entries will remain active for the duration an entry remains on the Top500 list. As entries are no longer kept in the Top500 list, they will be dropped from the Green500.
5. **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 reward those manufacturers who have made efforts to further improve energy efficiency.
6. **Recommended Power Measurement Equipment:** Watts Up Pro: Digital multimeter with PC interface. See <http://www.doubleed.com> for details. **Add some illustrations using Watts Up Pro here...**