

RESPONSE TO CNIS DRAFT

APRIL 2018

1. General Observations on the Draft Server Energy Efficiency Regulation

THE IMPORTANCE OF THE TEST METHOD CHOICE FOR DATA COLLECTION:

Ultimately, the success of any metric is going to depend on the robustness, repeatability and accuracy of the test method used to measure and collect server performance and power data. For the purposes of this paper, the ‘test method’ refers to the SERT testing procedure which generates raw, measured data on power, transaction counts, elapsed worklet test time and other measurements completed by the SERT test. A test metric refers to the equations and calculations used to convert the measured data into a final metric value including reference system values and aggregation computations and algorithms. Without a test metric that uniformly stresses the server to its full capacity and then accurately sets the performance interval levels, the test results cannot be reliably used to establish a server energy efficiency regulation that will be used to allow or deny placing product on the market. The critical nature of the test method and accurate data has been emphasized in the analysis work we have done for this paper. Where we observed changes in the relative value of the test results between configurations there have been significant changes in the outcome – where a given server is ranked based on the overall metric result. This leads to three specific concerns:

It is not realistic to manage a server energy efficiency regulation with two separate and distinct test methods.

Test methods that use different control parameters and workload algorithms will, by definition, lead to different results in transaction count and power demand. When properly constructed, two separate test suites can generate acceptable methodologies for testing servers – but they are extremely unlikely to generate metric results which will be comparable. This fact is illustrated by our comparison of the SERT and CNIS SEE metrics. While built on the same data set, the differences in the number of worklets between the two metric results in a different relative ranking of server products and configurations. Given two different sets of test data using distinct and unique test methods, there will be a similar difference in relative ranking between results. We conclude that it is important for CNIS to select and designate a single test metric for the generation of performance and power measurements for the “minimum allowable values of energy efficiency and energy efficiency grades for servers”.

The SERT metric offers the best test methodology for a server energy efficiency metric

The SERT metric was developed by SPEC, a recognized industry organization that “aims to “produce, establish, maintain and endorse a standard set of performance benchmarks for computers”. The metric has undergone thousands of hours of testing over a 6 year period, the worklets, workloads and overall metric results have been extensively analyzed and the calculation methods have been modified to improve the effectiveness of the overall metric to differentiate servers based on their ability to maximize work for each unit of energy consumed. Earlier, USITO and the Green Grid SERT Analysis Working Group have done extensive assessment of the SERT and SEEB metrics. USITO provided CNIS an analysis of the SEEB test method, “Preliminary China SEEB Analysis, The Green Grid February 2017”, to CNIS in March 2017. The document provided CNIS USITO’s assessment of the technical limitations of the SEEB test and aspects of the test method that required modifications to enable SEEB to fully stress a server to ensure test to test and server product to server product repeatability that could effectively differentiate server energy efficiency. Subsequently, we observed significant progress in conceptually aligning the latest CNIS Servers Energy Efficiency metric with SERT approach, however, as it will become evident from the analysis below, significant gaps remain to achieve equivalent server efficiency results with current SEE and SERT approaches.

Given CNIS’ stated desire to have a draft server energy efficiency standard published by year end 2018, USITO recommends that CNIS designate the SERT as the required test method and metric for the first version of the standard in the absence of an equivalent alternative. The SERT test and metric has been rigorously tested and validated by SPEC, The Green Grid, US EPA, JEITA and METI and others and a large data set of SERT results have been collected. The data set should facilitate CNIS’ efforts to set server energy efficiency thresholds based on a statistically significant set of results. Given the importance of server products to today’s digital economy and the operation of manufacturing and services industry across the full breadth of the economy, the SERT metric has been validated as an effective server energy efficiency metric.

There is also a benefit for all server manufacturers using the common test method and metric, as governments and international standards bodies are harmonizing on the SERT metric for measuring server energy efficiency. The metric is being designated as the server measurement method and metric in the current drafts of USEPA ENERGY STAR computer server requirements, EU Energy related Products Lot 9 for server and storage product requirements, the Japan Energy Law, the ETSI Energy Efficiency measurement methodologies and metrics for servers Standard and the ISO 21836.1 Server Energy Effectiveness Metric. It is worthwhile to reiterate that the manufacturers design and market the same server products for deployment in data centers globally. The use of SERT to collect server performance and power data will enable server manufacturers to test their server products once to comply with the global suite of server energy efficiency requirements. The common test method and metric provides a predictable energy efficiency framework against which server design engineers can develop future generations of server products.

