A Comparison of Amazon Web Services and Microsoft Azure Cloud Platforms for High Performance Computing

Charlotte Kotas
Computational Sciences and Engineering Division
Oak Ridge National Laboratory
Oak Ridge, TN, USA
kotascw@ornl.gov

Thomas Naughton
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN, USA
naughtont@ornl.gov

Neena Imam
Computing and Computational Sciences Directorate
Oak Ridge National Laboratory
Oak Ridge, TN, USA
imamn@ornl.gov

***Abstract*—Advances in commercial cloud computing necessitate continual evaluation of the cloud's performance on a variety of applications. This work looks at compute oriented instances from Amazon Web Services and Microsoft Azure cloud platforms and evaluates them with several high-performance computing benchmarks, including HPCC and HPCG. These benchmarks illustrate that the most cost competitive solution depends on the application to be run.**

Keywords—Amazon AWS, Microsoft Azure, cloud computing, cloud costs, scientific computing

I. Introduction

Previous work on high performance computing (HPC) on cloud platforms has suggested that cloud platforms were not able to compete with dedicated clusters for HPC applications, primarily due to slower network performance [1] [2]. However, cloud computing platforms are now offering more instance types geared toward high performance computing, including features like faster networks, larger amounts of RAM, and faster processors. These instances are marketed toward users with computationally intensive workloads like data analytics, video encoding, and engineering applications [3]. This study evaluates these new instance types for HPC applications that utilize the message passing interface (MPI) library for communication.

Recently, Sadooghi et al. [4] considered several benchmarks on a wide range of Amazon Web Services (AWS) instance types, including the high-performance Linpack (HPL) benchmark for compute performance. However, the study mostly considers single instances rather than clusters. In particular, while the cost analysis considers the cluster performance of two instance types, this performance appears to be based only on the HPL benchmark, which does not stress the network bandwidth at these sizes, and thus the analysis may not apply to communication intensive applications. The present study attempts to find a more complete look at cluster performance by considering a range of benchmarks with a more varied compute to communication load. It also expands on the previous work by considering a new AWS instance type as well as the Microsoft Azure cloud platform.

This study looks at the performance of the AWS and Azure cloud platforms on two HPC benchmarks: the HPC Challenge (HPCC) and the High-Performance Conjugate Gradient (HPCG). These benchmarks test several aspects of computer system performance, including computational speed, memory bandwidth, and network bandwidth. The results presented here focus on the cluster level application performance for various computation and communication patterns.

II. Cloud Instance Types

A. AWS

This study used the Amazon Web Services (AWS) elastic compute cloud (EC2) system's c4.8xlarge Linux instance types. During the time these tests were run (Summer 2017), these instances represented the largest and fastest of the compute-optimized instances offered by AWS. These instances specify Intel Xeon E5-2666 v3 (Haswell) processors with a nominal 2.9 GHz clock speed, 36 vCPUs (essentially Intel Hyper-Threads), 18 cores, and 60 GB of random access memory (RAM). They allow 10 Gigabit networking between the nodes utilizing a single root input/output virtualization (SR-IOV) technology. This instance type has a theoretical peak performance of 46.4 GFLOPS/core at 2.9 GHz (or 835.2 GFLOP/instance), however the clock speed is not constant in these tests as Intel's Turbo Boost technology is enabled by default. Linux versions of these instances are available for $1.591 per hour per instance in the US East region.

The instances were configured to use the Open MPI library version 1.10 and GNU compilers version 4.8.3. Instances were launched into the same virtual private cloud (VPC), placement group, and availability zone in order access the high-performance network.

B. Azure

The following tests used the Microsoft Azure H16r instance type, which provides Intel Haswell E5-2667 V3 processors, 16 cores, and 112 GB of RAM. This processor has a nominal clock speed of 3.2 GHz. At 3.2 GHz, this processor has a theoretical peak performance of 51.2 GFLOP/core or 819.2 GFLOP/node.

This work was sponsored by the U.S. Department of Energy's Office of Advanced Scientific Computing Research.

