Efficient Usage Of Amazon Web Service Platforms Oriented To Computation Of Cloud

D. Monica Seles | M. Caroline Viola Stella Mary
[1](PG Scholar, Department of IT, Francis Xavier Engineering College, Tirunelveli)
[2](Professor, Department of IT, Francis Xavier Engineering College, Tirunelveli)

Abstract— *Amazon Web Services (AWS) is a secure cloud services platform, offering compute power, database storage, content delivery and other functionality to help businesses scale and grow. Explore how millions of customers are currently leveraging AWS cloud products and solutions to build sophisticated applications with increased flexibility, scalability and reliability. AWS can provide many of the popular services to view documentation such as Amazon EC2, Amazon CloudFront, Amazon Simple Storage Service, AWS Identity and Access Management, AWS lambda, etc. AWS Management Console is a web application for managing Amazon Web Services. AWS Management Console consists of list of various services to choose from. It also provides all information related to our account like billing. The AWS Console mobile app, provided by Amazon Web Services, allows its users to view resources for select services and also supports a limited set of management functions for select resource types. This paper can show a review of amazon web services in a detailed manner.*

Keywords— Elastic Compute Cloud(EC2), Simple Storage Service(S3), Amazon Web Service(AWS), AWS lambda.

1. INTRODUCTION

In recent years, AWS play a vital role in all kind of developing platform using amazon products and services. Amazon Web Services (AWS) started to offer IT services to the market in the form of web services, which is nowadays known as cloud computing. With this cloud, we need not plan for servers and other IT infrastructure which takes up much of time in advance. Instead, these services can instantly spin up hundreds or thousands of servers in minutes and deliver results faster. We pay only for what we use with no up-front expenses and no long-term commitments, which makes AWS cost efficient. Cloud computing is emerging as a new important trend in distributed computing, and several vendors are awaiting feasible solutions to optimize the usage of their own data centers. The National Institute of Standards and Technology (NIST) said that Cloud computing can be defined as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (including networks, servers, storage, and services) that can be rapidly acquired, provisioned, and released with minimal management effort and service provider interaction. The rapid evolution of systems infrastructure technologies, specifically virtualization and cloud, present new opportunities and challenges in software design. In the evolution from bare metal hardware provisioning, to traditional virtualization, to cloud IaaS, to container technology and now cloud events, the software architect has a myriad of IT infrastructure tools upon which to design and deploy his/her software in optimized configurations (price, performance, scale, etc.). The evolution seems to have accelerated with tech hype and maturity cycles condensing from decades (physical infrastructure) to years (virtualization/cloud) to months (containers/microservices/cloud events). The most recent addition to this continuum is a set of services we will classify as cloud events. Specifically we are referencing new commercial services such as Amazon Web Service's Lambda, IBM, Google Cloud Platform's Cloud Functions, and Microsoft Azure's Functions. Cloud computing undoubtedly provides very good service to the user but still many organizations do not support cloud computing because of the security issues. The main security issues are data security and privacy protection. These security issues hinders the managers or customer to support the services provided by cloud computing. This is the reason why cloud computing not gaining the expected market size. Cloud security is the responsibility of both the cloud provider and cloud consumer. There should be the relationship of trust between them before availing any cloud service.. It is the responsibility of management to take care of security risk so as to protect data. There are many risks associated with the cloud computing. Some of them are: Security, service providers, Management and Control, Laws and regulation, virtualization risks, lack of standards and auditing, uncontrolled cost etc. Content Delivery Networks (CDNs) represent another critical service of the contemporary Internet and are responsible for moving around large shares of traffic served by thousands of datacenters to the end-users scattered worldwide. A particular feature Clouds and CDNs have in common is that, from the network point of view, they both separate the web- service or content from the actual server serving it. This dramatically complicates the task of monitoring the web- services and the provided quality. Indeed, for instance, in the case of the CDN, the originating server offering a given content can delegate its distribution to surrogate servers, or caches. Similarly for Clouds, the separation between content and the server generating it is brought by the virtualization layer.

2. Review of Literature

2.1 Tracking Causal Order in AWS Lambda Applications

Serverless computing is a new cloud programming and deployment paradigm that is receiving widespread uptake.

