Redfish Green500 Benchmarker (RGB): Towards Automation of the Green500 Process for Data Centers

*

Elham Hojati
Texas Tech University
elham.hojati@ttu.edu

Jon Hass
Dell Inc.
jon.hass@dell.com

Alan Sill
Texas Tech University
alan.sill@ttu.edu

Yong Chen
Texas Tech University
yong.chen@ttu.edu

***Abstract*—The goal of this research study is to develop and demonstrate methods to automate benchmarking the energy efficiency of high performance computing systems based on the Green500 methodology via direct measurements obtained from the baseboard management controllers using Redfish standard, instead of via an external power meter and a manual process. To achieve this goal, we have designed and developed an automatic Green500 benchmark tool based on Redfish, called RGB (Redfish Green500 Benchmarker). This tool also evaluates implementations of the Redfish standard to determine their ability to meet the requirements of the Green500 benchmarking protocols. We further develop another tool, RGB Checker, which checks the level of the accuracy and the precision of the RGB tool based on the Redfish standard. We have performed validation tests using these tools on a production cluster, on several different testbeds with various types of servers, and on a simulated environment.**

***Index Terms*—Redfish Standard, Green500, Performance, Power Usage, Energy Efficiency, Energy Consumption, Data Center, Supercomputer, Cluster, High Performance Computing**

I. Introduction

The primary goals of the work presented in this paper are to evaluate the degree to which current implementations of the Redfish standard from DMTF [1] can meet the requirements of Green500 measurement protocols, which define methodologies to calculate energy efficiency of HPC systems, and to identify areas in which the standard and its implementations can be improved to make such measurements conform better to these requirements. To achieve these goals, we have designed and developed a benchmarking tool called RGB (Redfish Green500 Benchmarker) and a corresponding validation tool called RGB Checker and tested the functionality of these tools by running them in a variety of real hardware and simulated settings. An additional goal of this project is to provide feedback to DMTF in areas in which the design and current implementations of the Redfish standard can be improved so as to produce the data required to meet the requirements for Green500 calculations.

Supported by the High Performance Computing Center and the Cloud and Autonomic Computing site at Texas Tech University (CAC@TTU), the National Science Foundation through award number 1362134, and the membership and collaboration with Dell Inc. and DMTF.

The contribution of this research is two-fold. First, we designed and developed a Green500 checking tool using Redfish technology through the integration of Redfish and Green500 methods. Second, we evaluated and demonstrated Redfish usage, and provided feedback from the HPC community to DMTF and Green500.

The rest of the paper is organized as follows. Section II introduces the background of this study. Section III presents the methodology of the project. Tests and results are presented in Section IV. Also, the current limitations of RGB are explained in this part. Section V provides the summary of the study. Section VI explains future work.

II. Background

Performance and energy usage are key factors for benchmarking supercomputers and data centers. To run more workload and increase the performance of a supercomputer, it needs to use more energy, which shows a trade-off between these two important benchmarking factors, both of which are vital and need to be considered. The Green500 measurement protocols have been integrated into the benchmarking requirements of the Top500 supercomputing list for the past several years, allowing HPC data center operators to consider the effects of both performance and energy efficiency in evaluating supercomputing clusters. In practice, however, the methodology for gathering data and performing the calculation for the Green500 has been difficult, as it generally requires external power measurement equipment and detailed experimental procedures.

Power consumption can be measured based on the Green500 protocols by one of the following methods: a power management solution, a mix of multi-meter and current probe, or an external power meter. These need to be applied either to the entire HPC system or a specific subset spelled out in the requirements. RGB automates this process based on the internal capability of Redfish-enabled devices (for example, through the baseboard management controllers of the HPC servers in the system). Therefore, RGB does not need to use any external power management solution and thus makes the Green500 calculation process easier and cheaper to carry out.

978-1-7281-5017-8/20/$31.00 ©2020 IEEE

This project proposes and then tests as an alternative both the value of integration and implementation of the methodologies of the Green500 project with the Redfish standard. Redfish is a new standard API for monitoring and managing data center components. Use of the Redfish standard, which has been implemented in the baseboard management controllers for HPC equipment from major manufacturers for the past several years and is increasingly becoming available in a variety of other data center equipment, can in principle simplify and allow automation of the Green500 measurement process.

