EMERSON
Network Power

Pasi Vaananen
Systems Architect
Embedded Computing

Emerson Network Power
46 Lizotte Drive, First Floor
Marlborough, MA 01752 U.S.A.

T + 1 (508) 357 8237
F + 1 (508) 357 8204
pasi.vaananen@emerson.com

March 20, 2009

Rebecca Duff,
ICF International

Dear EPA,

Emerson Network Power, Embedded computing division, designs and manufactures computer server equipment, specifically targeting the needs of Communications, Military, Aerospace, Government, Medical and Automation market segments.

We recognize that many of these market segments utilize server type equipment in significant quantities, and share the common goal of deploying Energy Efficient equipment to reduce associated system operational expenses and carbon footprint. As such, we do support EPA Energy Star specification for servers.

We would like to thank EPA for the opportunity to provide comments on the draft specification.

As a manufacturer of both AC and DC powered server category equipment, many of which can be qualified under EPA Energy Star requirements for “Blade Systems”, we have reviewed the present draft (Draft version 4) of the specification and associated set of normative documents.

We provide the enclosed comments to help ensure that the typical server equipment used by these markets can also be qualified in the context of the Energy Star for Servers specification program requirements and specified test processes.

This document may be published on the Energy Star for Computer Servers specification web site.

We hope that this feedback is useful for finalizing the EPA Energy Star Specification for Servers, and successful inclusion of “Bladed Systems” within the first release of the specification. Please contact me, if you have any questions with respect to our comments.

Sincerely,

Pasi Vaananen

EMERSON
Network Power

Summary

These comments are provided to:

- Help ensure that EPA plan to include DC-powered servers, and particularly typical DC-powered “Blade Systems” utilizing distributed power conversion architecture can be successfully tested for compliance with Energy Star for Servers Specification.
- That the EPA test requirements for the Energy Efficiency are aligned with other relevant Energy Efficiency and installation environment specifications and closely associated with the characteristics of installation environments of these systems during normal, long term operation.

These comments are arranged by issue, identifying the associated document(s) and affected specification sections / line numbers, as applicable.

Clarifications needed for the tests for Blade Systems with distributed powering architectures

Figure 1, below is a high level representation of the typical “Bladed System” server, employing the centralized DC-power distribution scheme. By “centralized DC-power distribution”, we mean that the “48V” power enters to one or more of the power supply modules, typically configured as a Field Replaceable Unit (FRU), which then converts the 48V to at least one intermediate voltage which is used to power a set of one or more “Blade Servers” hosted within the “Blade Chassis”. “Blade Servers” are connected to one (or two in case of redundant feed support) of the internal intermediate distribution voltages generated by the 48V/IVB power supply. The salient features of such system are exactly like the AC powered systems, particularly with the respect to the test requirements.

The power supply efficiency associated with bladed systems utilizing variations of this power distribution topology can be easily tested utilizing the specified test procedures in the present EPA draft specification and associated EPRI test procedures document, which is incorporated by reference. This is due to the fact that the power supply efficiency requires the access point to both equipment input power feed, and power supply output(s) for the load testing, and as the power supply is configured as a FRU, it can be relatively easily removed and tested as an individual assembly.

