Big Data Benchmarks on Bare Metal Cloud

Hyungro Lee, Geoffrey C. Fox,
School of Informatics, Computing and Engineering
Indiana University Bloomington
{lee212, gcf}@indiana.edu

May 31, 2018

1 Abstract

High performance computing requires to deal with a large number of applications running on different environments, and bare metal cloud is promising to enable new hardware features but in a easier way than traditional HPC systems. Data that we need to deal with is growing exponentially although many big data software support processing them at scale.

We perform big data benchmark on public bare metal clouds to demonstrate compute performance with direct hardware access and block storage. The results indicate that Amazon and Oracle are competitive supporting high throughput and low latency with operations in parallel. We investigate further on storage options available on Oracle bare metal with different data sets and anticipate to evaluate petabyte-scale workloads on cluster configurations in the future.

2 Introduction

We started using public bare metal clouds to evaluate big data applications because we wanted to see how they actually perform with various big data workloads. Oracle Cloud Infrastructure offers mainly two types of bare metal cloud servers, one is a standard server type with Intel E5-2699 v3 @ 2.30GHz resulting in a total of 52 Oracle virtual CPUs (OCUs), 104 logical processors (BM.Standard1.52) and the other is a I/O dense server type with eight of local NVMe SSDs in a total size of 51.2 TB (BM.DenseIO1.52) [1]. Every server types allows to attach up to 512 TB block storage with 25,000 IOPS maximum. Amazon EC2 provides one bare metal server type (i3.metal) with two Intel Xeon E5-2686 v4 processors running at 2.3 GHz, with a total of 36 hyperthreaded cores, 72 logical processors [2] along with 15.2 TB size of local NVMe SSDs.

Fast and power hardware are expected to accelerate dynamic data processing and we have observed that various big data applications generate different workloads such as compute intensive, i/o intensive, memory intensive and combined these which is ideal to estimate actual performance with effective behavior in practice. We find that there are big data benchmark suites [3, 4] and actively used them for the bare metal cloud evaluation.

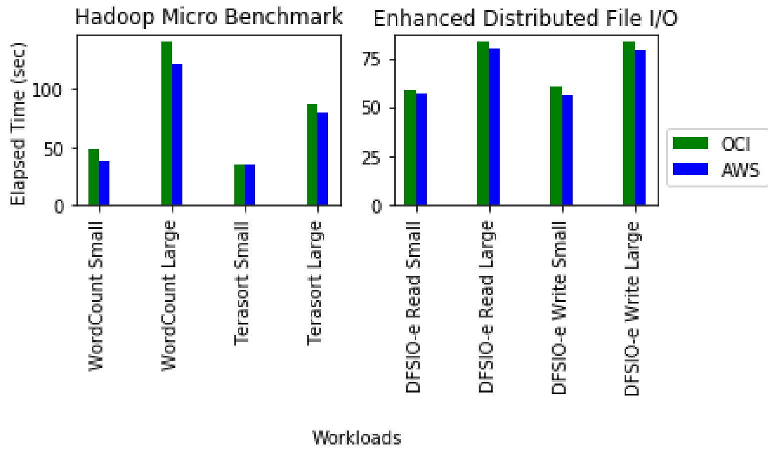

Figure 1: Big Data Benchmark on OCI and AWS

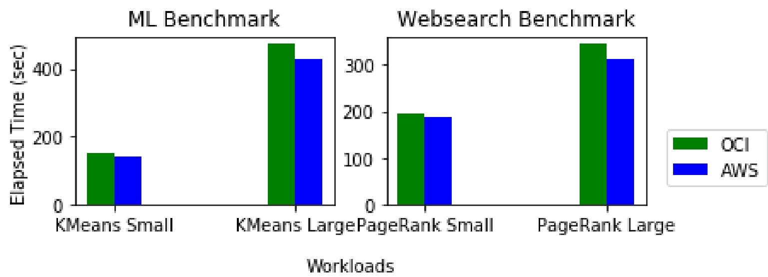

Figure 2: Other Benchmark on OCI and AWS

3 Result

3.1 System Setup

Amazon allows to specify IOPS (upto 32000 IOPS) through the Provisioned IOPS (SSD) option. We configured two volumes with a size of 50GB and 2500 IOPS and 1TB and 32000 IOPS on the i3.metal server type. BM.Standard52 and BM.DenseIO52 are chosen on the Oracle bare metal cloud with multiple storage configurations i.e. local storage, 50GB with 3000 IOPS and 1TB with 25000 IOPS and 16 TB. Detail descriptions will be added later.

3.2 Public Bare Metal Clouds

The first benchmark result is three workloads running on Hadoop systems. WordCound and Terasort workloads create both CPU stress and I/O where as DFSIO-E focuses on HDFS read/write intensive tasks. Enhanced Distributed File System I/O (DFSIO-e) is a part of HiBench workloads which generates aggregate performance curve between mapper and reducer functions while it's processing bulk data. Fig 1 2 shows comparisons towards two bare metal cloud offerings regarding to the six big data workloads.