THE CURRENT CNIS SEE METRIC WILL BE MORE EFFECTIVE USING THE FULL SUITE OF SERT WORKLETS

As the data analysis was performed on the different combinations of the worklet performance (transaction divided by elapsed time) and power measurements, it became clear that once a data set is created by manipulating the SERT data it is relatively easy to combine worklets to create an overall SEE score. We were able to fairly easily assess different combinations of two, four and seven CPU worklets and one and two memory and storage worklets. Creating the different overall SEE scores could be done quickly while the analysis of variance between the different SEE workload scores and the overall SEE score took much more time. Given the ease of combining the data, USITO recommends that CNIS seriously consider using the full suite of CPU, memory and storage worklets in the CNIS efficiency metric. This is important for the following reasons:

1. The SPECPower Committee designed the metric to stress a full range of CPU, memory and storage workloads to provide an overall assessment of the server's energy efficiency. Removing a subset of these worklets focuses the energy efficiency assessment on a less diverse set of workload types, reduces the robustness of the metric and may require reconsideration of the 65/30/5 CPU/memory/storage weighting.
2. While USITO encourages CNIS to use all seven CPU worklets, should CNIS decide to use a smaller number of CPU worklets the analysis that follows demonstrates that a minimum number of 4 CPU worklets should be used. However, the system ranking change analysis demonstrates changes from the baseline, even with 4 CPU worklets.
3. Reducing to a single memory worklet that measures bandwidth fundamentally changes the impact of the memory worklet on the overall server performance due to the exclusion of the capacity worklet. The analysis that follows shows that the change in the memory worklet score is responsible for up to 80% of the variation between the SEE and SERT metric server rankings, with significant rank increases for configurations with high performance, high core count processors and high memory capacity and significant rank decreases for configurations with low performance, low core count processors and lower memory capacity. USITO recommends that both memory worklets be used to calculate the memory workload score.
4. On the storage workload score, the use of only the random read/write worklet does not materially change the storage workload score and an argument can be made that a server product largely performs only read/write operations. Again, for the consistency reasons stated above, USITO sees a benefit to including both scores in the calculation of the overall efficiency metric.

Based on the SERT Working Group review of the Draft CNIS Server Efficiency standard, we performed analysis on the ITI/TGG dataset to understand how the changes in calculation method and the differences in configurations caused differences in the scoring and ranking of the configurations between the SERT metric and the draft CNIS metric. We were able to construct the CNIS metric from the SERT data set by using Equation 3 performing the following demonstrations or comparisons.

1. The SEE metric is the mathematical inverse of the SERT metric. (Section 2)
2. Analysis conducted using the normalized and un-normalized performance scores. (Section 3)
3. The sensitivity of the overall score to a 100%metric improvement in a single worklet score, where the overall score is apprised of 2, 4 and 7 CPU worklets. (Section 4)
4. The calculation of the memory worklet score using the Flood data. (Section 5)
5. Comparison of the CPU, storage and memory workload (aggregated worklet) scores and configuration rankings using 2 and 4 worklets for CPU and 1 and 2 worklets for storage and memory. (Section 6)
6. An analysis of how product families and individual configurations would be graded in a 3 tier grading system. Three thresholds were set at the 25%, 50% and 75% quartile points of the ITI/TGG dataset to perform the analysis. (Section 7)

2. CNIS SEE calculation methodology vs SERT efficiency calculation

SEE energy efficiency metric for any utilization level is calculated according to Equation 1.

Equation 1 SEE efficiency equation

$$SEE_i = \frac{E_i}{N_i} = \frac{Pwr_i * Time_i}{Transaction Count_i} = Pwr_i * \left(\frac{Time_i}{Transaction Count_i} \right)$$

Where:

SEE_i = SEE energy efficiency metric for interval i

E_i = SEE energy measurement: (Power * Elapsed Measurement Time) for interval i

N_i = SEE Number of transactions over measurement period for interval I (Transaction Count)

Time = Duration of the measurement interval

The SERT reported un-normalized performance score is calculated per Equation 2.

