TGG/USITO COMMENTS ON CNIS SERVER ENERGY EFFICIENCY REGULATION MAY 2020 DRAFT

Authors

Gary Verdun, Dell

David Reiner, AMD

Shahid Sheikh, Intel

Ricardo Gonzalez Llera, IBM

Jeff Doolittle, HPE

1 Introduction and Executive Summary

ITI, The Green Grid (TGG) and USITO have been collaborating with Chinese National Institute for Standardization (CNIS) over the last several years to review proposals for a new China National Standard for server energy efficiency. Since the computer servers are designed and manufactured for global markets, it has been industry's strong desire to harmonize China's server energy efficiency standards with ISO/IEC 21836:2020 international standard and the SPEC Server Efficiency Rating Tool (SERT™). Industry appreciates CNIS decision to align its latest standard draft with ISO/IEC 21836, especially on terms, definitions and test methods. However, the proposed CNIS server energy efficiency metric, which is expressed in Watt hours is different from the SPEC SERT metric. CNIS proposed metric is different, in that it fixes the number of transactions for each worklet and load levels, while varying time to complete the fixed number of transactions. SPEC SERT is not designed to complete a fixed number of transactions, and there is not a way to configure SPEC SERT to run until a specified number of transactions are completed. Current CNIS proposed metric will lead to new scoring method with implementations beyond the current scope of SPEC SERT. TGG is recommending not to pursue the new metric for the rationale provided below. Should CNIS decide to pursue its current metric approach, TGG has developed formulas to take data from the SERT reports and derive new values to represent the CNIS metric, in a way to ensure server energy efficiency ranking between two approaches remain the same. This additional step has added to the complexity of the new metric and without a tool to automate final results, this approach will be prone to errors and very difficult to implement.

TGG also recommends that unless there are exceptions or a reduction in scope, CNIS should simply reference the terms, definitions, scope, product categories, product family, tested configuration, test conditions, etc. as outlined in ISO/IEC 21836, and using the same terminology to the extent possible. TGG recommends not to have energy labeling requirements for servers. In regard to energy use data reporting for each grade, we recommend that server configurations that are part of the product family, their energy use (Wh) range should be reported, based on the low-end and high-end tested configurations. All values within the range should be considered to comply with the energy use (Wh) criteria for that grade. For single server configurations, the reported tested configuration results should comply with the energy use (Wh) criteria for that grade.

TGG is proposing the percentage threshold range of passing systems for each of 3 grade levels, based on shipping systems dataset, within each product category. TGG has extensive analysis experience working with government agencies globall, and looks forward to assisting CNIS during the server standard development process. TGG is offering to assist in determining the appropriate energy limits for each grade after the metric is finalized.

2 Elements of Regulation TGG Strongly Supports

ITI and The Green Grid appreciate China National Institute of Standardization's (CNIS) continuing due diligence to refine the requirements of the National Standard for Minimum Allowable Values of Energy Efficiency and Energy Efficiency Grades for Servers. We value the opportunity to continue to collaborate and the opportunity to provide additional input as part of its standard development process.

We welcome CNIS' decision to limit the scope of this standard to single- or dual-processor tower and rack servers designed with an x86-architecture and value CNIS' decision to align the testing methodology and requirements and the terms and definitions of this National Standard to those of the ISO/IEC 21836:2020 international standard and the SPEC Server Efficiency Rating Tool (SERT) Design Document 2.0.2. This action of CNIS will allow equipment manufacturers to comply to requirements more quickly and efficiently. ISO/IEC 21836 is an international standard written for the purpose of standardizing all of the best practices for server energy efficiency regulations and ECO Label programs, and was developed and approved by 22 countries (including China), and 11 observing nations, and in collaboration with The Green Grid (TGG), and SPEC.

Terms and Definitions: We value CNIS' decision to align the terms and definitions section of this National Standard to those provided in ISO/IEC 21836:2020 Section 3.1. This decision by CNIS will improve understanding of testing methods and requirements and will allow equipment manufacturers to comply to requirements more quickly and efficiently.

The Green Grid Server and Storage Energy (SSE) Working Group, has completed significant analysis on this current draft and have developed the following recommendations for your consideration.

3 Scope, References, and Test Config

Scope

While it is clear from the document that this standard is not applicable to other types of servers such as elastic servers, computer servers equipped with high-performance computing system and integrated auxiliary processing accelerators, thin servers, full fault-tolerance servers, storage devices (including blade storage) and network equipment, we request clarification to the applicability of the standard to blade, multi-node and resilient servers. Please confirm these are also not included in the scope of this regulation. We also seek clarification on the definition of elastic and thin servers.

Test Conditions and Configurations

Ambient Temperature

According to the CNIS Server EE National Standard, under the “Test Conditions” section, the tests of energy consumption of servers need to be conducted at an ambient temperature of 25 ± 10 °C. This is different from ISO/IEC 21836 and SPEC SERT which require the ambient temperature to be equal or greater than 20 °C but no higher than the documented maximum operating temperature of the system under test (SUT).

