International Journal of Management, IT & Engineering
Vol. 8 Issue 6, June 2018,
ISSN: 2249-0558 Impact Factor: 7.119
Journal Homepage: http://www.ijmra.us, Email: editorijmie@gmail.com
Double-Blind Peer Reviewed Refereed Open Access International Journal - Included in the International Serial Directories Indexed & Listed at: Ulrich's Periodicals Directory ©, U.S.A., Open J-Gage as well as in Cabell's Directories of Publishing Opportunities, U.S.A

AMAZON WEB SERVICES

Kalyan Sudhakar*

Abstract

This paper looks at the various products that are offered by the cloud computing platform AWS. It starts of by providing a brief overview of Amazon Web Services (AWS) itself and how it became the most profitable platform for Amazon the e-commerce platform. Most of the products that AWS provides as its name suggest come as services. These services are designed to offer solutions to storage issues, reliability of servers, latency problems, database management, business intelligence, data analysis, processing speeds, management challenges of having data warehouses among many other problems.AWS provides to its users control and security at a low cost. This together with easing the burden of users when it comes to the challenges mentioned allows the same users to concentrate of their core business. This paper looks at the main products that make this possible in greater detail.

Keywords:
AWS;
Redshift;
Cloud;
Web Services;
Amazon.

* **Master of Science, Electrical and Computer Science,Fairleigh Dickinson University, New Jersey, USA**

Introduction

Amazon Web Services is a cloud computing service that is a subsidiary of Amazon the e-commerce platform. It is used by individuals, businesses and even governments on a demand basis. The platform offers virtual computing hardware that lets the users choose their hard drive and RAM sizes as well as CPU or GPU processing capabilities (Amazon Web Services, n.d). The choice of a user is limited to their computing needs and financial abilities. In addition to hardware, users are offered the option to choose the operating systems to run this hardware, the databases, web servers and other software that they may need. The configuration behind all these can be done through the browser.AWS came about as a result of the company's effort to attempt to standardize and automate their activities such as storage through the web. In the end the people at Amazon saw there was a revenue potential in outsourcing their infrastructure to external entities who would want to be free of the worry of building their their own server firms. As such, AWS would take care of the storage, software and security concerns of such parties who would then move on to concentrate on their core businesses.AWS sells their products in form of services. Such services include networking, analytics, database, computing, storage, Internet of things and many more. Some of these services are not directly accessible to human users. Their usage is done through application program interfaces (APIs) that have to be integrated with developers over HTTP by using REST and SOAP protocols.

The technology itself is implemented on server firms that are maintained by AWS personnel. These server firms are available in locations all over the world. This architecture is advantageous in the sense that it provides availability and redundancy.

Storage

Amazon S3 (Simple Storage Service) is a storage platform that is offered by amazon via the use of REST and SOAP APIs. It is also offered through the BitTorrent interface. In S3 files and images are stored as objects in a scalable fashion based on your usage. Amazon itself uses S3 to store its own data and media that is found on the Amazon e-commerce network. Unlike most platforms that store data as files, amazon chose to store their data as objects. An object is composed of metadata, the data itself and a unique identifier to fetch it or update it. Object

storage highly advantageous because it is suitable for storing large amounts of unstructured data (Jian et al., 2014).

Amazon markets S3 as a highly available option with low latency. Like all of its other services, Amazon provides it on a scalable basis depending on your demands and needs. An application using the S3 services gets to store its objects under buckets.According to amazon, Glacier is a low cost storage service hosted on the cloud that is used for data archiving and backup. It provides security and durability at a low cost. Data can be retrieved from a few minutes to several hours. Its cost is way cheaper than having to store it on your own premises. Just like the rest of the infrastructure offered by Amazon, Glacier is highly available. This is owed to the distributed infrastructure which is found across the globe. Users are usually served by servers that are closest to them while the others are used for replication. Amazon logs all calls to the glacier storage for auditing purposes. These data is encrypted to ensure that it remains private to the users of the information and no one else. Amazon cloud front is a web service that is designed to have your static and dynamic web content delivered to the clients with the lowest level of latency. Cloud front delivers content through data centers that have been distributed all over the world. The server that is best suited to serve you is refered to as the edge location. According to a typical client server setup, when on requests a file like an image from a server, this file jumps from one location to another until it gets to your computer. This jumping from one computer/server to another is referred to as hoping. The sole purpose of cloud front is to reduce this hopping. The end result of which is to increase latency, reliability and availability of whatever files that are being served.Like all other AWS services Coudfront is provided on a pay as you go basis. That is, you do not have to pay for a set capacity up front. You only pay for what you have used over a given time period.