Equation 2 SERT Score Calculation

$$Score_i = \frac{Transaction Count_i}{Time_i}$$

The portion of Equation 1in parenthesis is seen to be the inverse of Equation 2. We can therefore represent SEE_i using Equation 3.

Equation 3

$$SEE_i = \frac{Pwr_i}{Score_i}$$

The SERT Interval efficiency is calculated using Equation 4.

Equation 4

$$SERTEFF_i = \frac{Score_i}{Pwr_i}$$

Therefore SEE_i is the inverse of SERT EFF_i or SEE is the inverse of the SERT efficiency score calculated using un-normalized performance values. This is true for all interval and worklet efficiency score calculations. It is at the workload level where SEE and SERT differ.

1. SERT uses 7 CPU worklets while SEE uses 2: Compress and LU (or equivalent).
2. SERT uses 2 Storage worklets while SEE uses 1: Unspecified, SERT WG used random read/write
3. SERT uses 2 Memory worklets while SEE uses 1: Unspecified, SERT WG used Flood (bandwidth)

3. Normalized vs Un-normalized efficiency score comparison

The normalized performance values reported in SERT are the raw performance values in transactions per second divided by the score of a baseline reference system. This is a common practice in the benchmarking realm intended to shift the magnitude of different test results such that they have similar overall magnitudes. Table 1 shows an example of this from one system. The compress and Hybrid SSJ tests have significantly different raw magnitudes with Compress being around 230 thousand and SSJ around 11 million. After applying the SERT normalization factors both worklets have magnitudes around 30 putting them in the same relative magnitude.

Table 1 Normalized and un-normalized example

Compress Raw Performance	Compress Reference Score	Compress Normalized Performance	Hybrid SSJ Raw Performance	Hybrid SSJ Reference Score	Hybrid SSJ Normalized Performance
229,329	6,924.18	33.12	10,963,739	354,125.94	30.96

If efficiency scores were to be combined using an arithmetic average, the use of normalized performance values would be critical to creating a metric that equally weighted the contribution of the individual worklets to the overall metric. Use of the geometric mean to combine performance, power, worklet, and workload scores minimizes the importance of the baseline reference system as the geometric mean minimizes the impact of high relative scores on the combined value of the individual measured or calculated values. Comparing the ranking of the SERT scores calculated from normalized and un-normalized performance values, it was found that the relative rankings of server product/configurations for either performance value were the same and that the overall SERT metric score using the normalized and un-normalized scores varied by a constant factor of 3.137.

4. Sensitivity to changes in a single CPU worklet

The SERT tool uses 7 CPU worklets to evaluate the CPU performance with each worklet measuring the performance of the CPU and system while doing work that is common in servers in real world applications. One of the driving factors in using 7 CPU worklets is the history of CPU vendors occasionally implementing hardware accelerators for certain common tasks seen in the real world. The most recent is encrypt/decrypt accelerators that were seen to drive drastic performance improvements in the Crypto AES worklet. In this case the system power stayed virtually constant but the performance running the Crypto AES worklet was as much as 5 times that of the same system without the accelerators enabled. This one hardware/software enhancement, while valuable, should not drastically improve the overall server efficiency scores.

In order to demonstrate the sensitivity of the magnitude of change in the workload score driven by a large change in an individual worklet score, an analysis was performed on a theoretical CPU workload as changes were made on one of the worklets. We did this evaluation using the geometric mean function to combine 2, 4 and 7 CPU worklets. Figure 1 shows the results of this analysis. We plot the percentage change in output (CPU aggregated Workload) vs a percentage change in one of the component worklets. Here we see that a 100% improvement in the overall worklet efficiency score of one worklet yields a 41% increase in the CPU workload score for the 2 CPU worklet case, a 19% increase for the 4 CPU worklet case and a 10% increase for the 7 CPU worklet case.

If the workload score were calculated using the arithmetic mean, the impact of the 100% increase in a single worklet score would be much higher. Table 2 shows the results of this analysis for 100% SEE increase in Compress and Lu. The size of the percentage change will depend on the magnitude of the worklet being changed relative to the magnitude of the other worklet values in the average. Large magnitudes will have larger impact than smaller magnitude values. The average calculation results were computed using the mean worklet SEE efficiency values in the current ITI data set.

