Overall comments:

I. My basic concern about the current document is that it assumes it is possible to have a single benchmark for meeting ENERGY STAR classification for all systems. The server market is very diverse. Servers do not run just one single workload like a monitor. Servers are designed with different workloads in mind...more disks for a database server; lots of memory for a scientific server; lots of processor cores for a web server. Servers designed for ISP providers are very different from those designed for database systems and those designed for high performance throughput.

Allow me to illustrate this problem with an example: consider a server designed for web-based workloads with a good SPECPowerjbb score in terms of energy efficiency. Unfortunatly, "efficiency" is not transferrable across workloads.
Someone running web-based workloads (like an ISP provider) would do well to compare, shop and purchase this machine or others with good SPECPowerjbb scores. Now consider a competing server designed for database transactions with lots of disks that has a slightly better SPECPowerjbb score than the first server. The ISP provider still does well to choose this second server. However, a large database warehouse chooses this second server for a 1000-machine installation. Unfortunately, the workload commonly used by this customer accesses disks constantly while the SPECPowerjbb benchmark did not. Alas, the server designer used cheap, power-hog disks to keep costs down. Thus, with the disks running all the time, the 1000-machine installation is terribly energy inefficient. This is an extreme example, but not unlikely to occur, especially if vendors attempt to design machines to meet the energy star requirement...itself a very likely scenario. This type of "cheating" with benchmarks is precisely why organizations like SPEC are attempting to create many different benchmark suites to classify systems based on workload.

My suggestion is that the EPA consider segmenting the server specification to reflect the marketplace. For example, the EPA could identify major markets that cover say 98% of all servers such as business class, web-server class and scientific class. Then the EPA could select a benchmark that best captures each market. Initially, using SPECjbb as proposed would be fine, but the EPA should classify this as representative of say business class machines (like those used by ISP providers). To me this would allow the EPA to make some progress on this effort, cover a large part of the sizeable server market, yet leave room for improvement and expansion.

II. Specific comments:

--On page 4 you describe the various configurations under consideration.
My advice here is to keep things small. As a co-founder of the Green500 List and the first researcher to measure component-level runtime power consumption of large scale machines, I can tell you first hand the methodology for measuring larger machines is quite different from that of smaller machines. For example, some larger machines use different types of multiphase power and measurement is difficult. Consider a single multiphase power connection that powers a full rack of systems, how do you separate out the power consumption of the network interconnect hardware?
We've come up with a workable methodology for Green500 (see http://www.green500.org/), however it is very different from the SPEC methodology as you can imagine. In this case, my suggestion is you simply focus on single system nodes for Energy Star since those are the basic commodity building blocks of large clusters of machines. For larger scale efforts, you could mention the Green500 as an approach and methodology to measure and encourage energy efficiency in larger systems. Here is some language you could use if so inclined, "We refer those interested in measurement methodologies designed to encourage energy efficiency in large scale systems to the efforts by Virginia Tech to establish a Green500 List of the world's most powerful and energy efficient systems. See http://www.green500.org for further information." If the ENERGY STAR program would like to endorse our Green500 List efforts, we would be honored.

--I agree with excluding DC power-based systems. There is no consensus whether these systems will become a large enough part of the market to warrant such inclusion.

--I agree with the focus on power supply efficiency. Our early research exposed terrible inefficiencies in power supplies that will critically improve energy use.

--On page 6, you mention "proposals or different approaches to determining efficiencies". Since this is only a note, its not critical to change the language, but I was not sure what was meant. If you are referring to measurement methodologies for server class systems, then feel free to mention Green500. However, Green500 is more of a measurement and benchmark approach.

--On page 6, RE: idle. This is a tough one. There is little agreement on what constitutes idle in a server. Though I think it is unfortunate not to tackle this now, I think I agree it would be best to table the idle discussion until another draft.

--On page 7, RE: Voltage Testing Requirement. This is back to my earlier suggestion that ENERGY STAR confine itself to servers for now since measuring larger systems is especially problematic as we have found in our experience.

--On page 7, RE: Utilization. Untilization is an overused term. It needs to be clearly defined here. I think you are using utilization = percentage of maximum load without reduced throughput. To me, it is very important that this is clearly defined since the term means different things to different readers.

--On page 7, RE: low power states. I think ENERGY STAR should absolutely recognize lower power states, but only in context. If the workload based metric used captures these power states, it should be used. If not, it should not be used. I think this is inextricably tied to the idle discussion and thus probably something else that should be tabled in this draft.

--The requirement to share SPECPowerjbb information is clearly in violation of the run rules established by SPEC. It's my understanding that the EPA would need SPEC approval to allow such reporting.

Thanks for considering my comments. Please contact me if you have any questions and let me know if I can participate further. I truly believe in this effort and am happy to play a part. 540-231-4238 (W)

Kirk W. Cameron
Associate Professor
Computer Science