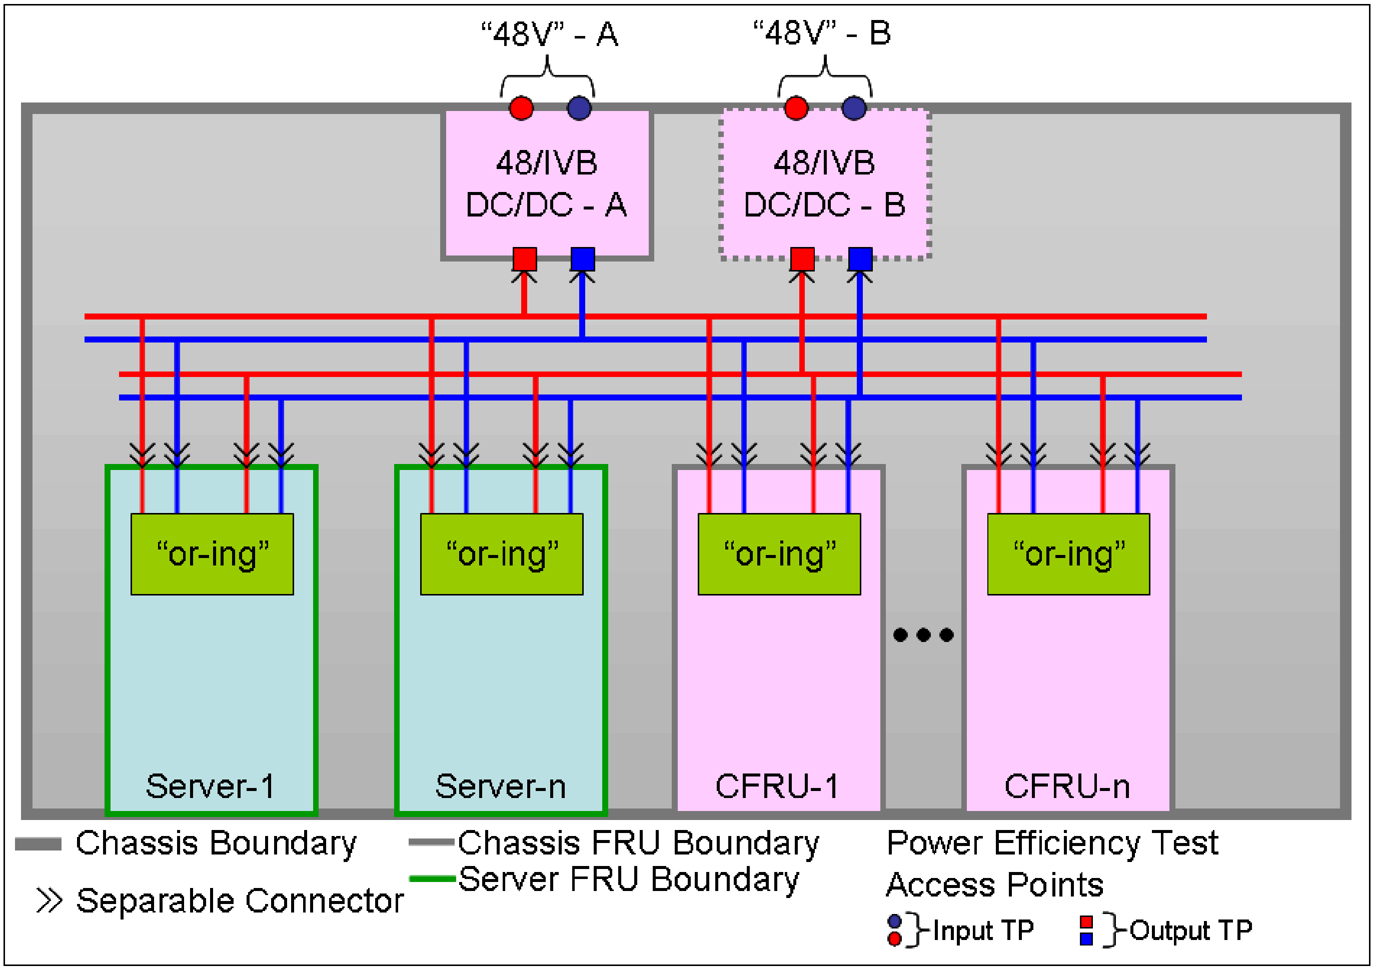

Figure 1 – Centralized DC powering architecture in Blade System

Figure 2, below is a corresponding high-level representation of the power distribution architecture that is typical on 48V DC powered systems. As compared to the centralized DC-power distribution architecture described above, and depicted in Figure-1, this topology has certain important differences that are relevant to the power supply efficiency test procedures:

- 48V power enters the system through power feed(s) connected to one or more Power Entry Modules (PEM), which typically (but not always) are configured as FRUs
- Power entry module may provide some limited functionality, such as shelf level over current protection and/or power feed filtering
- Power entry module does NOT convert the voltage to any intermediate voltage, i.e. the output voltage of the PEM is directly proportional to the input voltage (possibly with some small losses associated of protection/filtering circuitry within the PEM).
- 48V feed from PEM is directly distributed to the units that consume energy, such as "Blade Servers" and possibly any number of supporting FRUs enclosed within the "Blade Chassis".
- "Blade Servers" contain the DC/DC converters that convert the nominal "48V DC" to at least one intermediate voltage (typically, but not always 12V DC), which then feeds a set of the Point-Of-Load converters or other loads within the "Blade Server" FRU.
- "Blade Server" (and other, if applicable) FRUs may connect to two independent 48V DC feeds to support the power supply infrastructure redundancy.