Table 2: Sensitivity to 100% CPU worklet increase using Average Function

	2 CPU	4 CPU	7 CPU
Compress	77.66%	46.94%	12.83%
LU	22.34%	13.50%	3.69%

The analysis highlights the importance of using the geometric mean function to combine the worklet scores and using a larger number of CPU worklets to assess the efficiency of the server:

1. The use of the geometric mean reduces the undue impact of a large relative value for one worklet as compared to the other worklets being combined to calculate the aggregated CPU workload value. The geomean prevents a server tuned for performance on a specific worklet to dominate over servers constructed for general purpose service.

2. The use of a larger number of worklets further minimizes the ability of a single worklet to unduly influence the aggregated CPU workload score. As shown by the sensitivity analysis, the use of 7 rather than 2 CPU worklets reduces the impact of the doubling of one worklet score by a factor of four and prevents gaming of the test by accelerating a single worklet. In addition, more CPU worklets better enable the full testing of the broad capability of the server to run the range of workloads typically found in a data center environment.




Figure 1: Sensitivity of CPU workload efficiency to single worklet changes when worklet scores are combined with the Geometric Mean Function

5. Calculation of the SEE Memory worklet value

The SERT memory worklets results are reported and calculated differently in the SERT output files than the CPU and storage results. Those worklets are the calculated for SEE by dividing the measured power by the measured transaction counts per elapsed time (performance, equation 4). For the Flood worklet, the transaction count is set to match with the number of virtual Logical Partitions (LPARs), which is what is required to drive use of the full memory bandwidth in the server. The actual “workload” measured in the Flood test is the bandwidth in GB/sec. The bandwidth is indicative of the system performance – the amount of data moved per second, and is the equivalent measurement to the transaction counts per elapsed time in the CPU and storage worklet tests. Table 3 details the key measurements in the SERT test for the 25% and 50% Compress and Hybrid ssj worklets and the 100% and 50% Flood worklets. For the CPU worklets, the performance is the transaction count divided by the elapsed time. In the case of the Flood worklets, the

performance measures are independent of the value of the transaction count divided by the elapsed time. This value is not the value to use in assessing the Flood worklet. The transaction count, performance and the power measurement are the same (within measurement accuracies) for the two intervals, indicating that it is the bandwidth that is driving power demand. As discussed, for the Flood worklet the transaction count is indicative of LPARs driving the use of the memory bandwidth.

Table 3: Worklet Score measurements

Worklet	Load Point (Interval)	Performance	Performance Unit of Measure	Transaction Count	Elapsed Time (sec)	Power (Watts)
Compress	25%	31,593.62	transaction/sec	3,793,939	120.088	783.1
Compress	50%	63,176.06	transaction/sec	7,586,401	120.101	876.4
Hybrid ssj	25%	1,198,628.00	transaction/sec	287,762,954	240.0078	677.4
Hybrid ssj	50%	2,400,282.53	transaction/sec	576,336,276	240.121	759
Flood	50%	3,963.89	GB/sec	192	18.622	879.3
Flood	100%	4,040.33	GB/sec	192	35.482	878.9

Because the transaction count and elapsed time are not direct indicators of work performed, the memory worklet for the SEE test should be calculated taking the geometric mean of the power measurements divided by the geometric mean of the measured performance (bandwidth in GB/sec) values (equation 5) – which gives you units of energy over GB matching the CPU and storage worklet units.

Equation 5 Formula for SEE calculation of Flood3 worklet

$$SEE_{Floodi} = \frac{Pwr_i}{Performance\ Score_i}$$

6. Comparisons of the CPU, memory, storage worklet scores and overall metric scores with reduced number of workloads.

CNIS' proposed SEE metric uses 2 CPU worklets, compress and LU (integer calculation), one memory worklet, which we have assumed to be equivalent to the SERT Flood worklet for the purposes of this analysis, and one storage worklet, which we assumed to be equivalent to the SERT random read/write worklet. Equation 3 was applied to the SERT performance and power interval measurements to calculate the SEE CPU and storage interval worklet scores (equation 3). The 8 interval worklet scores for the 2 CPU worklets were combined using the geometric mean to calculate the SEE CPU workload score (equation 6).

