DIGITALEUROPE Analysis on ErP Lot 9 Idle Power Framework and Allowances for Enterprise Servers

Brussels, 31 March 2017

Table of Contents

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

Summary

DIGITALEUROPE has serious concerns with the proposed base idle state power allowances and the additional idle power allowances for extra components. These allowances come directly from an initial draft proposal made by EPA for ENERGY STAR server specification which have no strong foundation prior to wider stakeholder detailed comments and analysis including that of Industry. As proposed in the Draft Lot 9 working document, the idle power threshold proposal will drive an estimated 10-30% net increase in data center energy use and consumption (Section 4).

DIGITALEUROPE, in conjunction with the Green Grid, has completed an analysis of the relationships between idle power, server active efficiency[1], deployed power[2] and system launch year. The Green Grid SERT Analysis Working Group (TGG) has a data set of 71 one and two processor socket rack server products, with each server product having reported SERT data for three to five configurations. 262 separate configurations have measured idle power values and SERT metric scores. The data set also has 26 blade and 6 resilient two socket product families. The data analysis shows that:

1. Average server idle power has stayed largely flat by configuration over the past 6 years (since 2011) while average performance has steadily increased with each new generation of products. Server manufacturers have been increasing the performance capabilities of server products while maintaining a relatively constant power envelope (Section 1).
2. Higher performance servers, using more capable processors as measured by core count and frequency, will result in lower power consumption in the data center under both active and idle conditions since the single higher performing server has replaced several lower performance servers (Section 2).
3. Many higher performance servers have higher idle power than lower performing servers. A regulation based on the idle power limit and associated allowances will be biased against higher power, higher performing servers and further "tightening" of an idle limit will largely increase the number of high performing servers being excluded from the market (Section 3).

The proposal by DG GROW to define an energy efficient server that's solely based on its idle power measurement will exclude many high performance servers, as measured by workload delivered per unit of energy consumed, from the EU market. Excluding higher performing servers will mean deployment of lower performing servers in the data centers which will result in net increase of data center energy use, the opposite of the desired outcome of the ErP Lot 9 directive. Several use studies have been cited to justify the use of idle power as an efficiency indicator, citing the low utilization of many servers, particularly in offices and remote

---

1 As measured by the SPEC Server Efficiency Rating Tool weighted geomean active efficiency metric. (SERT metric).

2 Deployed power is the power of the number of servers required to deliver a specified workload in the data centre. A server's performance capacity is the weight geomean of the geomean of the interval performance measurements for each SERT worklet. The weighted geomean of the performance is divided into a specified performance value to determine the number of servers needed to do the work. The deployed power is calculated by multiplying the number of servers needed to do the work by the weighted geomean of

2 Deployed power is the power of the number of servers required to deliver a specified workload in the data centre. A server's performance capacity is the weight geomean of the geomean of the interval performance measurements for each SERT worklet. The weighted geomean of the performance is divided into a specified performance value to determine the number of servers needed to do the work. The deployed power is calculated by multiplying the number of servers needed to do the work by the weighted geomean of the worklet power measurements.

"data closets". Focusing on the individual idle power of an individual server is not an appropriate way to assess how much energy will be used in the data center.

Servers are sized to manage an expected peak workload, recognizing that there will be periods of time when servers are only partially utilized. With the use of virtualization technologies, workloads can be combined onto a single server to increase their utilization, thereby reducing the time they are at low utilization or idle and consequently reducing the number of servers required to do work. Higher performance servers typically have better utilization characteristics, enabling them to consolidate a greater number of individual applications onto a single server. This situation is applicable to both enterprise servers managed in a data center environment and servers installed in an office environment.

Further, tower servers, which are the server class largely expected to be installed in an office environment, represented 20.4% of the market in 2015 and are projected to represent only 16.3% of the market in 2020.[3] Rack server sales are declining as companies increasingly choose to purchase software services through platform or software as a service offerings from cloud computing companies. It is not logical to establish server requirements based on perceived concerns –i.e. servers sitting idle in office environments - when (1) total energy use is dictated by the number and size of the machines deployed to manage the workloads and (2) the servers installed in office environments is trending downward and will represent 16.3% of the market in 2020, around the time the regulation is expected to become effective. Contrary to the comments made by ECOS that "many servers still spend most of their time in an idle state" and "particularly in small and medium server rooms and colocation data centres", the data collected by DIGITALEUROPE confirms that the market in the scenario described by ECOS represents a minority of the high volume servers sold into the EU market. Colocation and enterprise data centres are complex and support a broad scope of applications and services and a generic statement as described would be inappropriate and not representative.

For these reasons, DIGITALEUROPE continues to recommend that server efficiency should be assessed using the DIGITALEUROPE Server Efficiency Metric (the SERT™ metric or weighted geomean efficiency metric).[4] DIGITALEUROPE and the Green Grid (TGG) have collected a large dataset of SERT™ results and have shown that the metric is effective in measuring server efficiency. DIGITALEUROPE believes that sufficient data has been collected to establish an effective SERT metric threshold for server energy efficiency to remove the least efficient servers from the market and driving reduced energy consumption in the data center (we estimate an energy reduction of 33-78% in power consumption compared to a business as usual case for worst performing servers, see Section 4).

If DG GROW proceeds with its plan to set a server efficiency threshold based on an idle limit, the information presented here demonstrates how higher performance servers would be unfairly and unjustifiably excluded with the potential net impact of higher energy consumption by data centers. Although there is no approach based on idle power limits that would entirely avoid this issue, DIGITALEUROPE has set out below several important changes that need to be made to the idle allowances to lessen the impact:

1. There needs to be an idle allowance for system performance. A low end processor with minimal cores and a lower frequency and a less energy intense supporting infrastructure – system chips, circuitry etc.

---

3 IDC Worldwide Quarterly Server Tracker 4Q2016, March 1, 2017. Referenced with permission from IDC.

4 "DE PP on server efficiency metric from SERT worklet results 20161214.pdf"; Posted to DIGITALEUROPE website on December 14, 2016.

- generates a lower idle power than a high end processor with maximum cores and frequency and its supporting infrastructure. The higher idle power of approximately 20 watts will deliver performance increases of 200% or more while reducing data center energy use because few higher performing systems will need to be deployed to deliver a given workload. TGG has developed a performance based idle power allowance for each server category identified in the Lot 9 regulation, which when combined with other allowances, eliminates more of the lower efficiency servers as measured by the SERT metric, than the Draft Lot 9 proposal. Incorporating the system performance adder enables a reduction of the base idle allowance to 15 W for a one socket server and 30 W for a two socket server, reduced from the 37 and 85 watts in the Draft Lot 9 document (Section 4).

2. The memory idle allowance can be reduced to 0.175 W/GB. Unlike the ECOS evaluation on one data point from a manufacturer's on-line product power calculator, TGG has evaluated over 25 data points from two technologies (DDR3 and DDR4), 3 process technology nodes and three manufacturers to justify the lower adder. A summary of the data is provided below (Section 5).
3. The allowances for the storage devices need to be broken into 6 categories by drive type, drive speed, and connection/controller type. TGG has evaluated 284 measured data points to identify the categories and the recommended idle limits. A summary of the data is provided below (Section 6).
4. Idle power allowances for I/O devices need to be added for 25 to 50, 50 to 100, 100 to 200, and greater than 200 Gb/s devices and the allowance for devices with 10 to 25 Gb/s needs to be increased to capture new technologies being released into the market over the Tier 1 period (Section 7).

These changes will better represent the idle power characteristics of the server systems and their associated components than the Lot 9 proposal. They are also based on an extensive set of measured or publically available data for components and systems. The suggestion by ECOS that "there is no contradiction between low-power idle and high-performance" is incorrect. Following that, the statement "the proposed Ecodesign framework is performance-based, with categories and adders that scale power allowance with performance" is also incorrect and it is exactly for this reason that DIGITALEUROPE proposes these changes. **DIGITALEUROPE notes that even with these changes, the revised idle power limit will still result in the exclusion of some high performance servers from the EU market even though these servers would yield a lower energy use than some systems that** would pass due to lower idle power (see Figs 11 and 12 in Section 3 below).**For this reason, DIGITALEUROPE maintains that idle power limits are not an effective or justified means of assessing the energy efficiency of servers.**

In addition, DIGITALEUROPE is proposing to modify the HPC systems definition to include servers optimized for artificial intelligence and machine learning. (Section 8).

The specific data analysis performed and the associated conclusions are presented below.

Section 1: Analysis of the change in idle power and the SERT metric over the past 7 years

IDLE POWER: Figures 1 and 2 depict the average idle power for each configuration type graphed against the year in which the rack server product was put on the market. The graphics illustrate that server idle power has remained largely flat from 2010 to 2015. Server manufacturers have continued to innovate to maintain a relatively constant power envelope while increasing generation to generation performance. As such, idle power is a poor characteristic against which to assess server efficiency as it is constrained by the inherent characteristics of the server components and their need to be on-line and communicating with the system with some minimum power draw in order to maintain system functionality and desired response times. Rather than tracking a server's ability to do more work for each unit of energy consumed, a tightening of the idle power limit will tend to increase the number of high performing servers excluded from the market forcing end users to purchase larger numbers of lower performing servers and increasing the energy demand of servers in the EU.