- DC/DC power converters themselves are typically not redundant, the independent input feeds (if available) are combined utilizing various possible circuit designs to essentially providing or'ing function to two feeds, providing the valid "48V" input voltage for the converter in the presence of at least one of the input feeds.
- DC/DC power converters are hosted by "Blade Server" FRU, and may be specified according to the specific requirements (voltage, power levels and efficiency curve) of the given "Blade Server". As such, these may be the same or different within a family of the "Blade Servers", depending on the selections and tradeoffs made during the "Blade Server" FRU design.

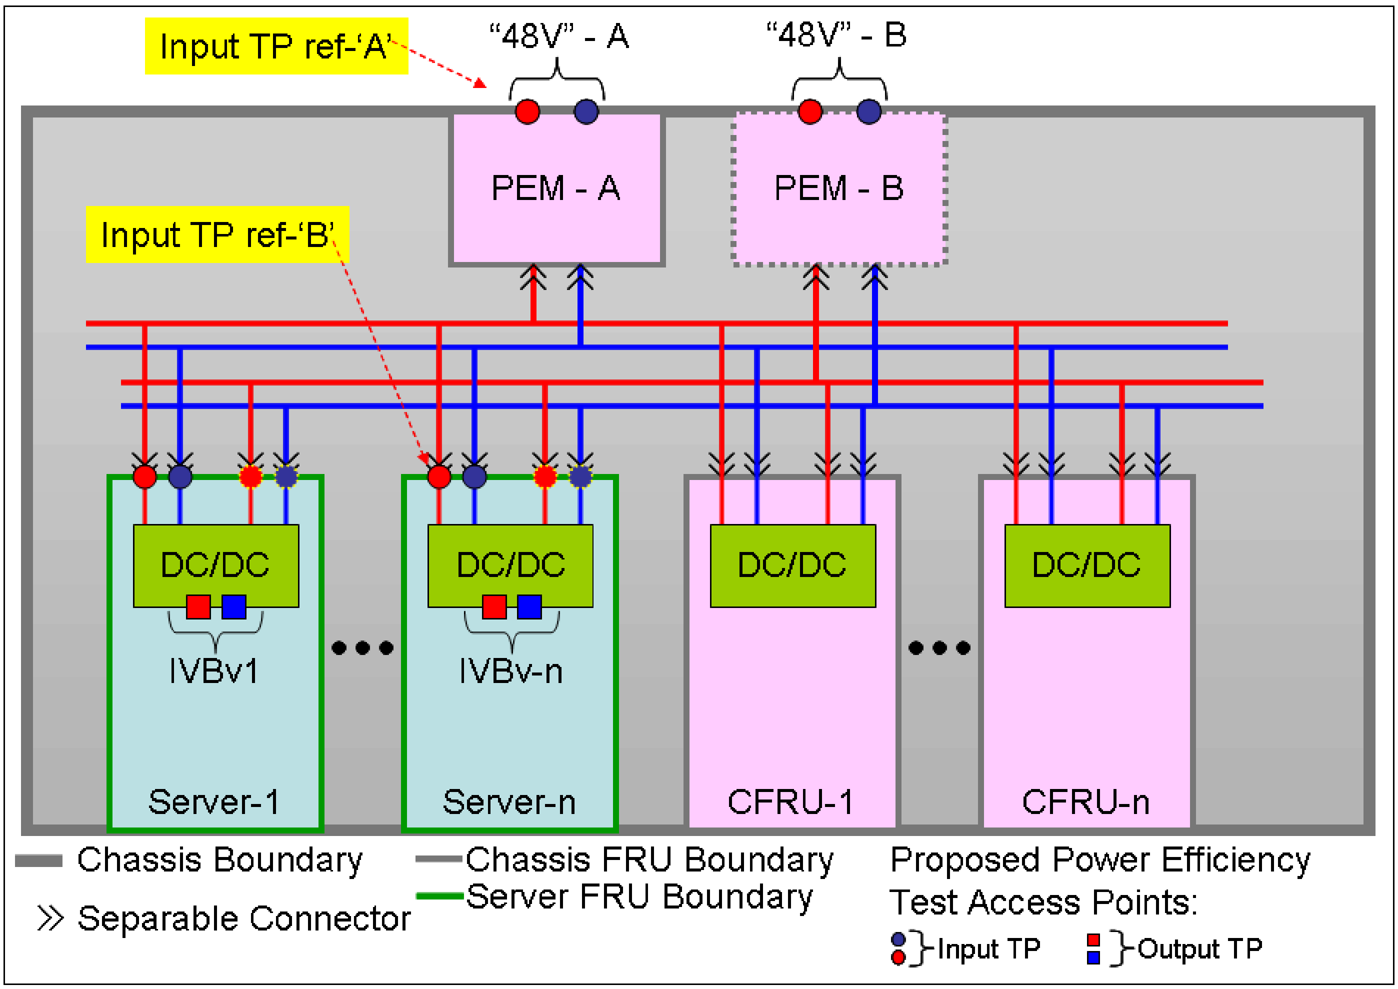

Figure 2 – Distributed DC powering architecture in Bladed System

This distributed conversion architecture has certain benefits over the centralized architecture described above with respect to its energy efficiency characteristics:

- Power supply capacity modular and increases naturally with load due to architectural coupling of the power supply with the load (vs. design for max case in centralized architecture). This can lead to better match of the power supply capacity to actual load, potentially resulting in higher energy efficiency, particularly for partially populated "Blade System" configurations, as the architecture allows the converter efficiency optimization around the known, more narrowly varying load range.

- As the power supplies themselves are typically not redundant, the load point can also be more easily matched to the range of the efficiency curve that yields to highest sustained efficiency (vs. 50% or lower point in centralized system utilizing two centralized 48V/intermediate voltage supplies).

Unfortunately, this commonly used DC distributed powering architecture represents some challenges with respect to power supply efficiency test procedure prescribed in the present EPA draft. The following comments are provided to identify these challenges, and provide input on how each of these can be addressed in the context of the language provided in the present draft specification.

- There are two potential input test access points: equipment input (points in PEM in Figure 2, above), and "Bladed Server" FRU inputs. EPA should clarify whether the converter needs to be tested from equipment input (reference point 'A'), or "Blade Server" input (reference point 'B'). It would be significantly simpler to test from reference point 'B', and we would therefore prefer this to be specified as reference point.
- If EPA decides that the distributed converters MUST be tested from the input reference point 'A', then EPA needs to clarify that all other loads than the converter under test are to be disconnected from the system prior to the test. Otherwise, it is impossible to meaningfully determine the efficiency of the converter under test.