Equation 6

$$SEE_{CPU} = \sqrt[8]{SEE_{Com25} * SEE_{Com50} * SEE_{Com75} * SEE_{Com100} * SEE_{LU25} * SEE_{LU50} * SEE_{LU75} * SEE_{LU100}}$$

The same calculation was performed using the geometric mean function to combine the two storage worklet interval scores to calculate the storage workload score.

Equation 7

$$SEE_{Storage} = \sqrt[2]{SEE_{RandomFull} * SEE_{RandomHalf}}$$

The memory worklet score was calculated by taking the geometric mean of the 100% and 50% memory worklet scores to calculate the memory workload score.

Equation 8

$$SEE_{Memory} = \sqrt[2]{SEE_{MemoryFull} * SEE_{MemoryHalf}}$$

The overall SEE score was calculated using a weighted geometric mean of the three workload scores.

Equation 9

$$SEE_{Server} = EXP(0.65 * \ln(SEE_{CPU}) + 0.3 * \ln(SEE_{Memory}) + 0.05 * \ln(SEE_{Storage}))$$

An overall SEE score was also calculated using 4 CPU worklets: Compress, LU, CryptoAES and Hybrid ssj. The TGG SERT Analysis Working Group had recommended that these 4 CPU worklets be used to calculate a 4 CPU workload score in an earlier comment document. The calculation of a 4 CPU workload score was done to determine what, if any, benefit would be achieved by using 4 CPU worklets for the CPU workload and overall metric scores as compared to the SEE and SERT scores.

All analysis discussed below was performed with a subset of the ITI/TGG data set consisting of 32 product families and 136 configurations, most of which were released in 2016 and 2017.

Comparison of the Rank Differences in the Workload and Overall SEE Scores to the SERT Scores


In order to assess the impact of reducing the number of worklets in each of the three workload types, we calculated and ranked the 2 and 4 CPU worklet, one memory and one storage SEE workload and overall metric scores using the SERT interval level measurement data. The SEE scores were compared to SERT workload and overall server metric score rankings. The SEE and SERT rankings were generated by calculating each of the workload and overall server metric scores for each configuration and then using the Microsoft Excel ranking function to rank each score from highest (rank of 1) to lowest (rank of 136) rank. We used the rankings of the calculated workload and overall server scores to perform two comparisons: (1) we calculated and plotted the rank difference for the workload and overall scores for each configuration (figures 2 to 5) to see the numeric quantity and direction (higher or lower) by which the ranking changed and (2) we assumed a grading system where the 136 configurations were divided into 4 equal size ranking groups of 34 configurations and compared the number of configurations that moved to different rating quartiles when comparing the SEE and SERT. (Tables 4-10. The raw data is found in the attached file "CNIS to SERT Ranking compare with spotlight.xls"). These comparisons of the ranking of each configuration's relative efficiency

under the SEE and SERT workloads and overall scores provide data and analysis to determine which server characteristics are causing the differences between the SEE and SERT metrics, and potential impact on server qualification under each scenario

Memory and Storage Workload Scores

Figure 2 shows the actual values of the rank differences in the memory and storage worklet scores. The data shows that there are significant ranking differences in the memory workload scores between the SEE workload using only the data from the Flood worklet test and SERT workload which uses both the Capacity and Flood workload tests. These rank differences, in turn, have a significant effect on the overall workload scores as well (Figure 5 and Figure 6). The storage worklets have small variations in the ranking, which are driven by the fact that the SEE workload score is calculated using only the Random Read/Write worklet (which is most relevant to server operations) while the SERT score uses both the Random and Sequential worklets.

Figure 2: Actual value of the rank difference between the SEE and SERT memory and storage workload scores



This analysis is assisted by assessing how configurations move within a 4 category grading system when changing from the SEE to the SERT score. Configurations were divided into 4 percentile groupings for both the SEE and SERT memory and storage workload scores. Comparisons were then made to how the configurations in the each percentile grouping for SERT were graded by the SEE score. If the SEE and SERT overall score ranking groups matched perfectly, then the comparison table would look like Table 4. All configurations would be in the same quartile for both metrics.

