ECODESIGN REQUIREMENTS FOR SERVERS – FROM SINGLE PRODUCT GROUPS TO EXTENDED SYSTEM APPROACH

Nils F. Nissen [1], *Lutz Stobbe* [1], *Thibault Faninger* [2], *Anton Berwald* [2], *Klaus-Dieter Lang* [1,3]

[1] Fraunhofer-Institut für Zuverlässigkeit und Mikrointegration, Gustav-Meyer-Allee 25,
13355 Berlin, Germany, corresponding author: nils.nissen@izm.fraunhofer.de
[2] BIO by Deloitte, 185 Avenue Charles de Gaulle, 92200 Neuilly-sur-Seine, France
[3] Technische Universität Berlin, Forschungsschwerpunkt Technologien der Mikroperipherik,
TIB 4/2-1, Gustav-Meyer-Allee 25, 13355 Berlin, Germany

Abstract: The currently ongoing ENTR Lot 9 preparatory study for ecodesign of enterprise server and storage equipment raises a number of methodological questions as the investigation requirements shift further from of a pure product scope to extended system analysis. The paper will give insight into the current status of the study. However, the core part of the paper will focus on the methodological aspects of addressing ecodesign on an individual product level while at the same time reflecting the extended system. The paper will explain the ecodesign implications of the various system levels, i.e. the sub-system (IT product) and extended system (support infrastructure) relationship. We then explore the options on how to cope with extended system aspects on a product scope level. The analysis will then examine existing and needed test standards, which are especially crucial for sub-system requirements.

1. INTRODUCTION

The environmental performance of complex products cannot be assessed completely by just analyzing and measuring a single piece of hardware in the lab, because the interaction with other system parts such as network elements or with infrastructure installations such as air conditioning has considerable impact on the product's performance including energy consumption. This was already the case with all networked devices or with modular installations such as machine tools or electric motors. In a generalized way we call this the ecodesign dilemma of sub-system scopes, where assessments focus on a partial system without modelling all influencing system parts.

The ENTR Lot 9 preparatory study for implementing measures of the Ecodesign Directive (2009/125/EC) started in June 2013 [1]. The product scope of this study was defined as "enterprise servers, data storage and ancillary equipment" according to the latest working plan document of the European Commission [2]. This product scope covers a rather broad spectrum of different products in terms of functionality and performance.

Moreover, enterprise servers and storage are typically operated in dedicated server rooms or data centers in order to fulfill quality of service (QoS) requirements with respect to availability, safety, and security. The term "ancillary equipment" in the full study title suggests this extended system character under which enterprise server and storage equipment is deployed. Ancillary equipment may indicate very different products in that respect. For example ancillary equipment could be the standard 19-inch server rack cabinet or a power distribution unit (PDU) sold separately from the actual server. During the first phase of the study the "ancillary equipment" has effectively been narrowed down to "network equipment" used directly in the server vicinity following clarification discussions with the representative of the European Commission supervising the study. The "extended system" in contrast may include the racks or the cooling systems not delivered by the server or storage system providers.

Against that background, the ENTR Lot 9 preparatory study has also been met with a number of challenges deriving from the required implementation of the new methodology for the ecodesign of energy-related products (MEErP) [3]. This method is an updated version of the MEEuP used in earlier studies, and again needs to be followed precisely to make results from different studies comparable. The MEErP consists of a logical structure of seven subsequent tasks defined as:

1. Scope (product definition, functional unit)
2. Market (analysis of sales and stock)
3. Users (use patterns, extended system)
4. Technology (HW/SW properties, BAT)
5. Environmental & Economics (eco-impact)
6. Design options (implementing BAT)
7. Scenarios (quantifying the improvement)

One set of problems derives from the purely sequential work plan as iterations are often necessary. This is especially true for the scope definition – unless the study deals with a product group which is already well defined and has a broad consensus. A considerably more complex problem is the so called "extended system approach", which has been introduced with this new methodology and has been incorporated as early as Task 3 on use patterns and user behavior. The challenges are on the one hand setting boundaries to that "extended system" and on the other hand making a clear distinction between what we like to call a "product functional unit" and an "extended system benefit".

In this paper we will discuss the product-system-interaction and its implications of eco-design measures on the example of ENTR Lot 9 product scope. We are first going to shortly describe the functionality and properties of current enterprise servers and storage equipment. Then we will be analyzing the developments towards interconnected product systems that create a "system benefit" under certain conditions, which are not reproducible under other boundary conditions. Finally, we will discuss the implications of the user's need for a system benefit with respect to the ecodesign requirements of an individual product.