- EPA should clarify the test procedure for equipment supporting redundant power feeds. We suggest that the efficiency testing for the redundant, distributed powering architecture should be only mandated with only one power feed active, as this represents the worst case. Testing with two active feeds should be left optional.
- As briefly described above, the converters in "Blade Servers" may be equivalent or different. EPA should clarify that the particular power converter needs to be tested <u>only once, even if used in differing</u> "Blade Server" designs. This is essentially equivalent to how the converters in the centralized architecture are tested, with the exception that if different converters are used, each configuration needs to be independently tested.
- In addition to distributed converters on "Blade Servers", there are typically additional converters hosted within "Blade Chassis" FRUs, such as in fan modules, management modules, etc. EPA should clarify if these converters need to be also independently tested for efficiency, and if EPA requires them to be independently tested, we suggest that EPA would limit the efficiency requirements to apply to only converters over some prescribed wattage (such as 20W) to reduce the associated test effort needed. Depending on a design, there may be quite many small converters to support powering of management infrastructure, which do not consume much power in the context of the overall "Blade System", and are not easily accessible for testing as they typically constitute of components directly soldered to FRU PCBs. As the efficiency of these small converters does not represent major energy efficiency improvement opportunity, testing of such converters should be exempted from the requirements.
- There is no clearly defined accessible standard output test access points for efficiency testing of the distributed converters associated with the "Blade Servers" on distributed conversion architecture, as these converters are typically not designed as FRUs. However, nothing in this architecture would preclude the system to meet the power supply efficiency requirements specified by EPA if the test access can be accomplished. As such, EPA should clarify that the distributed converters can be tested according to the EPRI procedure, as long as the associated converter output(s) can be isolated from the "Blade Server" load for the purpose of testing, and the converter can be brought to its normal operation state using any accessible interface, as needed. The specific methods to accomplish this may include use of sockets, connectors, etc. but should be left out of the normative test procedure specification.