Table 4: Example of SEE to SERT Ranking Comparison Table where all configurations are ranked in the same quartile for both SEE and SERT

		SEE Memory Worklet Ranking			
	Quartile	Top	Second	Third	Bottom
SERT Memory Worklet Ranking	Top	100%			
	Second		100%		
	Third			100%	
	Bottom				100%

Table 5 provides the percentile group changes for the memory workload. The memory workload score changes substantially on some configurations with the removal of the capacity workload, thereby significantly changing the rankings of some configurations on the memory workload. The systems with the highest improvement in rank score (best score represented by 1), an increase of more than 40 places, have high performance processors with high core counts (E7 and Platinum) and generally higher memory capacity. The configurations with the highest numeric reduction in ranking from SERT to SEE, a decrease of more than 40 places tend to be 1 or 2 socket systems with E3 processors with 4 to 8 cores and low memory capacity. The systems with high performance/high core count processors will have much higher relative flood scores than those with low performance/low core count processors, because of the higher number of memory channels and potentially higher bandwidth. There will also be a difference between systems with DDR3 and DDR4 memory. The differences in the Flood scores become more prominent when the Capacity worklet is removed. The capacity workload provides a valuable balancing on the memory scores for systems with lower performance processors and its absence in the SEE metric results in higher memory workload scores for higher performance systems. You can look at the rank changes by workload and overall metric in the spreadsheet.

Table 5: Quartile ranking comparison SERT and SEE Memory Score using Flood workload score

		SEE Memory Worklet Ranking			
	Quartile	Top	Second	Third	Bottom
SERT Memory Worklet Ranking	Top	50%	24%	24%	3%
	Second	26%	35%	26%	12%
	Third	18%	29%	29%	24%
	Bottom	6%	12%	21%	62%

The grading comparison for the storage worklet are shown in Table 6. The storage workload score changes resulting from using only the Random Read/Write worklet are relatively small with only a few configurations having ranking changes of around 20. Most other rank changes are small – only a few points up or down.

Approximately 80% of the systems retain their ranking quartile between SEE and SERT and that movement between the quartiles is confined to the boundaries between the quartiles.

Table 6: Quartile ranking comparison SERT and SEE Storage Score using Random Read/Write workload score

	Quartile	SEE Storage worklet ranking			
		Top	Second	Third	Bottom
SERT Storage Worklet Ranking	Top	88%	12%		
	Second	12%	85%	3%	
	Third		6%	77%	17%
	Bottom			18%	82%

2 CPU and 4 CPU Workload Scores

The SERT WG calculated the SEE metric with both 2 CPU worklets (Compress and LU) and 4 CPU worklets (Compress, LU, CryptoAES, and Hybrid ssj) to assess both the differences between the draft SEE metric and the SERT metric and to determine to what extent the use of 4 CPU worklets in the SEE CPU Workload value reduced the rank differences between the SEE and SERT metric while reducing the ability of having the acceleration of one worklet overly influence the SEE metric. Figure 3 details the rank differences between the 2 CPU worklet and 4 CPU worklet overall SEE metrics and the overall SERT metric, using a 2 and 4 CPU worklet and 2 memory and 2 storage worklets to focus the differences in the rank on the difference in the number of CPU worklets (2 versus 4 versus 7). Overall the 4 CPU workload has a higher average absolute value of rank change than the 2 CPU workload – 8.0 versus 6.2. In general, the maximum rank variation is approximately 20-30 for both the 2 and 4 CPU workloads; the 4 CPU workload shows more, higher variation in the rank change than the 2 CPU workload to drive the higher average rank change.

Looking at Table 7 and Table 8, which show the movement of configurations in the 4 quartile grading system, the grading stays largely the same for 2 and 4 worklet CPU workloads only with both workload scores having movement only between adjacent grading groups. Despite the larger average rank change, the 4 CPU worklet CPU workload score has a closer rank matching to the SERT CPU Workload score, with only 3 to 7 configurations changing ranks with all changes are within in one quartile change. Looking at those configurations that changed quartiles, in general the configurations with reduced SEE ranking had lower capacity or performance processors with 4 to 8 cores for E3 or E5 processors or silver level Skylake processors and often had higher memory capacity. Those configurations that moved up a rating quartile had higher capacity processors.

