IJIRCCE

e-ISSN: 2320-9801 | p-ISSN: 2320-9798

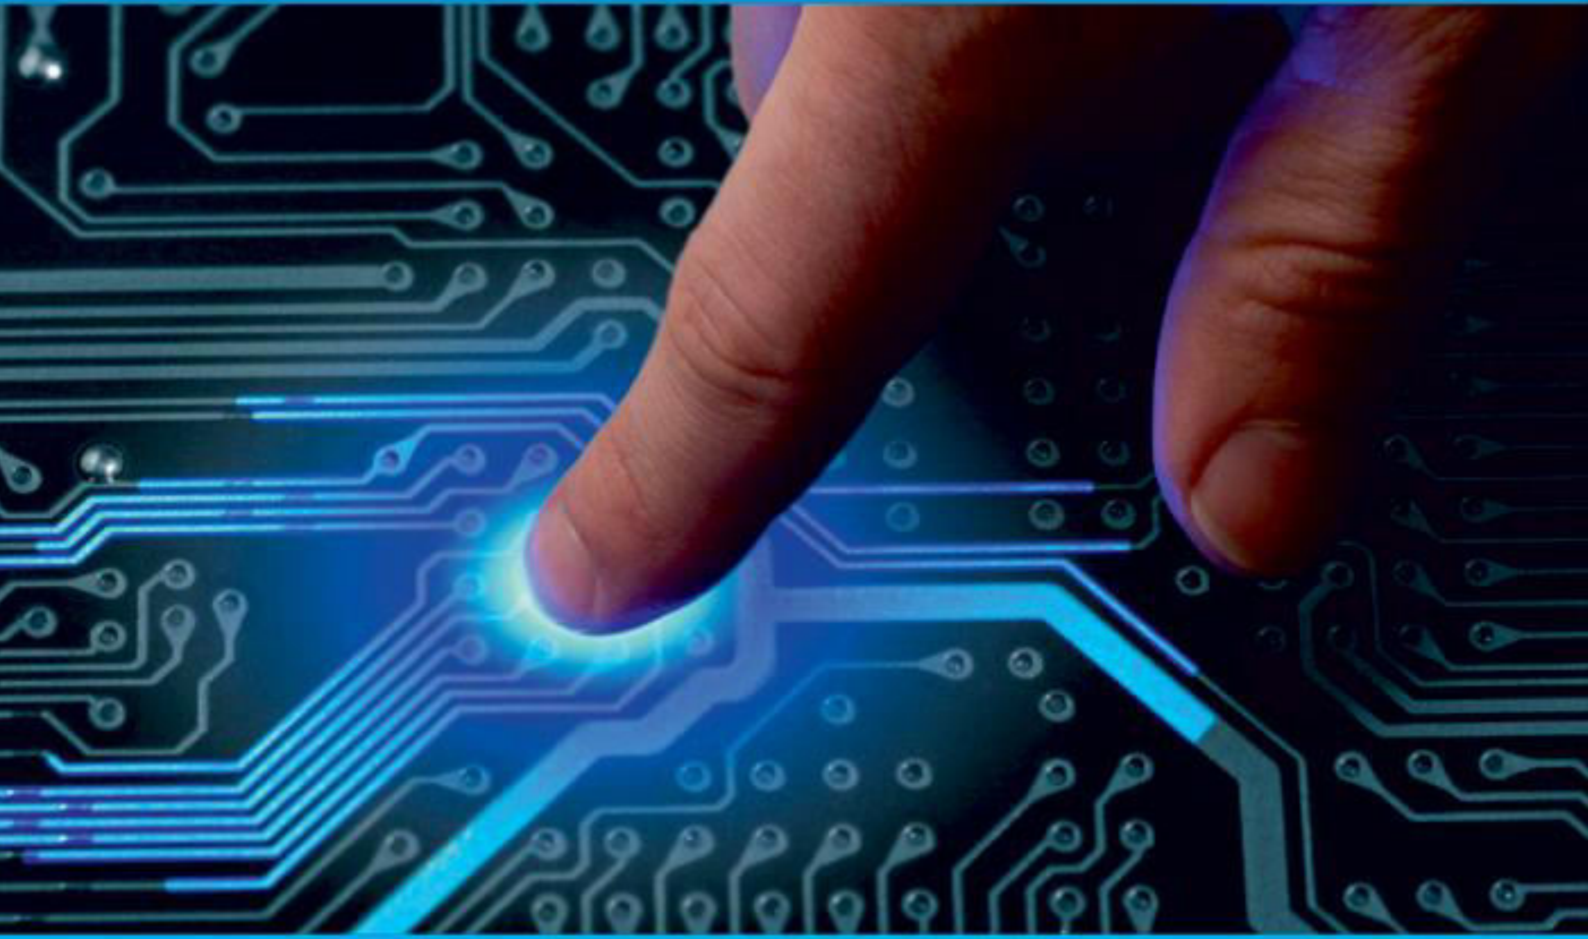

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

Volume 9, Issue 11, November 2021

INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA

Impact Factor: 7.542

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542

|| Volume 9, Issue 11, November 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0911027 |

