RGB (Redfish Green500 Benchmarker): A Green500 Benchmarking Tool Using Redfish

Elham Hojati
Texas Tech University

Yong Chen
Texas Tech University

Alan Sill
Texas Tech University

Jon Hass
Dell Inc.

ABSTRACT

Performance and energy are important factors for supercomputers and data-centers with a trade-off between them. Energy efficiency metric considers both of these properties. The Green500 is a branch of Top500 project which provides a list of supercomputers based on energy efficiency. It has a manual methodology for this calculation. Redfish is a new generation of management technologies for the hardware layer of data-centers. Our project focuses on designing and developing an automated Green500 benchmark tool using Redfish, called Redfish Green500 Benchmarker, or RGB in short. It offers the integration of Redfish and Green500, and automates Green500 benchmarking process with leveraging the internal capability of redfish enabled equipment. It also enhances the Redfish standard to make sure it addresses the requirements of Green500 calculations. This research will also conduct validation and evaluation of RGB on real-world clusters for small-scale to medium-scale tests, and on the data-center simulator we have developed.

KEYWORDS

Redfish, Green500, Energy Efficiency, Power Consumption, Data Center

ACM Reference Format:
Elham Hojati, Yong Chen, Alan Sill, and Jon Hass. 2018. RGB (Redfish Green500 Benchmarker): A Green500 Benchmarking Tool Using Redfish. In *SC '18: The International Conference for High Performance Computing, Networking, Storage, and Analysis.* ACM, Austin, TX, USA, 2 pages. https://doi.org/10.1145/3147213.3149212

1 INTRODUCTION

The contribution of this research is to design and develop a Green500 benchmark tool using Redfish technology (the integration of Redfish and Green500), broaden Redfish usage, and provide a feedback from HPC community to DMTF and Green500.

Firstly, RGB would benefit Green500 project. The current issue is that Green500 provides a manual methodology to find the level of energy consumption,which makes the process complicated. To address this problem, RGB automates Green500 methodology. Also, RGB accomplishes what green500 needs to achieve using Redfish standard technology. It uses the internal capability of Redfish enabled devices, instead of using some external expensive power meter or power management solutions, which reduces the cost of running Green500 process. In addition, RGB would benefit supercomputer stakeholders. RGB makes Green500 calculations easy for clusters. It gives them an automatic tool to run Green500 process, therefore they do not need to go through this methodology manually. Furthermore, RGB would benefit DMTF. It provides a feedback to DMTF which shows if Redfish API, and DMTF instrumentations provides all the properties and necessary data to calculate the level of energy efficiency in a cluster. DMTF could use the result to add needed data structure to the Redfish schema to make sure that Redfish API comes up with enough power consumption information, and DMTF instrumentations provide sufficient data for Green500 calculations.

ACM HPDC, November 11–16, 2018, Tempe, AZ, USA

ACM ISBN 978-1-4503-5149-2/17/12.
https://doi.org/10.1145/3147213.3149212

2 BACKGROUND

This section provides a background of the study. It talks about data centers and the importance of energy resource for them, Green500 project, and also Redfish technology.

2.1 The Importance of Energy in Data Centers

Energy is one of the most important and costly resources in data centers. Hence, it is very important to find a way to reduce the consumed power. There are several metrics to calculate energy efficiency factor in data centers. One of the most important energy efficiency metrics is performance per watt (PPW) [3]. This metric is calculated by dividing the total performance for running a workload by the total consumed power.

2.2 Green500

Green500, which is a subset of Top500 project, is firstly introduced in 2005. It provides a methodology for measuring energy efficiency. This methodology is the result of the collaboration between the Top500, the Green500, the Green Grid, and EEHPC WG (the Energy Efficient High Performance Computing Working Group) [4]. The Green500 uses PPW metric to provide a list of supercomputers based on energy efficiency every six months. It provides a measurement procedure for consumed power calculation in data centers. This methodology provides three quality level (Adequate quality level named L1, Moderate quality level named L2, and Best quality level named L3) based on four aspects [4], [3].

2.3 Redfish

Redfish [1] is a brand new technology for monitoring and managing the hardware layer of clusters. Redfish technology was introduced by the Distributed Management Task Force (DMTF) [2] to provide suitable interfaces for managing and monitoring the hardware layer of modern scalable data centers, and about 30 vendors and companies are contributing in and supporting this project. It provides some hardware layer monitoring information such as temperature, fan speed, voltage, and power.