Relative Humidity & Temperature Requirements

The CNIS National Server EE Standard requires the test environment to have a relative humidity of between 15%~80%. Neither ISO/IEC Standard nor the SPEC SERT Design document specify any relative humidity value. Both require the relative humidity of test environment to be within the documented operating specification of the SUT. Per ISO/IEC 21836, "The intent is to discourage extreme environments that can artificially impact energy consumption or performance of the server, before and during the measurement."

In order to avoid confusion in the use and compliance to this national standard, we recommend harmonizing requirements for ambient temperature and relative humidity and power requirements to those specified by the ISO/IEC 21836:2020 Standard and SPEC SERT V 2.0.2 Design Document.

Testing and Family Configurations

We recommend that all tested configurations, whether a single configuration or a configuration as part of product family, should be aligned with the definitions provided in ISO/IEC 21836 (section 7.8). For server configurations that are part of the product family, the test configurations are determined by the server manufacturer. These test configurations should meet the minimum allowable values of energy efficiency in order to qualify for the given grade.

Recommendation: For server configurations that are part of the product family, test only the low-end and high-end configurations per ISO/IEC 21836 (section: 7.8.3), using 2 or more storage devices. For single configurations testing should align with ISO/IEC 21836 (section 7.8.2) using two or more storage devices.

Energy Label and Data Reporting

Energy Labeling: We recommend not to have energy labeling requirements for servers. Unlike consumer electronics products, servers are sold in B-to-B enterprise and data center segments where the server buyer is more sophisticated and procures server products based on specific requirements. Secondly, the proposed server energy use metric (Wh) derived from SERT data, only provides an indication of relative energy use of shipping servers and is not reflective of actual server energy use in the data center environment. As such, the information on the energy label will not be meaningful for the end customer.

Energy use reporting: For energy use data reporting to comply with a given grade energy use thresholds we recommend the following: 1) : For server configurations that are part of the product family, report the energy use (Wh) range, based on low-end and high-end tested configurations. All values within the range should be considered to comply with the energy use (Wh) criteria for that grade. 2) For single server configurations, report the tested configuration test results to comply with the energy use (Wh) criteria for that grade

4 Functional Concerns with CNIS Methodology

Introduction

China Server EE Standard Annex B.1 specifies to “Run” a fixed number of transactions per worklet and load level. SPEC SERT is not designed to complete a fixed number of transactions, and there is not a way to configure SPEC SERT to run until a specified number of transactions is completed. In this section, TGG outlines its concerns and provides some proposed solutions.

Key Concerns

Different Iteration Counts Per Worklet

If it is required for SERT to execute the target number of transactions listed in CNIS Annex B.1, and since transaction count varies by server, it would be required to complete a preliminary SERT run to determine each worklet’s transaction count. Then, since it is not possible to run each worklet for a different number of iterations, the number of iterations per worklet would need to be calculated, and SERT then be run for the maximum number of iterations needed. Every SERT run is ~2.5 hours long, so this could add tens of hours to every server test, making the required testing unmanageable. TGG proposes two solutions to this problem in the *Possible Solutions* section.

Specified Test Method Data Not Directly Available

The CNIS standard specifies in A.3.3 d) to record the SERT energy consumption results in order to calculate SEC_{cpu} with 6.2.2 equation (3), however the energy consumption of the server at a fixed transaction count is not included in the SERT report. Thus, the formulas to take data from the SERT reports and calculate the inputs needed for 6.2.2 equation (3) are missing and need to be derived. TGG has derived these formulas, used them for the analysis of Section 5, and provided the formulas in the *Result Calculation* subsection in this section.

TGG is happy to provide specific recommended updates to the CNIS standard that could include these formulas if requested by CNIS.

Meaning of Metric

SE_{server} (CNIS Section 6.2.1) is specified in Watt-hrs. As shown in Section 5, SE_{server} results compare and rank accurately and in alignment with SPEC SERT results. However, TGG wonders if the intent of using a Watthour (Wh) result is because it is believed that this Wh number represents how much energy a given server would consume under a typical load. Many servers have different ratios of worklet transactions, and it is unclear how the rounded-off values in Annex B.1 were selected or what exactly the absolute value of SE_{server} indicates. Since this metric has not been used before and TGG has not done an extensive investigation to characterize SE_{server}, TGG is concerned that SE_{server} may not indicate absolute server energy efficiency (only relative energy efficiency). TGG is happy to investigate this further if requested by CNIS.

Manual Result Collection Problem

Using a new scoring calculation method would require manual collection of data from the detailed SERT reports and recalculation of the SERT scores using the CNIS algorithm for every server. This would add test time and be error prone as mistakes may be made in the manual calculations.