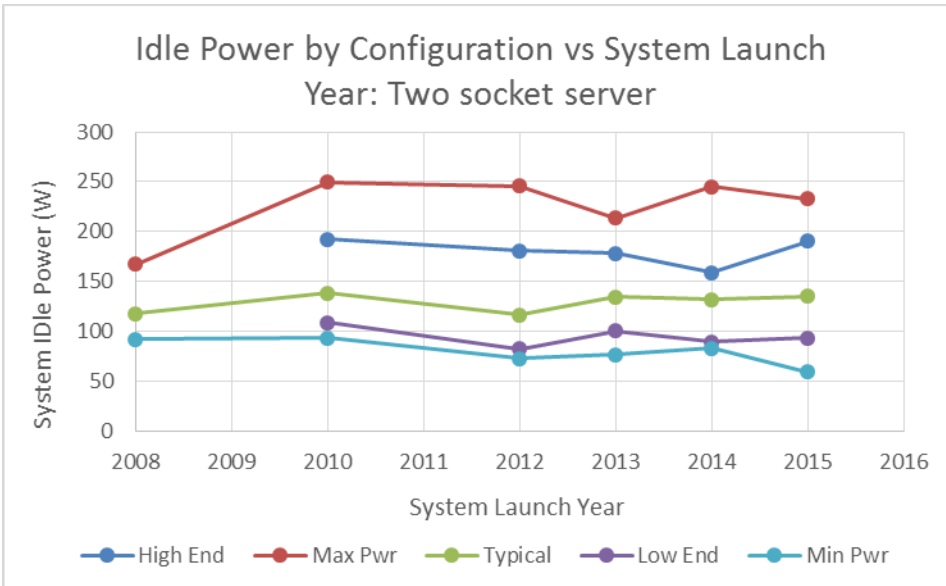

Figure 1: Average server idle power by configuration type for two processor rack servers.

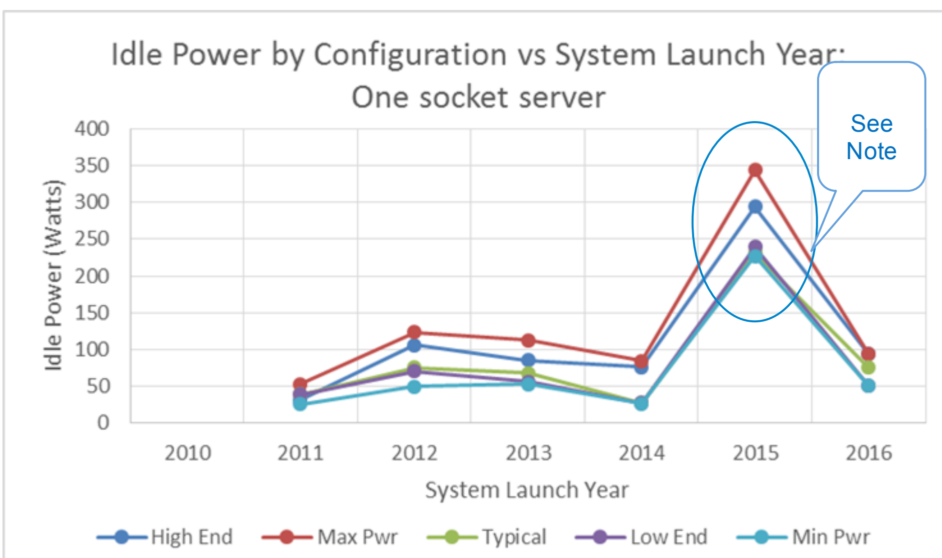

Figure 2: Average server idle power by configuration for one socket processor rack servers.

Note: The 2015 launch year has only a single product which is not representative of the typical x86 servers present in the other launch years, showing higher idle power and lower performance characteristics.

SERT METRIC: Figures 3 and 4 illustrate that average server active efficiency has steadily increased from 2010 to 2015. While manufacturers have largely kept power demand relatively constant with each new generation of products, they are increasing the performance capabilities of the server. Improving the performance in turn improves the active efficiency as measured by the workload delivered per unit of energy consumed. Figures 3 and 4 demonstrate that the active efficiency metric provides a means to assess real improvement in server efficiency. The active efficiency improvements with each new generation in the two socket rack server category are very clear. They are less clear in the one socket category, but the general trend is toward an overall improvement in workload delivered per unit of energy consumed with each new generation product.

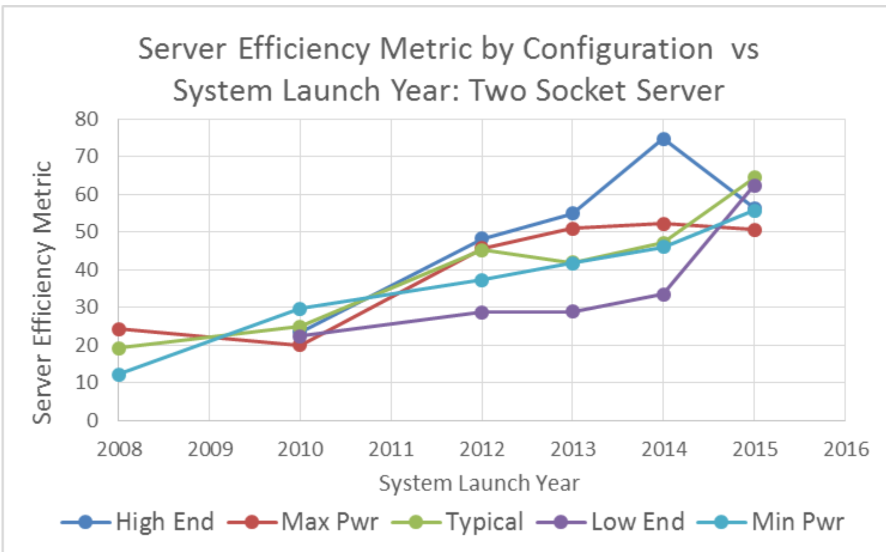

Figure 3: Weighted Geomean Efficiency by configuration and launch year for two socket rack servers.

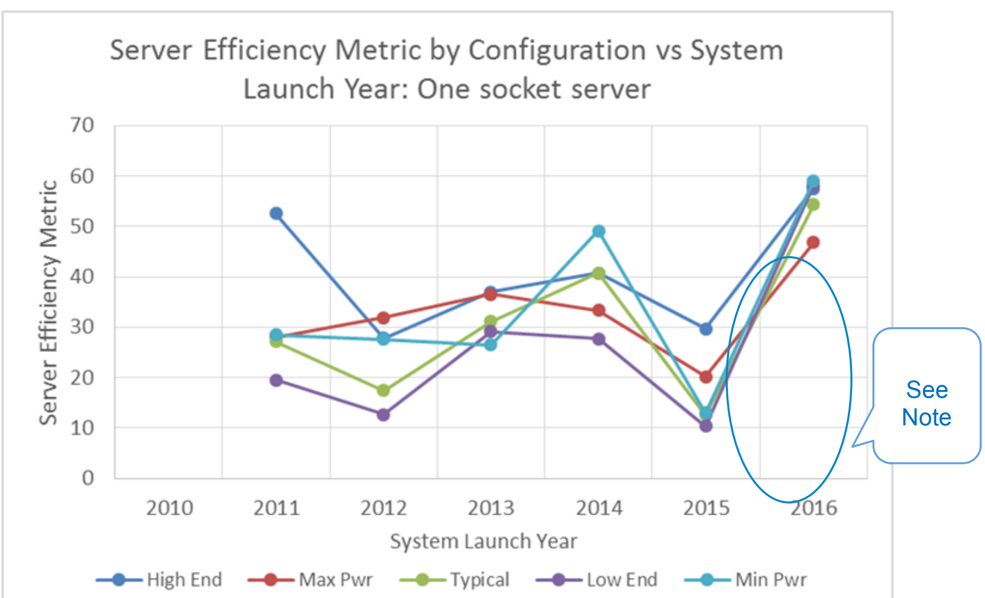

Figure 4: Weighted Geomean Efficiency by configuration and launch year one socket rack servers.

Note: The 2015 launch year has only a single product which is not representative of the typical x86 servers present in the other launch years, showing higher idle power and lower efficiency characteristics.

Section 2: evaluation of the idle power versus the SERT metric and deployed power

Given the power, performance and active efficiency metric trends discussed above, it is now appropriate to examine how server efficiency relates to idle power, i.e., do servers with lower idle power have higher active efficiency as measured by the SERT weighted geomean metric and do they result in lower deployed power in the data center?

Figure 5 provides a graphical comparison of the SERT weighted geomean active efficiency and the idle power of all of the one and two socket rack server configurations in the dataset. Examination of the figure offers the following observations:

1. Servers with an idle power in the middle of the idle power distribution generally have the best active efficiency score.
2. There are a set of server configurations which have low active efficiency across the full range of measured idle powers, indicating that an idle power limit will allow a significant number of low efficiency servers, as measured by workload delivered per unit of energy consumed, into the EU market.

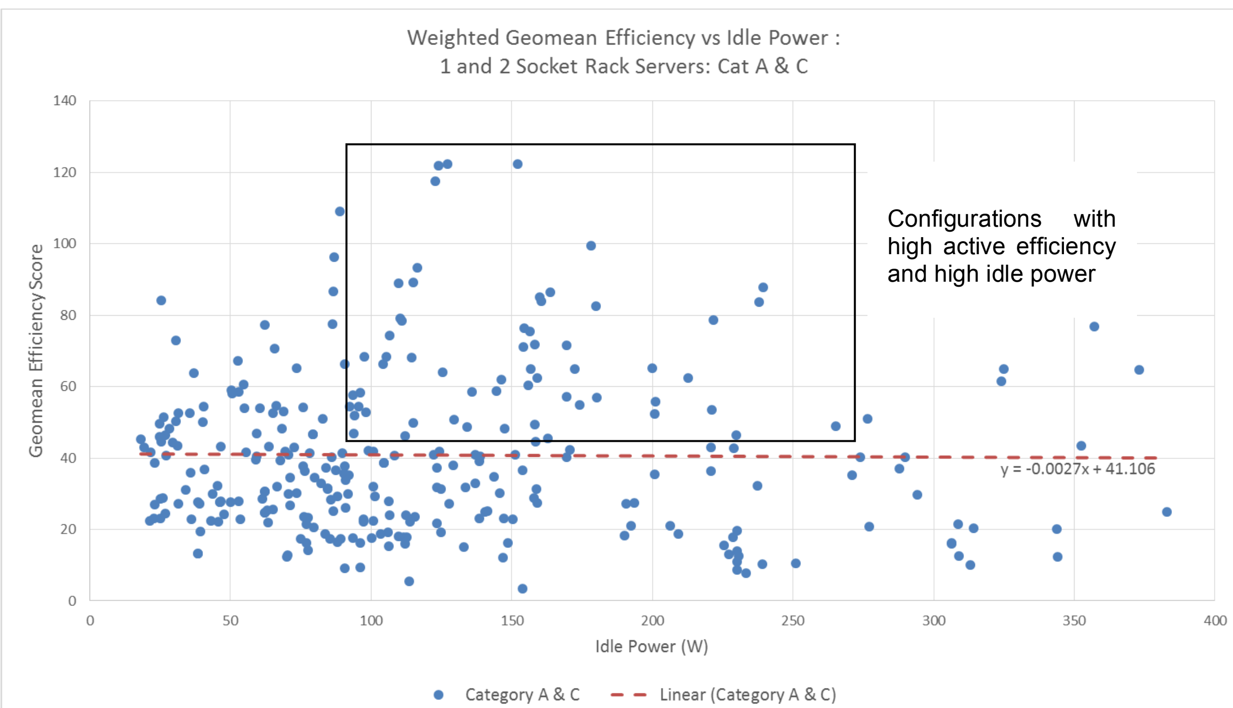

Figure 5: Comparison of the Weighted Geomean Efficiency Score to the configuration idle power