A. *Significance of the Energy Resource for Data Centers*

The energy consumption in data centers has significantly increased in recent years due to the massive growth in the number of computational activities. Since energy is one of the most expensive resources of data centers and because of the huge growth in data computation demands, reducing the energy consumption is an important challenge for supercomputers and High Performance Computing centers [2].

There are two types of energy usage in clusters: dynamic, related to the dynamic energy used by facilities to run workloads, and static, related to the fixed power consumption by equipment, which is not related to the running workload. To control and reduce waste of energy, there are active research studies focused on managing these two types of power usage in an efficient way [3], [4], [5], [6]. The focus of most studies is providing a new or optimized mechanism, like a new job scheduling method, to reduce dynamic energy consumption. Another example is a research study trying to control power leakage, a type of static power consumption, in order to improve the energy utilization [3].

A major step of energy optimization and reducing its costs is benchmarking and calculating energy efficiency factor in data centers. Several metrics have been proposed to measure the level of energy efficiency [7], [8], [9], [10]. Some of them are dedicated to specific facilities or systems such as cooling systems [7], [11], as well as network infrastructure and wireless communications [12]. The others consider the whole data center with all of its equipment, and calculate its energy efficiency , such as power usage effectiveness (PUE) [7], and performance per watt (PPW) metric [13].

Power Usage Effectiveness (PUE) is one of the defined metrics in this area. It is calculated by dividing the value of the whole consumed energy in a data center by the value of the consumed energy in just servers and network facilities. The PUE metric shows whether servers and network instruments in a data center use electric power effectively. Another important energy efficiency metric is performance per watt (PPW),which is calculated by dividing the performance for running a workload by the total consumed power. This metric is important because it considers two effective properties for clusters and supercomputers at the same time, which are performance and consumed energy.

B. *Green500*

The Top500 project [14] uses the Linpack benchmark metric to provide a list of the top 500 fastest supercomputers every six months. The Green500 project [15] is similar ranking methodology initially introduced in 2005. It provides a list of the data centers and supercomputers in the world based on energy efficiency by considering the performance per watt (PPW) metric. Like the Top500, the Green500 provides its ordered list every six months. Publishing this list leads to a competition between data centers and encourages stakeholders to underline the energy optimization in their data centers and find the possible solutions to reduce the level of the energy consumption. The Green500 project defines a precise methodology for calculating the energy efficiency. This methodology is the outcome of a group effort between EEHPC WG (the Energy Efficient High Performance Computing Working Group), the Top500, the Green500, and the Green Grid [16].

Based on the Green500 methodology, it is possible to use different workloads to run the Green500 test and measure PPW metric for a data center. But the default workloads of the Green500 methodology are the High Performance Linpack workloads (HPL) [13]. Green500 provides a measurement procedure for consumed power calculation in data centers. This methodology is based on three quality levels (Adequate quality level named level 1, Moderate quality level named level 2, and Best quality level named level 3), and four aspects. The first aspect provides detailed information for measurement, granularity, and timing of each level. The second aspect gives detailed information for machine fraction of each quality level. The third aspect talks about subsystems that need to be considered for each quality level. The forth aspect delivers detailed information about measurement location, and acceptable meter accuracy for each quality level. If the submission of supercomputers/data centers fulfills all the requirements for all the aspects of a quality level, then they can achieve Green500 quality for that level and the levels lower than that [16].

To measure PPW, Green500 calculates the maximum value of performance achieved by running the Linpack benchmark on the whole selected system based on the quality level requirements and divides that value by the average of consumed power in the same system obtained by a power meter tool, or combination of multi-meter and current probe, or power management solution [16], [13]. The Green500 methodology has a manual process, and there is no benchmark tool to automate this process.

C. *Redfish*

Redfish [17], [18] is an open standard which contains a collection of specifications and human-readable schemas to monitor and manage the hardware layer of data centers. This standard has been introduced by the Distributed Management Task Force (DMTF). About 30 hardware vendors and organizations are supporting and contributing in this open industry project. It describes a protocol using RESTful interfaces to get access to the hardware monitoring, and management data and operations using a schema based data model. Using this technology, it is possible to get hardware monitoring data, such as power and voltage information, which are useful for calculating the energy consumption of running a workload in a High Performance Computing center, and implementing the Green500 methodology. Redfish schema describes the Redfish resources using HTTPS protocol, in JSON format [19], based on OData v4 [20]. Redfish protocol works with the vast range of compute nodes, servers, and environments such as stand-alone servers, rack-mount environments, bladed environments, large-scale data centers, and cloud environments. The other strength of Redfish standard is that it is (and will continue) growing to cover the hardware layer of data center entirely. Redfish V1.0.0 just focused on compute nodes and servers, and Redfish V1.8 has been expanded to cover most of the equipment and facilities in data centers. This feature helps RGB project to use the internal capability of the Redfish enabled equipment to automate the Green500 process.

