Analyzing the Network for AWS Distributed Cloud Computing

Anshul Gandhi, Justin Chan
Stony Brook University *

ABSTRACT

Cloud computing is a global service. Cloud Service Providers, such as AWS, allow users to launch VM instances on multiple data centers (regions) around the world. However, the network connectivity and bandwidth between these different geographically distributed regions varies significantly depending on the user's location. In this paper, we analyze the network performance between pairs of AWS instances hosted on all available regions. We leverage our analysis to derive the optimal hosting region for web service providers depending on the customer locations.

1. INTRODUCTION

Cloud computing enables users in any part of the world to access IT services on a pay-as-you-go basis. Such services, including VMs and storage resources, are managed by Cloud Service Providers (CSPs), such as AWS [1] from Amazon, Azure [2] from Microsoft, and Cloud Platform [3] from Google. CSPs typically host such services on multiple, *geographically distributed data centers*, allowing users to launch VMs on different physical locations. As a result, CSPs offer a truly Distributed Cloud Computing (DCC) experience to users. For example, AWS offers users with nine different physical "regions" in which VMs can be launched [4]. These nine regions, listed in Table 1, are spread across five different continents.

In such a DCC environment, users have a choice as to *which region(s)* their VM(s) should be launch on. While the efficacy of compute and storage resources does not depend on the location of the VMs, the networking capability critically depends on the distance between the user and the VM hosting location. This is especially important for today's data-centric workloads and applications. Typically, users pick the region closest to them for launching VMs. For example, students in our cloud computing class at Stony Brook University (New York) were assigned the US East (N.Virginia) region by default when using AWS's EC2 service to launch VMs. However, this is not always the optimal choice. Consider a Web Service Provider (WSP) who is interested in leveraging the AWS cloud to host her website and associated data sets. While the WSP would benefit from lower latency by hosting the VMs in a region closer to her, the latency of the customers of the website will depend on their location relative to the host region. Thus, the WSP must take into account the location of her potential customers. For example, a WSP that exports products from New York to Ireland might benefit from hosting the product catalogues on the EU (Ireland) region as opposed to the US East region to provide lower latency, and thus a faster browsing experience, for customers. We refer to the above example as the *WSP problem*, and will revisit it later in the paper as a case study.

Index	Name	Location	Code
1	US East	N.Virginia	us-east-1
2	US West	Oregon	us-west-2
3	US West	N.California	us-west-1
4	EU	Ireland	eu-west-1
5	EU	Frankfurt	eu-central-1
6	South America	Sao Paulo	sa-east-1
7	Asia Pacific	Tokyo	ap-northeast-1
8	Asia Pacific	Sydney	ap-southeast-2
9	Asia Pacific	Singapore	ap-southeast-1

Table 1: AWS data center regions and locations.

In order to make the right VM placement decisions in the case of a WSP, information about the network performance between the WSP (source) and the CSP, and the CSP and the customer (target) is required. In this paper, we analyze the network performance between pairs of VMs hosted on different AWS regions. In particular, we analyze the *ping times, the download times, and upload times* between all possible pairs of VMs launched in different regions. The collected data and its analysis is the major contribution of this paper. To the best of our knowledge, this is the first paper to analyze the network performance between pairs of geographically distributed VMs launched on a commercial public cloud.

Our experiments reveal that there is significant variation in network performance (typically, *factor of 10*) between pairs of VMs launched in different regions. This variation provides interesting tradeoffs for the WSP problem in terms of the choice of the host CSP. Surprisingly, we find that there is little variation in performance for a given pair of regions with respect to time of day and day of week.

The rest of the paper is organized as follows. We describe our experimental setup in Section 2. We then present our network measurement and analysis results in Section 3, including ping times (Section 3.1), download times (Section 3.2), and upload times (Section 3.3). Based on the measured data, we analyze the WSP problem in Section 4. We discuss related work in Section 5, and conclude the paper in Section 6.

2. EXPERIMENTAL SETUP

We set up multiple EC2 micro instances on all nine regions listed in Table 1. For file download and upload tests, we create multiple, randomly generated 1MB files, and use `scp` for download and upload across regions. For ping times, we make use of the linux `ping` command.

*Supported by an AWS in Education Grant Award.

Copyright is held by author/owner(s).