3 METHODOLOGY

The goal of this study is developing a tool to benchmark power efficiency based on green500 methodology using Redfish standard, and test it in two different environments, a real cluster and a simulated one using containerization. This evaluator tool can be used for submission of HPC clusters for consideration for inclusion in Top500/Green500 annual ranking. This section explains the RGB process, and the steps of running it in a real environment.

3.1 RGB Process

RGB is an application written in Python language. It gets two inputs. The first input is a text file which contains all the required information of a supercomputer with Redfish enabled instrumentations to be submitted to the Green500 List. This file may contain information like 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 L1, L2, or L3. RGB tool returns the average value for two variables: $GFLOPSPerWatt$, and $\overline{P}(R_{max})$.

After getting input, RGB runs the initialization step. In this step, the application selects granularity method, timing method, measurements method, machine fraction, subsystems, and meter accuracy based on the input quality level. After getting input and initialization step, RGB goes to the measurement step. In this step, first it launches the benchmark tool to run a workload and find the performance. At the same time it starts recording the power measurements samples using Redfish commands. After providing enough samples based on the quality level requirements, it stops data gathering process, and saves the performance result, and calculates the unit average power by repeating the above steps. After providing all the required data, it calculates the output. RGB repeats the whole process at least three times, and reports the average of output values.

3.2 Using RGB in a Real Environment

We run our tool against a real cluster. This cluster, named Quanah, is commissioned in 2017, and located at the High Performance Computing Center of Texas Tech University. It contains 467 compute nodes, 36 cores per node, overall 16,812 Cores. The cluster has 87.56 TB Total RAM, 192 GB per node. The nodes have Xeon E5-2695 v4 Broadwell Processors. Quanah uses Intel Omnipath Fabric with the speed of 100 Gbps for networking process. The nodes in the cluster were installed from the scratch with CentOS 7.3. Quanah was used as a real testbed to check the capability of RGB.

4 IMPLEMENTATIONS AND RESULTS

RGB tool developed using Python language, and run against the Quanah cluster for the first quality level, and it returned the Green500 result for that level, L1, successfully. Testing how RGB works for the second and third quality levels is ongoing.

The result of the study indicated some limitations, which were reported to DMTF as a feedback. It showed that the current version of Redfish API interface does not provide sufficient information to perform different quality levels of Green500 precisely. The first limitation is that there is no underline timestamping for reading sensors in Redfish, therefore in some cases, and some especial scenarios, the current version of RGB does not show completely accurate results. To have precise output in all tests and for all devices, it is necessary to have Timestamps for reading some Redfish information such as PowerConsumedWatts, and AverageConsumedWatts. The other Limitation of the current version of Redfish API 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.

5 CONCLUSION

This project helps to enhance the Redfish Standard to make sure it is sufficient to meet the requirements of Green500 calculations, and it has all the necessary properties to calculate Green500 metrics, performance per watt (PPW) and energy efficiency to consider the reliability property), develop a green500 checking software based on Redfish standard, and run the automated green500 benchmark tool in a real data center situation and get the real results.

6 FUTURE WORKS

The results showed that the current Redfish model as published today, is not wholly sufficient and precise for these types of measurement and analysis, because of the lack of timestamp for sensor data samples, which leads to the conclusion that the current version of RGB is not completely accurate for all test cases and scenarios, and it is necessary to design a more precise tooling that provides more deterministic and higher confidence output.

As a future work, we are going to use the current version of Redfish specification as a basis, and design a new Redfish model which conveys necessary information, and provides the data necessary to perform the analysis. Also, we are going to run RGB against a simulated cluster which uses our new Redfish specification.

ACKNOWLEDGMENTS

This research is supported by the Cloud and Autonomic Computing site and High Performance Computing Center at Texas Tech University and the collaboration with Dell Inc, and DMTF.

REFERENCES

[1] 2017. *Redfish Scalable Platforms Management API Specification.* Technical Report.
[2] DMTF 2018. Distributed Management Task Force, Inc. (2018). https://www.dmtf.org
[3] H. Pyla K. Cameron W. Feng. R. Ge, X. Feng. 2007. *Power Measurement Tutorial for the Green500 List.* Document. Green500l.
[4] EEHPC WG. 2017. *EEHPC WG: Power Measurement Methology (version 2.0 RC 1.0).* Document. Green500.