978-1-5386-3025-9/18/$31.00 ©2018 IEEE

It should be noted that, as with the AWS instances, Turbo Boost is enabled by default. These instances offer an InfiniBand FDR network with remote direct memory access (RDMA) enabled. CentOS and Ubuntu Linux instances are available for $2.136 per hour per instance in the East US region, though prices may vary by region. (In the interest of full disclosure, compute time on the Azure system was donated by Microsoft for these benchmarks.)

The benchmarks were compiled against the Intel MPI library with GNU compilers version 4.8.5. For Azure, access to the high-performance network requires that the instances be launched in the same availability set and virtual network. Currently, Azure's high-performance network requires either the Intel MPI or Platform MPI library.

III. Benchmarks

A. High Performance Computing Challenge (HPCC)

The HPCC benchmark suite presents several tests to measure a range of performance parameters related to network bandwidth, memory bandwidth, local computational performance, and global computational performance [5]. Here, results are reported from:

- the HPL benchmark which measures global performance on solving a linear system of equations via LU decomposition
- the MPI FFT benchmark which measures global performance computing a distributed fast Fourier transform
- the PTRANS benchmark which measures global bandwidth during array operations (embodied by a parallel matrix transpose)

These benchmarks focus on the global system performance rather than the performance of individual nodes (or instances), as it has already been established that the individual node performance on cloud platforms can be equal to more traditional HPC environments [2].

The reference HPCC benchmark version 1.5.0 was compiled with the Intel Math Kernel Library (MKL) BLAS library. As the recommended problem size is based on the amount of RAM available, the input parameters were adjusted to use 80% of the available RAM for the HPCC benchmark (defining the N parameter). The (P, Q) parameters describe how the HPCC benchmark should internally divide the HPL and PTRANS benchmarks. For these tests, P was set to the number of instances and Q to the number of cores per instance. The NB parameter (block size) was set to 192 for all tests. Otherwise, input parameter values were as in the default "_hpccinf.txt" file.

B. High Performance Conjugate Gradient (HPCG)

The HPCG benchmark was developed as a counterpoint to the HPL benchmark and is meant to illustrate system performance on more communication intensive sparse matrix patterns [6]. The key computational kernels of the conjugate gradient benchmark include a global dot product, sparse matrix-vector multiplication, vector updates, and a "local symmetric Gauss-Seidel smoother."

The reference HPCG benchmark version 3.0 was used for these tests. The HPCG benchmark input parameters include (NX, NY, NZ), which defines the size of the sub-matrix to be computed by each MPI process, and the run time in seconds.

IV. Experimental Data

The experiments were run by creating a cluster of AWS or Azure instances, running the benchmarks associated with that many instances, and then destroying the instances before moving on to the next set of benchmarks. In general, only one set of tests was run for each cluster size, but a couple of unusual data points were retested, with both results marked on the following graphs. The cluster sizes were scaled from 1 to 32 nodes.

A. Repeatability

The HPCG benchmark was used to test the repeatability of the cloud configuration by launching 2 instances and running 5 iterations of the benchmark on 3 instance sets. The parameters for these tests were NX=NY=NZ=64 and a run time of 120 seconds for AWS and 300 seconds for Azure. For AWS, the performance on these 15 runs ranged from 18.2 to 18.8 GFLOP/s, a difference of 3.1%. For Azure, the performance ranged from 15.0 to 15.2 GFLOP/s, a difference of 1.3%. Based on these measurements, the cloud performance is assumed to be relatively stable.

B. HPCC PTRANS

The HPCC PTRANS benchmark determines the global bandwidth for a parallel matrix transpose operation. Fig. 1 compares the measured bandwidth observed with AWS and Azure, which illustrates that Azure's faster network performs better on this benchmark, as expected. Azure's eight instance test case was rerun because it seemed to produce an unexpectedly low bandwidth, and on the second run performed better. It should be noted that due to Azure's larger available RAM, HPCC was configured to allow it to handle larger $N \times N$ matrices. For instance, at 1 node, AWS used N of 38,688, while Azure used 52,896. As the number of instances increases, the matrix sizes increase nearly linearly with N^2.