The authors of this paper are the consultants that conduct the ENTR Lot 9 preparatory study. However, the contents of this paper are in no way representing the view of the European Commission. The intention behind this paper is to provide a scientific foundation for the further development of methodologies for the assessment of ecodesign in highly complex systems such as networked information and communication technology.

2. PROPERTIES OF ENTERPRISE SERVER AND STORAGE EQUIPMENT

Enterprise server and storage equipment are configured to provide a defined computing service or task specific application via a network connection to one or many remote clients. The storage system is basically only an extended memory for the server. Sometimes the storage is directly attached to the server (and therefore only accessible via the server) and sometimes connected via a more complex network.

What is important to understand is the fact that both product categories – server and storage – provide not only an individual functionality but through collaboration and deployed control a larger "system benefit". Connectivity and virtualization are essential elements in this functional system of servers and storage equipment.

System is the keyword that will be discussed throughout this paper. A "system" in the general sense can be defined as a set of independent elements (components) which are deliberately structured and interact to form an integrated whole or functionality. A system is usually defined by its structure, interrelationship of elements, and its boundaries and respective surroundings.

With respect to enterprise server and storage equipment the term "system" in practice is applied to various technical levels including:

- Component configuration on product level – including e.g. the type of processor and chipset or the storage media and capacity
- Modularity – a solution for optimizing server and storage scalability on the base of prefabricated network, power and cooling capacity
- Connectivity – defining not only bandwidth and latency but the location and efficiency of the interface control
- Software on product level – including type of operation system, respective performance of application software
- Virtualization – including the option of creating virtual entities, shifting loads and improve utilization of existing resources
- Climate on rack and room level – including thermal conditioning in order to ensure reliable operation
- Power supply – including power distribution and power conversion AC/DC as well as DC/DC

Keeping in mind that the ENTR Lot 9 preparatory study has the objective to investigate and assess the environmental impacts and improvement potential on both the product level and the extended system level, the first task is to clearly differentiate between product and system.

3. DEFINING THE PRODUCT

Enterprise servers have been defined as fully functional computing equipment, offered as a server product for professional use in a business (enterprise) environment. Enterprise servers are a complex product category. Individual server products are defined by various aspects including:

- Compute architecture and respective operation system – x86 architectures are dominating the volume market
- Form factor and number of nodes – e.g. rack-optimized, blade system, multi-node, micro-server
- Number and type of processor – power consumption of the complete server is directly scaling with the number and type of the CPU, as the CPU is often the single largest power drawing component
- Memory and storage configuration – performance and power consumption is scaling with memory
- Resilience and level of redundancy – this translates into additional hardware intensity and respective resource consumption

A screening regarding the main environmental impacts of enterprise servers has indicated that the power consumption in the use phase is responsible for more than 80% of the total product carbon footprint over an average four year product life time. This understanding sets the primary focus of ecodesign on the issue of optimizing the energy consumption in the use phase. The second priority is the efficient resource utilization and a feasible recovery of the large spectrum of valuable metals and other materials.

Data storage equipment shows a similar prioritization of environmental improvement necessities – energy efficiency before resource efficiency. The vast majority of enterprise storage is “always on” or as the storage networking industry association (SNIA) defines it online, near-online or removable (off-line) [4]. SNIA defines product categories by reflecting on the response time (e.g. online with a maximum response time of <80ms), storage capacity, application and indirectly the cost factor of a product.

Enterprise servers and storage equipment are assemblies that mainly contain the active and passive electronic components (e.g. CPU, RAM) on a printed circuit board (PCB), storage devices such as hard disk drives (HDD) or solid state devices (SSD), network interfaces and connectors (e.g. LAN, FC), as well as integrated cooling elements (e.g. heat pipes, fans) and power supply units (PSU). All this is fitted into a chassis or enclosure for different mounting options. The form factor describes dimension, modularity and mounting options of the product.

The boundary of the chassis (enclosure) seems to be sufficient for distinguishing individual products. Unfortunately the reality is more complex and poses some challenges with respect to a precise product (scope) definition. The challenge is the modularity and the multiple configuration options in which the products are either delivered as a unit or are combined with products delivered by other vendors.

