Modeling Cloud Storage: A Proposed Solution to Optimize Planning for and Managing Storage as a Service

Anita L. Timmons, Pavel Fomin, James Wasek

Department of Engineering Management & Systems Engineering School of Engineering and Applied Science, The George Washington University, Washington, DC

Abstract. Cloud-computing service providers are currently viewed as the best solution to the global need for massive data systems because of their superior flexibility, scalability, and cost benefits. Cloud computing that is enabled by virtualized services is still constrained, however, by the capacities of the underlying physical systems that are combined into sharable pools of resources. The next challenge for computation systems will arise when even the cloud is not sufficient. What comes after cloud migration and adoption? In this paper, we examine how service providers can manage cloud storage resources and costs when the amount of collected data to be retained grows exponentially, to the point that it strains even virtualized resource capacities. We assess the analytical frameworks being developed to identify which storage architectures can best accommodate the specific needs of large data storage consumers. We also investigate the areas in which these fail to fully address the problem, and propose solutions. We argue that a cloud storage framework that addresses data volume, data growth trends over time, and requirements for storage management will enable service providers to manage cloud storage resources and costs in such a manner that the cloud will continue to offer the greatest benefits for the storage of massive data systems.

1. Introduction

Many organizations are faced with the challenge of how to operate and manage information technology (IT) in an age of exponential data growth. As different types of data are generated, processed, and stored, data increase in terms of size (terabytes to petabytes) and velocity (the rate at which new content is added [1]. The accumulation of such large volumes of data over time is referred to as "big data," and represents a potentially valuable source of information for organizations [2]. The use of big data is changing the ways that organizations operate and do business. Companies use big data to optimize process flow, make predictions based on statistical analysis, and a myriad of other business, scientific, and engineering analyses [2, 3]. Over time, managing data can present many challenges when developing a data infrastructure to accommodate big data [2] . Several issues, such as scaling to meet increased user requests for existing and new data, performance demands and continuous availability as the amount of data grows, workload diversity as methods for combining and analyzing data increase, and data security as data and analysis products become desirable targets for hackers, must be addressed as part of big data analysis [4].

Cloud computing, which is a solution to the issues that arise from big data [4], has emerged to provide data-intensive computing environments with a cost-effective way to manage elastic (both growth and shrinkage) requirements for data and data processing [5]. Cloud architectures specify virtualized IT services and resources that are provided over a network by a cloud service provider [6]. Cloud adoption is driven by the appeal of lower participation costs, easier scaling, and effective provisioning [7]. Cloud consumers can also take advantage of the economic impacts of a scaled-down infrastructure cost (e.g., labor, facilities, and utilities) while relying on cloud providers to provide better performance and the guaranteed service quality that satisfies requirements for security, availability, and reliability.

With the exponential growth of data collection that has evolved over time to become big data and the growing dependence on IT in almost every field of human endeavor, cloud computing storage systems are widely considered to be the most viable solution for managing and preserving data [1]. Cloud consumers require storage solutions that are sufficiently heterogeneous, scalable, and flexible to allow for storage reuse and data sharing [8]. Systems architects and engineers grapple with ways to design storage architectures that deliver the service-quality metrics essential to meet these cloud consumer requirements [9].

Because no two cloud service providers are the same, cloud consumer systems architects and engineers

require a method to model, test, and evaluate their storage needs to determine the cloud provider that will best fit their business, research, or mission needs. This has created the need for a system that can (1) predict big data patterns and trends that will affect how big data is stored and (2) identify limitations and necessary improvements in current cloud storage infrastructure. In this paper, we examine how service providers can manage cloud storage resources and costs as the amount of collected data to be retained continues to grow exponentially.

Our study's contribution is to propose a solution to data growth. The optimal cloud storage framework (CSF) will address the volume of data growth based on performance requirements for storage management. Additionally, the framework should measure and report the behavior of the cloud environment, based on the size of the data set. In this paper, we construct and run cloud simulation models to collect data to answer the following research questions:

- Based on the collected performance data, can an optimal minimal point be identified when using the CSF?
- What improvements can be made to the framework that will assist engineers in increasing cloud storage performance?

Using cloud simulation tools in the CSF to generate the performance data needed for statistical analysis can benefit service providers by allowing them to identify methods for managing cloud storage resources, costs, and big data. This framework could also assist consumers and consumer system architects and engineers to determine which cloud storage provider's services best fit their needs. The CFS is based on general systems engineering principles that will assist service providers and cloud consumers who face challenges with the migration and adoption of large data sets to the cloud environment.

1.1.   Related work

Cloud storage environments are not new. Many researchers and engineers have addressed the problems and benefits of cloud storage environments, and numerous articles have discussed the challenges of data migration, quality of service, security, and unexpected hardware and software failures. Deciding whether to move to the cloud requires that businesses and organizations clearly assess their current infrastructure needs. Many articles have provided comprehensive reviews of cloud-based infrastructures and suggestions about how to manage different data types in the storage environment [5]. Kolodner et al. [5] propose a storage architecture that raises the abstraction level of storage using middleware to modify and manage data as a service. The authors discuss current offerings for cloud storage and present conceptual architectures for two scenarios. However, no quality of service or cost data are evaluated to determine how the storage solution performs for the cases presented. Another practical storage strategy is to examine the trade-offs between computation and storage usage. Yuan et al. [10] introduced a cost model that uses data sets for storage-producing algorithms to compute costs based on storage and computation. A cloud simulation tool called SwinCloud was used to test data sets with random sizes ranging from 100 GB to 1 TB. This approach—of using different data types of random sizes—takes into account the importance of size when analyzing storage functionality and performance. Unfortunately, the cost data that were used for the model and simulation came from only one service provider. Identifying trends, patterns, and limitations requires larger data sets and the comparison of several service providers' costs.