SERT partial load levels are not steady state

If CNIS requires the target number of transactions must be run and cannot be calculated, then the fact that SERT worklets below 100% are not steady state causes inaccuracies. SERT issues transactions randomly, but in a fixed total number during the measurement period. This means that using the power data from a portion of a SERT run to hit a transaction target may over or underestimate the energy consumption. TGG has not assessed the impact of this concern and does not know if it makes an important difference in the scores.

Proposed Solutions

Use SPEC SERT and ISO/IEC 21836 Scoring (strongly preferred)

SPEC SERT was designed with the scoring methodologies and formulas outlined in ISO/IEC 21836, and these values are pre-calculated and used in all other regulations and programs that use SERT. TGG recommends using the international standardized scoring calculations for the following reasons.

- This scoring methodology has been deeply analyzed by many parties, including SPEC, TGG, and EU consultants, is adopted by EU Lot 9, Japan Top Runner, and EPA ENERGY STAR Server, and so has been time tested for accuracy and real-world applicability.

- Testing all models of servers is an extremely expensive, time-consuming task and using a unique scoring method could mean servers that have been tested for other geographies must be retested for the Chinese market, resulting in huge incremental costs (see *Result Calculation* recommendation to avoid this).
- The official SERT scores are widely used and understood and using a new scoring method will cause confusion and make it difficult to compare server energy scores across geographies.
- Can use the pre-calculated SERT results instead of manual data collection and recalculation

CNIS Result Calculation Formulas

If CNIS decides to retain the current scoring methodology, TGG recommends using the below calculation method based on one SERT result to calculate the energy consumption of the fixed CNIS transaction count.

Calculation of SE_{server} CPU and Storage Workloads:

In order to analyze the SEC proposal using the existing TGG/ITI data set of SERT results it is necessary to calculate the energy that would be used to execute the required number of transactions. We first assume that the average power running the SERT worklet at any utilization level would be the same if the test were performed for a longer or shorter period of time. In the SERT data set the raw performance score for each server CPU worklet utilization level is recorded. This number is a transactions per second number and can be converted to the transaction count by multiplying by the elapsed test time in seconds.

Equation 1

$$TransactionCount = Perf_{util} * Time_{Util}$$

Where:

$Perf_{Util}$ = Raw performance score for the worklet at desired utilization level

$Time_{Util}$ = Mean test time for SERT to execute the worklet

To calculate the test time to complete the required number of transactions we scaled the current test time by the ratio of the desired transaction count to the existing transaction count as shown in Equation 2.

Equation 2

$$Time_{SEC} = Time_{SERT} * \frac{TransactionCount_{SEC\_Target}}{TransactionCount_{SERT}}$$

The SEC efficiency value is then calculated as in Equation 3.

Equation 3

$$SEC = \frac{PWR_{Util} * Time_{SEC}}{3600}$$

The calculations described in Equation 1 through Equation 3 must be carried out at each utilization level for each worklet for each server in the data set. These new values are then combined according to section 6.2.2 of the CNIS proposal. The storage workload transaction counts and test times in the SERT output file can be combined in the same manner as the CPU worklets to obtain SEC values for Random and Sequential worklets as well as the Storage workload.

Calculation of SE_{server} Memory Workload

The memory workloads in SERT do not implement transactions and transaction counts in the SERT output files in the same manner as the CPU and storage worklets do. The SERT output report generators fill in these columns but they do not represent real transaction counts. The stored values do not maintain the performance ratio to the 100% performance values that the CPU worklet transaction counts and performance scores do.

Flood3 Worklet

The Flood3 SERT performance score is a measure of the memory bandwidth of the system and is not related to a transaction count. The SERT test calculates the memory bandwidth using two different time intervals. The full and half performance scores are generally almost the same bandwidth values. The test just runs for half the time for the half utilization level, but the performance score reported is practically the same value as the full test.

Since SERT uses both performance values to calculate the Flood3 worklet efficiency scores, these values are treated the same as the CPU worklet utilization levels. SEC worklet energy values are calculated using the method described in the CPU section above. A default test time of 30 seconds is assumed for all Flood3 data conversions.

Capacity3 Worklet

The Capacity3 worklet performance and transaction counts stored in SERT output files do not have the same relationship to transaction counts and performance values as the CPU worklets do. The actual Capacity3 transaction count is a function of the memory size and the SERT reported performance value per Equation 4.

Equation 4

$$Transcount_{Capacity3} = \frac{Perf_{Capacity3} * \sqrt{\frac{Memsize}{1024}}}{\sqrt{((\frac{Memsize}{1024} - 1) * 0.51)}}$$

This adjustment was applied to Capacity3 performance values in the ITI/TGG data set for the purposes of evaluating the relative performance of SERT and SEC metrics. This conversion was applied to both the base and Max utilization levels for each server in the dataset. SERT uses the Max value and not the base value to determine the Capacity3 memory workload efficiency values. We calculated the SEC memory worklet score using only the Max value to assure better correlation between SERT and SEC.

Removing the performance scaling with memory size per Equation 4 leads to the rank correlation for memory shown in Figure 1.