To give an example, server and storage systems most often provide multiple configuration options including type and capacity of the storage media, the processor and memory configuration, as well as the power supply configuration. All these aspects are influencing the performance and the respective resource and power consumption. Let’s take a closer look into the environmental implications of modular product types with various configuration options.

4. EXAMPLE: MODULAR PRODUCTS

Modular server systems have the advantage of shared resources and prefabricated connectivity. Due to the requirement of providing scalable server capacity with exchangeable server and storage modules new form factors with modular enclosures have entered the market. Modular server systems such as blade systems or micro server systems consist of two housing elements. One element is the individual chassis (enclosure) for the server modules. A server module is basically the populated server boards with the processors and memory as well as connectors for power and network. These server modules are called blade, cartridge, or book, depending on the manufacturer’s preference.

A certain number of server modules are inserted into a larger system enclosure, which is housing the shared resources including network interfacing, additional storage devices, active cooling and power supply devices. The system enclosure is designed to fit into the 19-inch server racks. The market shows a high diversity in terms of dimensions, form factors, and system configurations. Sub-systems such as the power supply units may be delivered by different companies and installed by the operators on site.

To make these configurations comparable (as far as possible) benchmarks for servers have evolved over time. While for comparing raw compute power obviously maximum configurations are measured (such as the Linpack benchmark used for the TOP500 supercomputer lists [5]), environmental benchmarks for smaller enterprise servers such as SPECpower [6] are often measured with stripped down configurations. Omitting auxiliary features and minimizing redundancy often gives a better ratio of computations per watt consumed. However, for the environmental assessments in the current study the benchmarking results must be typical across all products delivered. What can be measured by a suitable benchmark can then be the basis for environmental differentiation of products and potentially for the setting of ecodesign criteria.

Using a blade center as one specific example case, in effect ecodesign requirements could address the various system levels encountered before:

- the main components: processors and the integrated storage (RAM, HDD, SSD),
- the single blades,
- one blade center unit,
- a full rack (possibly including units, which are not blade units) or
- a number of fully or partially filled racks, depending on how the setup is sold on the market.

Whether two racks filled with four blade center units are considered one product or four is not the main obstacle (by normalizing statistics to the number of processor cores) but finding the right benchmarks to aggregate product impact and show improvement potentials for the products “as delivered” is near impossible.

5. ADDING THE EXTENDED SYSTEM

So far we established that enterprise servers and storage equipment are defined by a certain set of hardware and software elements that are integrated in an enclosure. This enclosure could stand alone (pedestal) or be mounted in a cabinet (19-inch rack). The rack and the server room (data center) in which the products are mounted create a second and third system level (boundary). The environmental issue here is not the description of the physical structure. The cabinet and server room are also providing a necessary infrastructure for a safe, secure, and reliable operation of the server and storage equipment. Thus evaluating the product level but omitting infrastructure preconditions for the actual intended use might lead to very biased comparisons. The more essential infrastructure elements are:

- power distribution with reliable (uninterrupted) power supply,
- cooling, and air conditioning.

These essential infrastructure elements have a direct interaction with the IT equipment primarily in terms of energy consumption in the use phase. One example that we will explore shortly in this paper is the air inlet temperature level and therefore the amount of resources and energy used for the thermal conditioning (cooling) of the server room on the one hand and the amount of resources and energy used for dealing with higher inlet temperatures on the server or storage (product) level. Conventional server and storage equipment are typically operated in air-conditioned rooms at an average range of 20°C to 27°C. Higher or lower temperatures and respective humidity levels have had negative effects on the reliability of electronic and electromechanical components. However, air-conditioning is energy and resource intensive. Many examples in the past years have demonstrated considerable energy savings on a system level (data center) when the temperature envelop was pushed somewhat further (e.g. free cooling with inlet temperatures >27°C).

In conclusion, the environmental benefit of extended system (e.g. overall reduction of energy consumption on the data center level) could result from an increased material and energy consumption of the server and storage equipment (product level). The individual products certified for higher temperature operation might have to use more or more resource intensive materials for the thermal management.

Vice versa, improvements on the product level could lower the efficiency of the extended system, as these trade-offs may be highly non-linear.