Quality of service (QoS) plays an important role in consumers' and engineers' decisions about which cloud service provider to choose [11]. Zhang et al. [11] propose an infrastructure for a large-scale cloud cluster and focus on storage components, but the size of the data set used for the experiment was relatively small (1 GB). This is a differentiated approach to providing services based on the data path and the priority of the class of user based on allocations made by a management and scheduling panel. Differentiated cloud storage resources assist with latency and throughput, but other aspects of storage resources, such as Central Processing Unit (CPU) utilization and Random Access Memory (RAM), are not addressed. Throughput is important for the user to access the data faster; however, to truly study data behavior, all resources that impact data should be evaluated and analyzed.

Other studies have provided insight into the evolution of the data center environment. Kant [12] puts into perspective why cloud consumers are adopting cloud-based storage solutions, namely, the costly challenges and issues involved in maintaining data centers [12]. With the growing use of data centers to store massive amounts of data, rapid data growth can pose significant challenges for data centers and their customers [13]. In addition, with the increase in application build-ups and Virtual Machine (VM) scale-outs, bandwidth (throughput) rates can potentially cause degradation of transfer rates, which in turn causes storage

performance problems—such as latency—that can require data centers to increase their capacity to satisfy customer requirements [13].

Once a consumer categorizes his or her data, the CFS framework helps to assess the data's behavior in a virtualized environment using the response surface method (RSM), which is an alternative variance reduction and optimization method used in a variety of fields [14]. We use RMS to examine multiple parameters that impact storage performance to determine the near-optimal settings for predictor variables [14]. Use of this technique for state-of-the-art information systems such as clouds allows researchers, cloud service providers, and consumers to identify the optimal performance parameters to transmit, compute, and store large data sets in the cloud. Understanding the cloud infrastructure and architecture allows systems engineers and other users to build and make changes to best optimize the system.

2. Cloud Storage Framework (CFS): A Proposed Solution

The report "How Fast Is Our Data Volume Growing?" [15] shows that data stored in enterprise data centers will increase by sixty percent annually. Such a high annual increase demonstrates the importance of finding better solutions to managing data growth. CFS is based on the design of experiments (DOE) methodology, which is a tool used by research engineers and scientists to simultaneously determine how the interactions of many factors will affect the output design of the system being modeled. CFS uses an academic cloud simulation tool called CloudSim to measure the performance of cloud resources as data grow and begin to accumulate in the cloud environment [16]. The results of the cloud simulation can be analyzed further using RSM to gain a better understanding of the observed behavior of data growing in the cloud simulated environment. For the purposes of this paper, we will model a two-factor interaction to determine what minimal optimal point(s) can be predicted for the output results for the cloud storage performance factors CPU and RAM.

2.1. Design Of Experiments (DOE)

DOE is an applied statistical method used to assist researchers in gaining knowledge about a system based on input and output factors. Figure 1 depicts the typical principles associated with conducting a DOE [17].
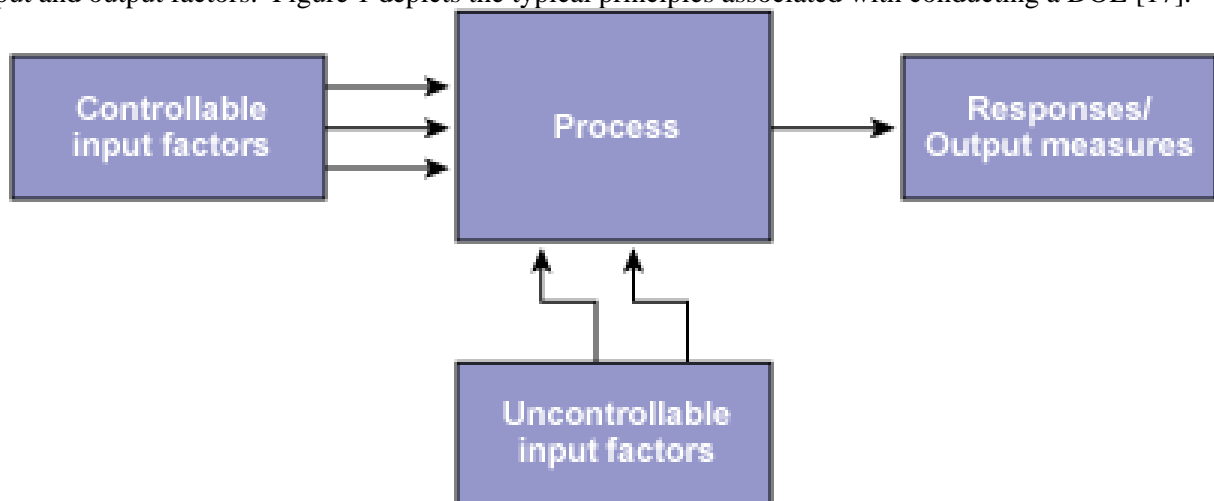

Figure 1: DOE Process

Using DOE, researchers are able to determine which individual factors interact with each other and how this interaction impacts the output results. After conducting the DOE, researchers should have gained a better understanding of their system's behavior based on results that will provide answers to the following questions [18]:

- What are the key factors in the experiment?
- What are acceptable performance settings for the system?
- What are the key system interactions noted, based on the output?
- Which settings would provide less variation in the output?