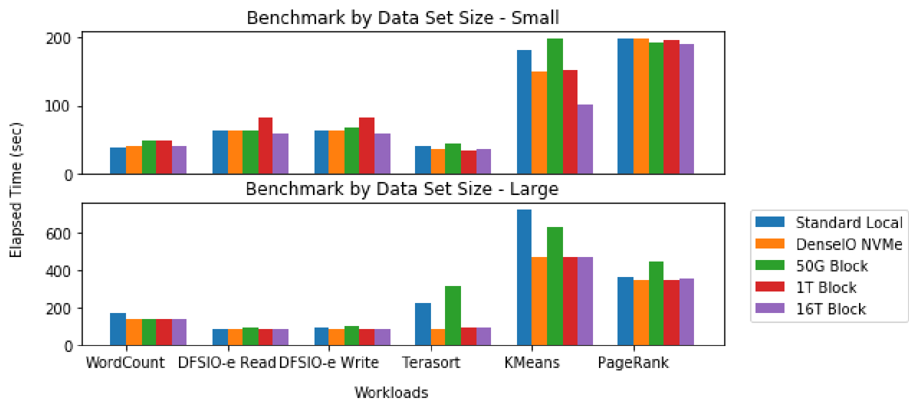

Figure 3: Benchmark Results between Small and Large size of Data Set

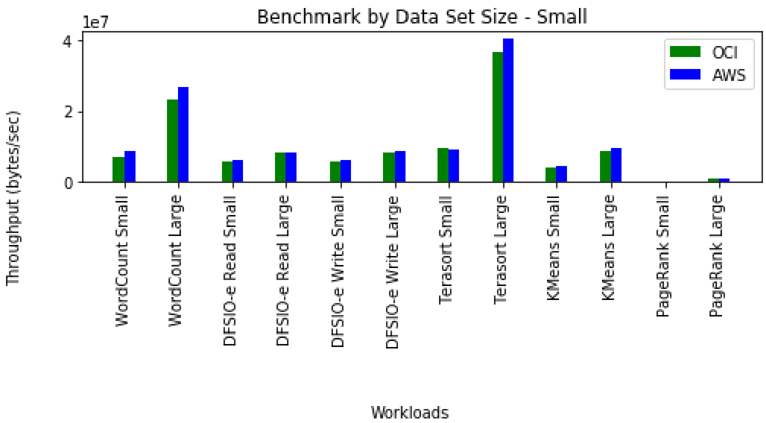

Figure 4: Benchmark Throughput on Bare Metal Cloud

3.3 Storage Choice

Our findings in this Fig 3 are competitive block storage performance compared to local NVMe storage. The ratio of 60 IOPS/GB is beneficial to small size workloads compared to 50 IOPS/GB by Amazon block storage. In addition, i/o intensive workloads with large data set perform effectively if we look at DFSIO-e workloads.

3.4 Throughput

Fig 4 shows results towards six workloads tested between AWS i3.metal and OCI BM.DenseIO52 with 1TB block storage attachment. Higher throughput (bigger bar) indicates better results where workloads with large data set perform effectively. Overall, two bare metal clouds produce similar throughput over these six workloads.

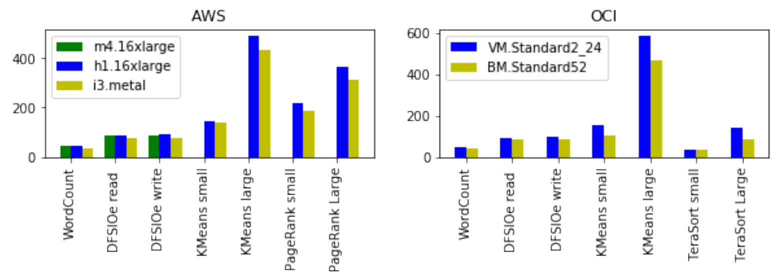

Figure 5: VM vs BareMetal (elapsed time - smaller is better)

4 Future Work

Bare metal cloud with direct hardware supports enables high performance compute and data processing for various workloads. We could not evaluate other workloads such as sql-based scan, aggregation and merge tasks using Hive and MongoDB. Streaming framework such as Flink, and Storm and deep learning toolkits such as TensorFlow, Caffe, and Torch are not included in this report. In addition, we anticipate to evaluate bare metal cloud systems with petabyte scale data sets on cluster configurations with a choice of job scheduling tools e.g. Mesos and Kubernetes. This will help us to design serverless computing architecture which is based on cluster management tools and containers.

5 Conclusion

We conducted big data benchmark on public clouds to provide a first run of bare metal compute performance. The results from two providers are competitive regarding to I/O throughput and compute performance through the benchmark results for six big data workloads. Additional test results are required to determine performance features and weak points on these platforms but we believe that the current results give you a chance to see performance benefits that bare metal cloud promises.

References

[1] Lee Gates. Predictable, consistent, high performance computing with oracle bare metal cloud compute service, 2017.

[2] AWS Amazon. Amazon ec2 bare metal instances with direct access to hardware, 2017.

[3] Shengsheng Huang, Jie Huang, Jinquan Dai, Tao Xie, and Bo Huang. The hibench benchmark suite: Characterization of the mapreduce-based data analysis. In *Data Engineering Workshops (ICDEW), 2010 IEEE 26th International Conference on*, pages 41–51. IEEE, 2010.

[4] Wanling Gao, Lei Wang, Jianfeng Zhan, Chunjie Luo, Daoyi Zheng, Zhen Jia, Biwei Xie, Chen Zheng, Qiang Yang, and Haibin Wang. Big data dwarfs:

Towards fully understanding big data analytics workloads. *arXiv preprint arXiv:1802.00699*, 2018.