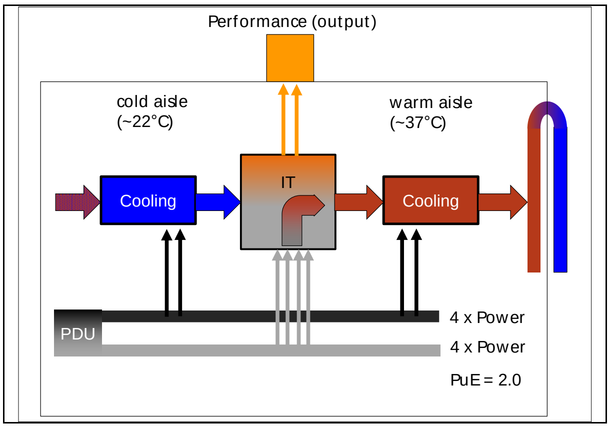

Figure 1: A simplified extended system view showing the PuE metric.

In the case of data centers the metric most often used to describe the efficiency of the infrastructure setup is the power usage effectiveness or PuE. The PuE of a data center operation is the total energy consumption divided by the energy allocated to the IT equipment, namely servers and storage systems in this study.

While this allows comparisons of basic data center layouts and of specific data centers of similar performance, it does not allow for improvement effects on the IT side. Lowering the IT power consumption while at the same time even increasing the computational output will lead to a worse PuE rating, if the cooling system is load dependent.

The efficiency and the adaptability of the infrastructure is thus essential for the extended system characterization, but cannot be measured with this metric, even though it is well established and ostensibly targets the same system level.

The extended system components are vital for the operation of the product, but add significantly more degrees of freedom to the problem description.

6. TEST STANDARDS

For comparing servers on a product level SPECpower has been the most prominent benchmark for a number of years. A closer look at the SPECpower test results and particularly the power consumption values in the partial loads reveal a certain mismatch. The measured power consumption is considerably lower in comparison to calculated power consumption based on individual component data such as the thermal design power of the CPU. This observation is confirmed by stakeholders. According to these sources, server devices under test have been specifically modified with a minimum hardware configuration for the SPECpower benchmark. The resulting power consumption is significantly lower than a typically or even maximally configured server.

This is one reason why the U.S. EPA Energy Star Program [7] in close collaboration with leading equipment manufacturers took the initiative and asked SPEC to create the SERT – server efficiency rating tool [8], a more comprehensive and by that more realistic set of tests. SERT considers the following basic aspects:

- Five different hardware configurations from minimum to maximum configuration and power (e.g. considering processor, memory, storage and interfaces).
- Thirteen different workloads (called worklets) for processor, memory, storage, hybrid (SPECpower equivalent) and idle.
- Distinction of partial loads including idle, 25/50/75/100 percent load.

Through this complex distinction of configuration and worklets it is now possible to get a more realistic understanding of the power consumption in active states. Eventually this should lead to better market information and to better distinction of Energy Star certified servers. First results of SERT data that have been made available to the study confirm that the power consumption e.g. in a typical configuration is higher in comparison to the SPECpower benchmark results. However, the detail discussion about the proper creation, documentation, transfer, validation and interpretation of SERT data is still ongoing.

Through the choice of the worklets best representing the actual or expected user workload, the SERT data offers the possibility to better match the most efficient hardware to the user’s needs. As has been introduced before, the system benefit is not fully covered by describing “computational power” and “electrical power”.

The user is additionally defining the availability of the product within a service environment. This determines the configuration of the product and the operating conditions as much as the expected workload. The following aspects have been identified to be critical (positive and negative) in that respect:

- Redundancy of components and potential over-provisioning of hardware elements,
- Additional reliability safeguards that add resource use on the product or extended system level,
- Wider thermal operating conditions (e.g. according to frameworks set by ASHRAE [9,10]),
- Sub-optimal utilization of available capacities.

The last point on utilization ratios is in fact the second biggest improvement lever of the last years via server virtualization and consolidation (the biggest still being the advances of processor efficiencies).

7. EFFICIENT UTILIZATION AND CONSOLIDATION BY VIRTUALIZATION

Efficient server utilization, hardware consolidation, and overall energy saving in a server room or data center are going hand in hand [11]. Efficient utilization of managed servers and storage equipment as well as associated network technology is effectively realized by a combination of system-level load management and to a smaller extent product-level power management. However, there is no clear distinction between both aspects.

Nowadays enterprise servers are fulfilling their intended purpose mostly in networked systems that combine multiple physical server and storage units in one virtual entity. Due to this development it is difficult to allocate the useful work (functional benefit of the energy-using product) solely to a single hardware element (server or storage product in a sense of a functionally complete device). To combine multiple (physical) server and storage devices in a large virtual entity is creating a considerably potential of energy and material reduction. In a virtual environment consisting of multiple networked devices it becomes possible to shift (virtual) servers and IT-loads between different hardware elements (physical servers).