Serverless offerings such as Amazon Web Services (AWS) Lambda, Google Functions, and Azure Functions automatically execute simple functions uploaded by developers, in response to cloud-based event triggers. The Serverless abstraction greatly simplifies integration of concurrency and parallelism into cloud applications, and enables deployment of scalable distributed systems and services at very low cost. Causality is an important tool employed in concurrent and distributed systems that facilitates reasoning about, analyzing, and drawing inferences from a computation. In particular, causal order is required for function design (to enable mutual exclusion, consistency, deadlock detection), for distributed debugging, failure recovery, and inconsistency detection, for reasoning about progress (termination detection, collection of obsolete data and state), and for measuring and optimizing concurrency. This lack of support in AWS Lambda limits the degree to which developers can identify the root cause of errors, performance bottlenecks, cost anomalies, and optimization opportunities for Lambda applications. To address these limitations, we present Gamma Ray, a cloud service for AWS Lambda applications that provides a holistic view of causal application behavior and performance end-to-end. Gamma Ray requires no developer intervention and works across AWS regions and AWS cloud services. Gamma Ray intercepts Lambda function entry points and calls to AWS services made by the application. It records these events synchronously using transactional database streams (to guarantee causal consistency) and processes them off line, in near real-time, to provide developers with service graphs and analysis data at both the function aggregate and instance level. As such, Gamma Ray precludes the need for developers to write their own Cloud Watch and X-Ray log parsing and aggregation tools for each application, and provides causal ordering for concurrent, multi-function Lambda applications. This paper investigates three implementation alternatives for Gamma Ray. Two of these alternatives are full X-Ray replacements that collect both performance data and causal relationships using static and dynamic instrumentation. The third is a hybrid approach that leverages X-Ray for performance monitoring (incurring some of its limitations) in exchange for lower runtime performance overhead. We investigate the overhead of each alternative using micro-benchmarks and multifunction serverless applications and find that the hybrid approach introduces the least overhead in terms of execution time and memory.

2.2Non-intrusive Anomaly Detection with Streaming Performance Metrics and Logs for DevOps in Public Clouds

Public clouds are a style of computing platforms where scalable and elastic IT-enabled capabilities are provided as a service to external customers using Internet technologies. Using public cloud services can reduce costs and increase choices of technologies, but it also implies limited system information for users. Thus, anomaly detection at user end has to be non-intrusive and hence difficult, particularly during DevOps operations because the impacts from both anomalies and these operations are often indistinguishable and hence it is hard to detect the anomalies. Our detection collects the log information of running operations and applications, and then uses the information to indicate the run-time system environments for choosing a statistical model, which has been trained beforehand. We analyze the metrics and their monitor data provided by Cloud Watch and then choose Support Vector Machine (SVM) to train the models under different system environments. At runtime, newly sampled metric data are fed to a classifier that has been indicated by the log information. In this paper, our work is specific to a successful public cloud, Amazon Web Service (AWS), and a representative DevOps operation, rolling upgrade, on which we report our anomaly detection that can effectively detect anomalies. Our anomaly detection requires only metrics data and logs supplied by most public clouds officially. We use Support Vector Machine (SVM) to train multiple classifiers from monitored data for different system environments, on which the log information can indicate the best suitable classifier. Moreover, our detection aims at finding anomalies over every time interval, called window, such that the features include not only some indicative performance metrics but also the entropy and the moving average of metrics data in each window.

2.3 On the Network Performance of Amazon S3 Cloud-storage Service

The advances in networking technologies and the increase in the need for storage resources have prompted many companies to outsource their storage needs. Cloud-storage providers offer clean and simple file-system interfaces, abstracting away the complexities of direct hardware management. At the same time, however, such services eliminate the direct oversight of performance that final users with high service-level requirements traditionally expect. While several works in literature have addressed security-related issues (such as privacy, integrity, availability, etc.) few of them have targeted the network performance of this kind of services. In this paper, we aim at improving the knowledge on the performance of cloud storage services, by focusing on the performance of the cloud-to-user network experienced when relying on Amazon Simple Storage Service (S3). S3 is the general purpose storage as a service provided by Amazon, where customer data is organized by means of objects stored in buckets. A bucket is a logical unit of storage uniquely identified and belonging to one of the locations in which the provider has deployed its storage infrastructures (hereafter cloud regions). Costs for the customer depend upon the storage class (standard, infrequent access, or long-term archive) and the cloud region in which the bucket is