Finally, using SAS JMP, the following consecutive steps were performed to develop a design model for the CFS framework [18]:

- Design screening to narrow the input variables for evaluation.

- Use the full factorial design to study the response combination of factors to determine the region of interest in which the process is close to an optimization maximum or minimum.
- Generate a response surface and analyze the response factors that contribute to the optimum region.
-

2.2. Response Surface Methodology (RSM)

RSM is part of the DOE statistical technique family. RSM uses the DOE process and statistical analysis to determine which input factors are significant interactions that affect the output measures [19]. These factors generate responses that develop the empirical model for the response surface. If the model can be defined as a linear model, then the first-order model is used. When using multiple input factors, a higher degree polynomial should be used. The general form of a model for a response variable y with linear terms and interaction terms up to order N is[20]

$$Y = \sum_{i=1}^{n} b_i X_i + \sum_{i<j}^{n} b_{ij} X_i X_k + \sum_{i<j<k}^{n} b_{ijk} X_i X_k X_j$$

Finally, the application of RSM can be employed to determine what factors can be used to pinpoint the optimum response. A design can be selected for fitting a second-order model. Central composite design is the method most commonly used [19]. SAS JMP will generate the analysis of variance to determine the goodness of fit for the model. Finally, contour and three-dimensional surface plots can be generated to help researchers visualize which factors of the model significantly relate to the response on the surface.

In the next section, we will discuss how data will be captured using DOE to develop the RMS empirical model for the response surface.

3.   Data and Experiments

3.1. Cloud Data Center Simulation

Cloud infrastructure services, or Infrastructure as a Service (IaaS), generally consist of computing resources such as processing, storage, and networks that can be simulated by the cloud using a layered architecture [16]. Along with the extension layer, the simulator supports the modeling of cloud behavior based on interfaces such as virtual machines (VMs), memory, storage, and bandwidth [16]. Using a CloudSim GUI extension, such as CloudReports, a cloud customer can study different VM provisioning based on their cloud workload by simply changing the input parameters to model a particular scenario [21]. In CloudSim, the data center manages a specific number of hosts. Each host is assigned to one or more VMs based on the policy that is created in CloudReports. The data center host can be preconfigured and assigned a process capability, which is stated in terms of millions of instructions per second (MIPS), allocated random access memory (RAM), storage sectioning, and processing elements (cores) to the VMs [22]. This layer can easily be used by a cloud service provider to determine the issues that are associated with provisioning the host to the VMs. Additionally, CloudSim allows the customer to study performance rates based on the applications and functionality of the VMs, based on the workload. The user can set virtualizations services in CloudReports, such as the number of VMs, the VM image size, scheduling priorities, and the VM [21]. Using CloudReports, we represent a virtualized data center to evaluate the increase in the amount of data.

In the next section, we demonstrate how this tool was used to observe the growth of data in a cloud environment.

3.2. Experimental Setup and Workload Generation

We consider a number of workloads by varying the size of the output files and other data center storage resources. By increasing the size of the input and output files, we are able to simulate the data growth in the cloud. A series of experiments was conducted while closely monitoring the performance input variables for Customer 1, as shown in Table 1.

In Experiments 1-5, the output sizes of the files being stored in the Data Center for Customer 1 ranged from 1 MB to 1.2 MB. Other customer resource configurations are set to manage the cloud workload. The virtualized resources of Experiments 1-5, such as the machine amount, VM image size, and bandwidth, are the baseline input parameters that are used to conduct the experiment. The parameters were chosen to resemble a virtual machine using cloud technology that can process large amounts of data. Setting a realistic VM environment at the onset of testing increases the likelihood that a machine that uses the same configurations will perform acceptably as the amount of data grows rapidly in the cloud environment. Experiments 6-19 used some of the same baseline configurations while increasing the VM image size, processing capability of the VM (MIPS), and output files. Increasing these configurations could potentially

degrade the performance of resources, such as storage in the cloud environment. The remaining experiments (20-48) increased the configurations to simulate saturating the resources so that extreme output values could be captured for optimization.