By now it has been shown that virtualized workloads can in extreme cases be shifted around the globe not only according to availability of compute power, but also according to lowest CO_2 emissions or best local cooling conditions. The actual utilization ratio of smaller servers without virtualization is still very low at around 10-20%. Consolidating four of such workloads into one physical server clearly shows the improvement potential still possible in many cases.

Returning to the central question of system assessment, the virtualization options are not trivial to integrate either in the “per product” calculation or in the “extended system” calculations.

8. CONCLUSION

The primary environmental impact derives from the actual utilization of the products and the resulting electrical energy consumption in the use phase. Existing product carbon footprint (PCF) of enterprise servers, conducted by leading manufacturers and environmental assessment experts, indicate that over the whole product life cycle the use phase is dominant with more than 80% of the total PCF. The following aspects can be considered in that respect:

- The products are typically operated 24/7 for a period of three to eight years (average four years).
- Server utilization is currently still quite low at about 20% load on average.
- Server utilization is slowly increasing on average, primarily driven by virtualization.
- The performance increase of processors according to Moore’s Law combined with parallel processing and more power management is still leading to substantial increases in energy efficiency.
- Energy efficiency has been addressed strongly by the industry over the past five years in conjunction with SPEC, SNIA and TEER benchmarks.

The infrastructural operating conditions including cooling infrastructure and energy supply are directly related to the power consumption of the products. It is necessary to consider energy trade-offs between the ICT equipment and the supporting infrastructure. The best energy savings result from a system approach.

The energy consumption in the use phase is influenced by the application and performance requirements of the equipment. In other words, the specific hardware and software configuration determines the power consumption. This includes not only the power draw of the electronics but also the dimensioning and conversion efficiency of the power supply unit, the cooling requirements, as well as energy and material overheads for the required cooling.

Reliability and QoS requirements so far can only be addressed through more complex product categorization (such as distinguishing resilient servers from normal servers), as the added energy and resource consumption would otherwise penalize these products compared to simpler products. Using the delivered application specific compute power as normalization for functional units is therefore not going far enough.

Adding the extended system scope for a diverse product group, such as enterprise servers and storage, is certainly an important step forward. Exploring the interdependencies systematically and with quantifiable data is still a big challenge and an open research topic of a broader scale than the current study.

9. REFERENCES

[1] Project website for ENTR Lot 9 Server Study: www.ecodesign-servers.eu ; accessed 28th July 2014.

[2] Establishment of the Working Plan 2012-2014 under the “Ecodesign Directive”, Final document, 7th December 2012. Available here: http://ec.europa.eu/enterprise/policies/sustainable-business/documents/eco-design/working-plan/files/comm-swd-2012-434-ecodesign_en.pdf

[3] Methodology for the Ecodesign of Energy-related Products (MEErP), Study report, December 2011. Available here: http://ec.europa.eu/enterprise/policies/sustainable-business/ecodesign/methodology/index_en.htm

[4] SNIA Storage Networking Industry Association., http://www.snia.org ; accessed 28th July 2014.

[5] Top500, The Linpack Benchmark http://www.top500.org/project/linpack/ ; accessed 28th July 2014.

[6] Standard Performance Evaluation Corporation, SPEC Power Committee http://www.spec.org/power/ ; accessed 28th July 2014.

[7] Energy Star, Enterprise Servers Specification Version 2.0 http://www.energystar.gov/products/specs/enterprise_servers_specification_version_2_0_pd ; accessed 28th July 2014.

[8] Standard Performance Evaluation Corporation, Server Efficiency Rating Tool (SERT) http://www.spec.org/sert/ ; accessed 28th July 2014.

[9] ASHRAE TC9.9, Data Center Networking Equipment – Issues and Best Practices http://tc99.ashraetcs.org/documents/ASHRAE%20Networking%20Thermal%20Guidelines.pdf ; accessed 28th July 2014.

[10] 2008 ASHRAE Environmental Guidelines for Datacom Equipment http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf ; accessed 28th July 2014.

[11] Schlösser, A; Stobbe, L.; Nissen, N. F.; Schäppi, B.; Lang, K.-D., “Energy Efficient Network Equipment for Data Centers”. *Electronics Goes Green 2012+*, Berlin, September 9-12, 2012. Available at: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6360440 ; accessed 28th July 2014.