placed, according to a pay-as-you-go model. In more details, cost is calculated as the sum of three quotas depending on the size of the stored object, the number of download requests, and the volume of the traffic transferred. Cloud Front (CF) is the global Content Delivery Network (CDN) service offered by Amazon and integrates with S3 in order to distribute contents to the end users with low latency and high data transfer speeds. Data is distributed to the users through the global network composed of the Amazon edge locations spread all over the world.

2.4Cloud Computing Security: Amazon Web Service

Cloud Computing is a recently emerged model which is becoming popular among almost all enterprises. It involves the concept of on demand services which means using the cloud resources on demand and we can scale the resources as per demand. Cloud computing undoubtedly provides unending benefits and is a cost effective model. The major concern in this model is Security in cloud. This is the reason of many enterprises of not preferring the cloud computing. Cloud computing has a focus on maximizing effectiveness of the shared assets. Cloud computing have certain features like they are agile, have reduced cost, easier maintenance, reliable, secure, scalable, etc. Cloud computing involves communication over a loose coupling mechanism which involves multiple cloud components. Certain security issues faced by cloud computing include sensitive data access, data segregation, privacy, authentication, bug exploitation, recovery, accountability, account control. Cloud computing is the just the virtual environment for the customer who are using the cloud services in which he will give its data to the cloud without knowing even the location of the data. The data can be there with thousands of other data on the cloud. So the most important facilities that storage provider should provide is Confidentiality, Integrity and Availability (CIA). A model should be developed that promotes CIA. CIA can be provided by encrypting the data, access control to prevent unauthorized user to access data and scheduled backup should be there to ensure availability. This paper provides the review of security research in the field of cloud security. After the security research, we have presented the working of AWS (Amazon Web Service) cloud computing.

2.5 A Distributed Architecture for the Monitoring of Clouds and CDNs: Applications to Amazon AWS

Clouds and CDNs are systems that tend to separate the content being requested by users from the physical servers capable of serving it. From the network point of view, monitoring and optimizing performance for the traffic they generate is a challenging task, given the same resource can be located in multiple places, which can in turn change at any time. The first step in understanding Cloud and CDN systems is thus the engineering of a monitoring platform. In this paper, we propose a novel solution which combines passive and active measurements, and whose workflow has been tailored to specifically characterize the traffic generated by Cloud and CDN infrastructures. We validate our platform by performing a longitudinal characterization of the very well-known Cloud and CDN infrastructure provider Amazon Web Services (AWS). By observing the traffic generated by more than 50,000 Internet users of an Italian ISP, we explore the EC2, S3 and CloudFront AWS services, unveiling their infrastructure, the pervasiveness of web-services they host, and their traffic allocation policies as seen from our vantage points. The solution provided in this paper can be of interest for i) developers aiming at building measurement tools for Cloud Infrastructure Providers, ii) developers interested in failure and anomaly detection systems, and iii) third-party SLA certificates who can design systems to independently monitor performance.

3. AMAZON WEB SERVICE(AWS)

Amazon Web Services (AWS) is a secure cloud services platform, offering compute power, database storage, content delivery and other functionality to help businesses scale and grow. Explore how millions of customers are currently leveraging AWS cloud products and solutions to build sophisticated applications with increased flexibility, scalability and reliability.

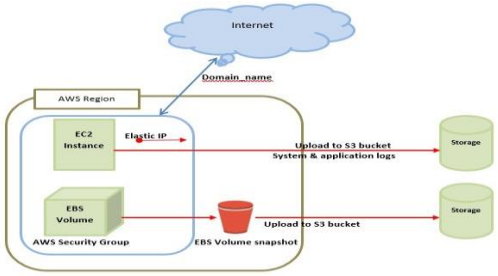

Figure 3.1 Basic Architecture of AWS Web Service

3.1 AWS LAMBDA

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume - there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app. AWS provides a portfolio of services to architect a Big Data platform without managing any servers or clusters. AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics.