Table 1: Performance Input Variables for Customer 1

Experiment	VM Amount	VM Imagine	VM BW	VM Priority	Virtual PE	MIPS	VM Ram	PE	Max Length	File Size	Output Size
1-5	1	1,000	155,000	1	4	4,000	1,024	8	65,000	25,000	1.0×10^6- 1.2×10^6
6-19	1	1,500	155,000	1	4	5,000	1,024	8	65,000	25,000	1.4×10^6- 2.1×10^6
20-24	1	2,000	155,000	1	4	6,000	1,024	8	65,000	25,000	2.2×10^6- 2.6×10^6
25-35	1	2,000	155,000	1	8	7,000	1,024	8	65,000	50,000	2×10^6- 3.1×10^6
36-48	1	3,000	155,000	1	8	8,000	1,024	8	65,000	75,000	3.3×10^6- 10.9×10^6

After running the experiments, the input and output values were captured in Table 2 for statistical analysis using SAS JMP software, which has a number of platforms for data analysis. For instance, the Fit Model platform allows the user to fit multiple output values against the input effects used to create the data set. The fitting platform offers the user a multiple fit regression model, models with complex effects, RSMs, and multivariate models for analysis. It also has special screening tools for cube plots, prediction profilers, and contour profilers. We will use the continuous response using the following x-input variables to represent customer performance parameters [23]:

- Virtual machine image (VMI): A single file that contains a virtual disk with a bootable operating system and other software applications.
- Virtual processing element (VPE): A virtual processor is a single logical processor that is exposed to a partition by the hypervisor.
- Million instructions per seconds (MIPS): This is a general measure of computing performance and, by implication, the amount of work a larger computer can do.
- Processing element (PE): A physical processor is effectively a hierarchically ordered collection of logical processors with some form of shared system resources.
- File size: Initial size of the file being stored in the simulator.
- Output size: Final size of the file being stored in the simulator.

The Y variables are the output responses generated by the CloudReport tool. The following Y output responses will be fitted to the model:

- CPU: Where most of the calculations are performed; considered to be the brains of the computer. This value in CloudReport represents usage by the multiple cores, which are independently processing computing instructions.
- RAM: A discrete amount of memory assigned to run the virtual machine. This value in CloudReport represents the memory usage required to execute a virtual machine instance.

In the next section, we will report the results generated to produce models for the output responses for the CPU and RAM and conduct analyses to determine the model's validity and whether an optimum minimum response was generated for the output responses.

3.3. Analysis and Results

3.3.1. Model Validity Assessment

We analyzed our results to determine whether the two-level factorial design could accurately model the response. Standard regression outputs for the model are shown in Table 2.

3.2.2. Prediction Profiler with Minimum Desirability Set for Response Analysis

After determining which variables are considered a good fit and significant to the second-order polynomial, we can begin the optimization process. Using JMP's Prediction Profiler, or simply Profiler, allows us to explore cross sections of predicted responses across multiple factors. Based on the values predicted to produce a minimal value for CPU and RAM, the Profiler tool yielded the results shown in Figures 2 and 3.

Table 2: Model Summary Of Fit Table For Output Response CPU And RAM

	CPU	RAM
R-squared	.7350319	.196599
R-squared adj.	.584863	.031751
Root mean square error	21.206326	.391705
Mean of response	105.4658	95.09479
Observations (sum of weights)	48	48

The R-squared value measures the proportion of variation around the mean explained by our polynomial model [24]. If R-squared is 1, the model is considered to be a perfect fit [25]; a zero value indicates that the fit did not improve from the simple mean model [25]. R-squared increases with the number of variables added to the polynomial. A model that has too many predictor variables will begin to model random noise in the data.   Therefore, adding more variables can overfit the model and lessen its ability to make extrapolations about the data [25]. Adjusted R-squared is a modification of R-squared that adjusts for the number of terms in a model. R-squared always increases when a new term is added to a model, but adjusted R-squared increases only if the new term improves the model [25]. Table 2 shows the values of R-squared and R-squared adjusted that best represent a fit model for further analysis.
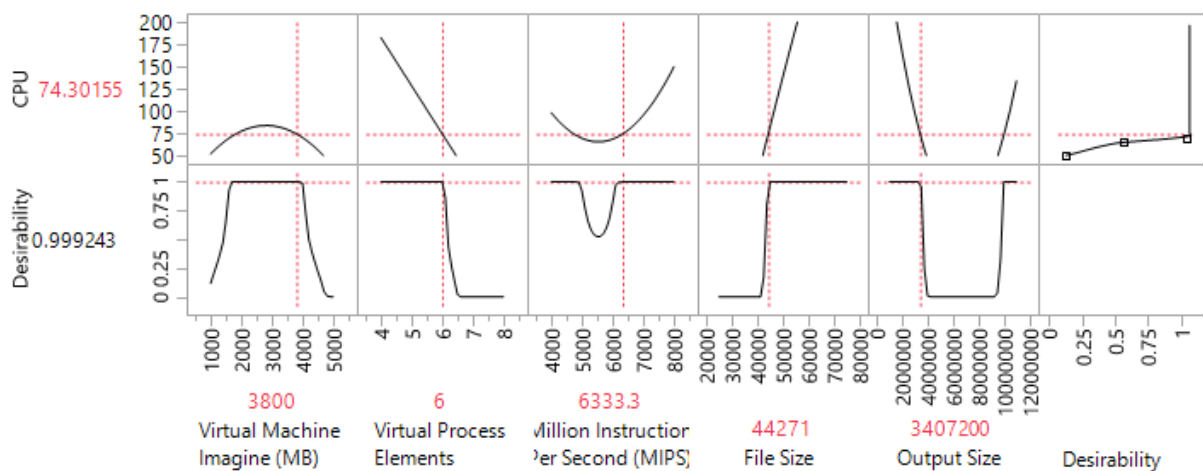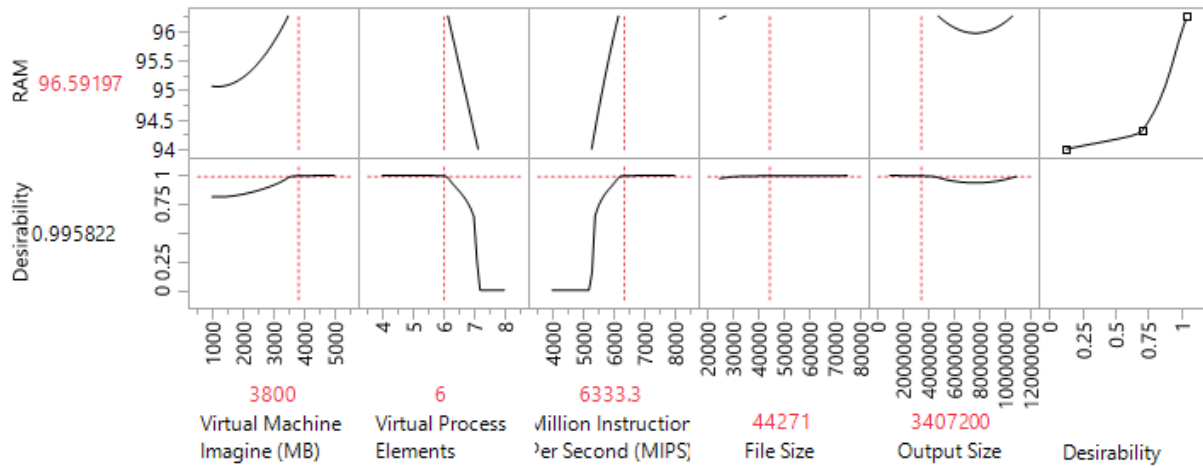