Databases

Redshift is a fully managed petabyte-level warehouse for data that is provided by Amazon web services. It is scalable from a few gegabytes to the petabyte that you will need to run your data needs. The sole purpose of redshift is to allow businesses to run the analysis that they need to understand their customers and their own activities. After you have provisioned a Redshift cluster by launching a set of nodes you are now in a position to run the queries that you need.

SQL-based tools can be used to run the needed queries on the dataset available on your cluster. The speeds of these queries depend on the data size that you have and how good you or your database administrate is.Management of the Redshift cluster can be done through a variety of ways. You can use the Amazon Redshift console, the AWS command line interface, APIS or a Software Development Kit provided by AWS.Aurora is a relational database that is offered by AWS that is compatible with both MySQL and PostgreSQL. It provides the reliability of all the services that are offered by AWS together with speed. According to Amazon, Aurora is five times faster than MySQL and three times faster that PostgreSQL in terms of throughput (Amazon Web Services, n.d). In a way it provides the speed and reliability of commercial databases at the cost-effectiveness of open source platforms.One can port their MySQL or PostgreSQL databases to Aurora and still use the same tools that they use on these platforms. Shifting to Aurora frees the enterprise and database administrator the stress of thinking about scaling, reliability and efficiency of their database. Tasks such as repair, patching and recovery of databases can still be done through the administration tools that are provided by Amazon for Aurora management.An Aurora instance is created when you create a DB cluster. The DB instance is composed of the several DB instances and the cluster volume responsible for managing the data for those instances. The cluster volume is replicated across different time zones to improve availability. The DB instances are divided into a primary instance and replica volumes. Read write operations are done on the one primary instance that exists in all Aurora instances. Replication is done across all the other instances which only support read operations.

Amazon EC2 and IAM

Amazon Elastic Compute Cloud (EC2) provides a flexible computing environment for the server needs of the entity in question. The EC2 instance that one gets from Amazon provides the concerned party with a scalable way of using the cloud. The user can configure the Linux or windows based operating system to run on their server as well as the storage, CPU or RAM needs that they feel is adequate for their endeavors. The size of these virtual computing resources can be increased or reduced depending on finances and use (Amazon Web Services n.d). Amazon provides web and command line tools to manage this. If one feels that manual configuration can lead to low latency during periods that they have high traffic to their server there is an option for automatic scaling and down scaling. This is done through what is refereed

to as AWS auto scaling.AWS Identity and Access Management is a tool that is used to control the use of the various services that Amazon offers. Amazon allows one AWS account to have users and groups. It is through the IAM console that one can decide the level of access and privileges that users and groups can have when it comes to the services that an entity AWS uses. What is surprising is that AWS offers IAM for free. The account holder only has to pay for the other services already in use.Further limits can be provided to users and groups via the use of fine-grained access control, multi-factor authentication for highly privileged users, managed access control for mobile applications and even integration with a corporate directory.

In fine-grained access control for example, access to specific resources can be limited. You can constrain users to have access to a resource depending on the time of day, their IP addresses and whether or not they are using SSL.