The graph above demonstrates that there are server products that will have much higher weighted geomean efficiencies but also higher idle power. Early work done by TGG demonstrated that idle power had a negative correlation of 0.2 to 0.35 to the CPU worklet efficiency score, indicating that there is a poor correlation between the two values.[5] Some of those servers will be excluded from the market with an idle power limit based solely on the criteria and allowances detailed in the Draft Lot 9 regulation. More data will be provided on this point in a discussion of Figures 7 and 8.

Having looked at how idle power relates to active efficiency, it is important to evaluate how idle power relates to the deployed power, the power demand of the number of servers of a given configuration required to deliver a defined workload. A server with a lower performance will require more servers to deliver a specified workload than a server with a higher performance capability. Figure 6 plots the deployed power against idle power for the one and two processor socket rack servers for a workload equivalent to 100 times the capability of the highest performing server in the data set.

1. Evaluating deployed power, Figure 6 demonstrates that there is a range of servers with idle power levels from 20 to 300 watts that can deliver that workload efficiently (below 2,000,000 W deployed power). The choice of the server size and configuration will depend on the workload type, the level of reliability desired and other policy requirements set by the data center operator. However, the data suggests that there are acceptable configurations in a range of performance capabilities to get the necessary work done.

2. Establishing a straight idle limit will indiscriminately exclude both low and high performance systems with higher idle power demand. It is likely that some servers used efficiently in highly virtualized or compute intensive environments will be excluded.

3. High-end configuration, higher performance servers generally have higher idle power but lower deployed power[6] when compared to lower performance servers.

5 “Analysis of the Server Efficiency Rating Tool: Implications of Server Configurations and Components on SERT™ Efficiency Results”, The Green Grid, November 2015, page 18.

6 Deployed power is calculated by (1) dividing the weight geomean of the SERT performance scores into a large workload value to calculate the number of servers need to perform that workload and then (2) multiplying the number of servers required to perform the workload times the weighted geomean of the power use at the 25% workload interval. See the document “Deployed Power Analysis Workload Descriptions” submitted to DG GROW on August 11, 2016.

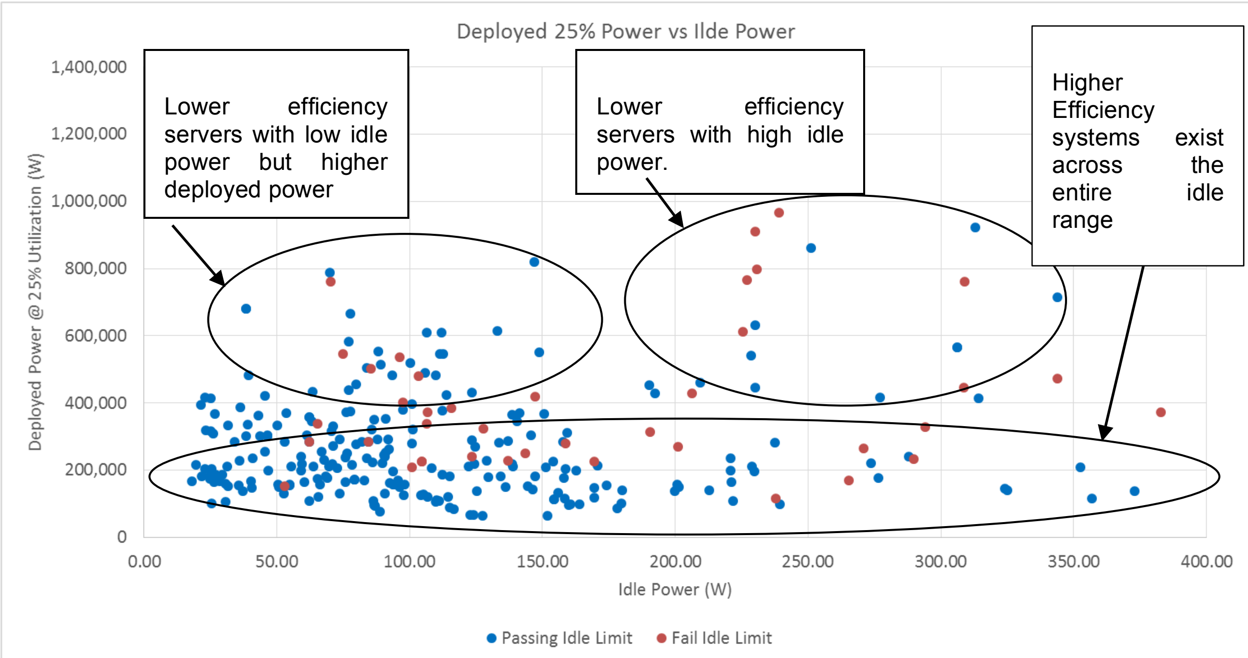

Figure 6: Idle power versus deployed power of the servers needed to execute a defined workload.

Section 3: servers eliminated by idle power and SERT metric thresholds

From a practical standpoint, creating a regulatory system based on idle power limits will result in the exclusion of high performing, more efficient servers from the market. These servers will be excluded at both the low-end and high-end configurations. Let's look at the data, using Figures 7 through 10. All evaluations are done using the one and two socket rack servers (21 and 44 product families respectively) from the ITI/TGG SERT metric data set.

Draft Lot 9 Idle Limit: First consider low-end configuration two socket servers. Using the idle power base and component allowances, Figure 7 illustrates which low-end two socket rack server configurations will fail the idle test (points marked in red). Under this scenario, 11 of the 44 server products have one or two configurations that fail the Lot 9 idle test. The base idle power allowance was adjusted to 70 Watts to achieve a failure rate of 25% of the systems, slightly higher than the 20% failure rate for the Draft Lot 9 85 watt limit. The higher failure rate was selected to allow comparison to the DIGITALEUROPE proposal for the addition of the system performance allowance and the modification of the storage and memory idle power allowances. Both low-end and high-end systems with much higher efficiencies and lower deployed power values fail, while less efficient systems with higher deployed power demands passed the test. A further reduction of the idle limits to fail 50% of the server products (not shown graphically), results in two thirds of the 11 additional failed systems having active efficiency scores of 20 to 50 (representing the more efficient systems).

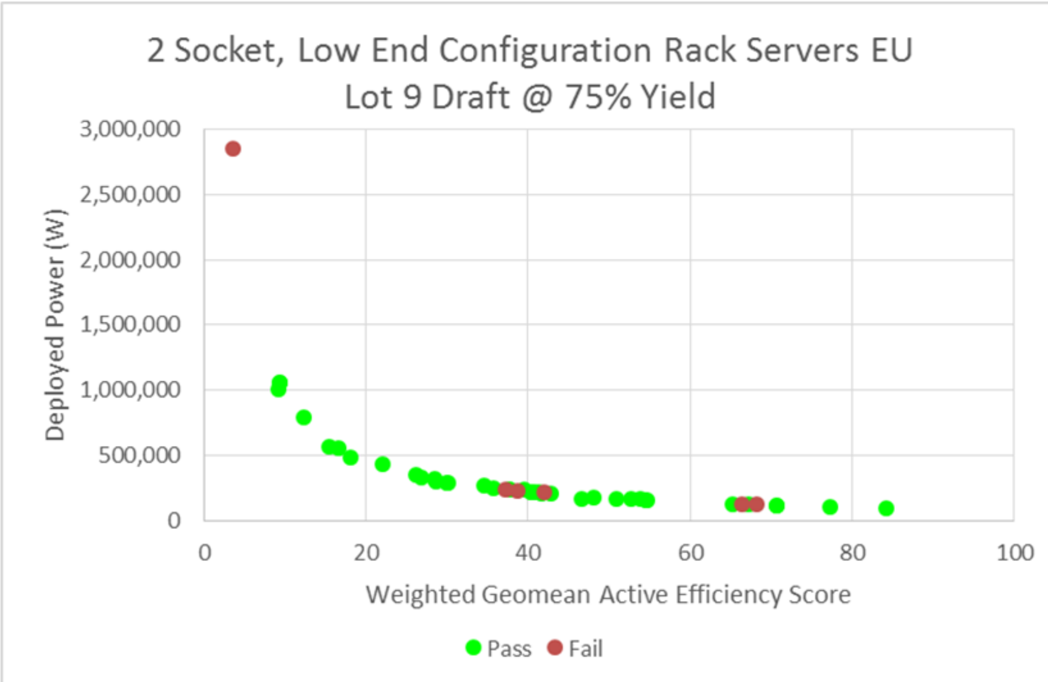

Figure 7: Low-end Configuration two socket rack servers which fail Lot 9 Idle power limits