Figure 2: CPU Prediction Profiler

Figure 3: RAM Prediction Profiler

The desirability function allows the researcher to choose an operating condition for an overall desirability-optimized value, such as a minimum or maximum [26]. The desirability function is a transformation of the response variable to a 0 to 1 scale [26]. A response of 0 represents a completely undesirable response, and 1 represents the most desirable response [11]. Based on the Prediction Profiler, the values shown below were considered the most desirable minimum responses for the CPU and RAM models.

Table 3: Most Desirable Minimum Prediction Profiler Values

Virtual Machine Imagine	3,800
Virtual Process Elements	6
MIPS	6,333
File Size	44,271
Output Size	3,407,200

The desirability function can be used to make the appropriate trade-offs when optimizing a system's performance in the cloud. Due to the dynamic and agile environment of cloud resources, it is difficult to find a global optimal minimum. When equipment is off or not in use, the optimum minimal will be zero. Because of this limitation, we can only define a feasible operating region for optimal performance [27]. By using response surface plots, the researchers are better able to pinpoint feasible operation regions.

3.3.3. Response Surface 3D Plot

A response surface is a graph of a response variable as a function of the various factors. One important aspect of RSM is visualizing the response. Response surface plots were used to assist us in:

- Observing how changes in variables affect a response of interest.
- Finding variable levels that optimize a response.
- Selecting operating conditions that meet specifications.

The response surface 3D plot for CPU and RAM are shown in Figure 4. The CPU response surface plots represent a minimum curvature pattern. The VMI variables depict a strong interaction with the other response variables (i.e., VPE, MIPS, File Size, and Output Size), with the optimal region of interest centered on the values predicted for desirability. For example, in Figure 4(a) the response surface shows that the region of interest for minimal optimum points lies on a slightly curved plane within the green areas; stronger responses are darker green. In Figure 4(b) the response is in the saddle surface, which shows increases and