C. HPCC MPI FFT

The HPCC MPI FFT benchmark computes a global FFT, with an array size that increases as the number of instances increases. It should be noted that the MPI FFT only utilizes a power-of-two number of processes and that it only computes a power-of-two length array. Because Azure has nearly twice the available RAM as AWS, it is computing an FFT that is twice the length of AWS for a given cluster size. The similarities in the 16 and 20 node cases are likely due to the power-of-two limitations, as these cases are actually computing the same size FFT with the same number of processes. It is unclear how these processes are distributed among the instances, but it is possible that they are distributed the same way for the two tests. Fig. 2 shows that Azure performs better than AWS for the MPI FFT, likely due to its faster network. Both of the eight instance Azure results are shown in the graph, with the better performance coming from the second test with the larger PTRANS bandwidth.

D. HPCC HPL

The HPL benchmark is designed to show the computational power of a cluster with relatively moderate communication overhead. Fig 3 shows that AWS and Azure scale very similarly

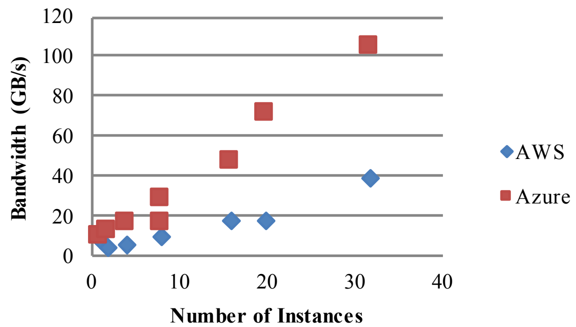

Fig.1 HPCC PTRANS bandwidth for AWS and Azure as a function of cluster size.

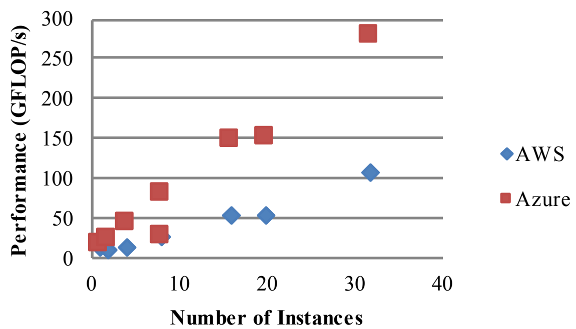

Fig. 2 HPCC MPI FFT performance for AWS and Azure as a function of cluster size.

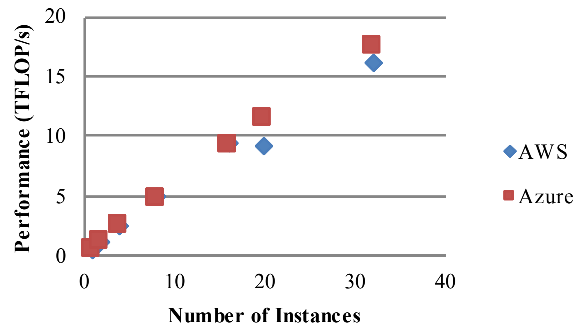

Fig. 3 HPCC HPL performance for AWS and Azure as a function of cluster size.

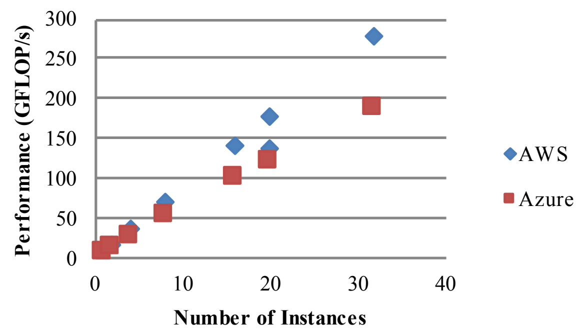

Fig. 4 HPCG performance for AWS and Azure as a function of cluster size.

for the HPL benchmark. Since the underlying processors have a similar amount of computational power, this is not a surprising result. (Both Azure eight node tests are shown in this figure, but produce nearly identical results as the computation is not communication bound.) AWS performance appears to "level off" between 16 and 20 instances, which is surprising. However, as this is a single result from a benchmark with many parameters, it is possible that this result is an anomaly, particularly as scaling continues as more instances are added.