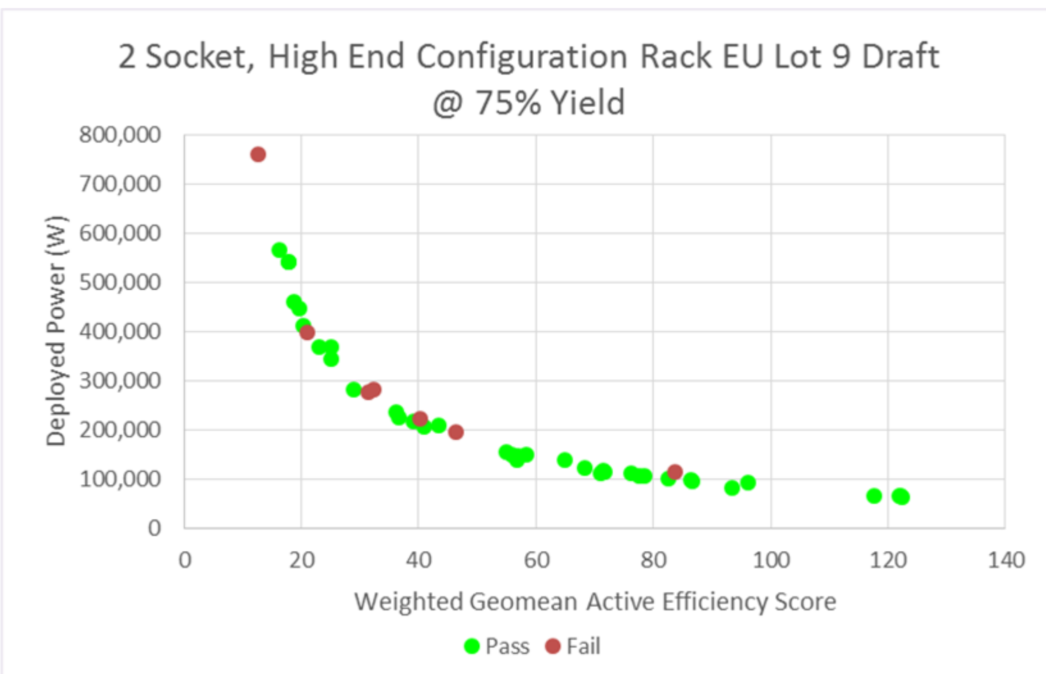

Figure 8: High-end Configuration two socket rack servers which fail Lot 9 Idle power limits

SERT Metric Threshold: DIGITALEUROPE has consistently advocated for the use of the active efficiency metric to assess the energy efficiency of servers. Setting a threshold of 20 for the active efficiency metric on the same dataset of 44 servers result in 75% of the systems passing. In this case, the failed systems occur at the lowest active efficiency score and the highest deployed power values – truly eliminating the products which will increase power use in the data center. It also does not discriminate against higher power, higher performance systems.

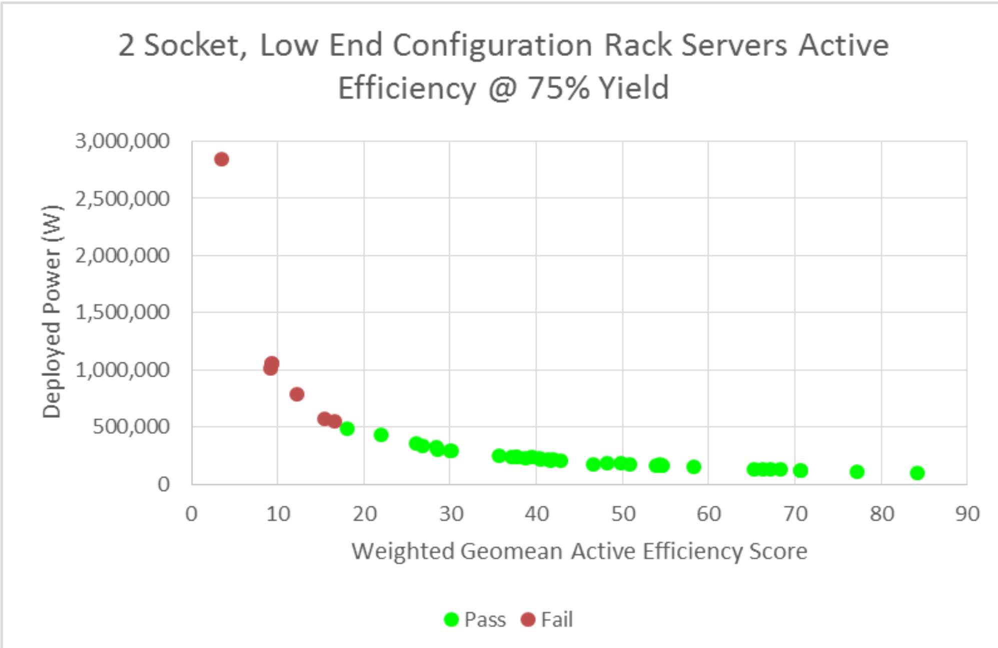

Figure 9: Low-End Configuration two socket rack servers which fail the SERT Metric Threshold

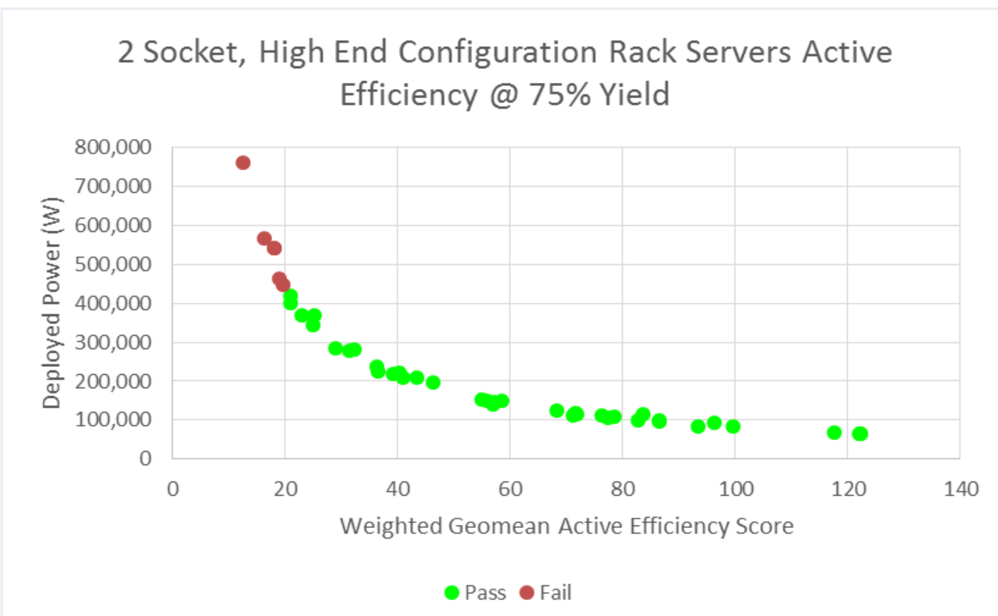

Figure 10: High-end Configuration two socket rack servers which fail a SERT Metric Threshold

Section 4: DIGITALEUROPE Recommendation for the Modified Base Idle Allowances

Though DIGITALEUROPE strongly opposes use of idle power to assess server efficiency if DG GROW insists on establishing idle power limits, DIGITALEUROPE is recommending several critical changes to the idle power base limit and allowances so that an idle power threshold better distinguishes the higher performance, higher efficiency systems available in the market. Table 1 shows the idle allowance values for the newly proposed performance based idle power allowance and those values that were changed from the Draft Lot 9 idle proposal.

		BASE IDLE POWER & COMPONENT ALLOWANCES			
	One socket rack server (Cat A)	One socket resilient server (Cat B)	Two socket rack server (Cat C)	Two socket resilient server (Cat D)	Blade server (Cat E)
Base Idle power (W)	15	There is insufficient data to develop an idle limit for this server category.	30	200	40
System Performance Multiplier[7]	4.7		4.0	0.9	8
Base Performance threshold excluded from multiplier (high-end/low-end)	0/1		2.1/0	26/0	3/3
System Performance Allowance (W)	4.7 * (CPU geo mean peak perf – base perf threshold)		4 * (CPU geo mean peak perf – base perf threshold)	0.9 * (CPU geo mean peak perf – base perf threshold)	8 * (CPU geo mean peak perf – base perf threshold)
Storage Devices (W)[8]	4.5		4.5	4.5	4.5
Memory (W/GB)	.175		.175	.175	.175

Table 1: Base Idle and Idle Allowances used for the DIGITALEUROPE System Performance Idle Analysis

Notes: See next page.

7 Appendix A: System Performance Idle Allowance Description

8 Industry proposal outlined in Table 8

(1) These allowances were used in the comparative calculations between the Draft Lot 9 and DIGITALEUROPE analysis.

(2) No idle limit is proposed for a one socket resilient server (Cat B); only one system is on the market. It should be removed from the idle power requirements due to the small market share and product volume.

(3) Details of the revised memory, storage and I/O allowances will be provided in subsequent sections of this paper.

(4) All other idle allowances are consistent with the allowances identified/provided in the Draft Lot 9 document.

Under the Draft Lot 9 Idle limit proposal, server idle power is affected by the performance capability of the processor as measured by the core count, frequency and performance capacity as measured by the SERT CPU worklets. A low end processor with minimal cores and a lower frequency and a less energy intense supporting infrastructure – system chips, circuitry etc. - generates approximately 20 Watts less idle power and a lower average SERT geomean CPU worklet score than a high end processor with maximum cores and frequency and associated supporting infrastructure. Failing to recognize the higher power demands of higher performance processors and their associated infrastructure results in the idle limit being biased against higher power, higher performance server products.