decreases in the response. Visualizing the region of interest is easier to see in this display plot, due to the descending curvature in the center represented by the highlighted green values. Figure 4(c) shows a response curvature similar to Figure 4(a), with VMI and File Size values to decrease CPU appearing toward the center of the plane. Figure 4(d) shows a flatter-plane response plot, indicating that this response has very little impact on the output response for CPU.
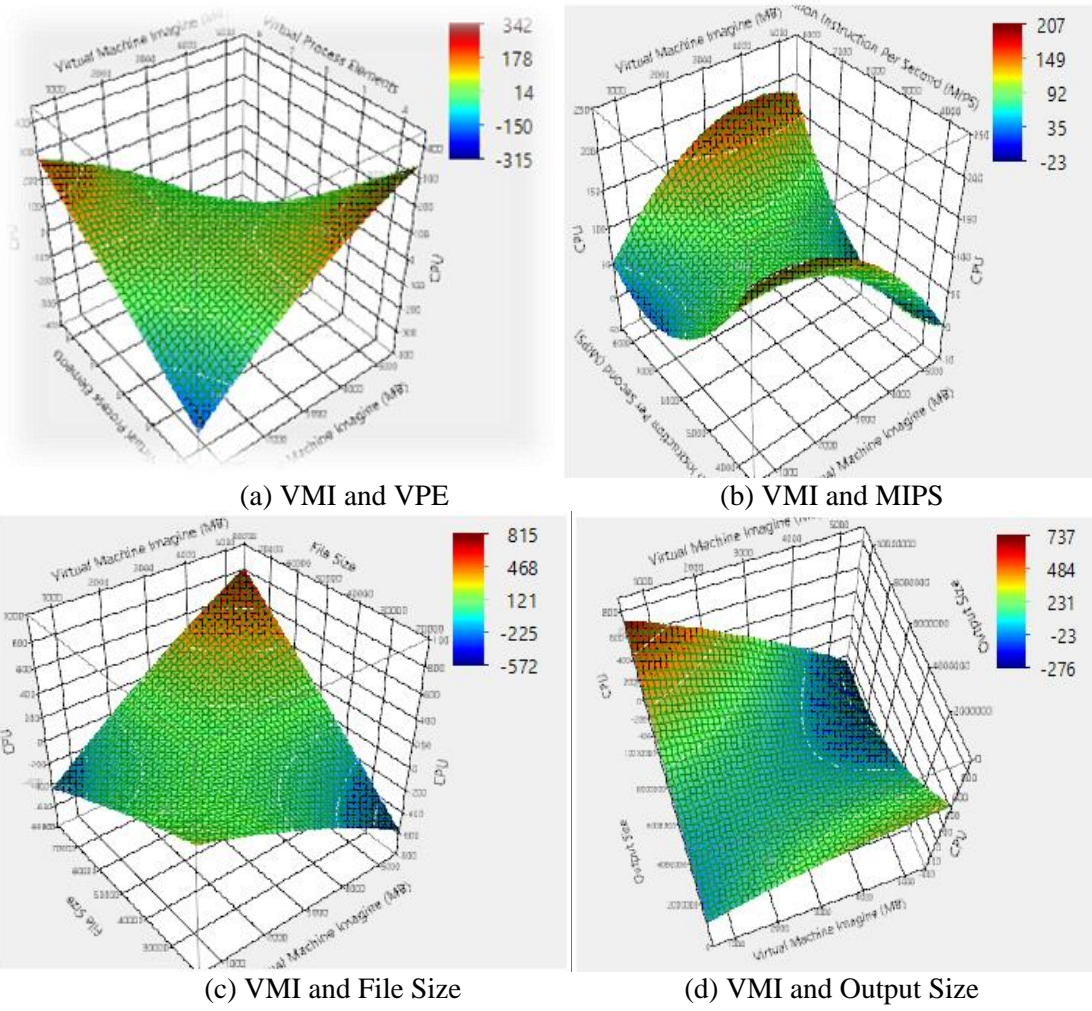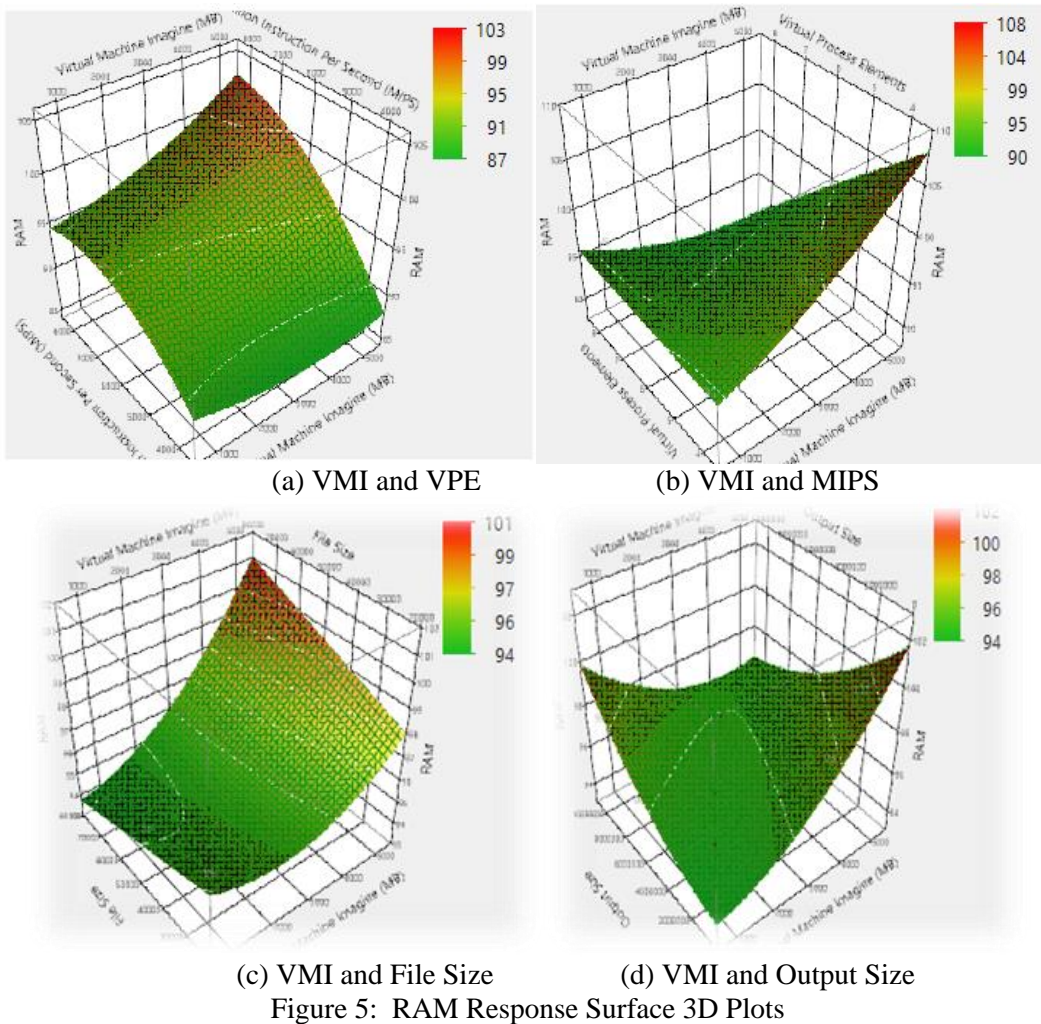

