Benchmarking Automated Hardware Management Technologies for Modern Data Centers and Cloud Environments

Elham Hojati
Department of Computer Science at Texas Tech University
elham.hojati@ttu.edu

Yong Chen
Department of Computer Science at Texas Tech University
yong.chen@ttu.edu

Alan Sill
High Performance Computing Center at Texas Tech University
alan.sill@ttu.edu

ABSTRACT

Traditional management standards are often insufficient to manage modern data centers at large scale, which motivates the community to propose and develop new management standards. The most popular traditional standard for monitoring and controlling the health and functionality of a system at hardware layer is Intelligent Platform Management Interface (IPMI). Redfish is a new hardware-based management technology designed as the next-generation management standard. The goal of this study is to investigate hardware management technologies and to find out if they are powerful enough to meet demands of modern data centers. Particularly, we focused on Redfish and IPMI, and we benchmarked and compared them from four different aspects: latency, scalability, reliability, and security. Our result shows that there is a trade-off between improving the performance of a system and increasing the security and the reliability of that. Our results show that Redfish is more secure and more reliable, but the performance of IPMI tends to be better.

KEYWORDS

Redfish, IPMI, cloud environments, data centers, benchmarking

ACM Reference Format:
Elham Hojati, Yong Chen, and Alan Sill. 2017. Benchmarking Automated Hardware Management Technologies for Modern Data Centers and Cloud Environments . In *UCC '17: 10th International Conference on Utility and Cloud Computing*. ACM, Austin, TX, USA, 2 pages. https://doi.org/10.1145/3147213.3149212

1 INTRODUCTION

In this research study, we focus on two hardware management technologies for data centers, IPMI (Intelligent Platform Management Interface) [1] and Redfish [2]. We compare them from four different aspects: latency, scalability, reliability, and security. Our result shows that there is a trade-off between improving the performance of a system and increasing the security and the reliability of that. Redfish is more secure, but the performance of IPMI is usually better than Redfish.

UCC '17, December 5–8, 2017, Austin, TX, USA

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

2 METHODOLOGY

This research mainly focuses on benchmarking and comparing two different hardware management tools for data centers, Redfishtool [3] and IPMItool [4]. We run Redfishtool in two different modes: Redfish in the raw mode; and Redfish in full tree traversal mode, we call it Redish in short. In the tree traversal mode, Redfish command traverses URIs until it finds asked information. But in the raw mode, it goes exactly through the specified URI, not all of them. [3].

2.1 Benchmarking Environment

This research was performed using the QUANAH cluster hosted at the High Performance Computing Center of Texas Tech University. The cluster was installed from the scratch with 64-bit version CentOS 7.2.1511 with kernel version 3.10. QUANAH cluster contains 224 Redfish enabled compute nodes from Dell Inc. The general specifications of the hardware of each node include: Intel Xeon CPU E5410 with 36 cores, 192GB RAM, and networks including 1Gbps Ethernet and 100Gbps Intel Omnipath. All required packages for running IPMItool and Redfishtool were installed on the controller node of the cluster, named Charlie.

2.2 Benchmarking Methodology

To benchmark the latency, we selected a list of management and monitoring actions, and for each action in the list, we provided the appropriate commands in IPMItool and Redfishtool (by considering which action in Redfishtool is equivalent to the action in IPMItool). For each tool, we developed a bash script that checks the latency of the action for each of the nodes of the cluster using that tool. By running the scripts, the latency of each action was evaluated and compared.
The list of actions selected to benchmark latency contains: Temperature Checking, Fan Speed Checking, Sensor Data Record (SDR) Checking, Chassis Status Checking, System Event Log (SEL) Checking , Field Replaceable Unit (FRU) Checking , and User Account Checking. To benchmark the scalability, we varied the size of the cluster under test, and observed the impact of the changes on the performance of running hardware monitoring commands on the cluster. To compare reliability and security, we considered network protocols used for transferring and securing information in Redfish and IPMI technology.