In order to limit the impact of this bias, TGG has developed a system performance multiplier allowance for each server category identified in the Lot 9 regulation, which when combined with other allowances, eliminates more of the lower performance servers as measured by the SERT metric than the Lot 9 proposal. A common performance multiplier is proposed for each server category with adjustments defined for high and low end configurations. Systems with performance values greater than the base performance threshold will get a system performance based idle allowance equal to the multiplier value times the delta between the base system performance value and the threshold. The system performance allowance calculation is explained in detail Appendix A. Figures 11 and 12 illustrate the systems which fail when using the DIGITALEUROPE proposed method for the Idle Power threshold calculation.

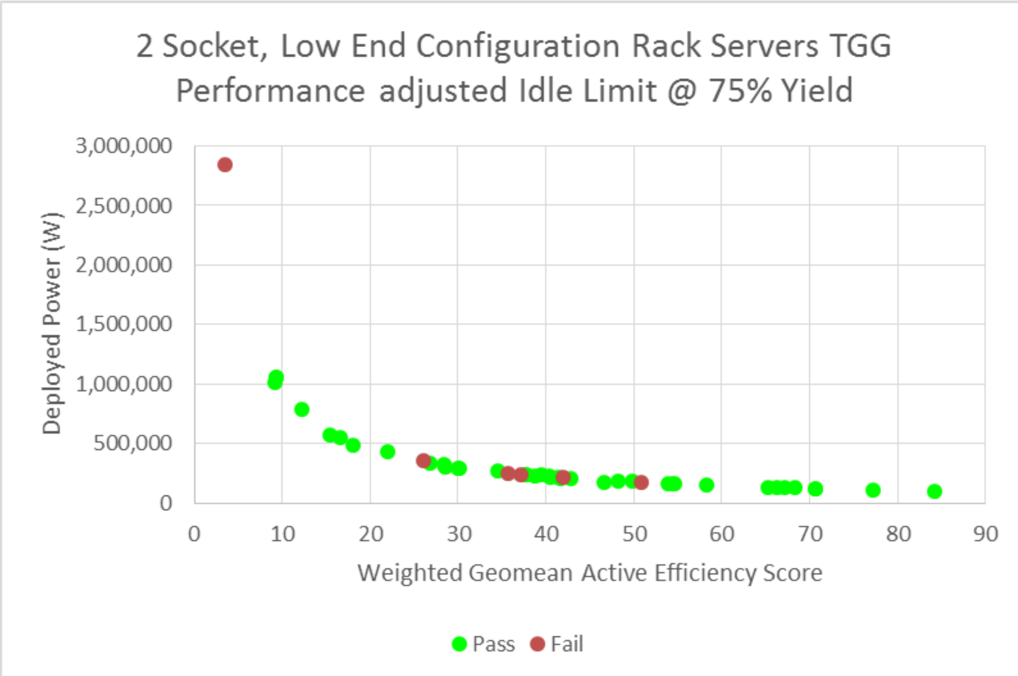

Figure 11: Low-end configuration two socket systems that fail the DIGITALEUROPE proposed idle allowances.

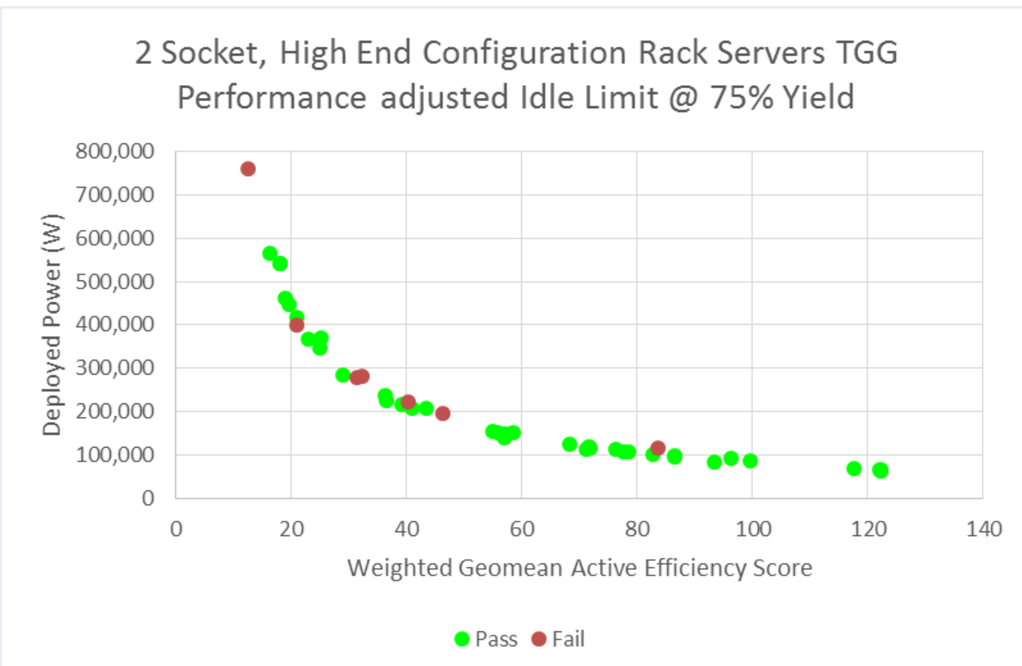

Figure 12: High-end configuration two socket systems that fail the DIGITALEUROPE proposed idle allowances.

Using the system performance allowance calculation, the low-end configuration failed systems move to the left on the weighted geomean active efficiency score scale as compared to the Draft Lot 9 Idle limit while the failed systems on the High-end configurations stay largely the same. Examining the individual idle limits, it was found that the idle power limit for higher performance systems, both low-end and high-end configurations, was increased using the system performance allowance, better matching the threshold to system performance. DIGITALEUROPE wishes to emphasize this proposal is only an improvement to a flawed proposal: the SERT metric is the better approach to setting server energy efficiency requirements.

Section 5: Comparison of the Average Deployed Power of the Servers Allowed by the Three Thresholding Proposals

In order to assess the relative effectiveness of each of the three thresholding methods in reducing data center power use and consumption, TGG calculated the average deployed power for the server configurations which passed each of the three thresholding methods discussed above. While this is a rough assessment of the impact of the three thresholding methods on available servers, recognizing that data center operators and IT professionals selecting servers for office and data center installations will typically select the server which most efficiently manages their workload(s), it indicates the average power use of the servers allowed into the market. Assigning the SERT Metric a value of 1, figure 13 below illustrates that the DIGITALEUROPE Modified Idle Power proposal results in 0.95 and 1.1 times more average deployed power use for one and two socket rack servers respectively and the Draft Lot 9 proposal results in 1.1 and 1.3 times more energy use. Looking at the data by comparing the ratio of the average deployed power of the failed systems to the average deployed power of the passed systems in Figure 14, it is shown that the SERT metric offers the biggest improvement of 1.5 to 4.5 higher deployed power in the failed systems. This indicates that on average the failing systems would consume 1.5 to 4.5 times the energy of the passing systems. The Lot 9 proposal offers the lowest ratio of .9 to 1.75 times the energy use of the failing systems versus the passing systems. Both graphical depictions of the deployed power benefits of the DIGITALEUROPE SERT Metric and Modified Idle Power thresholding approaches show that the approaches are superior to the Draft Lot 9 proposal in their ability to contribute to reduced energy consumption in the data center.

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

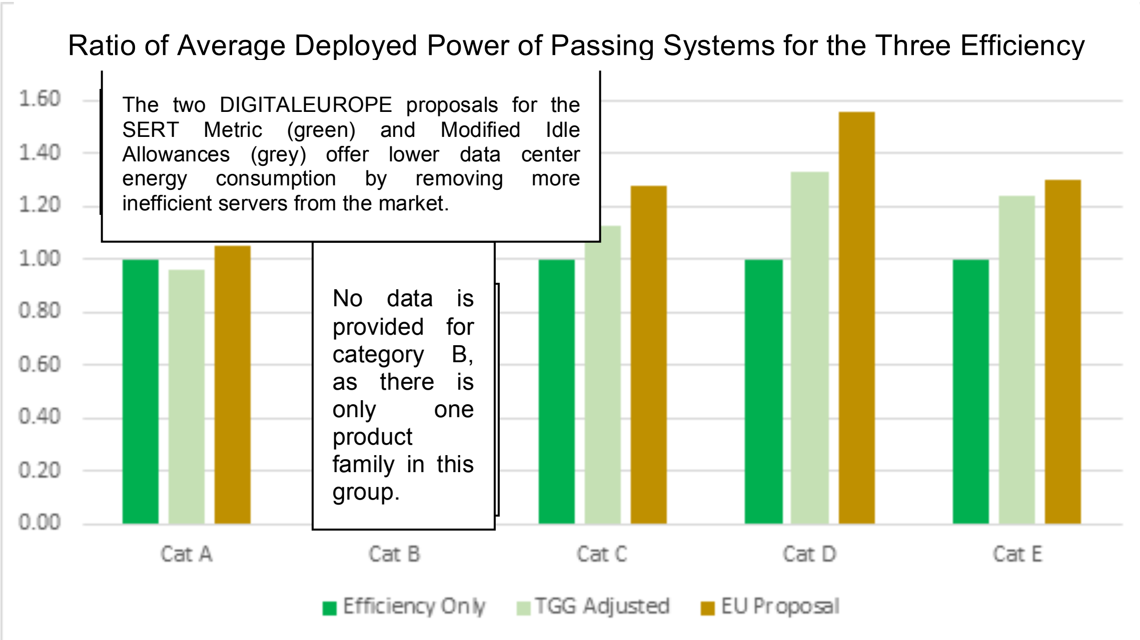

Figure 13: Ratio of the average Power Use of Servers passed by the three server efficiency thresholding methods.

