



September 28, 2012

To: Energy Star Product Team

From: Chris Bergen, AppliedMicro

AppliedMicro appreciates the opportunity to submit comments for Energy Star's specifications for Computer Servers.

In October 2011, ARM® announced the ARMv8 instruction set and AppliedMicro revealed X-Gene®, the world's first high performance, 64-bit ARM powered, server-on-a-chip product family purpose-designed for the cloud. Our objective in creating X-Gene was to dramatically lower energy consumption in the large-scale data center. We applaud Energy Star's efforts to provide a consistent, transparent basis to evaluate energy consumption under reality-based server workloads and to provide a framework to encourage conservation.

Section 1 Definitions

Lines 41-48. Section A) 3) attempts to define a blade system. While we agree with the definition and agree that sections a) (1) through a) (4) represent the state of the industry today, the same sections may be construed as constraining the definition of a blade system to these specific form factors. We recommend that they are qualified as examples of prevalent blade server form implementations.

Lines 79-84. Section A) 6) attempts to define the emerging server form factor, where a large number of relatively smaller and lower-power processing units in a scalable architecture. This section could benefit from a clear definition of a "node." We would like to propose that a compute node be defined as a processor (or cluster of coherent processors) with its directly attached memory (DRAM) and network interface. Supporting logic (e.g., voltage regulators, boot flash, sensors, and/or management entities such as a BMC) may be per-node or shared across nodes. However, each node represents an independent functional compute server and, as such, may have a directly attached storage device or could be "diskless." This definition of a node can apply to any form factor. Once a node is defined, we agree with the definition that a "multi-node server" is thus two or more independent server nodes that share a single enclosure or blade and one or more power supplies, and that, in a multi-node server, power is distributed to all nodes through shared power supplies. However, the statement that "Server nodes in a multi-node server are not design to be hot-swappable" is not pertinent to the definition of a multi-node server and may unnecessarily constrain the definition.

Line 119. Section C) should also reference multi-node servers as a separate form factor.

Line 298. Section 2.1.1. should include reference to multi-node servers.

Lines 421 – 499. It is not clear if section 3.6 or section 3.8 is intended to apply to multi-node servers.

Line 458. Table 4 provide an allowance for additional IO devices. It is not clear if those additional devices must be pluggable (e.g., PCIe slot) or can be permanently soldered “down” on the board, or integrated into the processor itself. We believe that the budget for additional devices should be the same regardless of the implementation.