EMERSON
Network Power

“48V” DC Power Supply Test Voltage

Emerson Network Power, Embedded Computing is aware of the prior comments (for draft 3) that have been made with respect to the test voltages used in “48V” DC power systems, and presently stated EPA position on the test voltage.

We would like to suggest that EPA should re-evaluate their stated position with respect to this issue. This is due to the majority of the installed “48V” systems operating from / optimized for battery plants, with “Float” voltage significantly different that the presently specified 48VDC +- 1% test voltage.

As the goal of the EPA Energy Star specification is to work to *improve* the energy efficiency of the computer server equipment, due consideration should be given to the typical equipment installation environments. Otherwise, on the worst case, this efficiency test specification might lead to optimization of the power supply efficiency at the voltage that is *different* than the intended equipment typical installation environment; with reduced efficiency in real installation as compared to if the test specification would be aligned with the typical installation environment. This would obviously not serve anybody’s interest, nor would this be good for the goal of the improving the energy efficiency.

The existing energy efficiency specifications all recognize the typical installation environment to be around 53V. The overview of some of other associated efforts with references is given below:

- ANSI ATIS-0600315-2007 defines the typical voltage as 53.0, nothing that “The value is a compromise between equipment operating with VRLA and flooded technology battery.”
- Verizon VZ-TPR-9205 (http://www.verizonnebs.com/TPRs/VZ-TPR-9205.pdf#page=9) defines the voltage as -53.0 +/- 0.25V to be used when measuring energy efficiency.
- The -53 +/- 0.25V value was used in early drafts of the ATIS energy efficiency standard, but that was later widened to -53.0 +/- 1V. This is now reflected in the ANSI family of standards for telecom energy efficiency (ATIS-0600015.2009, ATIS-0600015.01.2009, etc.); see https://www.atis.org/docstore/product.aspx?id=24547 and https://www.atis.org/docstore/product.aspx?id=24548.
- AT&T's energy efficiency requirements are to follow the ANSI/ATIS requirements; this is indicated in AT&T's TP-76200 document (https://ebiznet.sbc.com/sbcnebs/Documents/ATT-TP-76200.pdf#page=99), which indicates the -53V +/-1V test condition.

Therefore, we propose to align this specification with the characteristics of the majority of typical installation environments, like most other relevant test specifications do, and propose that EPA specifies the test voltage for “-48V” DC systems as -53V ±1V.

This proposed change should be applied to:

- EnergyStar® program Requirements for Computer Servers Draft4, Section 4A, table after line 1033, 2nd column, line #4
- EnergyStar® program Requirements for Computer Servers Draft4 Appendix A, table after line 1266, 3rd column, line #4

Recognizing that EPRI efficiency test document is not under the direct control of EPA, we suggest that EPA would pass the resolution on test voltage as input to EPRI. Regardless of EPRI resolution,

EMERSON
Network Power

EPA should write a statement within ESTAR Draft4, section 4A that incorporates the EPRI document by reference. Suggested example wording for such statement is as follows:

“Nominally 48V DC powered systems SHALL be tested according to EPRI Generalized Test Protocol for Calculating the Energy Efficiency of Internal AC-DC and DC-DC Power Supplies, Revision 6.4.2, with the exception that the test voltage used should be 53VDC±1V instead of 48VDC +-1%.” Alternatively, language that <u>allows</u> testing at 53V level would represent a good compromise – this could be accomplished either by expanding the % tolerance band associated with the test voltage or providing the alternate test voltage specification.

