

AN EXAMINATION OF CLOUD COMPUTING'S SECURITY RISKS AND PERFORMANCE EVALUATION

Subodh Kumar

Research Scholar, Glocal School of Science
The Glocal University, Mirzapure Pole , Saharanpur (U.P).

Dr. Uma Shanker

Research Supervisor, Glocal School of Science,
The Glocal University Mirzapure Pole , Saharanpur(U.P).

Abstract:

Cloud computing offers many advantages, including improved collaboration, increased storage capacity, mobility, and accessibility. However, cloud computing also poses security risks. The purpose of this study is to analyze various security risks to cloud computing and assess cloud performance. Due to the inadequacy of conventional methodologies for assessing cloud performance, performance analysis of clouds is a common undertaking. This study does a performance analysis of a cloud based on many features, including fault tolerance, scalability, elasticity, and pay per usage. This report also examines a number of cloud computing security risks.

Key Words - Mathematical Models, Cloud computing, performance metrics,

I. INTRODUCTION

Performance testing in cloud computing is different from that of traditional applications. The traditional performance testing focused on the performance metrics for applications that are under a particular workload for a fixed configuration. Cloud test need to measure the performance metrics related to the workloads that run in a distributed fashion on multiple virtual and real machines. Computing and computer systems are becoming more complex but easy to use due to advent in network technology such as cloud computing. Performance testing is a type of testing intended to determine the responsiveness, throughput, reliability, and/or scalability of a system under a given workload [1]. The enormous growth of cloud computing created a demand for mathematical models that can measure the performance characteristics of cloud applications. This paper is organized as follows. Cloud computing and its characteristics are described in Section II. Section III describes various security threats in cloud computing. Section IV discusses and explains performance analysis of cloud computing followed by conclusion (Section V).

II. CLOUD COMPUTING

Data and programs are being swept up from desktop PCs and corporate server rooms and installed in "the compute cloud." The Greek Myths tell of creatures plucked from the surface of the Earth and enshrined as constellations in the night sky. Something similar is happening today in the world of computing. Whether it's called cloud computing or on-demand computing, software as a service, or the Internet as platform, the common element is a shift in the geography of computation. [2] Cloud computing is a model for pooling IT resources to provide real-time on-demand self-provisioned services to business users on an as-needed basis (Fig 1). [3]




Fig.1 Idea behind Cloud Computing

2.1 Cloud computing characteristics

Essential Characteristics: The cloud model is composed of five essential characteristics, three service models, and four deployment models.

It also lists three "service models" (software, platform and infrastructure), and four "deployment models" (private, community, public and hybrid) that together categorize ways to deliver cloud services. The National Institute of Standards and Technology's (NIST) definition lists five essential characteristics of cloud computing: on-demand self-service, broad network access, resource pooling, rapid elasticity or expansion, and measured service. The definition is intended to serve as a means for broad comparisons of cloud services and deployment strategies, and to provide a baseline for discussion from what is cloud computing to how to best use cloud computing. [4]

On-demand self-service: A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access: Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

Resource pooling: The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

Rapid elasticity: Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Measured service: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

2.2 Service Models

Software as a Service (SaaS): The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user specific application configuration settings.

Platform as a Service (PaaS): The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services,

and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Infrastructure as a Service (IaaS): The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

2.3 Deployment Models

Private cloud: The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Community cloud: The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public cloud: The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Hybrid cloud: The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds) [5]

III. SECURITY THREATS IN CLOUD COMPUTING

The most typical security risks associated with cloud computing are discussed below (Fig 2):

3.1 Data Loss

Data loss is the most common cloud security issue connected with cloud computing. It is often referred to as data leakage. The act of having data rendered, damaged or destroyed unusable by a user, application or programme is known as loss of data. In the environment of cloud computing, in the hands of a third party using the sensitive data, one or more data pieces cannot be used by the data owner, the hard drive is not functioning properly, un-updated software, data loss happens.



Fig.2 Security Threats in Cloud Computing

3.2 Hacked Interfaces and Insecure APIs

Because cloud computing is completely reliant on the Internet, it is critical to secure the interfaces and application programming interfaces (APIs) used by external users. APIs are the most convenient way to access most cloud services. Few cloud computing services are available to the general public. These services are available to third parties, hence there is a risk that they could be corrupted or compromised by hackers.

3.3 Data Breach

When confidential information is accessed, viewed taken by an authorization of the third party, an hacking of data in organization (a process known as a "data breach").

3.4 Vendor Lock-in

Vendor lock-in is one of the most serious security risks in cloud computing. Transferring an organization's services from one provider to another can be difficult. It may be difficult to transition from one cloud to another because different vendors providedifferent platforms.

3.5 Increased complexity strains IT staff

For the IT team, operating, integrating, and migrating to the cloud services is complicated. IT workers must have new competencies and skills in order to preserve, integrate and manage data in the cloud.

3.6 Meltdown and Spectre

Programs can observe and steal data that is presently being processed on computers thanks to Spectre and Meltdown. It can function on desktop PCs, mobile phones, and cloud computing systems. It has the ability to save the password as well as your emails, business papers, and other private data in the memory of other open programmes.

IV. PERFORMANCE ANALYSIS IN CLOUD COMPUTING

Although performance analysis is a complicated process, finding a problem's core cause is its straightforward objective. To do this, we take a systematic, step-by-step approach to narrow in on the fault domain holding the fundamental cause. We'll talk about how cloud computing and virtualization affect this process as those two topics are the focus of this chapter. But let's first look at a root-cause analysis' high-level flow.

We start with a common issue—slow application response time—and gradually eliminate the possibilities as follows:

Determine whether the issue only impacts a certain transaction or transaction type, or if it affects the entire programme. If the latter, then identifying the slow-moving transaction type is our first step. This will establish the setting for additional research.

Decide which tier, if any, is troublesome. The issue could exist between two tiers (i.e., due to network latency). Check whether the environment itself has a negative impact, such as CPU depletion or memory limits, to further isolate the issue. This would include everything not directly related to the application, including suspensions of garbage collection.

Finally, narrow down the problematic application component, method, or service call. From here, we may establish whether the underlying issue is algorithmic, CPU-centric, the result of frequent VM suspensions, or brought on by external bottlenecks like I/O, the network, or locks (synchronization). This final action is the starting point for a prospective remedy.

This process can be used with any system, including cloud-based ones, even though it might differ slightly in your environment. However, there is a crucial caveat: virtual systems are susceptible to the time-keeping issue caused by VM suspensions, which can lead us to mistakenly identify the wrong tier as problematic. As a result, real time measurements are required whenever we examine reaction or communication times (as opposed to apparent time). To do this, we may either use virtualization-aware timers inside the VM or measure reaction time outside the VM in an unaffected system (e.g., an appliance).

Increased transaction volume or overuse of the underlying network infrastructure may be the cause of increasing inter-tier time. (If the latter is the case, it might be best to pay the administrator a nice visit.) Instead, if we've only found one troublesome layer, we'll start by looking for excessive VM suspensions. We may accomplish this by comparing the reported stolen time to our time-based transactions. We can identify which single transaction or even method was suspended if the granularity of our metric is fine enough. We simply have a vague idea that a transaction might have been halted. We cannot be certain of our results, for instance, if our granularity is measuring in 20-second increments (the default period of VMware's vCenter). Greater assurance results from finer granularity.

Be aware that there will always be some stealth time in a particular virtual system, although it shouldn't be excessive. It should take a few milliseconds per 20-second cycle, around the same as a garbage collection.

The tier itself is examined next for method-level trouble locations. Here, we deal with some well-known issues:

The overhead that virtualization-aware timers impose on precise measurements is more expensive than that of the standard OS timers and may be detrimental.

Even for a quick approach, VM suspensions can provide the appearance of slowness, which could direct our study in the wrong direction. Determining real time (which includes suspensions) and actual execution time is therefore necessary (which does not).