Figure 3: Rank difference (SEE-SERT) for the two and four CPU worklet CPU workload score




Table 7: Quartile ranking comparison SERT and SEE Workload Score for the SEE 2 CPU Workload

		SEE 2 CPU Worklet Ranking			
SERT CPU Worklet Ranking	Quartile	Top	Second	Third	Bottom
	Top	91%	9%		
	Second	9%	79%	12%	
	Third		12%	68%	21%
	Bottom			21%	79%

Table 8: Quartile ranking comparison SERT and SEE CPU Workload Score for SEE 4 CPU Workload

		4 CPU SEE CPU Worklet Ranking			
SERT CPU Worklet Ranking	Quartile	Top	Second	Third	Bottom
	Top	97%	3%		
	Second	3%	85%	12%	
	Third		12%	74%	15%
	Bottom			15%	85%

Overall SERT and SEE Metric Scores:

Figure 4: The rank difference (SEE-SERT) by configuration between the overall metric scores Figure 4 and Figure 5 show the rank differences by configuration when comparing the SEE and SERT overall metric scores. The rank differences for the Overall Metric scores largely match the rank differences in the memory workload

score. In general, the overall SEE workload scores should increase for the servers using higher capacity processors based on the fact that configurations increased in the CPU and memory workload rankings where they had higher performance processors and decreased in the rankings where they had lower performance processors. This behavior is shown in Table 9 and Table 10, where the presence in grade category shifts by roughly 25% up and down between the SERT and SEE overall scores, with a slightly smaller shift in the bottom quartile. Looking at the stoplight plots in the file “CNIS to SERT Ranking compare with stoplight.xls”, noting which configurations shift ranking ranges, and then looking at the component characteristics of the configurations which shift ranking the impact of the processors is shown.

Figure 4: The rank difference (SEE-SERT) by configuration between the overall metric scores
SEE metric uses 2 CPU, one memory and one storage worklet(s).




Figure 5: The rank difference by configuration between the SEE and SERT overall metric scores
SEE metric uses 4 CPU, 1 memory and 1 storage worklet workload




Table 9: Quartile ranking comparison SERT and SEE Overall Metric score using 2 CPU Workload Table

		SEE 2 CPU Overall Ranking			
	Quartile	Top	Second	Third	Bottom
SERT Overall Score Ranking	Top	76%	24%		
	Second	24%	56%	21%	
	Third		21%	50%	29%
	Bottom			29%	71%

Table 10: Quartile ranking comparison SERT and SEE Overall Metric score using 4 CPU Workload

		SEE 4 CPU Overall Score			
	Quartile	Top	Second	Third	Bottom
SERT Overall Score	Top	74%	26%		
	Second	26%	53%	21%	
	Third		21%	59%	21%
	Bottom			21%	79%

The data for the SEE 7 CPU overall metric score are not presented here. It is assumed that if CNIS chose to use the full complement of CPU worklet scores, both memory and storage worklet scores would also be used and the SEE overall metric score would be the inverse of the SERT overall metric score. When using the same worklet compliment as SERT and using an inverse of the SERT efficiency score the ranking compared to SERT is almost identical.

The analysis above illustrates several key points:

1. In summary, we see that the change in the differences in the SEE and SERT memory workload score ranking has over 70% of the impact on the relative ranking of the SEE and SERT overall results. Based on the analysis above, the SEE ranking places greater emphasis on the CPU performance capabilities than the SERT test and the SEE metric will give server configurations with higher performance processors higher grades.
2. Using just the Flood worklet, or an equivalent bandwidth based memory metric, without including a capacity worklet will markedly change the relative ranking of server products under the overall SEE energy efficiency metric. This is the largest single cause of relative score and rank changes between the SEE and SERT overall metrics. Given the increased emphasis that the single memory workload brings to the overall SERT metric, maintaining a

more balanced memory workload by using both the Flood and Capacity worklet scores to calculate the workload score is recommended.

3. The use of 2 or 4 CPU worklets does not result in a significant change in the relative ranking of products. However, it does reduce the robustness of the analysis and the use of only 2 CPU worklets to create the overall SEE metric almost ensures that efforts will be made to accelerate the Compress and/or LU worklets to improve the overall metric score. As the data for all of the SERT worklets are available, it makes sense to use all seven worklets in the SEE metric.