3 TESTS AND RESULTS

In the latency test, we primarily focus on benchmarking the latency of Redfish and IPMI by running various commands to query information of a node in the cluster. This process uses a black box approach for benchmarking and comparing these two technologies. Figure 1 reports one of these results: the latency of retrieving

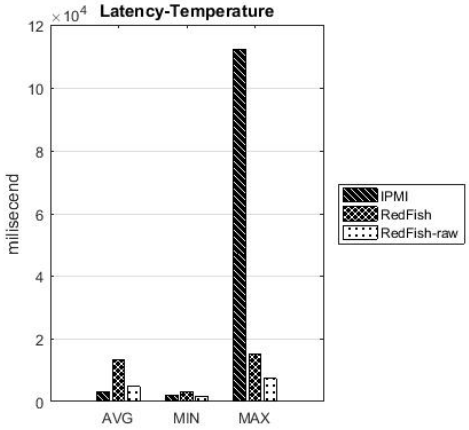

Figure 1: Latency comparison of querying temperature using IPMI, Redfish, and Redfish-raw

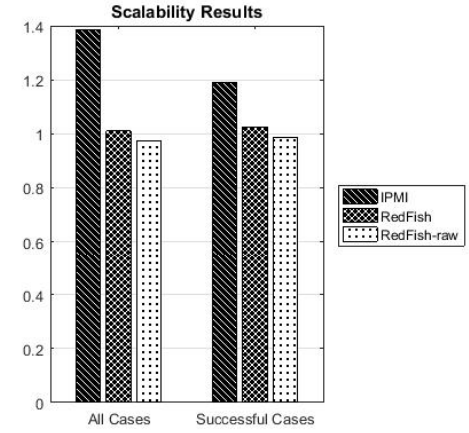

Figure 2: Scalability comparison of querying temperature using IPMI, Redfish, and Redfish-raw

temperature data using IPMI, Redfish (Redfish in full tree traversal mode), and Redfish-raw. Those results show that the performance of IPMI is generally better than Redfish based on the latency of the response. Figure 2 plots the scalability in querying temperature data using IPMI, Redfish, and Redfish-raw. The best scalability was achieved for the Redfish in the raw mode. from reliability aspect, IPMI uses UDP protocol for communication which is not a reliable protocol. But Redfish uses TCP protocol as its transport layer protocol, which is a reliable protocol with error checking and error recovery. Therefore, the reliability of Redfish is better than IPMI. Furthermore, Redfish uses HTTPS and TLS protocol to achieve security in its connections, which is stronger than RMCP+ used by IPMI to make its connections secure. Therefore the security of Redfish is better than IPMI.

Table 1: Overall Result

Tool	Latency	Scalability	Reliability	Security
IPMI	√√	√	×	×
Redfish	×	√√	√	√
Redfish-raw	√	√√√	√	√

Table 1 shows an overall comparison based on our benchmarking results. The traditional standard, IPMI, has better performance (considering the latency of running management and monitoring commands) than Redfish. Redfish is more reliable and more secure than IPMI. Redfish also has better simplicity of usage from our evaluations and experiences. The scalability result shows that Redfish in the raw mode has the best scalability, which is often critical for large-scale data centers nowadays.

4 CONCLUSION

Data centers and cloud environments are complex and dynamic. There is a strong desire for benchmarking and evaluating technologies and standards for dynamic cloud environments and data centers compared to traditional static systems. In this research, we focused on one of the most important traditional management standards, IPMI, and the new standard, Redfish. The goal of this research is to benchmark these hardware management technologies based on four different properties, latency, scalability, reliability, and security. We intended to develop quantitative comparison of these technologies and standards for modern scalable clouds and data centers.

This current research has also identified that the latency of Redfish can be further improved by issuing commands at once in a scatter/gather fashion. Moreover, the performance of a particular individual command will not matter significantly if many of them can be issued simultaneously and collectively. Even a large number of commands can be issued within a small amount of time. It would be difficult to perform such collective optimizations with IPMI as it uses a UDP transportation protocol and can have possible message losses. On the other hand, Redfish adopts a TCP transportation protocol and such collective optimizations are possible. We plan to further study these optimization strategies in the near future.

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.

REFERENCES

[1] 2013. *IPMI -Intelligent Platform Management Interface Specification Second Generation v2.0.* Document Revision 1.1. Intel Hewlett-Packard NEC Dell.
[2] 2016. *Redfish Scalable Platforms Management API Specification Document.* Technical Report.
[3] DMTF. 2017. Redfishtool. (2017). https://github.com/DMTF/Redfishtool
[4] BSD License. 2017. IPMItool. (2017). https://sourceforge.net/projects/ipmitool/