The amount of detail required to address the first two problems is almost never reported for VM suspensions! There are two approaches to this apparent conundrum:

We don't need to employ single-transaction diagnostics, so we can rely solely on guest timers, which you might assume might distort our results. The rule of large numbers, however, eliminates skewing when we examine aggregated data from thousands of transactions, so long as we do not experience significant VM suspensions.

Similar to GC suspensions, we can relate VM suspensions to a thorough transaction analysis. Our analysis avoids the timing issue by omitting transactions that are more significantly impacted (greater than a few milliseconds).

The amount that the VM suspension impacted the transaction response time is depicted in this hot spot graph. This allows us to disregard the suspended area from our analysis. Finally, we've narrowed down the likely root cause to a certain code area, service call, or function. We now understand the cause of the delay. The most frequent reason is excessive CPU usage. Since we're discussing virtual environments, we should once more question the veracity of the reported CPU time. The surprising answer is generally yes.

The majority of hypervisors manage stolen time to have the least amount of CPU effect on any user programmes. This implies that an estimate of the average CPU usage for a transaction expressed in CPU time can be made. (Remember that CPU utilisation is an exception.) This measurement can be trusted to indicate that the application code is indeed using more CPU time if it consistently increases.

The opposite is also accurate. It must be the increased latency in the virtualization layer if the time of our transaction increases in a particular area of the application (as opposed to throughout) and cannot be attributed to CPU, sync, steal, wait, or garbage-collection time. But we don't have to only speculate.

We are aware of which application methods in ours access the network or the hard drive. We know that when these measured times increase, there is either a real latency issue or a suspension because we are employing virtualization-aware timers for these access I/O locations. On any case, the next step is to explore for solutions in the virtualization layer and discuss them with the administrator.

Nearly everything is covered in that. As was already established, the objectives for performance management in private vs. public clouds differ significantly. The performance analysis has certain technical variances as well. That brings up our next subject.

A metric is used to measure and understand the behavior of software. Cloud metrics can be used to measure the behavior of cloud which utilizes the resources from the computers as a collective virtual computer, where the applications can run independently from particular computer or server configurations [6].

Cloud delivers its services through internet and provides the full user functionality of a software application by the web sites which provide Software as a Service. Dynamic web sites provide regularly changing information to users and utilize dynamically generated pages and maintain data for display in a database [7]. Cloud uses the dynamic web sites to deliver the web applications on demand. Cloud metrics should follow some characteristics which help to evaluate cloud on each and every parameter which is necessary for a good quality cloud, so that a client can rely on it to choose the best cloud.

The main advantages of cloud computing are elasticity, scalability, reliability, availability, pay-per-use and fault-tolerance [8].

a. Elasticity: is one of the major factors for the success of the cloud as an IT infrastructure [9]. For a DBMS deployed on a pay-per-use cloud infrastructure an added goal is to optimize the system's operating cost. Elasticity, i.e. the ability to deal with load variations by adding more resources during high load or consolidating the tenants to fewer nodes when the load decreases, all in a live system without service disruption, is therefore critical for these systems. Even though

Elasticity is often associated with the scale of the system; a subtle difference exists between elasticity and scalability when used to express a system's behavior.

b. Scalability: is a desirable property of a system, which indicates its ability to either handle growing amounts of work in a graceful manner or its ability to improve throughput when additional resources (typically hardware) are added. A system, whose performance improves after adding hardware, proportionally to the capacity added, is said to be a scalable system.

c. Reliability: is the probability that a product or part will operate properly for a specified period of time (design life) under the design operating conditions (such as temperature, volt, etc.) without failure [10]. The outcome of the measurement process is reproducible that is similar to results over time for some different inputs and across many different situations. Cloud gets many requests simultaneously and will also give the similar results for some requests in a period of time so clouds have to be reliable.

d. Availability: Cloud Services should be available maximum time [11]. The on demand, elastic, scalable, and customizable nature of the cloud must be considered when deploying cloud architectures. Many different clients might be accessing the same back-end applications, and many providers are providing the cloud services has the expectation that only their application will be properly delivered to users. In cloud computing it is essentially required to gather the information instantly without making a user to wait and the gathered information should be related to each other.