E. HPCG

The HPCG benchmark was run with the same parameters for AWS and Azure, 1 MPI process per core, NX=NY=NZ=64, and a run time of 300 seconds. Because Azure has fewer cores (16 vs 18), it is running a slightly smaller problem size than AWS, but the work per core should be the same. The HPCG result shown in Fig. 4, with AWS outperforming Azure, is a little surprising, as Azure's faster processors and network should give it an advantage on this benchmark. Looking at some of the internal timing measurements shows that Azure spends a significantly longer time in the global dot product section of the benchmark. For the reference version of HPCG, this routine is implemented using the MPI_Allreduce function, and it may be that the difference is due to differences in the Intel MPI library and the Open MPI library implementations or configurations.

V. Cost Analysis

The HPCC and HPCG benchmarks can provide measures of the raw performance of the different cloud systems, but a more complete picture should also include the operating cost for running these benchmarks. This section discusses a few ways to quantify the operating costs based the various benchmarks above. However, these measures can only capture the dimensions considered by the benchmarks. For instance, the HPL benchmark is designed to produce a large FLOP count with a moderate communication cost. It will not reflect the advantage of having a faster network as strongly as a communication intensive benchmark like the MPI FFT. Although the desired applicationIn order to predict how a given application will behave, it is important to have benchmarks with similar scaling characteristics.

In terms of theoretical compute power, AWS's c4.8xlarge instance types can provide 1.889 PFLOP/Dollar (based on a cost of \$1.591/hour and 835.2×10^9 FLOP/s), while Azure's H16r instances rate 1.381 PFLOP/Dollar (based on \$2.136/hour and 819.2×10^9 FLOP/s).

Of course, the theoretical measurement for a single instance does not reflect the performance that is actually achievable across multiple nodes for a realistic workload. By using the performance data from the HPCC HPL, HPCC MPI FFT, and HPCG benchmarks, the FLOP/Dollar can be computed across a range of workloads. As shown in Fig. 5 for the HPL benchmark,

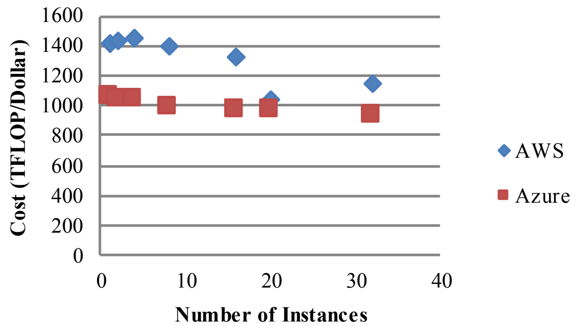

Fig. 5 Cost analysis based on the HPCC HPL benchmark as a function of cluster size for AWS and Azure.

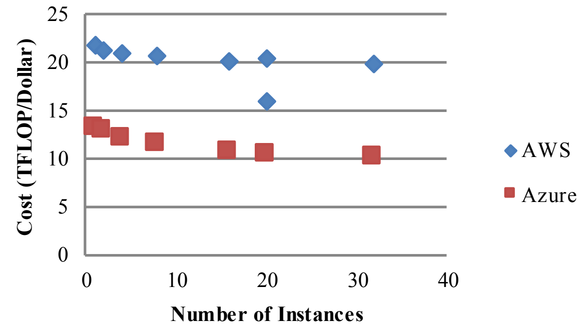

Fig. 7 Cost analysis based on the HPCG benchmark as a function of cluster size for AWS and Azure.