Test and Reporting Temperatures

Again, recognizing that the goal of this specification is to promote more energy efficient designs, the test and reporting temperatures should be aligned with the “normal” operation conditions, not the worst case conditions, where the equipment may be subjected during certain abnormal, but typically short-term operation conditions.

Test report template, which is incorporated as a normative reporting by the EPA request for comments cover sheet document, and specification (lines 828-833) requests the ASHRAE Delta-T, and minimum and maximum airflows to be reported at the “Peak Temperature”, which is stated to be 35°C.

We have identified several issues with these reporting requirements:

- There are no supporting test procedures or reference to test procedures for these reporting requirements within the specification. If this information is supposed to be supported by test data for compliance, then the associated test procedures need to be specified. Note that the ASHRAE does not provide test specification for these reporting requirements.
- 35°C is only one of the temperatures specified by relevant ASHRAE documents. This temperature is specifically associated with ASHRAE Class-2 maximum operation temperature. Is it the goal of EPA to preclude the inclusion of the systems that are specified to operate on more controlled data center environments (i.e. ASHRAE Class-1 with 32°C maximum operation temperature)? If not, allowances need to be made to provide data for the class that system is actually designed to comply with.
- In virtually all present systems, whether “Bladed Systems” or “Volume Servers”, the fan speed is dynamically controlled based on the function of the environmental conditions and often also depends on the offered system computational load. This is done to reduce both fan energy consumption and acoustic noise emissions. As such, it is not possible to meaningfully determine the “minimum fan speed at peak Temperature”, as requested.
- The 35°C is maximum temperature for ASHRAE Class-2, which is almost never representative of the equipment inlet temperature of the equipment. The representative input temperature that most facilities are trying to maintain for equipment is to keep temperature within the “recommended” operation temperature range. In addition for this temperature not reflecting the intended long-term operation point that is typical to the facilities, elevated temperature will generally yield to increased power dissipation due to increase on the fan power dissipation due to need for increased airflow to sustain operation at the elevated temperatures (fan power required is proportional to cube of the airflow), and/or the increase of the power dissipation of the silicon chips operating at increased junction temperatures due to increase in the leakage power. EPA should recognize these facts, and preferably adjust the test and reporting requirements to reflect the “recommended” system ambient climatic conditions as specified by ASHRAE (for either class-1, or class-2), or at least add the corresponding reporting requirements in addition to maximum.

Note that our position with respect to reporting temperatures is aligned with the ASHRAE position. The accurate reference to ASHRAE document in question, which is referenced on the reporting template as “Information (as reported in ASHRAE thermal guidelines)” is “Thermal Guidelines for Data Processing Environments”, ASHRAE, 2004, ISBN 1-931862-43-5. This document, in section-5, which contains the equipment manufacturer reporting requirements, specifically requires that the information is provided in ambient temperature between 20°C and 25°C (representing the recommended range at 2004 version, it should now be aligned with extended range referenced below). The referenced ASHRAE document requires both nominal and maximum temperature airflow and power dissipation data for “minimum”, “full” and “typical” configurations. EPA should either reference the above ASHRAE document directly or provide the equivalent level of clarity on the reporting requirements.

There are also other differences between EPA and ASHRAE reporting requirements you may want to address, including:

- ASHRAE requires airflow pattern reporting, EPA does not
- ASHRAE requires dimensions and weight reporting, EPA does not
- EPA requires Delta-T reporting, ASHRAE does not
- ASHRAE provides additional clarifications on the reporting data and conditions

As a summary, the reporting temperatures in “Information” section should either be aligned with, or amended with the updated AHSRAE specification, i.e. System ambient dry-bulb temperature used for airflow and Delta-T reporting should be within range of 18°C to 27°C.

For reference, the recently updated ASHRAE climatic specifications with extended “recommended” temperature range can be found at:
http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf