

We had some further internal discussion regarding the new concept papers for servers (inofficial versions). Below you find three rather general aspects already briefly mentioned in the comments I have sent to Jan which again are summarised here as you might find them relevant for your discussions tomorrow.

I. Type and number of benchmarks:

The idea to work with adapted existing benchmarks (at least for a next step) sounds reasonable. However as we all know there is already a number of benchmarks available covering different workloads (web, CPU, database, email etc.) and we probably could imagine that industry will not be happy with declaring many benchmarks for all their products to meet Energy Star compliance (you might get a first impression on that tomorrow.....may be we are to pessimistic). Thus it may be fruitful to use the potential subsequent stage of benchmark evaluation and adaptation for a filtering of benchmarks in order to arrive at a minium number of most representative benchmarks. Maybe 2-3 benchmarks would be sufficient to cover the most important aspects which actually could mean 1-2 additional benchmarks to SPECpower for example. It might be possible to limit the set to a CPU oriented benchmark (e.g.SPECpower), an I/O oriented benchmark and may be a virtualization oriented benchmark (speculative statement to stimulate discussion :-))? We made the experience that manufacturers do not like to use established performance benchmarks (like SpecCPU) for general information purposes covering all their products. However this problem might be avoided by a slight modification of the existing benchmarks.

II. Consideration of virtualisation:

To take relevant market and technology development into account it may be also important to consider virtualization (as a rather important approach to save energy) as far as possible. Including virtualization into the tier 2 concept basically could mean to include an additional virtualization benchmark showing how efficient the server hardware is regarding operating typical workloads in a virtualised environment. Thus the benchmark would be a somewhat more "complex workload" covering the hypervisor and several standard workloads running in parallel. To our knowledge SPEC is currently also working on a virtualization benchmark. It may be worthwhile to check if this or other benchmarks already available in this context could be used for Energy Star (at least options could be considered for a tier 3).

III. Coverage of blades:

One other question may be if it is important to make blades 1:1 comparable to rack server technology in the next step. Blades are mainly bought for specific reasons regarding management etc. (scalability/easy extension, place requirements, easy wiring, lower management costs etc.). Blade systems are less appropriate for small companies operating only very few servers. Overall there are areas of application where blades make more sense than rack servers and vice versa. Overall it would probably be adequate and easier for a next step to keep the two technologies separate concerning mandatory quantitative Energy Star criteria but to allow a comparison via the information/reporting criteria (energy consumption per blade/per chassis etc.).

more comments on the current papers will follow next week.

good luck for the meeting and

best regards
bernd

Dr. Bernd Schäppi
Austrian Energy Agency