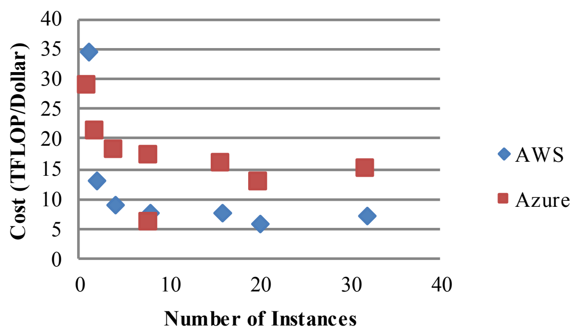

Fig. 6 Cost analysis based on the HPCC MPI FFT benchmark as a function of cluster size for AWS and Azure.

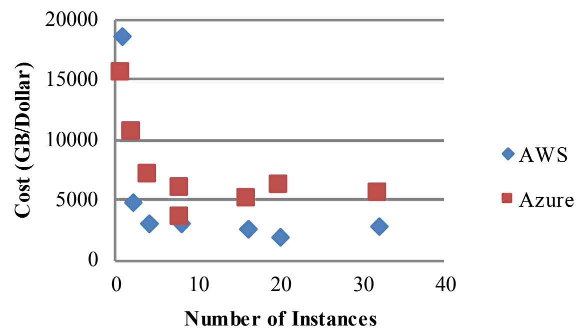

Fig. 8 Cost analysis based on the HPCC PTRANS benchmark as a function of cluster size for AWS and Azure.

Fig. 6 for the MPI FFT benchmark, and Fig. 7 for the HPCG benchmark, the best performance to cost ratio depends on the algorithm. AWS provides a larger number of floating point operations per dollar across node sizes for HPL (Fig. 5) and HPCG (Fig. 7), even for the 20-node case that had unexpectedly poor performance. Azure is better for most of the MPI FFT configurations (Fig. 6), with the exceptions being the single node case where MPI is communicating with local memory and the anomalous eight node cluster.

Similarly, the HPCC PTRANS benchmark can be used as a measure of the communication capability of the networks by normalizing the observed GB/s with the cost of the system. As shown in Fig. 8, Azure provides a better performance to cost ratio than Amazon, due to its faster network infrastructure.

VI. Conclusion

This study looks at the performance of several HPC benchmarks on the AWS and Azure cloud platforms, focusing on the compute oriented c4.8xlarge and H16r instance types. However, which cloud platform is the cheapest for a given use case depends on the computation and communication patterns of the application. This study finds that at the point in time that the tests were run, the AWS c4.8xlarge was cheaper in terms of raw computing, while Azure's H16r had cheaper bandwidth. This suggests communication intensive applications may benefit from the Azure H16r's faster network and larger RAM, resulting in an overall cheaper solution. Considering that cloud providers are continually improving their offerings, the best way to determine how any given application will perform in the current cloud environment is to test it on the prospective system.

Acknowledgments

Author C. K. thanks the supportive staff at AWS and Azure for many helpful discussions.

References

[1] K. R. Jackson, L. Ramakrishnan, K. Muriki, S. Canon, S. Cholia, J. Shalf, H. J. Wasserman and N. J. Wright, "Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud," in *2nd IEEE International Conference on Cloud Computing Technology and Science*, Indianapolis, IN, USA, 2010.

[2] P. Mehrotra, J. Djomehri, S. Heistand, R. Hood, H. Jin, A. Lazanoff, S. Saini and R. Biswas, "Performance Evaluation of Amazon EC2 for NASA HPC Applications," in *Proceedings of the 3rd Workshop on Scientific Cloud Computing*, Delft, The Netherlands, 2012.

[3] Amazon Web Services, "Amazon EC2 Instance Types," 2017. [Online]. Available: https://aws.amazon.com/ec2/instance-types/. [Accessed 25 August 2017].

[4] I. Sadooghi, J. H. Martin, T. Li, K. Brandstatter, K. Maheshwari, T. P. P. de Lacerda Ruivo, G. Garzoglio, S. Timm, Y. Zhao and I. Raicu, "Understanding the performance and potential of cloud computing for scientific applications," *IEEE Transactions on Cloud Computing,* vol. 5, no. 2, pp. 358-371, 2017.