This leads to poor correlation and breaks scaling with memory size for the CNIS proposal.

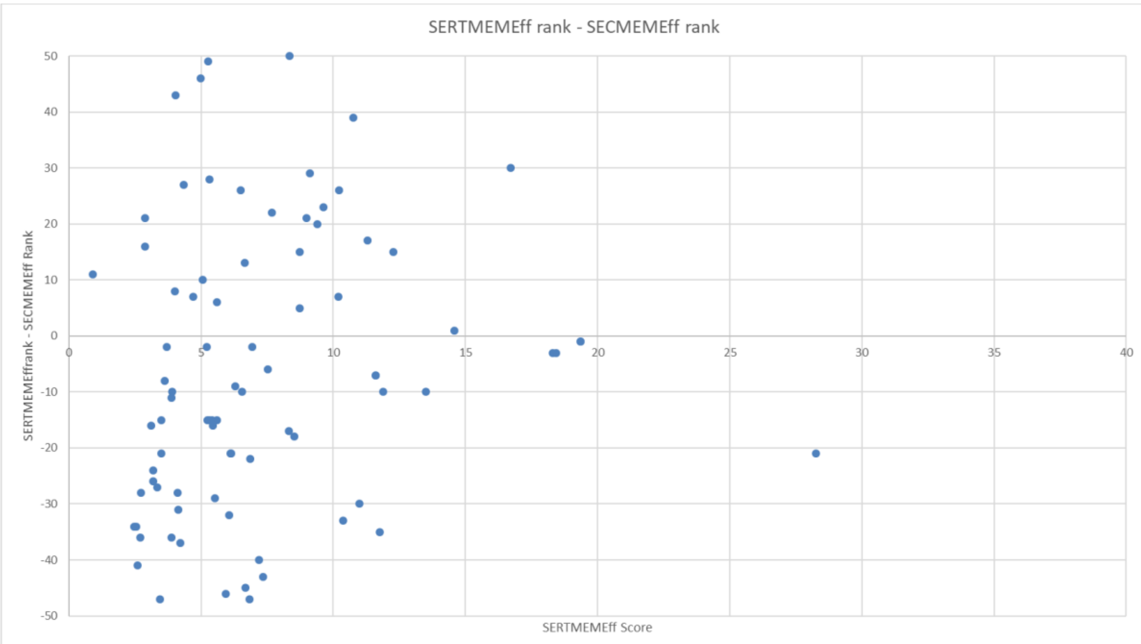

Figure 1

If we do not convert to obtain a true transaction count but instead treat the capacity worklet as we did the CPU and storage worklets by assuming the performance score is the transaction count divided by the elapsed transaction time, then we get the correlation shown in Figure 2. This is a much better scenario and provides the SEC metric with scaling due to memory size just like is done in SERT.