e. Cost: Cloud Computing allows an organization to pay by the hour of computing resources, potentially leading to cost saving even if the hourly rate to rent a machine from a cloud provider is higher than the rate to own one. This is essentially preferable when demand for a service that arise over time.

f. Fault Tolerance: is one of the key issues of cloud computing. There are many fault tolerance techniques in parallel computing [12]. Fault tolerance is concerned with all the techniques necessary to enable a system to tolerate software faults.

Resource allocations, workloads and system behaviors can fluctuate widely, and performance can further vary based on system configurations. Despite the assorted issues with performance analysis of virtualized environments, researchers have made much headway in recent years developing tools and techniques for measuring, modeling and simulating virtualized environments.

Various research groups have developed performance metrics and benchmarks specifically to measure the performance of virtualized systems such as cloud computing. In some cases, research groups have been able to use custom implementations of existing benchmarks when analyzing virtualized systems. For example, a particular implementation of SPECweb2005 benchmark, which is the standard benchmark used to determine the performance of Web servers, can be successfully used in performance testing of virtualized applications [13]. However, most benchmarks that were originally developed for physical machines cannot be used for reliable performance testing of

virtualized environments due to the diverse and complex nature of these virtualized systems.

V. CONCLUSION

The physical host is necessary because we need to know about other VMs operating on the same host, not because we need to assess the physical host's utilization specifically. Clouds bring a few particular concerns in addition to the performance issues we encounter in every virtualized system. First, a VM could relocate or retire at any time. In most cases, we need to reconstruct the path of a transaction across its VMs, including information on the original physical hosts, in order to assess performance concerns after the fact. The programs that can affect ours can change at any time depending on how many and which other VMs are running on the same host as ours. Without this information, we might be aware that we don't have enough resources, but we wouldn't know why.

The conventional performance metrics for performance analysis of cloud computing is not enough because of virtual environment. The exponential growth of cloud computing has resulted in a dire need for metrics that can measure the performance characteristics of cloud applications. In this paper, the concept of cloud computing and its essential characteristics along with security threats has been discussed. Performance of a cloud has been analyzed on the basis of their characteristics such as pay per use, elasticity, fault tolerance and scalability. Traditional approach for performance testing in cloud computing is also being discussed. The importance and benefit of using mathematical models for performance testing in cloud computing are also explained in this paper.

REFERENCES

- [1] Gurdev Singh, Shanu Sood, Amit Sharma CM- Measurement Facets for Cloud Performance. International Journal of Computer Applications (0975 – 8887) Volume 23– No.3, June 2021
- [2] David Cleary “Web Based Development and Functional Size Measurement” IFPUG Annual conference.
- [3] M. Armbrust et al. Above the clouds: A Berkeley view of cloud computing. Technical Report UCB/EECS-2019-28, 2019.
- [4] Team Sardes, Inria Rhône-Alpes, Elasticity in Cloud Computing, June 23, 2021
- [5] Kareim M. Sobhe, Ahmed Sameh “Multi-Channel Clustered Web Application Server”
- [6] “Making Cloud Service Continuity a Reality” NetPrecept Software Ltd.
- [7] Chunye Gong, Jie Liu, Qiang Zhang, Haitao Chen and Zhenghu Gong The Characteristics of Cloud Computing.
- [8] J.D. Meier et al. Performance Testing Guidance for Web Applications. Microsoft Corporation, United States, 2022.
- [9] B. Hayes, “Cloud computing”, Communications of the ACM, vol. 51, no. 7, pp. 9–11, Jul. 2018.
- [10] Md Shamshoddin Altamash and Prashant Y Niranjana “A Survey of Identifying Key Challenges of Performance Modeling in Cloud Computing” International Journal of Computer Science and Information Technology Research (IJCSITR), Vol.1, Issue1, pp(33-41), Month: October-December 2023.