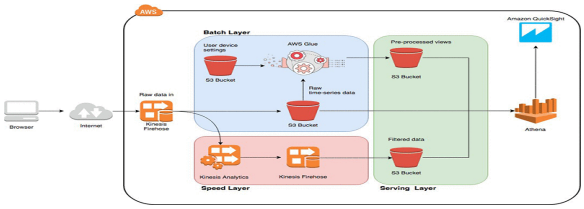

Figure 3.2 Working of Services in 3 different layer

3.2 AWS SES

Amazon Simple Email Service (Amazon SES) is a cloud-based email sending service designed to help digital marketers and application developers send marketing, notification, and transactional emails. It is a reliable, cost-effective service for businesses of all sizes that use email to keep in contact with their customers.You can use our SMTP interface or one of the AWS SDKs to integrate Amazon SES directly into your existing applications. You can also integrate the email sending capabilities of Amazon SES into the software you already use, such as ticketing systems and email clients.

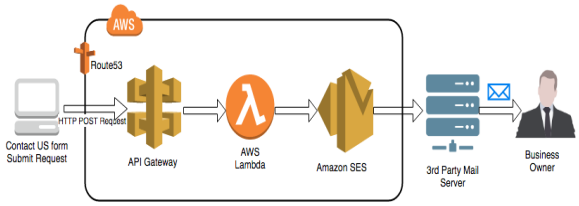

Figure 3.3 Architectural flow of Amazon SES

3.3 AWS S3

Companies today need the ability to simply and securely collect, store, and analyze their data at a massive scale. Amazon S3 is object storage built to store and retrieve any amount of data from anywhere from websites and mobile apps, corporate applications, and data from IoT sensors or devices. It is designed to deliver 99.999999999% durability, and stores data for millions of applications used by market leaders in every industry. S3 provides comprehensive security and compliance capabilities that meet even the most stringent regulatory requirements. It gives customers flexibility in the way they manage data for cost optimization, access control, and compliance. S3 provides query-in-place functionality, allowing you to run powerful analytics directly on your data at rest in S3. And Amazon S3 is the most supported cloud storage service available, with integration from the largest community of third-party solutions, systems integrator partners, and other AWS services.

3.4 AMAZON API GATEWAY

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. With a few clicks in the AWS Management Console, you can create an API that acts as a “front door” for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon Elastic Compute Cloud (Amazon EC2), code running on AWS Lambda, or any web application.Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. Amazon API Gateway has no minimum fees or startup costs. You pay only for the API calls you receive and the amount of data transferred out.

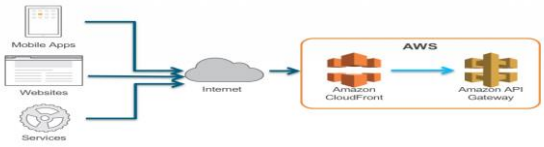

Figure 3.4 An Overview of API Gateway

Edge-optimized endpoints helped you reduce latency to clients accessing your API on the internet from anywhere; typically, mobile, IoT, or web-based applications. Behind API Gateway, you could back your API with a number of options for backend technologies: AWS Lambda, Amazon EC2, Elastic Load Balancing products such as Application Load Balancers or Classic Load Balancers, Amazon DynamoDB, Amazon Kinesis, or any publicly available HTTPS-based endpoints.

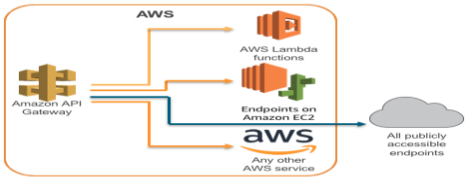

Figure 3.5 AWS lambda functions to access services in public endpoints

API Gateway launched regional API endpoints, which are publicly available endpoints without any preconfigured CDN in front of them. Regional endpoints are great for helping to reduce request latency when API requests originate from the same Region as your REST API

4. Conclusion And Future Work

The effectiveness of our architecture by presenting a characterization of Amazon Web Services (AWS) traffic from passive measurements. By digging into the traffic generated by EC2, S3 and CloudFront over a period of two years, we show some interesting findings. The concern in the cloud computing is Security around data, access and privacy protection. Cloud computing should be secure and

504 robust and should mitigate the risks. According to the analysis of cloud computing it was found that security should be the core operation rather than an add on operation. AWS(Amazon Web Service) has an outstanding performance in cloud computing because of the its excellent work in the area of Security of data.