III. Methodology

Two benchmark tools have been developed in this research, including the RGB tool based on the Redfish standard and the RGB Checker tool [21] that can be run in a real data center to go through the steps of Green500 methodology and validate the accuracy and the precision of the RGB tool.

One of the reasons that Redfish is suitable for implementing Green500 methodology is that it is (and will continue) growing to cover the hardware layer of data centers entirely. Redfish started with baseboard management controller functionality on individual servers, and has since expanded in scope to cover an increasing variety of data center equipment. The use of simulation techniques helps us to simulate a Redfish enabled version of some of the cluster instruments that are not Redfish-enabled yet, and this technique helps us to test our tool at scale. For real hardware, our RGB tool based on the Redfish standard can be used for measurements of HPC cluster efficiency for inclusion in the Top500/Green500 annual rankings. This section explains the process of RGB measurement based on the Redfish standard, and the steps of running it in a data center.

A. RGB Process

Figure 1 and algorithm 1 show the workflow of the RGB tool. The tool takes two inputs. The first input is a text file which contains all the required information of a supercomputer with Redfish enabled instruments to be submitted to the Green500 List. This file may contain information including the IP address of the compute nodes, switches, routers, PDUs, and UPSs. The second input is the requested Green500 quality level, which can be selected between numbers 1, 2, and 3. RGB tool returns the average of two values, $GFLOPSPerWatt$, and $\overline{P}\left(R_{max}\right)$.

In the initialization step, RGB selects granularity method, timing method, measurements method, machine fraction, subsystems, and meter accuracy based on the input quality level. After finishing the initialization step and gathering the required inputs, RGB goes to the measurement step. In this step, first it launches the Linpack benchmark. Then it starts recording the power measurement samples using Redfish. After getting enough samples based on the quality level requirements, it stops the data gathering process, and saves the Linpack performance result, and calculates the unit average power by repeating the above steps. After getting all required data, it calculates the output. RGB repeats the whole process at least three times, and reports the average of the output values. The process explained in a poster summary presented in SC18 [22].

Algorithm 1: RGB Process

Input: 1- A supercomputer with Redfish enabled instruments to be submitted to the Green500 List.
2- Requested Green500 level (1,2, 3).
Output: $average\left(GFLOPSPerWatt, \overline{P}\left(R_{max}\right)\right)$

Step A) Initialization Step
- Step 1: Select Granularity Algorithm based on the input level.
- Step 2: Select Timing Algorithm based on the input level.
- Step 3: Select Measurements Algorithm based on the input level.
- Step 4: Select Machine Fraction based on the input level.
- Step 5: Select Subsystems based on the input level.
- Step 6: Select Meter Accuracy based on the input level.

Step B) Measurement Step
- Step 1: Launch the Linpack benchmark
- Step 2: Start recording the power measurements samples using Redfish command.
- Step 3: Stop recording the power measurement samples based on selected algorithms in the initialization step.
- Step 4: Save the Linpack performance.
- Step 5: Calculate the unit average power by repeating the above steps based on selected algorithms in the initialization step.
- Step 6: Derive the output.
- Step 7: Repeat the above measurement procedure at least three times and find the average of each output.

B. Using RGB Checker in a Real Testbed

RGB Checker is an additional tool that we developed to find the accuracy and the precision of running RGB in a real data center against Redfish-enabled equipment. To check the accuracy and the precision of RGB using current implementations of the Redfish standard, we ran the Redfish Checker tool against individual instances of ten types of servers: Dell-XR2, IDRAC 14g, Intel, Supermicro, HPE, PowerPC (Supermicro), Insyde, and Dell PowerEdge C6320. We also ran our tool in a cluster named Quanah located at the High Performance Computing Center of Texas Tech University. Commissioned in 2017, this cluster contains 467 compute nodes, 36 cores per node, with a total of 16,812 cores.