(a) VMI and VPE   (b) VMI and MIPS

(c) VMI and File Size   (d) VMI and Output Size

Figure 4: CPU Response Surface 3D Plots

(a) VMI and VPE   (b) VMI and MIPS

(c) VMI and File Size   (d) VMI and Output Size
Figure 5: RAM Response Surface 3D Plots

RAM responses all generated planes with slight curvature plots. Additional shading was needed to clearly see the optimal area of interest. Due to lack of curvature, we are reevaluating the second-order polynomial for the RAM. Figure 5(a)'s plot shows the plane responding with descending values for RAM for VMI and VPE. This plot shows the most significant correlation of the variables based on the output response for RAM. Plots in Figures 5(b) and (c) show the plane responses increasing for VMI, MIPS, and File Size with respect to RAM. Figure 5(d)'s plot depicts a gradual increase of RAM values toward the center of the plane, as VMI and Output Size increase.

4. Conclusions and Future Work

The purpose of this research was to observe exponential data growth in the cloud storage environment and develop an approach to determine how to assess and determine current storage needs for the cloud. Due to the strain that large data workloads place on cloud storage architecture, the CFS—which uses current cloud-simulation tools—is limited in its capabilities and cannot simulate a storage environment when bombarded with random IOPS requests. Today, cloud designs for resources such as networks, computation, and storage hinge on choosing the right size of virtual instance [28], which is an operating system (OS) or application environment installed on software that imitates dedicated hardware [29]. Applications and data hosting in the cloud environment can quickly consume the vast amount of resources assigned to a virtual instance. Understanding the behavior of the applications, data, and cloud resources is vital to understanding how to optimize a customer's performance in the cloud.

CPU usage is the amount of actively used virtual CPU as percentage of available CPU. Windows users can access CPU usage via the Task Manager. Short spikes in CPU usage are considered normal; however, extended usage values above 90 percent could impact performance by introducing latency. RAM usage is similar to CPU, as it also represents the virtual hardware memory being used by the guest physical host. Generally, users would want a slightly larger memory size to accommodate workload spikes. If the host feels that a machine needs more memory, it can be swapped or increased based on what is available in the host's resource pool. If virtual machines begin demanding greater memory than available, this could potentially degrade performance of the cloud storage system. These two resources produce values that are unique to their workloads. Understanding RAM and CPU utilization are crucial for optimizing a resource-management strategy for the cloud storage system.

We will continue to examine areas in the response surface plane and conduct tests to identify ways to improve optimal response regions. Future research will seek ways to lower overuse of resources such as RAM, CPU, and power as data continue to grow in the cloud environment. When considering the adoption of any cloud service, a consumer must have a clear understanding of and the technical knowledge necessary to manage their data in the cloud. Resource utilization for RAM, CPU, and power can be expensive based on consumption. The purpose of the CFS is to assist the consumer in understanding how large data needs would behave in the cloud environment and how to reduce these resources while ensuring the storage system's best performance. Study results indicate that large data sets have the potential to overcommit resources, and this could impact how data are accessed and stored.

In conclusion, this article presents an approach that can assist cloud storage consumers in managing their storage resources by using a framework that consists of cloud simulation modeling and optimization using RSM. The CFS is a new approach that will assist researchers in future work on methods for handling exponential data growth in the cloud storage environment. We will investigate ways to improve the cloud simulation tool that models the IaaS environment, adding server inputs for data access patterns associated with saturating the amount of I/O request before latency increases and degrades server performance. As data continue to be valued and grow, continual research for cloud storage services is needed to deliver consistent, reliable, and efficient cloud storage performance—and, enhanced by ongoing research, CFS will help meet current and future needs.

5. References

[1] Mackie, S., How Fast Is Our Data Volume Growing?, in Storage Strategies. 2009: http://c.ymcdn.com/sites/www.theasbc.org/resource/dynamic/blogs/How_Fast_is_Our_Data_Growing.pdf.

[2] Alexander, F.J., A. Hoisie, and A. Szalay, BigData. IEEE Computing in Science and Engineering, 2011. 13(6): p. 10-13.