7. Impacts on differences of the overall SEE metric values within CNIS energy grading system.

TGG believes that the range of SEE scores possible for configurations within a given server family is sufficiently large, which makes it extremely difficult to establish energy efficiency criteria for three grades and four categories of servers as is proposed in the current draft. Most servers could meet multiple grades depending on the customer configuration with many servers having configurations across all grades. For the purpose of this evaluation, we used the data for the 2 socket rack servers from the ITI/TGG dataset.

The difficulty with a 3 level, 4 category grading system is that a subset of the product family, potentially over 50%, will have different grades within a 3 tier, 4 grouping grading system. For illustration purposes, a grading system has been established using the quartiles of the calculated SEE metric for the 2 socket rack servers, with product introduction dates between 2012 and 2017 and plotted the range of SEE scores for each family in Figure 6. The families are sorted in ascending date of introduction with 2012 systems on the left and systems introduced in 2017 on the right side of the graph. Each family is represented by a vertical line that goes from the minimum SEE score for that family to the maximum SEE score for that family using the 3 to 5 configurations for each family represented in the data set.

The horizontal lines in Figure 6 labeled A, B and C are located at the 25%, 50% and 75% quartiles of the maximum SEE score for each family. Line A is the SEE score for which 25% of the 2 socket rack server configurations have maximum SEE scores above this level. Line B is the SEE score for which 50% of the server configurations have a maximum SEE score that is above this level. Line C is the SEE score for which 75% of the server configurations have a maximum SEE score that is above this level. Looking at the distribution of scores across the configurations in each family, we can make the following observations:

1. 3 families have SEE scores that indicate at least one configuration will fail the highest (least efficient threshold, A) and others will be in grades 1, 2, and 3.

2. 9 families have SEE scores that would put at least one configuration in each of the three “passing” grade levels.
3. 27 families have configurations that would put them into 2 “passing” grade levels.
4. 17 families have SEE scores that would put them into a single grade, with 5 of those product families having all of their configurations in grade 1.

It is very important to (a) designate the specific limits on the configuration characteristics that will be tested to qualify a given product family and (b) designate how the product family and its individual, tested configurations will be assessed against the grading thresholds. The work that the TGG SERT Working Group has done with the SERT data set has shown that SEE or SERT scores are very dependent on the selection of components – configurations with higher performance processors and higher memory capacity will tend to have higher SERT scores or lower SEE scores. Specific product families will have overall more efficient SERT or SEE scores based on the technology level of their components and their firmware and operating system characteristics. As Figure 6 shows, many product families will have configurations that will span the grading levels of a multi-threshold grading system.

USITO recommends that only 2 energy grading levels be implemented with one intended to remove the worst efficiency servers from the market and the other intended to indicate the highest efficiency servers in the market.




Figure 6: SEE overall metric ranges for the individual configurations for each two socket rack family

8. Conclusions

Based on the analysis performed by USITO and the TGG SERT Analysis Working Group, USITO makes the following recommendations:

1. Using different test methods and worklet algorithms will result in different measurements that will create server energy efficiency metric scores which are not comparable. Given the large SERT data set, building SEE based on the SERT data provides a solid foundation for the metric and provides a data set with sufficient size to support setting thresholds for the server energy efficiency metric.
2. With the available data set, there are benefits to using all of the available worklets to calculate the SEE metric. The broader set of worklets provides a better overall assessment of a server configuration's workload delivered per unit of energy consumed and limits the opportunity to bias the metric by accelerating one worklet to improve the overall score.
3. Should CNIS chose to continue to use a reduced set of workloads, USITO encourages CNIS to use a 4 CPU, 2 memory, one storage metric to provide better representation of the server CPU and memory characteristics.

4. Only 2 energy grading levels be implemented with one intended to remove the worst efficiency servers from the market and the other intended to indicate the highest efficiency servers in the market.

USITO appreciates the opportunity to review the Draft Server Energy Efficiency regulations and provided comments and recommendations. We are prepared to continue to work with CNIS to develop a final regulation that effectively assesses server energy efficiency while enabling industry innovation to drive more efficient computing solutions.