Route 53

Amazon Route 53 is a scalable and reliable Domain Name System (DNS) that is cloud based. It is AWSs way of making developers have an easier time of mapping the IP addresses of their servers to human readable form such as www.example.com instead of number like 19.169.0.34. This service is provided at a relatively cheap price. If you have other services such as EC2 or S3 buckets, Amazon Route 53 can effectively route the traffic that needs this services. It is also used to route users to servers that are outside of the AWS infrastructure. The other advantage provided by Amazon Route 53 is that it can be used to perform health checks on your application endpoints. Amazon Route 53 comes with a simple editor that can be used to manage the end users that are routed to registered applications. In addition to this can be used for Domain Name Registration. Route 53 will then automatically configure DNS settings for the domains that you purchased from AWS.

Kinesis and Lamda

In the present day world where the amount of data comes in at a fast pace and needs to be analyzed with equal speed, Anazon Kinesis does this for you. Kinesis is used to collect process and analyze data as it being streamed in at real-time. It can be used for data such as video, audio and website logs. In a traditional setting a company would need to set up its own data warehouse

infrastructure to do this. A lot of thought would have to go into the design of the data and hardware layers that would form such a big data processing farm. Small businesses that can not afford the equipment or find the expertise to do this can leverage the financial muscle and experience of Amazon to do this via the use of Kinesis.Therefore AWS gives users the ability to process data at real-time instead of days and hours after the events have occurred. This data will come from thousands of sources and at large amounts and yet it will still be processed at a very low latency.AWS Lamda on the other hand is a service that runs code that is written in an AWS supported language. Instead of having to worry about the operating system, processors or CPUs of the server you want your code to run in, AWS allows you to upload your code the way you would do a file and it will just run.Lamda will take care of scaling your code according to the workload that is needed at any given point in time. Charges are made to the AWS account for every 100ms the code that is uploaded to Lamda runs.What makes your code run is what are refereed to as triggers. Triggers can come from endpoints to your application or AWS services such as S3, Dinamo DB and Kinesis to mention but a few.

VPN

AWS Virtual Private Connections VPCs can be done to remote networks through the aid of Virtual Private Network VPN connections. Amazon allows its users to do this through a variety of ways.

The through VPNs managed by AWS itself an IPsec VPN connection can be made between a VPC and a remote network. In this setting there will be a Virtual Private Gateway with two endpoints that come in handy for failover purposes.AWS VPN Cloudhub is another product from AWS that can help users with separate remote networks connects. The Virtual Private Gateway mentioned above is what will be used to ensure secure communication between your separate networks.AWS can still work with third party VPN software from other vendors or from the open source communities. This can be done when you have your VPC running an EC2 instance.

Quicksight

Business intelligence is an idea that is now available to small and big businesses through quicksight with the aid of a few clicks. It can be used to analyze data from excel sheets, MySQL, PostgeSQL and any other source of interest in a matter of minutes.Once you upload this data you

can get graphical visualizations of whatever metrics you want to understand without having to be an analytics expert. It works well with S3, Aurora and other AWS sources of data. The service is highly scalable as the rest of its AWS siblings. One no longer has to invest in the hardware and software that was traditionally needed for business intelligence.

Conclusion

I have highlighted few services provided by AWS. Amazon is adding more services and features every day. Companys can choose the tools and services based on the business needs and the budget.

References

- Amazon Web Services (n.d). *Cloud Computing with Amazon Web Services*. Retrieved April 30. 2018 from https://aws.amazon.com/what-is-aws/
- Jiang, L., Da Xu, L., Cai, H., Jiang, Z., Bu, F., & Xu, B. (2014). An IoT-oriented data storage framework in cloud computing platform. *IEEE Transactions on Industrial Informatics*, *10*(2), 1443-1451.
- Amazon Web Services (n.d). *What is AWS Cloudfront?* Retrieved April 30, 2018 from https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html
- Amazon Web Services (n.d). *Overview of Amazon Aurora?* Retrieved April 30, 2018 from https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Aurora.Overview.html
- Amazon Web Services (n.d). A*mazon s3?* Retrieved April 30, 2018 from https://aws.amazon.com/ec2/