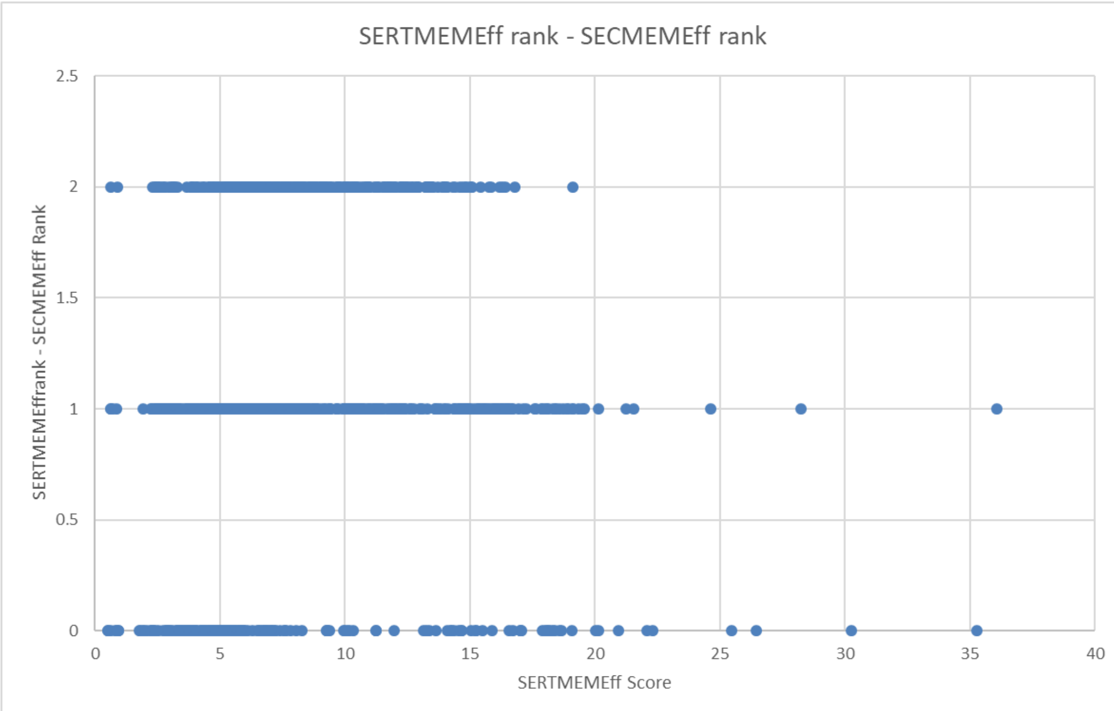

Figure 2

Result processing tool development

TGG experience with the TGG/ITI data set has shown significant issues with manually manipulated data. Significant effort has been required to find and correct all forms of data entry and copy and paste errors as each new set of servers is added to the dataset. It is highly recommended that an automated tool be created to post process the SERT results files in order to obtain the necessary values to calculate SE_{server} efficiency. Use of a manual method for these data conversions and calculations will have significant risk of assigning incorrect values to servers.

5 Server Ranking Assessment of CNIS Calculation Methodology

Background

The May 2020 draft of server energy efficiency in many ways leverages the pending international standard ISO/IEC 21836. There are some key differences that cause considerable variation in whether or not the two implementations identify the same servers as efficient or not efficient.

How SERT Works

TGG has done a preliminary analysis and believes it has identified the key elements of the differences between the SERT efficiency scores and efficiency scores proposed in the CNIS draft.

The SERT tools runs a set of worklets for a fixed time interval. For example, the CPU worklets each run for 30 seconds. At 100% utilization transactions are applied to the server by the test system at the rate determined during the calibration run as the maximum execution rate of the server. For each lower utilization level (75%, 50% and 25% for CPU worklets) transactions are sent to the test unit at random time intervals and transaction sizes such that the total number of transactions over the 30 second time interval will be the indicated percentage of the 100% transaction count value. This methodology is essential to properly ascertaining the energy efficiency of the server. The varying levels of work and varying pauses between units of work allow the server to exercise the built-in power management capabilities thus lowering the energy needed to execute the workload. This method also more closely matches how work would be applied to the server in real life situations than a constant application rate would.

Scoring calculation differences

The CNIS proposal collects energy for the same number of transactions at the same utilization levels identified in the SERT worklets. Applying a similar quantity of work at half the rate of application of work will take twice the amount of time to execute. This section looks at how SERT and SEC worklets run comparing 100% utilization levels to lower utilization levels and making both mathematical and graphical comparisons of the two methods.

The ratio of SERT efficiency score at lower utilization levels will be approximately the fractional power at the lower utilization level divided by the utilization level percentage of the 100% level. This is shown in Equation 5 and Figure 3 below. We see that the efficiency at a lower utilization level is related to the efficiency at the 100% utilization level by a factor of the current % utilization level divided by the ratio of average power at the lower utilization to the average power at 100% utilization.

Equation 5

$$SERTEff_{CPUj} = \frac{Perf_{CPUj}}{Power_{CPUj}} = \frac{0.75 * Perf_{CPU100}}{pm_j * Power_{CPU100}} = SERTEff_{CPU100} * \frac{0.75}{pm_j}$$

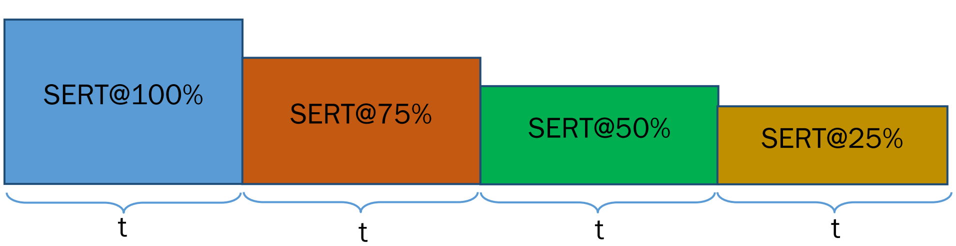

Figure 3

$SERTEff_{CPUj}$ = SERT CPU Efficiency score at j utilization level

$Perf_{CPUj}$ = Server CPU performance at j utilization level

$Power_{CPUj}$ = Server CPU average power at j utilization level

Pm_j = Average power level of server at j utilization level divided by average power of server at 100% utilization level

Performing a similar analysis for the SEC proposal yields Equation 6 and can be visualized using Figure 4. This shows that efficiency values for utilization levels lower than the 100% utilization level show an inverse relationship when comparing SERT to the SEC proposal. Efficiency at lower utilization levels for SEC are the efficiency at 100% utilization times the inverse of the ratio described for SERT.