AMAZON EC2 A CLOUD’S COMPUTER

Saurabh Bhate[1], Nikhil Soni[2], Nitin Kamble[3]

Student, School of Computer Science and Engineering, Ajeenkya D Y Patil University, Pune, India[1]

Student, School of Computer Science and Engineering, Ajeenkya D Y Patil University, Pune, India[2]

Professor, Ajeenkya D Y Patil University, Pune, India[3]

ABSTRACT: Cloud Web Services is emerging today as a commercial infrastructure that eliminates the need for maintaining expensive computing hardware. Through the use of virtualization, clouds promise to address with the same shared set of physical resources a large user base with different needs. Thus, clouds promise to be for scientists an alternative to clusters, grids, and supercomputers. However, virtualization may induce significant performance penalties for the demanding scientific computing workloads. In this work we present an evaluation of the usefulness of the current cloud computing services for scientific computing. We analyze the performance of the Amazon EC2 platform using micro-benchmarks, kernels, and e-Science workloads. We also compare using long-term traces the performance characteristics and cost models of clouds with those of other platforms accessible to scientists. While clouds are still changing, our results indicate that the current cloud services need an order of magnitude in performance improvement to be useful to the scientific community.

KEYWORDS: Virtual computers, Amazon Machine Image (AMI), Persistent storage, Cloud monitoring,

I.INTRODUCTION

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) Cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.Amazon EC2 runs instances on its physical infrastructure using the open-source vir-tualization

middleware Xen. By using virtualized resources, a cloud can serve with the same set of physical resources a much broader user base; configuration reuse is another reason for the use of virtualization. Sci-entificsoftware, compared to commercial mainstream products, is often hard to install and use. Pre- and incrementally built virtual machine (VM) images can be run on physical machines to reduce deployment time for scientific software.The term cloud computing refers to a method through which information and programs can be stored and accessed without storing or accessing it on any physical media. This is highly advantageous to companies that require large amount of disk space. It is a modern means to save internal IT resources because data is not stored in-house. Rather, data is stored in a “cloud” from where it can be retrieved anytime. In today’s modern world, using cloud computing helps large corporations to make huge savings. This is because they do not need to worry about financing the required software or hardware. Instead, they simply choose a cloud service and obtain the required software or hardware in a few clicks. This is overall a far more economic and fast process compared to traditional methods of storing and accessing data [1]. Cloud Computing provides three kinds of services: i) Private cloud: This type of cloud owned by the organization is meant to provide services to its own users. ii) public cloud: Third party are providing the services. Examples include Amazon Web Services (AWS), Microsoft Azure, IBM/SoftLayer and Google Compute Engine. iii) Hybrid cloud: This is a combination of services provided by private and public clouds. The main goal of this kind is to achieve scalability. Cloud computing has three categories of services: i) Infrastructure as a service: It helps users to transfer work from one machine to another, usually a virtual machine. ii) Platform as a Service: PaaS is used for general software development. Common PaaS providers include Salesforce.com’s Force.com, Amazon Elastic Beanstalk, and Google App Engine. iii) Software as a service: SaaS delivers software applications over the Internet; these are often called Web Services. Microsoft Office 365 is an example of a SaaS. There are many examples of cloud computing services, including Google drive, Apple iCloud, Dropbox, SugarSync and AWS.

II. PROBLEM STATEMENT

I In some cases, you might have underlying resources that you want to upgrade incrementally. For example, you might change to a higher performing instance type in your Auto Scaling launch configuration so that you can reduce the maximum number of instances in your Auto Scaling group. If problems occur after you complete the update, you might need to roll back your infrastructure to the original settings. To do this manually, you not only have to remember which resources were changed, you also must know what the original settings were.