C. Using RGB in a Simulated Environment

Use of simulated data center equipment allows us to develop the protocols for measurement in general without being limited by the current capabilities of equipment or implementations of the Redfish standard. This allows us to develop methods that

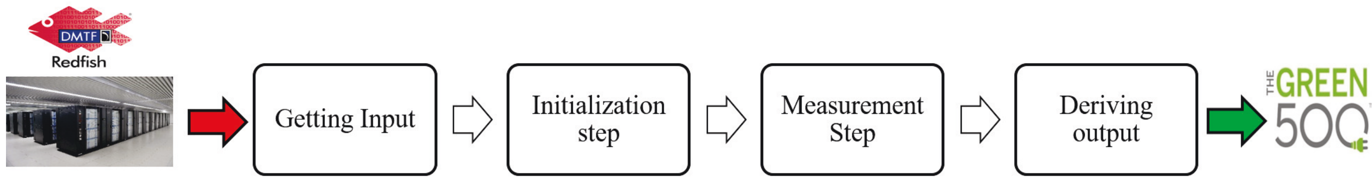

Fig. 1. Overview of RGB process.

test all aspects of the Green500 methodologies and explore the capabilities of the standard to meet these measurement needs. We used Docker [23] to simulate a Redfish enabled cluster and tested the RGB tool on that environment [23].

To consider the third aspect for the first quality level of Green500, adequate level named L1, it is required to measure the power consumption of compute nodes of the selected cluster, and measure or estimate the power consumption of inter-connecting facilities. Redfish supports compute nodes, therefore it is possible for RGB to get the required information just by communicating with the BMC of compute nodes, and sending Redfish requests to them. To consider the third aspect for the second quality level of Green500, moderate level named L2, it is required to measure the power consumption of compute nodes of the selected cluster, and measure or estimate the power consumption of whole participating subsystems. Redfish supports compute nodes, therefore it is possible to measure the power consumption of compute nodes. To estimate the power consumption of whole participating subsystems, we have used a simulation methodology.

To consider the third aspect for the third quality level of Green500, the best level named L3, it is required to measure all participating subsystems. THis requires fetching information at very high rates, which are difficult to achieve for large-scale clusters with individually polled measurements. We use simulation to show how RGB can provide the output for the quality level L3 using the future versions of Redfish API which is going to support all cluster facilities. More specifically, we use docker Swarm [23] to provide a simulated cluster.

IV. Implementations and Results

In this section, we present the results, the outcomes, and the constraints of running the RGB tools.

A. Validating RGB using RGB Checker

The RGB Checker script goes through the steps of the RGB process and checks if the data collected by Redfish satisfies the needs of Green500 methodology by calculating the energy consumption. It also provides the accuracy and the precision level of the RGB results. We ran the tool against the Quanah cluster, and also ten types of testbeds with various servers explained in the methodology section.

Redfish provides two types of authentication: Basic Authentication and Session Authentication [24], [17]. Therefore, We provide two implementations of the RGB Checker, based on these two types of authentication methods. Basic authentication uses the TLS protocol for data transfer between a client and a server. In a session based authentication process, the server creates a session for a connection, stores the client state and the session information in its memory, and assigns a session authentication token to that session log. Using that token, the client can send the next Redfish requests to the server without need for any extra authentication process.

The results of running the RGB Checker in the Quanah cluster shows that the average time duration for gathering a power consumption sample using RGB based on basic authentication is 3.57 seconds. Therefore, the RGB tool based on basic authentication method is far from the granularity requirements of the first aspect of the Green500 methodology for the sampling rate, which should be at least one sample per second for levels 1 and 2 of Green500 and it needs a higher rate for level 3 [16]. The second implementation of RGB Checker was based on Redfish session authentication. In this implementation, the average time duration for gathering a power consumption sample is less than one third of the basic authentication implementation. We ran ten tests and ran the RGB Checker against ten different testbeds containing servers from different brands (Dell-XR2, IDRAC 14g, Intel, Supermicro, HPE, PowerPC (Supermicro), Insyde, and Dell PowerEdge C6320), and gathered 10,000 power samples to find the accuracy and the precision level of the RGB tool for those testbeds. Figure 2 shows the accuracy and the precision percentage of the RGB tool in ten different tests against ten different testbeds. Figure 3 shows the average time for getting power consumption samples (milliseconds), and the number of power reading requests that took more than 1 sec in those ten different tests. Table I shows the average results of all the tests. The table shows the total number of the power reading requests, the average number of the requests that took more than one second, the average, maximum, and minimum time to get power samples, the precision percentage, and the average of the total energy usage for a test.