The future efforts in this area should also center on issues of ease and accessibility encountered during use of cloud event services. Significant potential work exists in the development of robust versioning systems, test frameworks that account for unpredictability in cloud event environments, clearer optimization paths, and deployment tooling that understands cloud event application composition across many different services. Furthermore, application development frameworks that are effective and performant across cloud events and IaaS services could be very powerful, and allow cloud events to be viewed as a way to deploy applications and not just a way to develop applications. Additionally, experimenting with application domains could find new uses for cloud event technologies.

REFERENCES

[1] Ambeth Kumar .V.D, Ashok Kumar .V.D, Divakar .H, and Gokul .R, "Cloud Enabled Media Streaming using Amazon Web Service", IEEE International Conference on Smart Technologies and Management for Computing, Controls, Energy and Materials (ICSTM), pp. 195-198, 2017.

[2] Christof Ebert, Gorka Gallarda, Josune Hernantes, and Nicolas Serrano, "DevOps", IEEE Software, published by the IEEE Computer Society, pp. 94-100, 2016.

[3] Daniel Sun, Min Fu, Liming Zhu, Guoqiang Li and Qinghua Lu, "Non-intrusive Anomaly Detection with Streaming Performance Metrics and Logs for DevOps in Public Clouds: A Case Study in AWS", IEEE Transactions on Emerging Topics in Computing, pp. 1-12, 2016.

[4] Daeli, Jacopo. "Writing AWS Lambda Function in Go," 9 January 2016, retrieved: May 2016. Available: http://jacopodaeli.com/writing-aws-lambda-functions-in-go/

[5] Fabio Bracci, Antonio Corradi, Luca Foschini, "Database Security Management for Healthcare SaaS in the Amazon AWS Cloud", IEEE, pp. 812-819, 2012.

[6] Garrett McGrath, Brenden Judson, Paul Brenner Jared Short, and Stephen Ennis, "Cloud Event Programming Paradigms Applications and Analysis", IEEE 9th International Conference on Cloud Computing, pp.400-406, 2016.

[7] http://aws.amazon.com/what-is-aws/

[8] Ignacio Bermudez, Stefano Traverso, Maurizio Munafo, Marco Mellia, "A Distributed Architecture for the Monitoring of Clouds and CDNs: Applications to Amazon AWS", IEEE Transactions on Network and Service Management, pp. 1-13, 2014.

[9] Narasimhan and Nichols .R, "State of Cloud Applications and Platforms: The Cloud Adopters View", IEEE Computer, Vol. 44, no.3, pp. 24-28, Mar. 2011.

[10] Rajkumar Buyya, James Broberg, Andrzej Goscinski, "Best Practices in Architecting Cloud Applications in the AWS Cloud", https://doi.org/10.1002/9780470940105.ch18, 03 January 2011.

[11] Saakshi Narula, Arushi Jain, Prachi, "Cloud Computing Security: Amazon Web Service", Fifth International Conference on Advanced Computing & Communication Technologies, pp. 501-505, IEEE, 2015.

[12] S.Subashini, V.Kavitha, "A survey on security issues in service delivery models of cloud computing", Journal of Network and Computer Applications, vol.34, pp.1-11, january 2011, https://doi.org/10.1016/j.jnca.2010.07.006.

[13] Sumit Khurana, Anmol Gaurav Verma, "Comparison of Cloud Computing Service Models:SaaS, PaaS, IaaS", IJECT Vol. 4. April-June 2013.

[14] Valerio Persico, Antonio Montieri, Antonio Pescape, "On the Network Performance of Amazon S3 Cloud-storage Service", 5th IEEE International Conference on Cloud Computing, pp. 113-118, 2016.

[15] Wei-Tsung Lin, Chandra Krintz,Rich Wolski, and Michael Zhang, "Tracking Causal order in AWS Lambda Applications", IEEE International Conference on Cloud Engineering, pp.50-60, 2018.