[3] Villalpando, L.E.B., A. April, and A. Abran, Performance Analysis Model For Big Data Applications In Cloud Computing. Journal of Cloud Computing, 2014: p. 1-20.

[4] Tole, A.A., Big Data Challenges. Database Systems Journal, 2013. 4(3): p. 31-40.

[5] Kolodner, E.K., et al. Data Intensive Storage Services on Clouds Limitations, Challenges and Enablers. in Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference. 2011. Athens: IEEE.

[6] Bojanova, I., J. Zhang, and J. Voas, Cloud Computing. IEEE Computer Society, 2013. 15(2): p. 12-14.

[7] Bapna, R., et al., A Market Design for Grid Computing. INFORMS Journal on Computing, 2008. 20(1): p. 100-111.

[8] Zeng, W., et al., Research on Cloud Storage Architecture and Key Technologies, in International Conference on Information Systems. 2009. p. 1044-1048.

[9] Huo, Y., et al., A Cloud Storage Architecture Model for Data Intensive Applications, in 2011 International Conference on Computer and Management (CAMAN). 2011, IEEE: Wuhan. p. 1-4.

[10] Yuan, D., et al., A Highly Practical Approach Toward Achieving Minimum Data Sets Storage Cost In The Cloud. IEEE Transactions On Parallel And Distributed Systems, 2013. 24(6): p. 1234-1244.

[11] Zhang, Y., et al., DifferCloudStor: Differentiated Quality of Service for Cloud Storage. IEEE Transactions on Magnetics, 2013. 49(6): p. 2451-2458.

[12] Kant, K., Data Center Evolution: A Tutorial on State of the Art, Issues, and Challenges. Computer Networks, 2009. 53(17): p. 2939-2965.

[13] Leong, Y.K., M.M.A. Khin, and P.S. Alexopoulos, Storage System Architecture for Data Center of the Future. International Journal of Advancements in Computing Technology (IJACT), 2012. 4(9): p. 184-192.

[14] Chang, K.-H., L.J. Hong, and H. Wan, Stochastic Trust-Region Response-Surface Method (STRONG)—A New Response-Surface Framework for Simulation Optimization. INFORMS Journal on Computing, 2013. 25(2): p. 230-243.

[15] Mackie, S., How Fast Is Our Data Volume Growing. 2009, Storage Strategies: Columvbia, MD.

[16] Calheiros, R., et al., CloudSim: A Toolkit For Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms. Software Practice and Experience, 2011. 41: p. 23-50.

[17] Sundararajan, K. Design of Experiments – A Primer 7/28/15; Available from: http://www.isixsigma.com/tools-templates/design-of-experiments-doe/design-experiments-%e2%90%93-primer/.

[18] Quality, A.S.f. What Is Design of Experiments (DOE)? ; Available from: http://asq.org/learn-about-quality/data-collection-analysis-tools/overview/design-of-experiments.html.

[19] Proust, M., JMP Release 8 Design of Experiments Second Edition. 2009, SAS Institute Inc.: Cary, NC. p. 273.

[20] Veli-Matti, Tapani, and Taavitsainen, Experimental Optimization and Response Surface, in Chemometrics in Practical Applications, K. Varmuza, Editor. 2012: Finland. p. 92-138.

[21] S á T. CloudReports. 2011 2011 [cited 2014 September 2014]; Software Download]. Available from: https://github.com/thiagotts.

[22] Kumar, R. and G.Sahoo, Cloud Computing Simulation Using CloudSim. International Journal of Engineering Trends and Technology (IJETT), 2014. 8(2): p. 82-86.

[23] Eh, B. Hyper-V Concepts - vCPU (Virtual Processor). 2011 [cited 2015 6/28/2015]; Available from: http://social.technet.microsoft.com/wiki/contents/articles/1234.hyper-v-concepts-vcpu-virtual-processor.aspx.

[24] Proust, M., JMP Statistics and Graphics Guide, Release 7. 2007, SAS Institute Inc: Cary, NC.

[25] Frost, J. Multiple Regression Analysis: Use Adjusted R-Squared and Predicted R-Squared to Include the Correct Number of Variables. 2013 [cited 2015 August 5, 2015]; Available from: http://blog.minitab.com/blog/adventures-in-statistics/multiple-regession-analysis-use-adjusted-r-squared-and-predicted-r-squared-to-include-the-correct-number-of-variables.

[26] Shanock, L.R., et al., Polynomial Regression with Response Surface Analysis: A Powerful Approach for Examining Moderation and Overcoming Limitations of Difference Scores. Journal Business Psychology, 2010. 25: p. 543-554.

[27] He, Z., et al., Robust optimization for multiple responses using response surface methodology. Applied Stochastic Models in Business and Industry, 2010. 26(2): p. 157-171.

[28] MacVittie, L. CAPACITY IN THE CLOUD: CONCURRENCY VERSUS CONNECTION. 2012 [cited 2015 July 25, 2015]; Available from: http://datacenterpost.com/2012/05/capacity-in-cloud-concurrency-versus/.

[29] Rouse, M. Virtual Machine. 2014 [cited 2015 August 1, 2015]; Available from: http://searchservervirtualization.techtarget.com/definition/virtual-machine.