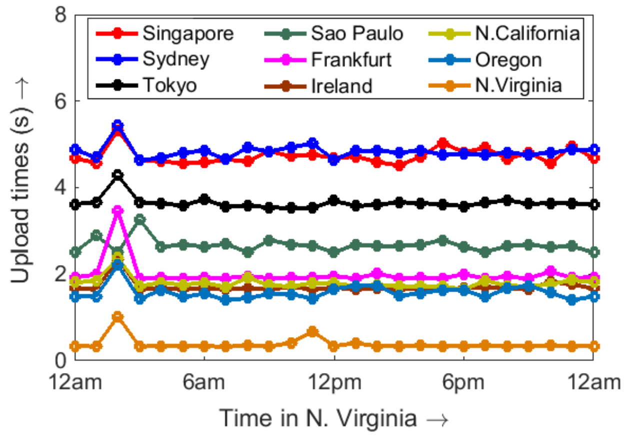

Figure 1: Upload times for a 1MB file from N.Virginia to other regions on 2/20/2015.

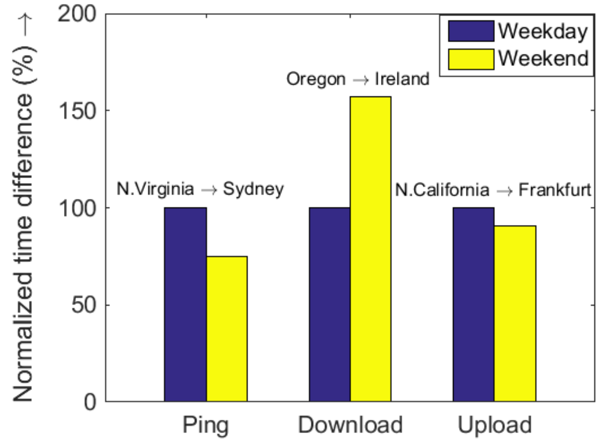

Figure 2: Normalized performance variation between certain weekday and weekend tests.

2.1 Measurement methodology

For each pair of regions, say source and target, we perform multiple ping, download, and upload tests between VMs of these regions. If the source and target region are the same, we perform the tests between two different VMs in the region. The tests are conducted over multiple days, including weekdays and weekends, and are repeated multiple times a day. In order to eliminate outliers, we discard any measurements that are outside of two standard deviations of the mean of the collected data. Using this heuristic, we discard approximately 2% of the data. The results in Sections 3.1-3.3 are averaged over the remaining (undiscarded) data.

3. NETWORK PERFORMANCE ANALYSIS

In our experiments, we find minor *temporal variations* in performance with respect to time-of-day and day-of-week. The most significant of these variations is shown in Figure 1, which illustrates our upload times, every hour, from N.Virginia to the other nine regions on 20th February, 2015. Observe the spike in upload times at 2am. Since the spike appears for upload times to all regions, we believe the cause has to do with the (common) source region, N.Virginia in this case. We also observed similar spikes at 2am for a few other days. The spikes were significant only for upload tests. Our best guess is that the N.Virginia data centers have some network maintenance scheduled at 2am.

We also found differences in performance between certain weekday tests and weekend tests. The most significant of these variations are shown in Figure 2. The measurements from the specific download test shown in Figure 2 had sufficient variation that they were discarded (based on the two standard deviations rule).

We now focus on *spatial variations* in performance with respect to source and target regions. Figure 3(a) shows the approximate distance, in miles, between the different regions. We concisely display the results in the form of a matrix. Here, the rows represent the source region and the columns represent the target region. In the matrix, the light shaded (white) regions indicate lower values and the dark shaded (red) regions indicate higher values. Note that, for the case of distances between two regions, the matrix is symmetric (transpose of the matrix returns the same matrix).

3.1 Ping times

Figure 3(b) shows the results of the ping time tests for all pairs of regions. We see that the ping times vary significantly. For example, ping times from N.Virginia (row 1), ignoring the self ping time (column 1), vary from 70ms (to Oregon) to 233ms (to Singapore). Interestingly, the ping times to Ireland are faster than those to N.California. The variation between minimum and maximum ping times is much larger for (the rows) Ireland and Frankfurt. Note that the ping times matrix is not truly symmetric, which is in agreement with prior studies on ping times [5]. For example, the largest ping time from N.Virginia is to Singapore, whereas the largest ping time to N.Virginia is from Sydney, although the difference is not too large. As expected, the lowest ping times are on the diagonal (pings between VMs in the same region).

Note the dark zones of the matrix corresponding to rows 4-6 and columns 7-9, and likewise, the symmetric zones corresponding to rows 7-9 and columns 4-6. These represent the large ping times between VMs launched in Europe and South America on one side, and Asia Pacific on the other. These large ping times are in agreement with the large physical distances between these regions as shown in Figure 3(a). Likewise, note the light zones of the matrix corresponding to rows 7-9 and columns 7-9. These represent the small ping times between VMs in the Asia Pacific group of regions. Again, these small ping times are in agreement with the (relatively) shorter physical distances between the Asia Pacific regions as shown in Figure 3(a). The shorter physical distance between Oregon and N.California, and Ireland and Frankfurt, likewise results in lighter shaded zones in rows and columns 2-3 and 4-5, respectively.