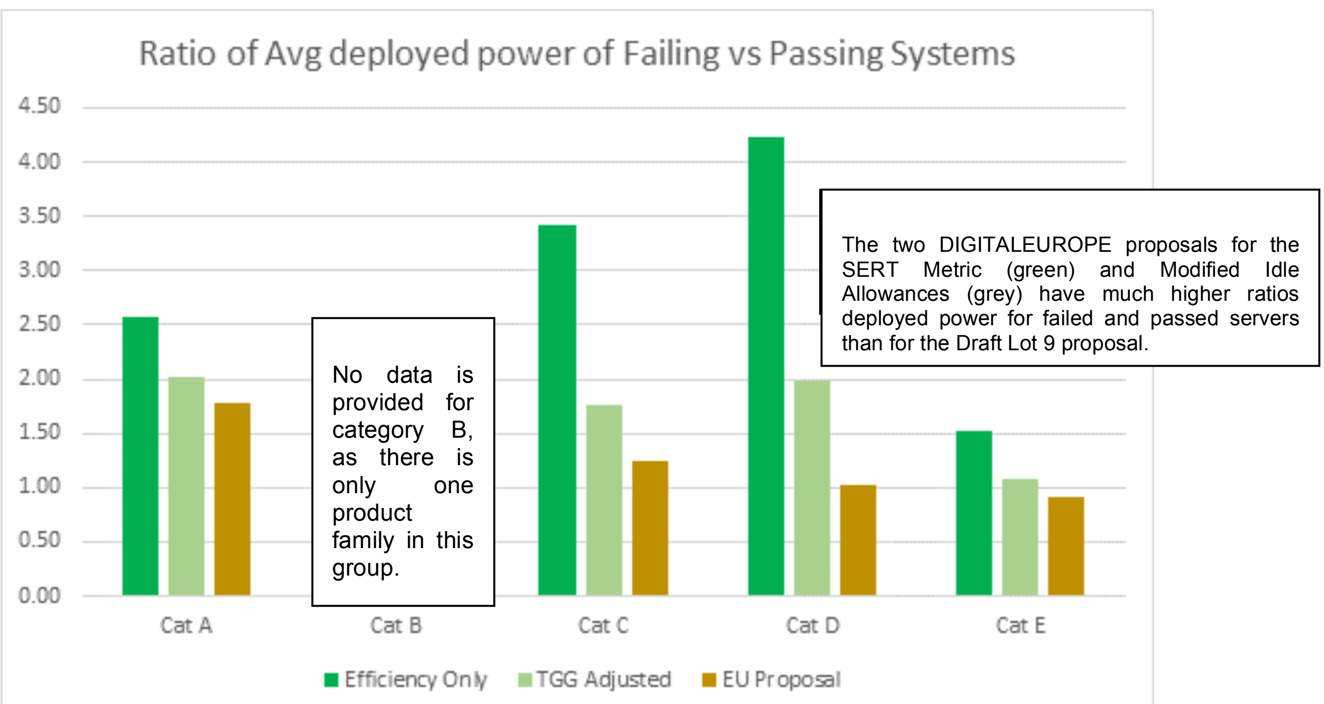

Figure 14: Relative Comparison of the Average Deployed Power of each group of servers passed by the three server energy efficiency thresholding methods.

The SERT metric offers the most efficient fleet of servers to the market, the DIGITALEUROPE Modified Idle Power proposal a slightly less efficient fleet. The Draft Lot 9 proposal results in a significantly less efficient approach as compared to either of the two recommended DIGITALEUROPE proposals for a server energy efficiency threshold.

DIGITALEUROPE strongly recommends that DG GROW implement a server energy efficiency threshold using the SERT Metric. Should DG GROW choose to implement an idle power threshold, DIGITALEUROPE recommends that the Modified Idle Power proposal and the modified memory, storage and I/O allowances be implemented. DG GROW should not implement the Draft Lot 9 proposal.

Section 6: Memory Idle Power Allowance

We have updated the memory idle power data set that we provided to the commission during the Task 1 to 7 study work. We have obtained data on a full set of idle power for the DDR4 DIMMs, augmenting the DDR3 data we provide for the task reports (table 2). This data illustrates several important points:

a. The Watts/GB value decreases with increasing DIMM size and newer process technologies (smaller line sizes).

b. There is a variability of 20 to 70% in the W/GB between memory technology, manufacturer and process technology node. Server manufacturers will typically source and randomly use memory DIMMs from two or three manufacturers for a given server product. This introduces a high degree of variability in the power levels based on the specific DIMM product used in the server.

c. There is not a definable linear relationship by which a W/GB value can be calculated for a given DIMM. The variability discussed in (b) above precludes determining a base power demand for the DIMM platform and then a W/GB value for the memory.

	Process Technology Node					
4GB Chipset	30 nm		25 nm		20 nm	
	Min (w/GB)	Max (w/GB)	Min (w/GB)	Max (w/GB)	Min (w/GB)	Max (w/GB)
DDR3	0.076	0.108	0.061	0.108	0.039	0.095
DDR4	0.166	0.223	0.07	0.218	0.065	0.082

Table 2: Watts per GB data by DDR3/4 and process technology node (Watts DC)Error! Not a valid link.

Table 3: Watts DC/GB and total DIMM watt use data for 4 DDR4 DIMM sizes from 3 manufacturers

As the current low-end and high-end configuration calls for the use of the lower DIMM sizes and current products will largely be using DDR4 memory technology, the likely W/GB for DIMMs used in current test servers will be based on the 4 GB and 8 GB DDR4 DIMM values between .15 and .22 WDC/GB. Assuming 90% power supply efficiency, this would be translated to .17 to .24 WAC/GB. As it is expected that manufacturers will largely move to 8 GB DIMMs by the time the Lot 9 requirements take effect, DIGITALEUROPE proposes that the memory idle power allowance be reduced to 0.175 W/GB.

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

Section 7: Storage Idle Power Allowance

DIGITALEUROPE has partnered with TGG to gather data on idle power use for storage devices. As Table 4 shows, there is a wide variation in idle power between and within device types as segregated by form factor, device speed or SSD, communications protocol used and capacity. Data was collected from 284 drives manufactured by 5 HDD manufactures and 5 SSD manufacturers.

					Device Idle Power	
Form Factor	**Speed**	**Comms**	**"G"**	**Count**	**Min (Wdc)**	**Max (Wdc)**
2.5	7.2	SAS	12G	3	3.2	3.6
2.5	7.2	SAS	6G	3	3.1	4.5
2.5	7.2	SATA	6G	3	2.6	2.9
2.5	10	SAS	12G	18	2.6	5.5
2.5	15	SAS	12G	14	3.9	6.3
2.5	10	SAS	6G	22	3.1	4.9
2.5	15	SAS	6G	4	4.9	5.4
2.5	SSD	NVMe		27	4.0	9.0
2.5	SSD	SAS	12g	22	1.8	5.4
2.5	SSD	SAS	6G	3	3.4	3.5
2.5	SSD	SATA	6G	66	0.5	2.0
3.5	5.7	SATA	6G	3	3.9	5.0
3.5	7.2	SAS	12G	23	4.5	8.8
3.5	7.2	SAS	6G	21	4.5	8.4
3.5	7.2	SATA	6G	52	3.3	8.8

Table 4: Summary of Idle Power data for storage devices (Allowance groups are color coded to Table 5)

Based on the collected data, TGG has proposed the creation of six categories of storage devices. We have tried to set allowances which are representative of the range of products available in the marketplace while minimizing the number of categories created. For each proposed category and allowance, we have excluded those devices at the upper edge of the idle power use those categories where there are obvious outliers.

Similar to the discussion about memory DIMMs, server manufacturers source storage devices from several manufacturers and randomly use those devices when building configurations. Because the low-end and high-end server configurations designated for testing are required to have 2 storage devices, the allowance will have a minimal impact on the idle limit for the tested configurations. A representative adder is important for surveillance testing where a procured product may have multiple storage devices of a type with a high idle power.

Form Factor	Speed	Idle Allowance (Wac)	Number of Pass	Number of Fails
2.5	7.2	4	8	1
2.5	SSD SAS, SATA	4.5	71	4
2.5	SSD SAS, SATA >1920 GB	8	16	0
2.5	NVMe per interface port	10	27	0
3.5	5.7 and 7.2	8	91	8
2.5	10 and 15 K SAS 6G, 12G	6	55	3

Table 5: Proposed Storage Categories and Idle Allowances

DIGITALEUROPE recommends that DG GROW create 6 categories of storage devices with the idle allowances detailed in Table 5.

Section 8: I/O Device Idle allowances

Communications technology is currently transitioning to a virtualized environment, allowing servers to dynamically partition and provision high performance capacity ports to multiple logical ports. The larger capacity ports support a virtualized communication environment referred to as Software Defined Networking (SDN) and Network Function Virtualization (NFV). Although systems equipped with these high capacity ports demand higher incremental idle and active power, a higher capacity port (e.g. 100Gbs) can be dynamically provisioned to replace multiple dedicated lower capacity ports (e.g. 10 x 10Gb/s or 5 x 20Gb/s) thereby decreasing overall energy use. The additional adders for the higher capacity ports are proposed in order to anticipate the integration and software based deployment of these higher capacity ports in server systems.

The industry is developing and releasing new network ports with data processing speeds of 25 Gb/s to 200 Gb/s. For ports up to 10 Gb/s, we are in agreement with the EU proposal. These network ports have a higher per port power use, which is offset from a performance/power standpoint by the increased data transfer speeds and a reduced number of ports due to the virtualization capabilities. Table 6 details the proposed idle allowances for these higher speed ports.

Network Port Speed	Proposed Idle Allowance (W)
10 to 25 Gb/s	15 W
>25 to 50 Gb/s	20 W
>50 to 100 Gb/s	26 W
>100 to 200 Gb/s	35 W
>200 Gb/s	45 W

Table 6: Proposed Idle Allowances for High Gb/s Network Ports

Like memory DIMM and storage device allowances, the primary purpose of establishing idle allowances for these higher throughput ports to accommodate these in the case of market surveillance activities. These higher output ports are unlikely to be installed in the two configurations tested to demonstrate compliance with the Lot 9 requirements.