Equation 6

$$SEC_{CPUj} = Power_{CPUj} * Time_{CPUj} = pm_j * Power_{CPU100} * \frac{Time_{100}}{0.75} = SEC_{CPU100} * \frac{pm_j}{0.75}$$

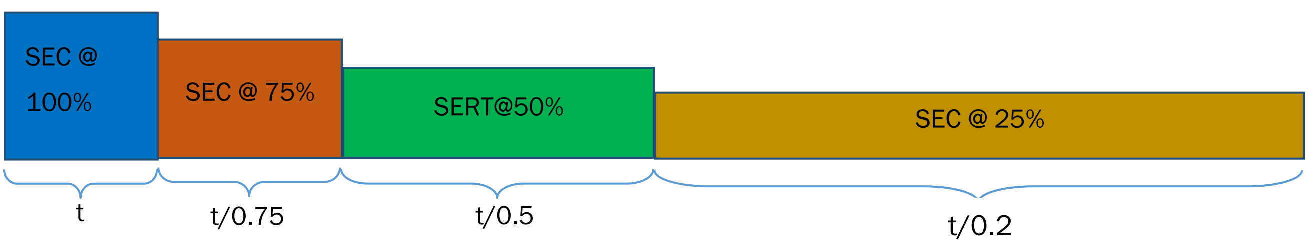

Figure 4

Where:

SEC_{CPUj} = SERT CPU Efficiency score at j utilization level

$Time_{CPUj}$ = Server CPU transaction time at j utilization level

$Power_{CPUj}$ = Server CPU average power at j utilization level

Pm_j = Average power level of server at j utilization level divided by average power of server at 100% utilization level

Ranking changes/efficiency assessment changes

In order to compare the proposed SEC method and evaluate its effectiveness relative to SERT using the ITI/TGG data set it was necessary to make some modifications to model or simulate the implementation of the SEC methodology. In our analysis we model this by using a scaled time interval to execute the required unit of work and multiply this value by the average power level of the indicated utilization level.

Analysis Details

The SERT metric is a performance divided by power metric and assigns increasing efficiency numbers to higher efficiency servers. The SEC metric is an energy consumption for a fixed unit of work and assigns lower numbers to higher efficiency servers. TGG decided to use a server efficiency ranking method to compare the two metrics. This method takes a set of servers and ranks them from most efficient being number 1 to least efficient being number N. This process is performed for the SERT efficiency score and for the SEC proposed efficiency metric on all 944 servers in the current TGG_ITI server dataset of servers introduced in the market from 2012 through 2020. For each server the rank score for the SEC metric is subtracted from the rank score for the SERT metric. This change in rank is a good indicator of how closely the two systems rank servers from most efficient to least efficient.

TGG performed the rank change analysis for the CPU workload, the Storage workload, the Memory workload and for the total server efficiency values. In this analysis servers are ranked from most efficient server (rank 1) to the least efficient server (rank n) in a set of n servers.

We calculated the numerical rank of the 944 server test configurations in the ITI/TGG data set for both SERT and the SEC/SE proposal. We then calculated the delta between the rank value for the two metrics by subtracting the rank using the SEC/SE metric from the rank using the SERT metric. A value of 0 indicates the two methods calculate the exact same rank for that server in the set. Figure 5 below shows a plot of the rank change for each server for the CPU workload. Figure 5 shows very good correlation in ranking of the servers in the data set between SERT and SEC for the CPU workload using the methods described above.

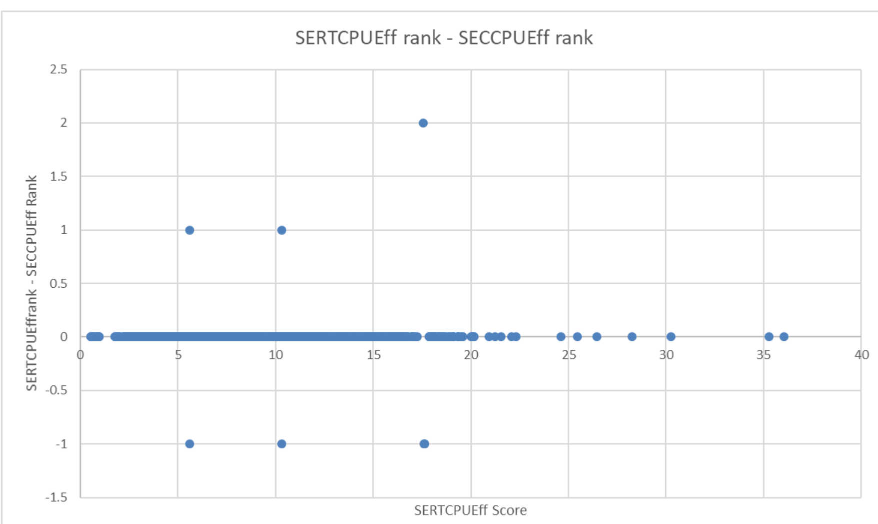

Figure 5

The next workload evaluated is the storage workload. The same ranking analysis was performed on the storage workload and Figure 6 shows the results of that analysis. SERT and SEC have a perfect correlation in ranking the TGG data set servers from most efficient to least efficient.