3.2 Download times

Our download times experiments involve downloading randomly (pre-)generated 1MB files from various source regions to the target region. Figure 4(a) shows the results of the download times from source regions (columns) to target regions (rows). For a given target, ignoring the self download times, we see significant variation in download times. While, qualitatively, this is to be expected, it is interesting to note the quantitative variations. For N.Virginia, there is a factor 3 difference between maximum (from Sydney) and minimum download times (from Oregon). For Ireland and Frankfurt,

	N.Virginia	Oregon	N.California	Ireland	Frankfurt	Sao Paulo	Tokyo	Sydney	Singapore
N.Virginia	0	2309	2282	3424	4168	4690	6825	9687	9735
Oregon	2309	0	593	4701	5274	6673	4909	7653	8201
N.California	2282	593	0	5103	5712	6294	5333	7491	8631
Ireland	3424	4701	5103	0	744	5757	6026	10763	7037
Frankfurt	4168	5274	5712	744	0	6093	5814	10240	6374
Sao Paulo	4690	6673	6294	5757	6093	0	11513	8312	9941
Tokyo	6825	4909	5333	6026	5814	11513	0	4842	3300
Sydney	9687	7653	7491	10763	10240	8312	4842	0	3918
Singapore	9735	8201	8631	7037	6374	9941	3300	3918	0

(a) Matrix for geographical distance (in miles).

	N.Virginia	Oregon	N.California	Ireland	Frankfurt	Sao Paulo	Tokyo	Sydney	Singapore
N.Virginia	1	70	84	76	90	123	172	219	233
Oregon	68	1	20	159	158	180	104	176	166
N.California	82	20	1	170	172	202	106	158	178
Ireland	81	157	172	1	22	193	283	308	349
Frankfurt	90	159	169	22	1	217	274	297	358
Sao Paulo	123	180	202	192	217	1	288	314	353
Tokyo	172	104	106	277	272	287	1	104	74
Sydney	239	177	158	309	298	314	104	1	175
Singapore	233	166	178	348	360	352	74	175	1

(b) Matrix for ping times (in ms).

Figure 3: Matrices showing (a) the distance, in miles and (b) ping times, in milliseconds, between VM instances located in different AWS regions.

the difference is almost a factor 10.

Note, again, the dark zones of the matrix corresponding to rows 4-6 and columns 7-9, and likewise, the symmetric zones corresponding to rows 7-9 and columns 4-6. Likewise, note the light zones of the matrix corresponding to rows 7-9 and columns 7-9. This is similar to our observations for the ping times.

3.3 Upload times

For upload tests, we upload randomly generated 1MB files from source regions to various target regions. Figure 4(b) shows our results for upload times from source regions (rows) to target regions (columns). The results for upload times are very similar to those for download times, and this similarity can be easily observed by looking at the heat map and numbers in Figures 4(a) and 4(b).

4. CASE STUDY: THE WSP PROBLEM

We now discuss a simple case study that highlights the benefits of the above network performance measurements. Consider a WSP that is located in a given region, R_w, and whose potential customers are located in a different region, R_c. For now, we assume that the customers are located in one region; we will soon relax this assumption. The WSP is interested in hosting the web service and associated data on AWS, and would like to optimize the placement of their VMs by finding the optimal hosting region. The location of the hosting region, R_h, is critical for the performance of the WSP. The "closer" (in terms of network connectivity and performance) R_h is to R_w, the smaller is the time taken by the WSP to upload new data. Likewise, the closer R_h is to R_c, the smaller is the access time for WSP's customers. Note that we are ignoring the access time between R_c and the customers; we will consider this explicitly in future work.

Let $0 \leq \alpha \leq 1$ be the relative importance of the WSP's upload times to that of the customers' access latency. Then, the optimal hosting region, R_h^*, can be derived by finding region R_i that minimizes the following objective function:

$$\mathrm{Obj}(R_i) = \alpha \cdot \mathrm{UL}(R_w, R_i) + (1 - \alpha) \cdot \mathrm{DL}(R_c, R_i)$$

Here, UL and DL are the Upload and Download matrices shown in Figure 4. Intuitively, one would expect that R_h^* is either R_w or R_c. In fact, if we rely on the Distance matrix in Figure 3(a) to optimize the objective function, our optimal hosting region would be R_w or R_c. Surprisingly, this is *not* always true. Consider the case where $\alpha = 0.5$ and the WSP is in R_w = Tokyo and the customers are in R_c = Ireland. In this case, the value of the objective function, $\mathrm{Obj}(R_i)$, for $R_i = R_w$ is about 2.95, and for $R_i = R_c$ is about 2.9. However, for R_i = N.Virginia, the value is about 2.55. This means that hosting the data in the N.Virginia region would provide 12% more value than hosting it locally in Tokyo or Ireland. Similarly, we get about 10% more value by hosting the data in N.Virginia when the WSP is in Ireland or Frankfurt, and the customers are in Singapore. Our above optimization strategy can be easily extended to the problem of finding the optimal hosting regions when customers upload data to an intermediate region, and the service provider downloads customer data from the intermediate region and provides some results (here, R_w and R_c are interchanged).

Now consider the more challenging scenario where the customers are spread over two regions, R_{c1} and R_{c2}, with relative weights $0 \leq \beta \leq 1$ and $(1 - \beta)$. In this case, the objective function is:

$$\mathrm{Obj}(R_i) = \alpha \cdot \mathrm{UL}(R_w, R_i) + (1 - \alpha) \cdot \{\beta \cdot \mathrm{DL}(R_{c1}, R_i) + (1 - \beta) \cdot \mathrm{DL}(R_{c2}, R_i)\}$$

We can again use the optimization approach to derive the R_i that minimizes the above objective function by leveraging the Download and Upload time matrices from Figure 4. In this case, the optimal region provides 5-12% more value than hosting at the source (WSP) region. Likewise, the objective function can be extended to consider multiple WSP locations and multiple customer locations.

Our empirical measurements allow us to easily optimize for the hosting regions. Understanding the causes behind these optimal hosting regions will be part of future work.

5. RELATED WORK

We now discuss related work that focuses on performance analysis and measurement of public clouds. CloudCmp [6] examines the performance of public cloud provider VMs. The authors analyze the network performance between VMs of the same data center, and between VMs of different data centers hosted in the US. Hajjat et al. [7] conduct a measurement study of EC2-deployed VMs to understand the impact of provider-specified policies and limits. The authors do analyze the network performance between VMs launched in the N.Virginia and N.California regions. Our study specif-

	N.Virginia	Oregon	N.California	Ireland	Frankfurt	Sao Paulo	Tokyo	Sydney	Singapore
N.Virginia	0.3	1.6	1.8	1.7	2.0	2.6	3.4	4.8	4.6
Oregon	1.6	0.3	0.6	3.2	3.2	3.6	2.2	3.6	3.3
N.California	1.8	0.7	0.3	3.3	3.6	4.0	2.3	3.2	3.5
Ireland	1.7	3.2	3.4	0.4	0.7	4.0	5.6	6.1	6.7
Frankfurt	1.9	3.2	3.4	0.7	0.3	4.5	5.3	5.6	6.8
Sao Paulo	2.5	3.5	3.9	3.8	4.3	0.7	5.6	6.1	6.8
Tokyo	3.5	2.2	2.2	5.4	5.3	5.7	0.6	2.2	1.6
Sydney	4.6	3.5	3.1	6.0	5.7	6.1	2.2	0.5	3.5
Singapore	4.6	3.3	3.5	6.7	6.8	7.0	1.8	3.6	0.4

(a) Matrix for 1MB download times (in s).

	N.Virginia	Oregon	N.California	Ireland	Frankfurt	Sao Paulo	Tokyo	Sydney	Singapore
N.Virginia	0.3	1.6	1.8	1.7	1.9	2.5	3.4	4.6	4.5
Oregon	1.5	0.3	0.6	3.2	3.1	3.5	2.2	3.5	3.3
N.California	1.8	0.6	0.3	3.3	3.6	3.9	2.2	3.1	3.5
Ireland	1.7	3.2	3.3	0.3	0.7	3.8	5.5	5.9	6.6
Frankfurt	1.9	3.1	3.4	0.7	0.3	4.3	5.3	5.7	6.9
Sao Paulo	2.5	3.5	3.9	3.7	4.2	0.3	5.5	6.0	6.8
Tokyo	3.4	2.1	2.2	5.4	5.3	5.6	0.3	2.2	1.6
Sydney	4.6	3.5	3.1	5.9	5.9	6.1	2.2	0.3	3.5
Singapore	4.5	3.3	3.5	6.8	6.8	6.8	1.6	3.4	0.3

(b) Matrix for 1MB upload times (in s).