DIGITALEUROPE recommends that DG GROW create 4 additional network port categories with the idle allowances detailed in Table 6 and increase the idle allowance for the 10 to 25 Gb/s Network port from 8 W to 15 W.

Section 9: Modifications to the HPC definition to include Artificial Intelligence and Deep Learning Systems

High Performance Computing (HPC) server systems are being applied to and modified to support Artificial Intelligence and Deep Learning applications. In order to incorporate these systems into the HPC server definition supplied with DIGITALEUROPE's comments to the Lot 9 Draft, several modifications need to be made.

High Performance Computing (HPC) System: A computing system which is designed and optimized to execute highly parallel applications. HPC systems feature a ~~large~~ number of clustered ~~homogeneous~~ nodes often featuring high speed inter-processing interconnects as well as ~~large~~ high memory ~~capability and~~ bandwidth. HPC systems may be purposely built, or assembled from more commonly available computer servers. HPC systems must meet ALL the following criteria:

A. Marketed and sold as a Computer Server optimized for higher performance computing or deep learning / artificial intelligence applications;

B. Designed (or assembled) and optimized to execute highly parallel applications;

C. Consist of multiple ~~a number of typically homogeneous~~ computing nodes, clustered primarily to

increase computational capability;

D. Includes high speed inter-processing interconnections between nodes.

The red text indicates edits from the previous definition submitted to DG GROW on March 17,2017.

CONCLUSION AND RECOMMENDATIONS

DIGITALEUROPE has thoroughly reviewed the ErP Lo9 draft regulation on servers and storage and provided its detailed assessment over the several years of regulation development process. The industry impact assessment is based on deployment of higher performance servers and storage equipment in the data center environment with the goal to increase the performance capabilities of these products while maintaining a relatively constant power envelope. Industry has demonstrated with data detailed above that its proposed approach will lead to higher data center equipment utilization and a net reduction in data center energy consumption in the EU, than that provided by the ErP Draft Lot 9 regulation approach.

The recommendation below is provides a summary of industry's overall concerns, both those raised in this detailed analysis of idle power limits and allowances and the recommendations from our written comments previously submitted on March 17 and March 31, 2017.

Summary of Recommendations:

In Scope/Definitions: Industry is proposing changes to definitions in the draft regulation including solid state drive (SSD), small storage product and changes to memory levels in the low and high-end configuration definitions.

Out of Scope/Definitions: Industry is proposing to expand out of scope list in the regulation. This should include defining server and storage products that are exempt from the regulation thereby removing any uncertainty regarding the regulation intention.

- Industry is proposing to modify HPC systems definition to include servers optimized for artificial intelligence and machine learning, add large servers, server appliance, and servers with direct attach auxiliary processing accelerators definitions to the list of out of scope products..
- Storage server and Hyper-converged storage systems are defined and assigned to on-line storage class of products.
- Network servers are defined and assigned to the network products category, which is out of scope in this proposed Draft.

Power Supply Efficiency: Industry proposes that separate efficiency requirements should be set for multi-output power supplies. The efficiency requirements for multi-output power supplies should start at 80plus silver and move to 80plus gold for Tier 2 and 80plus platinum for Tier 3.

Materials Efficiency requirements: The purpose of eco-design is to address market failures and act where significant improvements are possible. For the proposed material efficiency requirements, based on the EC's data, we do not see a market failure when it comes the product reuse of servers and storage products and the proposed requirements will not have a significant impact.

- Industry agrees with the Commission's assertion at the recent Consultation Forum that servers are best in class when comes to recycling rate. As such, Industry believes the costs of further increasing the recycle rate far outweigh any significant improvement in recyclability rates for servers and storage systems and propose that **dismantling requirements** be removed.
- Regarding the **proposed ban on gluing and welding**, the industry views prohibition of certain fastening techniques to be overly prescriptive, hampering future product innovation. Instead the focus should be on the outcome to achieve the regulatory requirements – ensure that components are removable and recoverable while clearly understanding and delineating the responsibilities of the manufacturer and the end-user. Industry recommends removing any specific sealing or joining techniques from the regulation.
- DIGITALEUROPE considers making available **built-in data deletion software** as highly inappropriate due to several risks, including security compromises, customer specific issues with

ensuring acceptable destruction of data, manufacturer's liability, SW licensing etc. and recommends that the Commission remove these requirements from the Lot 9 regulation.

Idle State Power (Base limits and allowances):

- DIGITALEUROPE continues to reiterate that server efficiency should be assessed using the SERT™ weighted geomean efficiency metric (SERT metric). DIGITALEUROPE and the Green Grid (TGG) have collected a large dataset of SERT results and have shown that the metric is effective in measuring server efficiency and identifying the servers that will result in higher energy use if installed in the data center.
- Should DG GROW proceed with its plan to set a server efficiency threshold based on an idle limit, DIGITALEUROPE is recommending several critical changes to the idle power base limit and allowances so that an idle power threshold better distinguishes the higher performance, higher efficiency systems available in the market.
  - Server Categories: Industry proposes 4 categories instead of 5. There is insufficient data to establish criteria for Category B (one socket resilient server) since there is only one system currently on the market.
  - Server Base idle power targets: As outlined above, the pure idle mode approach is not workable. There is a need for an idle allowance to account for system performance as discussed above (based on CPU geomean peak performance). This adder was developed assuming a 75% pass rate, with the objective to ensure that higher performing efficient server systems are not removed from the market. As described in Table 1 the system performance allowance is based on a category multiplier and a unique performance threshold for high and low configuration systems below which the system performance allowance is not permissible. The mechanics of the system performance adder approach is further outlined in a separate treatise (Appendix A).
  - Other Adders (allowances):
    - Memory: Based on measured data, Industry proposes to reduce the memory allowance from 0.25 W/GB to 0.175 W/GB, as outlined in Table 3.
    - Storage: The allowances for the storage devices need to be broken into 6 categories by drive type, drive speed, and connection type, with idle allowances of 4-10 W, as outlined in Table 5.
    - I/O devices: For ports up to 10 Gb/s, DIGITALEUROPE is in agreement with the EU proposal. However, the industry is developing and releasing new network ports with data processing speeds of 25 Gb/s to 200 Gb/s. These network ports have a higher per port power use. DIGITALEUROPE recommends that DG GROW create 4 additional Network port categories with the idle allowances detailed in Table 6 and increase the idle allowance for the 10 to 25 Gb/s network port from 8 W to 15 W.
    - DIGITALEUROPE supports other allowances as per the Draft Lot 9 Working Document.

Information Requirements: DIGITALEUROPE outlined the following key issues in March 17, 2017 submission with rationale.

- The ASHRAE table 6 needs to be modified to show the recommended and allowable temperature ranges for each ASHRAE environmental classes.

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

- DIGITALEUROPE strongly opposes any proposal to require idle testing at higher boundary temperature. Should DG GROW insist on this data, DIGITALEUROPE recommends that manufacturers be required to provide the fan power to temperature curves for the server product to enable data center operators to calculate idle power at higher temperatures.
- DIGITALEUROPE opposes the draft requirement on critical raw material (cobalt, neodymium, and palladium) disclosure, as the costs far outweigh the limited usefulness. Industry proposes to remove these requirements for servers and storage from the Lot 9.
- DIGITALEUROPE recommends that the draft requirement to **make the current firmware for servers publically available** be removed from the requirements.

Measurements and calculations: In absence of any adders for Auxiliary Processing Accelerators (APAs), DIGITALEUROPE proposes that a clarification be added to remove graphics and expansion cards before any testing occurs. Further, a definition of APAs needs to be included in the Lot 9 regulation and the servers with direct attach APAs should be excluded from the regulation, as they are expected to represent a small market share over the Tier 1 period and they are generally incorporated in servers with HPC characteristics.

Verification procedure for market surveillance: DIGITALEUROPE recommends that one of the two tested configurations be obtained for market surveillance and test values compared to the measured reported values.

Timetable for eco-design requirements: Industry recommends the regulative effective date to be set 2 years after its publication in the Official Journal of EU.

--

For more information please contact:
Sylvie Feindt, DIGITALEUROPE’s Sustainability Policy Director
+32 2 609 53 19 or Sylvie.feindt@digitaleurope.org

Appendix A

System Performance Idle Allowance description

The fundamental limitation of idle power limits is that they do not account for the work capability of the servers under evaluation. The deployed power analysis[9] takes into account the fact that servers in data centers and office environments get deployed in sets or groups of systems based upon the performance capabilities of the system being deployed and a number of units of a given server configuration required to execute the desired workloads. The net impact on power consumption in the data center is not the power characteristic of a single system but of the group of systems needed to perform the specified work activities.

The TGG SERT™ analysis work group determined that the Draft Lot 9 and ENERGY STAR idle power limit proposals do not adequately account for the fact that systems with higher overall performance capabilities also tend to have higher idle power values. High performance systems use processors with maximum power demand and have more infrastructure (e.g. memory bandwidth, PCIe lanes) and system chips resulting in higher overall power demand, including higher idle, as compared to lower performing systems. This fact can be seen by comparing idle and maximum power of low-end and high-end performance configurations from the same product family in the ITI/TGG SERT database[10]. In order to ensure the idle power limits do not unfairly penalize higher performance servers, TGG SERT™ WG determined that an idle power allowance, based upon the performance capabilities of the system, was needed to account for the additional power requirements in high performing systems.

The TGG SERT™ analysis work group evaluated several options for using a performance metric to calculate an idle power allowance to increase the idle power limit for more capable, higher power demand systems to minimize the current bias toward lower performing systems based on the Draft Lot 9 idle proposal. Some of the system attributes evaluated were Processor thermal design Power (TDP), CPU Capacity calculated by number of sockets times the number of cores times core frequency, and a Geomean of the normalized performance scores of the CPU worklets at 100% utilization level (GeoPeak). The CPU GeoPeak score was determined to provide the most useful value to calculate a performance based idle power allowance.