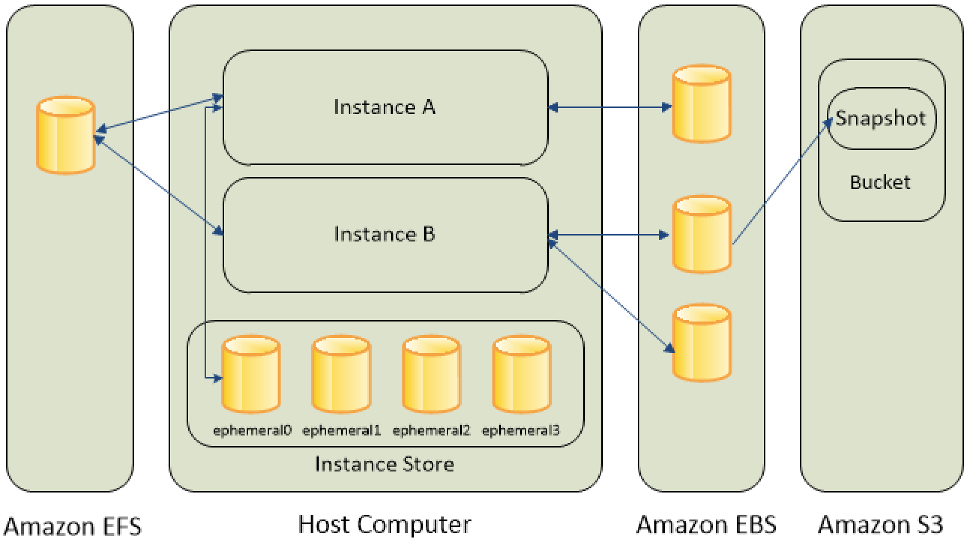

III. CONSEQUENCES

When you provision your infrastructure with CloudFormation, the CloudFormation template describes exactly what resources are provisioned and their settings. Because these templates are text files, you simply track differences in your templates to track changes to your infrastructure, similar to the way developers control revisions to source code. For example, you can use a version control system with your templates so that you know exactly what changes were made, who made them, and when. If at any point you need to reverse changes to your infrastructure, you can use a previous version of your template.

An EC2 instance may be launched with a choice of two types of storage for its boot disk or "root device." The first option is a local "instance-store" disk as a root device. The second option is to use an EBS volume as a root device. Instance-store volumes are temporary storage, which survive rebooting an EC2 instance, but when the instance is stopped or terminated, this store is lost. EC2 is a service whereby you can create virtual machines and run them on one of Amazon's data centers. These virtual machines take a slice of that center and simulate the hardware of a physical server. You can choose from a variety of machine configurations that have different processing powers, memory configurations, and virtual hard drive sizes. Once you have picked the class of your machine, you start it by loading an Amazon Machine Image (AMI). You can pick from one of the many public AMIs, most of them based on Linux, or any custom ones you may have created.One of the biggest security risks involved with cloud computing is the element of trust. This is because the service user is not buying, configuring, monitoring, and implementing their own personal computers and servers physically. These services are provided to them on request by a cloud computing service provider. This makes their configuration and internal settings vulnerable to the services on which they rely. This situation can lead to an attack on the confidentiality and/or integrity of the data present in the cloud [7]. In [8], the authors have elaborated on the significant risks involved with cloud services in particular. As a result, the authors have provided some solutions related to security concerns for industries and cloud computer service providers. The security concerns include secure data transmission, encryption, access controls, authorization and authentication, customer service, data privacy, log monitoring, intrusion detection, and auditing. Besides these, proper risk management procedures must be followed by cloud providers. Strategies, measures, and rules should be created, recorded, and executed. Training materials or courses should be produced that set a standard for giving basic security and risk administration abilities and education to the cloud computing suppliers, the security group, and their internal partners.

IV. LITERATURE SURVEY

In a study [13], the authors selected Amazon EC2 instances, and used them to show the methods of gaining access to and exploiting an instance. Also, in [14], the authors showed that cloud clients are not watchful when picking EC2 instances. By distributing a malicious instance, they observed that this instance was started a few times and that data about the use of the instances could be gathered. In addition, they demonstrated that it was possible to evade the payment mechanism of paid images by changing the AMI file. The authors in [15] used graph theory techniques to study security issues associated with Amazon EC2 with respect to the configuration of images. The authors provided various security recommendations that are applicable at the infrastructure level of AWS, rather than at instance level.

Establish trusted access to realize data value:
The authors in [16] researched security issues associated with images on the Amazon EC2 Service. An automated system was created by the authors. Various tests were done to achieve the desired results. The system tested if the programs used in images are up to date or not. The system was also used to test common vulnerabilities in various operating systems like Windows and Linux. Nessus [25] was used for this purpose. The results indicated that clients and providers of public images can both suffer from the dangers of potential security weaknesses present in EC2. Most of the software programs used were updated two or more years ago. 98% of Windows AMIs and 58% of Linux AMIs contained software with critical vulnerabilities. The associated security risks include loss of privacy, authority, and system infections through malware. Another test was performed to investigate the probability of compromise of cloud systems through malware. ClamAV [26] was the anti-virus software used for this purpose. The results indicated that Windows machines are more vulnerable to internet malware compared to Linux systems. The results also indicated that EC2 had no mechanism to differentiate connection of a legitimate source from a malicious source; also that, if clients using a particular image has not removed their credentials fully from that machine, there were ways to recover full credentials using various tools available online. Anyone renting the image subsequently could gain access to credentials, and then use AWS with the original client being billed. This study concluded that there must be appropriate vulnerability assessment before renting and using a cloud-based image.