Based on information in the Table I and the figures 2 and 3, in 80% of the test scenarios, the accuracy and the precision rate is more than 93%. In 80% of the test scenarios, the average time to get power samples is less than one second. Overall, the average time for getting power samples is 725.7 milliseconds, which is less than one second. It means that in the average case, the RGB tool satisfies the requirements of

TABLE I
THE AVERAGE RESULTS OF RUNNING RGB CHECKER IN TEN DIFFERENT TESTBEDS

Total power sample requests	Avg. # of req.s took more than 1 sec	Avg. time to get a power sample(ms)	Precision level (%)	Total energy usage (Watt)	Min time to get a power sample(ms)	Max time to get a power sample(ms)
10000	220.8	725.7	77%	97343	251	1900

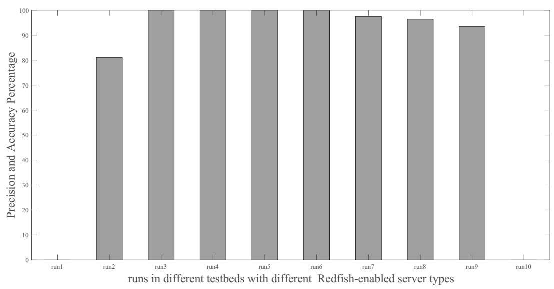

Fig. 2. The precision and the accuracy percentage of RGB tool for level 1 and 2 of Green500, in ten different testbeds with ten different Redfish-enabled server types.

the granularity aspect for level 1 and 2 of Green500. In 20% of the test scenarios, there was no sample gathering with the rate more than 1. It means that in 20% of the scenarios and testbed environments, level 1 and 2 of Green500 has been completely satisfied. But in 80% of test scenarios, there is at least one out of 10,000 samples which took more than one second to be obtained. The result in Table I shows that the RGB tool satisfies the Green500 requirements for the level 1 and the level 2, but it does not satisfy the requirements of the level 3 of Green500. The reason is that based on the requirements of the granularity aspect in the Green500 methodology, the power sample rate for level 1 and 2 are at least one sample per second, and the energy, voltage, and current sample rates for level 3 are at least 5 kHz for AC, and 120 Hz for DC [16].

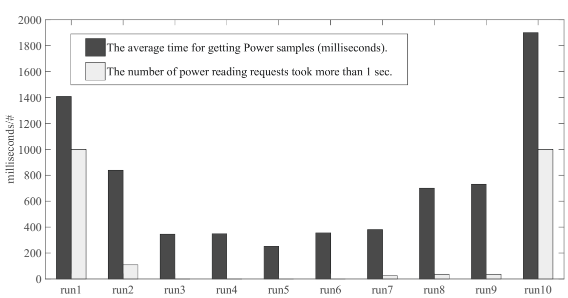

Fig. 3. The average time for getting power samples (milliseconds), and the number of power reading requests took more than 1 sec.

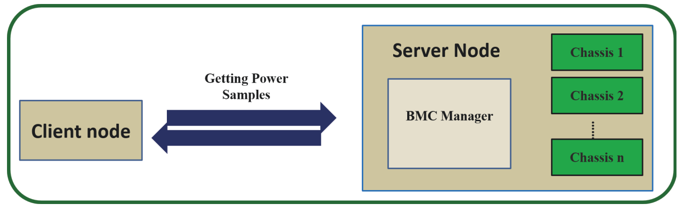

Fig. 4. The structure of a chassis collection.

B. Current Limitations of RGB

The current implementation of RGB is a Python script based on the Redfish standard. The tool was run in a simulated environment for the first quality level, and it returned the results successfully [21]. The results of the study showed that in some specific scenarios, Redfish API interface does not provide sufficient information to perform different quality levels of Green500 precisely. The indicated limitations have been reported to DMTF as a feedback of the research. The first limitation is that there is no timestamp for reading sensors in Redfish. To have precise outputs in all tests and for all devices, it is necessary to have timestamps for reading some Redfish information such as PowerConsumedWatts and AverageConsumedWatts.

Figure 4 shows a server with a collection of chassis. This server may receive several Redfish requests related to different chassis at the same time. Due to the lack of using multi-tasking or multi-threading techniques in a server with a chassis collection, and the lack of timestamps in Redfish reports, testing with individually polled measurements may not satisfy the requirements of the Green500 granularity aspect in this scenario. Another limitation of individually polled measurements is that the rate of reading energy consumption is not enough for the third quality level of Green500. To achieve the third quality level, it is necessary to be able to read voltage and current samples at the rate of 5 kHz for AC / 120 Hz for DC [16]. To achieve such rates will require further capabilities for polling within the servers and reporting of bulk results, for example through improvements in Redfish telemetry features.

V. CONCLUSIONS

The main objective of this project is to automate the process of Green500 benchmarking using the internal capability of data center equipment. The main achievements include automating the manual Green500 process, providing an open benchmark tool for data centers to go through the Green500 methodology easily with no cost, and using the internal capability of Redfish-enabled facilities instead of external expensive tools to gather power samples.

This project has introduced two benchmark tools, the RGB and RGB Checker. The RGB is based on the Redfish standard and can benefit data center stakeholders and the Green500 project by introducing an open-source benchmark tool for the Green500 calculation. This tool has been tested using a simulated data center as well. The RGB Checker tests the ability of Redfish implementations to provide the necessary information such as power, voltage and current samples with a sufficient rate to calculate the consumed energy. It also helps to find the limitations of such implementations to address the Green500 requirements and provides pointers to improve the Redfish standard project to support the energy consumption calculation part of the Green500 methodology. The RGB Checker also validates the accuracy and the precision level of the current version of RGB, by running that against several types of servers and testbeds, as well as a real data center with a large number of compute nodes.

VI. Future Work

As demonstrated above, the current version of Redfish standard as published today is not completely sufficient based on the implementations tested in this research to gather the necessary power, voltage, and current samples and calculate the consumed energy based on the requirements of Green500 methodology. We are able to draw a conclusion that the insufficiency is primarily due to three aspects, lack of timestamps for power, voltage, and current samples, lack of multi-tasking or multi-threading methodology in the case that a collection of chassis need to process several Redfish requests simultaneously, and inadequate sampling rate for the highest Green500 quality level. Therefore, RGB is not completely accurate in some test cases, and it is necessary to design a more precise tool to provide more deterministic and higher confidence outputs for all possible scenarios.

In the future, we plan to design a Redfish Green500 Benchmarker based on additional and recently proposed features of the Redfish standard, such as the telemetry model [25]. The telemetry model is a proposal to provide a Redfish service to support defining registries of metrics and specifying ways to retrieve measurements at intervals through a metric report. It also supports providing triggers for acquisition of a particular metric. Based on the Redfish telemetry model, it is possible to check the timestamp of the gathered samples. In this case, even if there is a delay in the network or other bandwidth limitations in retrieving individual results, it should be possible to provide the pre-defined reports that will meet the requirements for the Green500 process. We also plan to further elaborate on, implement, and test the design of a simulated data center that contains several different types of equipment that gathers results based on the Redfish telemetry model.

Acknowledgment

This research is supported by the High Performance Computing Center and the Cloud and Autonomic Computing site at Texas Tech University (CAC@TTU), which is supported by the National Science Foundation through award number 1362134, and the membership and collaboration with Dell Inc. and DMTF.

References

[1] DMTF. Distributed management task force, inc. website. [Online]. Available: {https://www.dmtf.org}(accessed:10/14/2019)

[2] F. Soppelsa and C. Kaewkasi, *Native Docker Clustering with Swarm*. Packt Publishing Ltd, 2016.

[3] I. Takouna, W. Dawoud, and C. Meinel, "Energy efficient scheduling of hpc-jobs on virtualize clusters using host and vm dynamic configuration," *ACM SIGOPS Operating Systems Review*, vol. 46, no. 2, pp. 19–27, July 2012, doi: 10.1145/2331576.2331580.

[4] L. Liu, H. Wang, X. Liu, X. Jin, W. He, Q. Wang, and Y. Chen, "Greencloud: a new architecture for green data center," *ICAC-INDST '09 Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*, pp. 29–38, June 2009, doi: 10.1145/1555312.1555319.

[5] K. Le, R. Bianchini, J. Zhang, Y. Jaluria, J. Meng, and T. D. Nguyen, "Reducing electricity cost through virtual machine placement in high performance computing clouds," *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2011.