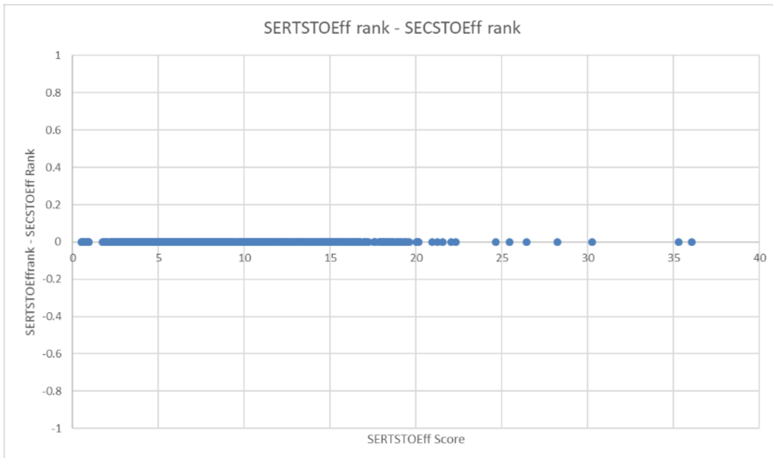

Figure 6

The memory workload ranking comparison is shown in Figure 7. Here we see a much larger number of servers with non-zero errors between the two methods in ranking servers in the TGG data set.

Note memory bandwidth tests are relatively short and the time value varies for each server. We cannot evaluate this against the TGG data set as we did not collect these total transaction time values for every worklet utilization level. For our analysis we assumed a time value of 30 sec for all Flood3 and Capacity3 utilization levels. This analysis calculated transaction count directly from the SERT performance score without applying the memory scaling factor of Figure 4 to the value.

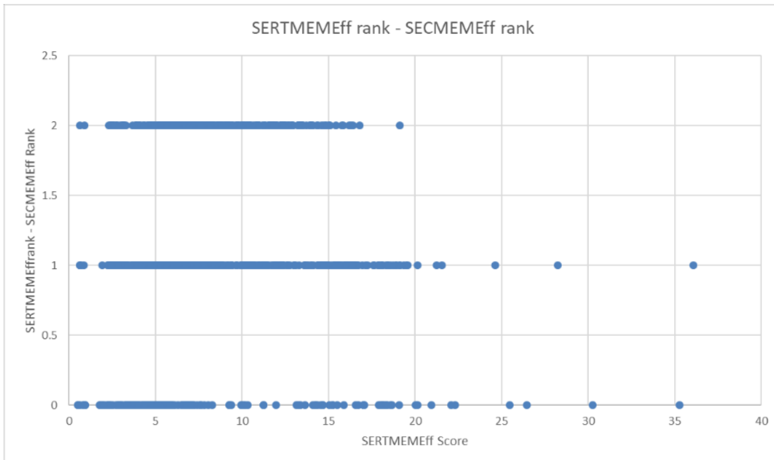

Figure 7

We next compare the rankings of the overall server efficiency ratings of the two metrics. Figure 8 shows a plot of the rank change for each server in the TGG/ITI data set. The proposed method of calculating SE values from the values available in the TGG/ITI data set show very good correlation between SERT efficiency scores and the proposed SE metric efficiency values. TGG believes the dominant source of ranking errors at the server level are the memory workloads and will investigate methods to improve the ranking correlation further.

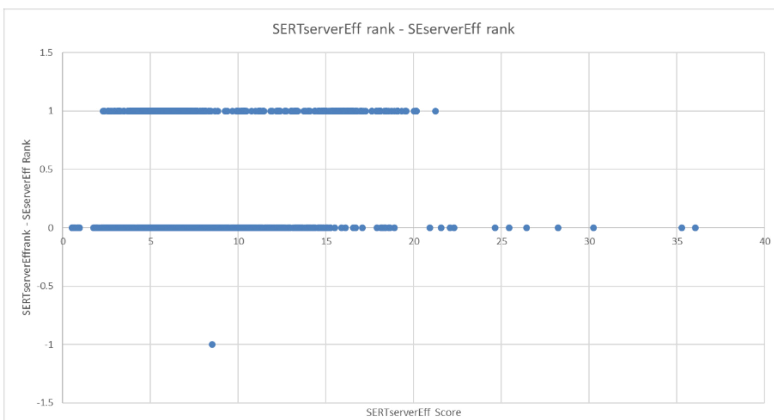

Figure 8

6 Categorization, Limits & Energy Efficiency Grading

Product Categorization

Before establishing minimum allowable values of energy efficiency and energy efficiency grades for servers, product categories for servers in scope of the standard, should be defined in alignment with ISO/IEC 21836 (section 8.1.2), using the same terminology to the extent possible. TGG recommends CNIS state that Blade/Multi-node and Resilient server categories are out of scope of the China national standard.