---

9 Deployed power is the power of the number of servers required to deliver a specified workload in the data centre. A server's performance capacity is the weighted geomean of the geomean of the interval performance measurements for each SERT worklet. The weighted geomean of the performance is divided into a specified performance value to determine the number of servers needed to do the work. The deployed power is calculated by multiplying the number of servers needed to do the work by the weighted geomean of the worklet power measurements.

10 ITI/TGG SERT™ database is a compilation of the server products, with configuration and SERT suite data, certified to the ENERGY STAR criteria through March of 2016. In some cases the database holds data from the full xml SERT results file, in others data from the results or results detail file are available where the files were to submitted to EPA or available through ITI and TGG member companies.

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

In order to provide a numerical value that indicates system processor complex performance capabilities it was decided to take the geomean of the normalized performance score at the 100% utilization level for the 7 CPU worklets and the hybrid SSJ worklet.

The current conformity assessment proposal for energy efficiency is to assess both a high end configuration and a low end configuration of the server against a pass fail criteria. In evaluating the server data set it became obvious that the performance threshold criteria would need to be different between the high end and low end configurations in order to fairly assess the server family. By the nature of the configurations we observed higher idle power and higher performance and weighted geomean efficiency scores on the high end configurations than what was observed on the low end configurations.

To keep the process manageable, it was determined to establish a common performance multiplier within each category type and that adjustments for high and low end configuration would be applied through the use of a performance adder adjustment threshold. This threshold value defined the highest performance below which no adder is applied. Systems with performance values greater than the threshold will get a performance based idle adder equal to the multiplier value times the delta between system performance value and the threshold value. The evaluation of whether a configuration receives a performance adder and the calculation for the performance adder are as follows:

Case 1: System Performance <= performance threshold

Performance adder = 0

Case 2: System Performance greater than Performance threshold

Equation 1

$$\text{Performance adder} = (\text{Perf}_{\text{GeoPeak}} - \text{Perf\_Threshold}) * \text{PerfMultiplier}$$

Where:

PerfGeoPeak = Geomean of normalized CPU and SSJ worklets at 100% utilization

Perf_Threshold = Threshold for application of performance adder. Only systems with performance values greater than this value will have a performance adder applied.

PerfMultiplier = Multiplication factor for performance based adder.

In order to accommodate an additional adder and to give the adder a sufficient impact to help higher performing systems to pass TGG WG reduced the base idle power limits in the Draft Lot 9 idle proposal. We normalized dataset to 75% pass rate, eliminating the bottom 25% of the systems in the data set according to the new performance adjusted idle power limit.

Table 1: Base and component idle allowances, with attendant failure results, for the DIGITALEUROPE proposal.

Configuration	Performance Allowance Property	Category A	Category B	Category C	Category D	Category E
	Base Idle Power	15	There is insufficient data to propose an idle limit for this category	30	200	40
	Performance Multiplier	4.7		4	0.9	8
Low End	Max Performance No Adder Threshold	1		0	0	3
High End	Max Performance No Adder Threshold	0		2.1	22	3
% Passing Idle Limit		76.2		75.0	76.9	66.7
Total Number of Systems		21		44	26	6

In order to evaluate the effectiveness of the proposed idle performance adjustments, the TGG working group compared the average 25% utilization deployed power[11] of the passing and failing systems for each category segmented by high and low end systems. An effective efficiency criteria should result in failing systems with higher energy use in the data center and passing systems with lower energy use in the data center (Higher ratio is better). Table 2 shows that the proposed performance based idle allowance delivers this result when comparing the Draft Lot 9 idle limit proposal with the DIGITALEUROPE proposed Performance compensated idle limit. The cells highlighted in red indicate that the average deployed power of the passing systems is higher than the average deployed power of the failing systems. No highlight indicates that the average deployed power of the failing systems is higher than the average deployed power of the passing systems.

In Table 2 we see that the Draft Lot 9 idle proposal will result in passing servers which will likely have higher average data center power levels for both low end and high end configurations of category D. This means that on average, the proposed Draft Lot 9 idle limit will increase energy use in the data center for both low and high end configurations in category D because the pure idle approach is failing higher performing systems. In each of these cases the DIGITALEUROPE proposed performance adjusted idle limit has higher average data center power levels for the failing systems than the passing systems as should be required of an efficiency metric.

For the Category A case, the Draft Lot 9 idle metric yields higher data center power for the failing systems when compared the ratio of failing system average power to the passing system average power. The Draft Lot 9

11 The 25% deployed power scenario is used, as discussions have center on the fact that many servers are identified as operating at 25% CPU utilization or less. For the purposes of this discussion it was determined this was the best average deployed power against which to assess the relative data center power demand of different configurations sized separately to execute a single, specified workload type. DIGITALEUROPE notes that the discussion of server capacity utilization is a function of CPU, memory, storage and I/O capacity utilization and must more complex than the very simple analysis of a 25% CPU utilization.

proposal yields a ratio of 1.06 for high end systems and 2.89 for low end systems. The DIGITALEUROPE proposal performs even better with a ratio of 1.21 for the high end and 3.6 for the low end configurations, showing that the DIGITALEUROPE proposal removes more low efficiency servers from the market. For the category C high end configurations the two methods perform very similar with the Draft Lot 9 proposal showing a ratio of 1.52 and the DIGITALEUROPE proposal showing a ratio of 1.53. The Category C low end configurations have a somewhat larger spread with the Draft Lot 9 proposal showing a ratio of 1.96 and the DIGITALEUROPE proposal showing a ratio of 2.17.

Category E shows both the Draft Lot 9 proposal and the DIGITALEUROPE proposal perform the same for the low end configuration and the Draft Lot 9 proposal has a ratio of 1.07 for the high end configuration while the DIGITALEUROPE proposal has a ratio of 1.38. Category E systems also have a 66.7% yield instead of the 75% yield due to having only 6 systems in the data set.

Please note that in the analysis, the Draft Lot 9 proposal was adjusted for a 75% yield (product families passing) by reducing the base idle allowance. The adjustment was made so that we were comparing like scenarios, i.e. 75% of the product families passing, for both the Draft Lot 9 and DIGITALEUROPE proposals.

Table 2: Deployed Power Comparisons for the Draft Lot 9 and DIGITALEUROPE proposals

		Category A		Category C		Category D		Category E	
		EU Idle	TGG Perf Adjusted	EU Idle	TGG Perf Adjusted	EU Idle	TGG Perf Adjusted	EU Idle	TGG Perf Adjusted
Percent Systems Passing Idle Limit		76.2	76.2	75	75	76.9	76.9	66.7	66.7
High End	AVG Dep Power of Failing Systems (kW)	308	261	322	322	143	398	195	246
High End	AVG Dep Power of Passing Systems (kW)	290	215	212	211	178	144	181	179
High End	Ratio Fail/Pass	1.06	1.21	1.52	1.53	0.8	2.76	1.07	1.38
Low End	AVG Dep Power of Failing Systems (kW)	838	784	632	679	201	407	209	209
Low End	AVG Dep Power of Passing Systems (kW)	290	218	323	313	317	306	218	218
Low End	Ratio Fail/Pass	2.89	3.6	1.96	2.17	0.63	1.33	.96	.96

Notes:

Red Highlight indicates failing systems deployed power is lower than the passing systems deployed power

Red text indicates option with the lowest ratio of failing deployed power to passing deployed power.

ABOUT DIGITALEUROPE

DIGITALEUROPE represents the digital technology industry in Europe. Our members include some of the world's largest IT, telecoms and consumer electronics companies and national associations from every part of Europe. DIGITALEUROPE wants European businesses and citizens to benefit fully from digital technologies and for Europe to grow, attract and sustain the world's best digital technology companies.

DIGITALEUROPE ensures industry participation in the development and implementation of EU policies. DIGITALEUROPE's members include 61 corporate members and 37 national trade associations from across Europe. Our website provides further information on our recent news and activities: http://www.digitaleurope.org

DIGITALEUROPE MEMBERSHIP

Corporate Members

Airbus, Amazon Web Services, AMD, Apple, BlackBerry, Bose, Brother, CA Technologies, Canon, Cisco, Dell, Dropbox, Epson, Ericsson, Fujitsu, Google, Hewlett Packard Enterprise, Hitachi, HP Inc., Huawei, IBM, Intel, iQor, JVC Kenwood Group, Konica Minolta, Kyocera, Lenovo, Lexmark, LG Electronics, Loewe, Microsoft, Mitsubishi Electric Europe, Motorola Solutions, NEC, Nokia, Nvidia Ltd., Océ, Oki, Oracle, Panasonic Europe, Philips, Pioneer, Qualcomm, Ricoh Europe PLC, Samsung, SAP, SAS, Schneider Electric, Sharp Electronics, Siemens, Sony, Swatch Group, Technicolor, Texas Instruments, Toshiba, TP Vision, VMware, Western Digital, Xerox, Zebra Technologies.

National Trade Associations

Austria: IOÖ
Belarus: INFOPARK
Belgium: AGORIA
Bulgaria: BAIT
Cyprus: CITEA
Denmark: DI Digital, IT-BRANCHEN
Estonia: ITL
Finland: TIF
France: AFNUM, Force Numérique, Tech in France
Germany: BITKOM, ZVEI
Greece: SEPE
Hungary: IVSZ
Ireland: TECHNOLOGY IRELAND
Italy: ANITEC
Lithuania: INFOBALT
Netherlands: Nederland ICT, FIAR
Poland: KIGEIT, PIIT, ZIPSEE
Portugal: AGEFE
Romania: ANIS, APDETIC
Slovakia: ITAS
Slovenia: GZS
Spain: AMETIC
Sweden: Foreningen Teknikföretagen i Sverige, IT&Telekomföretagen
Switzerland: SWICO
Turkey: Digital Turkey Platform, ECID
Ukraine: IT UKRAINE
United Kingdom: techUK

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20