[6] J. Panneerselvam, "A prescriptive analytics approach for energy efficiency in datacentres," PhD Thesis, University of Derby, 2018, http://hdl.handle.net/10545/622460.

[7] C. Fiandrino, D. Kliazovich, P. Bouvry, and A. Y. Zomaya, "Performance and energy efficiency metrics for communication systems of cloud computing data centers," *IEEE Transactions on Cloud Computing (Volume: 5 , Issue: 4 , Oct.-Dec. 1 2017)*, pp. 738 – 750, April 2015, doi: 10.1109/TCC.2015.2424892.

[8] S. Sharma, C.-H. Hsu, and W. chun Feng, "Making a case for a green500 list," *Proceedings 20th IEEE International Parallel & Distributed Processing Symposium*, pp. 738 – 750, April 2006, doi: 10.1109/IPDPS.2006.1639600.

[9] A. S. M. Uddin, A. A. Rahman, "Criteria to select energy efficiency metrics to measure performance of data centre," *International Journal of Energy Technology and Policy, vol. 8, no. 3*, pp. 224 – 237, 2012.

[10] P. Mathew, S. Greenberg, S. Ganguly, D. Sartor, and W. Tschudi, "How does your data center measure up? energy efficiency metrics and benchmarks for data center infrastructure systems," *HPAC Journal*, 2009.

[11] M. K. Herrlin, "Airflow and cooling performance of data centers: Two performance metrics," *ASHRAE Transactions, vol. 114, no. 2*, pp. 182 – 187, 2008.

[12] T. Chen, H. Kim, and Y. Yang, "Energy efficiency metrics for green wireless communications," *2010 International Conference on Wireless Communications & Signal Processing (WCSP)*, November 2010, doi: 10.1109/WCSP.2010.5633634.

[13] R. Ge, X. Feng, H. Pyla, K. Cameron, and W. Feng., "Power measurement tutorial for the green500 list," Green500, Document, June 2007, https://www.top500.org/files/green500/tutorial.pdf.

[14] Top 500 List. [Online]. Available: {https://www.top500.org/}(accessed:10/14/2019)

[15] The green 500. [Online]. Available: {https://www.top500.org/green500/}(accessed:10/14/2019)

[16] EEHPC WG, "EEHPC WG: Power measurement methology (version 2.0 rc 1.0)," Green500, Document, 2017, https://www.top500.org/static/media/uploads/methodology-2.0rc1.pdf.

[17] DMTF, "Redfish Scalable Platforms Management API Specification," DMTF, Document, 2019, https://www.dmtf.org/sites/default/files/standards/documents/{DSP}0266_1.8.0.pdf.

[18] Distributed Management Task Force, Inc., Redfish API. [Online]. Available: {https://www.dmtf.org/standards/redfish}(accessed: October14,2019)

[19] "Json schema: A media type for describing json documents draft-handrews-json-schema-01," Cloudflare, Inc., Tech. Rep., 2018. [Online]. Available: https://tools.ietf.org/html/draft-handrews-json-schema-01

[20] "Odata version 4.0 part 1: Protocol," OASIS Standard, Tech. Rep., 2014. [Online]. Available: https://docs.oasis-open.org/odata/odata/v4.0/os/part1-protocol/odata-v4.0-os-part1-protocol.html

[21] "Rgb-Checker," Available at https://github.com/nsfcac/RGB-Checker (1/17/2020) NSF CAC@TTU, 2020.

[22] E. Hojati, Y. Chen, A. Sill, and J. Hass, "RGB(Redfish Green500 Benchmarker):A Green500 Benchmarking Tool Using Redfish," Available at https://sc18.supercomputing.org/proceedings/tech_poster/poster_files/post254s2-file3.pdf , poster summary in SC18, 2018.

[23] E. Hojati, "RGB based on the Redfish Standard)," Available at https://github.com/nsfcac/RGB-Standard NSF CAC@TTU, 2020.

[24] "Redfish®reference guide revision 2.0a," Super Micro Computer, Inc., Tech. Rep., 2019. [Online]. Available: https://www.supermicro.com/manuals/other/RedfishRefGuide.pdf

[25] "Redfish telemetry model," DMTF, Redfish SPMF, Tech. Rep., 2018. [Online]. Available: https://www.dmtf.org/documents/redfish-spmf/redfish-telemetry-white-paper-010a