Minimum Allowable Values of energy efficiency for servers (Limits)

Once the product categories are defined, the next step should be to gather server energy efficiency data for products in the market. Since the servers are designed for global markets, the data for similar products within a given category, shipping in other jurisdictions may also be applicable for analysis to establish minimum allowable values for servers within each category. Analysis for establishing limits should be strictly based on comparing like products within each category. The data for analysis should include all shipping product, within a product category, from most energy efficient to least energy efficient servers. This is important, as any analysis based on selective or most energy efficient products (e.g., ENERGY STAR certified products only) will lead to removing many efficient server products from Chinese market.

Data Analysis Support

TGG SSE was one of the primary data analysis resources to advise the governments in the US, EU and Japan on SERT scores analysis resulting in the desired percentage of passing servers for EU ECO Design Lot 9, US EPA ENERGY STAR Server v3.0, and Japan Top Runner program for Servers. ISO/IEC 21836 (section 6.5) also addresses methodology for threshold selection within the product category. Once the CNIS SERT scoring calculations are finalized (see section on CNIS results calculations formula above), TGG is happy and ready to support with similar analysis to aid CNIS in setting appropriate limits.

Energy Efficiency Grades

Most jurisdictions either have a mandatory energy efficiency requirement on ICT products based on Minimum Energy Performance Standards (MEPs) or a voluntary labeling program like ENERGY STAR mandated for government procurement. TGG SSE understanding of China energy efficiency grades and labeling is based on how China PC energy efficiency standard was implemented and learning from other jurisdictions. Here’s the TGG SSE recommendation:

Grade 3

Grade 3 sets the minimum allowable limits of energy efficiency for servers. This is the minimum requirement to meet in order to place products on Chinese market. The grade 3 should correspond to China PC MEPs based approach, with the goal to pass 85-90% of the most efficient server products in the market (while removing the bottom 10-15% least efficient products from the market) within each product category.

Grade 2

Grade 2 sets the values of energy efficiency of servers higher than Grade 3. The grade 2 represents the minimum allowable values of energy efficiency for servers for government public procurement. This grade should be similar to the ENERGY STAR approach with the energy limits set to pass 30-35% of most efficient servers, within each product category.

Grade 1

Grade 1 sets the most aggressive values on energy efficiency for servers. This top grade is voluntary and intended to incentivize the manufacturers and to demonstrate best in class most energy efficient servers in the market. This grade should correspond to energy limits that pass 15-20% of most efficient servers within each product category.

Figure 9 below summarizes the above discussion

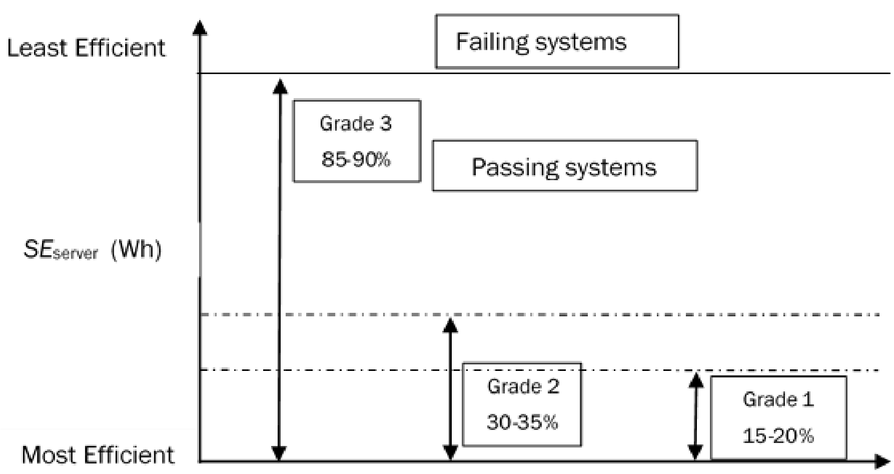

Note: As discussed earlier all tested configurations, whether a single configuration or configuration as part of product family, as defined in ISO/IEC 21836 (section 7.8), should meet the allowable values of energy efficiency for servers in order to qualify for the given grade.

7 Conclusions and Next Steps

Industry appreciates CNIS decision to align its latest standard draft with ISO/IEC 21836. However, there remain several key gaps. The proposed server energy efficiency metric expressed in Watt hours is different from the SERT metric. CNIS proposed metric fixes the number of transactions for each worklet and load levels. SPEC SERT is not designed to complete a fixed number of transactions, and there is not a way to configure SPEC SERT to run until a specified number of transactions is completed. Current CNIS proposed metric will lead to new scoring method to be implemented beyond the current scope of SPEC SERT. TGG is recommending not to pursue the new metric for the rationale stated in this paper. Should CNIS decide to pursue its current metric approach, TGG has developed formulas to take data from the SERT reports and derive new values to represent the CNIS metric. This additional step will add to the complexity of new metric and without a new tool to automate calculation of the final result, this approach will be prone to errors and very difficult to implement. TGG is proposing an energy efficiency threshold approach for each grade level. TGG has extensive analysis experience working with government agencies globally and looks forward to assisting CNIS during the server standard development process, and determining the appropriate energy limits for each grade, after the metric is finalized.