Figure 4: Matrices showing (a) the download times, in seconds and (b) the upload times, in seconds, for a 1 MB file size between VM instances located in different AWS regions.

ically analyzes the network performance of VMs hosted in all 9 regions (7 countries), and reveals some interesting results for such geographically distributed VMs, as discussed in Section 3.

Wang et al. [8] analyze the impact of virtualization on network performance of EC2 instances. Their study reveals that network performance is impacted by virtualization, and is often unstable. Schad et al. [9] benchmark EC2 with the objective of analyzing variance in the cloud. They empirically analyze EC2 from various standpoints - CPU performance, memory speed, disk I/O, network, instance startup, and S3 access time. Palankar et al. [10] study the Amazon S3 service and evaluate file transfer times between EC2 and S3. Gandhi et al. [11, 12] examine the compute capacity of cloud VMs. In all of the above studies, the authors focus on the performance of instances launched in a given region; our focus in this paper is on the performance of the network connectivity between instances launched in different regions.

6. DISCUSSION AND CONCLUSION

Today's geographically distributed cloud computing platforms provide users with a truly DCC experience. However, leveraging DCC requires an understanding of the trade-offs between different hosting locations. In this paper, we use AWS's EC2 as our case study of a DCC system, and measure the network performance between VMs launched in nine different EC2 hosting regions. Our measurements of ping times, download times, and upload times reveal that there is significant variation in network performance between different regions. By ignoring these variations, users' AWS-deployed applications can suffer severe performance losses. We demonstrate, using a simple web service provider example, that by leveraging our AWS DCC network performance measurements, users can benefit from the distributed nature of cloud computing platforms and optimize their application performance accordingly.

In future work, we will further expand our measurement study to include path information and delays between VMs in different regions. This will help us understand the bidirectional connectivity between different regions and will allow us to analyze the *causes* of network delays in more detail, as has been emphasized in recent work [13]. We will also consider more interesting and practical use cases such as content delivery networks provided by Akamai and Amazon, and examine the impact that the content hosting location has on end-to-end performance. Further, we will analyze the network performance between hosting locations and end-users by leveraging remote (Atlas [14]) probes.

7. REFERENCES

[1] "Amazon Web Services." `http://aws.amazon.com`.

[2] "Microsoft Azure." `http://azure.microsoft.com`.

[3] "Google Cloud Platform." `https://cloud.google.com`.

[4] I. Amazon Web Services, "Global Infrastructure." `http://aws.amazon.com/about-aws/global-infrastructure`.

[5] V. Paxson, "End-to-end Routing Behavior in the Internet," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 601–615, 1997.

[6] A. Li, X. Yang, S. Kandula, and M. Zhang, "CloudCmp: Comparing Public Cloud Providers," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, pp. 1–14, 2010.

[7] M. Hajjat, R. Liu, Y. Chang, T. S. E. Ng, and S. Rao, "Application-Specific Configuration Selection in the Cloud: Impact of Provider Policy and Potential of Systematic Testing," in *Proceedings of IEEE INFOCOM'15*, 2015.

[8] G. Wang and T. S. E. Ng, "The Impact of Virtualization on Network Performance of Amazon EC2 Data Center," in *Proceedings of the 29th Conference on Information Communications*, (San Diego, CA, USA), pp. 1163–1171.

[9] J. Schad, J. Dittrich, and J.-A. Quiané-Ruiz, "Runtime Measurements in the Cloud: Observing, Analyzing, and Reducing Variance," *Proceedings of VLDB Endowment*, vol. 3, no. 1-2, pp. 460–471, 2010.

[10] M. R. Palankar, A. Iamnitchi, M. Ripeanu, and S. Garfinkel, "Amazon S3 for Science Grids: A Viable Solution?," in *Proceedings of the 2008 International Workshop on Data-Aware Distributed Computing*, (Boston, MA, USA), pp. 55–64.

[11] A. Gandhi, P. Dube, A. Karve, A. Kochut, and L. Zhang, "Adaptive, Model-driven Autoscaling for Cloud Applications," in *Proceedings of the 11th International Conference on Autonomic Computing*, pp. 57–64, 2014.

[12] A. Gandhi, P. Dube, A. Karve, A. Kochut, and L. Zhang, "Modeling the Impact of Workload on Cloud Resource Scaling," in *Proceedings of the 26th International Symposium on Computer Architecture and High Performance Computing*, pp. 310–317, 2014.

[13] A. Singla, B. Chandrasekaran, P. B. Godfrey, and B. Maggs, "The Internet at the Speed of Light," in *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, pp. 1:1–1:7, 2014.

[14] "RIPE Atlas." `https://atlas.ripe.net`.