Quickly replicate your infrastructure:
The authors in [17] clarified the shortcoming of AWS regarding security services in 2013. Netflix leases space from Amazon Web Services (AWS) to operate membership administration to watch their movies and TV episodes. As indicated by the investigation of information leakage of 209 worldwide organizations in 2011, 37% of information leakage cases included malicious attacks. In 2012, a retailer owned by AWS, known as Zappos, was the victim of cyber theft [18]. The number of clients whose login information might have been leaked is up to 24 million. In [19], the authors evaluated an ordinary environment of the mainstream AWS cloud with a focus on security. The cloud security was surveyed by implementation of Dionaea honeypots for a couple of months in the given systems. In the experiment done by authors, three AWS EC2 instances ran in parallel in the regions of Singapore, US East Virginia, and Sao Paulo. An overall comparison was also provided among the three regions.

V. RELATED WORK

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment.
Amazon EC2 offers the broadest and deepest compute platform with choice of processor, storage, networking, operating system, and purchase model. We offer the fastest processors in the cloud and we are the only cloud with 400 Gbps ethernet networking. We have the most powerful GPU instances for machine learning training and graphics workloads, as well as the lowest cost-per-inference instances in the cloud. More SAP, HPC, Machine Learning, and Windows workloads run on AWS than any other cloud.

Increase or decrease capacity within minutes, not hours or days
SLA commitment of 99.99% availability for each Amazon EC2 region.
The AWS Region/AZ model has been recognized by Gartner as the recommended approach for running enterprise applications that require high availability

VI. CONCLUSION

The world congress on Internet Security survey [8] in 2013 indicated that there is always going to be a high demand for products providing security management. Since AWS gives its clients full control of an instance, it can be concluded that security is not only the responsibility of the cloud provider, but also of the client. Human beings are the weakest link, as they say. Another possible solution to the above-mentioned security issues can be that AWS should not allow services and clients to share account login information with each other. In addition, AWS users must read the tips and techniques for how to secure AWS [27] before starting to use the service.

REFERENCES

[1] K. Ren, C. Wang, and Q. Wang, “Security challenges for the public cloud,” IEEE Internet Computing, no. 1, pp. 69–73, 2012.
[2] Cloud Security Alliance, “Cloud control matrix CCM v3.0.1,” 2014, available at: https://cloudsecurityalliance.org/research/ccm/ [3] ISO Std IEC, “ISO 27017,” Information technology- Security techniques (DRAFT), 2012.
[4] OpenStack, “OpenStack open source cloud computing software,” 2015, available at: http://www.openstack.org.
[5] Open Data Center Alliance, “Open data center alliance usage: Cloud based identity governance and auditing rev. 1.0,” Tech. Rep., 2012.
[6] B. Tang and R. Sandhu, “Extending OpenStack access control with domain trust,” in Network and System Security, 2014, pp. 54–69.
[7] A. Gouglidis and I. Mavridis, “domRBAC: An access control model for modern collaborative systems,” computers & security, 2012, 31(4).
[8] K. Fisler, S. Krishnamurthi, L. A. Meyerovich, and M. C. Tschantz, “Verification and change-impact analysis of access-control policies,” in ICSE, 2005. [9] G.-J. Ahn, H. Hu, J. Lee, and Y. Meng, “Representing and reasoning about web access control policies,” in COMPSAC ’10, 2010.
[10] K. Arkoudas, R. Chadha, and J. Chiang, “Sophisticated access control via SMT and logical frameworks,” TISSEC, vol. 16, no. 4, 2014.
[11]. "Can AWS Microsoft reach 50 market share cloud", (2015), [Online]. Available: http://marketrealist.com/2015/11/can-aws-microsoftreach-50-market-share-cloud/
[12]. Cusumano M. 2010, “Cloud computing and SaaS as new computing platforms”. Communication. ACM 53, 4 (April 2010), 27-29. DOI: https://doi.org/10.1145/1721654.1721667
[13]. Ristenpart T., and Savage S., 2009, "Hey, you, get out of my cloud: exploring information leakage in third-party compute clouds", In Proceedings of the 16th ACM conference on Computer and communications security.
[14]. Haroon M. and Nick A., 2010, "Clobbering the cloud", defcon.org, Communication Inc, U.S.
[15]. Bleikertz S., Probst P. and Eriksson K 2010, "Security audits of multi-tier virtual infrastructures in public infrastructure clouds", In Proceedings of the 2010 ACM Workshop on Cloud Computing Security, p. 93
[16]. Balduzzi M., Zaddach J., Balzarotti D., Kirda E., 2012 "A Security Analysis of Amazon’s Elastic Compute Cloud Service", in Proceedings of the 27th Annual ACM Symposium on Applied Computing - SAC ’12, p. 1427.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING