



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 9 Issue: XI Month of publication: November 2021

DOI: <https://doi.org/10.22214/ijraset.2021.38764>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

A walkthrough of Amazon Elastic Compute Cloud (Amazon EC2): A Review

Anurag Choudhary¹, Pradeep Kumar Verma², Piyush Rai³

¹M. Tech Student, Dept. of Computer Science and Engineering, IET. Dr. RMLAU, Ayodhya, U.P, India

^{2,3}Assistant Professor, Dept. of Computer Science and Engineering, IET. Dr. RMLAU, Ayodhya U.P, India

Abstract: Cloud services are being provided by various giant corporations notably Amazon Web Services, Microsoft Azure, Google Cloud Platform, and others. In this scenario, we address the most prominent web service provider, which is Amazon Web Services, which comprises the Elastic Compute Cloud functionality. Amazon offers a comprehensive package of computing solutions to let businesses establish dedicated virtual clouds while maintaining complete configuration control over their working environment. An organization needs to interact with several other technologies; however, instead of installing the technologies, the company may just buy the technology available online as a service. Amazon's Elastic Compute Cloud Web service, delivers highly customizable computing capacity throughout the cloud, allowing developers to establish applications with high scalability. Explicitly put, an Elastic Compute Cloud is a virtual platform that replicates a physical server on which you may host your applications. Instead of acquiring your own hardware and connecting it to a network, Amazon provides you with almost endless virtual machines to deploy your applications while they control the hardware. This review will focus on the quick overview of the Amazon Web Services Elastic Compute Cloud which also containing the features, pricing, and challenges. Finally, unanswered obstacles, and future research directions in Amazon Web Services Elastic Compute Cloud, are addressed.

Keywords: Cloud Computing, Cloud Service Provider, Amazon Web Services, Amazon Elastic Compute Cloud, AWS EC2

I. INTRODUCTION

Cloud computing is becoming a more common method of gaining access to computing resources. Switching and performing on the cloud, It appears that everything these days is happening "in the cloud computing." However, how actually is this vague concept? The simple truth is that it's placed from the other ends of your internet connection — a point on which you can browse applications and services plus safely store your data. Amazon launched its subsidiary Amazon Web Services in July 2002, with the objective of "allowing developers to build creative and entrepreneurial products through their own." Amazon Web Services (or simply AWS) is a global cloud services platform that provides almost everything that an organization needs to develop sophisticated applications that are adaptable, scalable, and reliable. It is a billing model that seems to be "pay-as-you-go," without any upfront or capital expenditures. Amazon provides nearly 100 on-demand services, with the number rising on a regular basis. Amazon.com's offering Elastic Compute Cloud, launched in 2006, revolutionized cloud computing. Elastic Compute Cloud, or EC2, is the backbone of Amazon Cloud Services; it is a cloud computing framework that helps individuals to rent virtual servers on which to execute their own applications. Traditionally, you would be familiar with the following costs while setting up a server for the company: the cost of the server, the cost of the server's maintenance, and the cost of the server's upkeep, the configuration of the server, traffic distribution, resource allocation, privacy, and security. When users adopt EC2, Amazon offers you a service based on the "cloud computing model." On such a "pay as you go" approach, users may borrow a chunk of their desired size through the network infrastructure. We can use Amazon EC2 server instances at any moment, for as long as you need, and for any legitimate use.

II. OVERVIEW OF AWS ELASTIC COMPUTE CLOUD

A. What is Amazon EC2

In 2006, Amazon launched AWS EC2, which is a cloud-based virtual machine that has now dramatically changed the industry. AWS EC2 stands for Amazon Elastic Compute Cloud, which is an Amazon web service that offers highly customizable computational cloud-based services. It is also re-sizable since you can massively increase or decrease the number of server instances that deploy if the computational needs change. On Amazon's EC2, an instance is a virtual server that is being used to run applications. It may also be considered as a kind of aspect of a large computer, with its own hard drive, network connection, operating system, so on and so forth. However, this is all a virtual environment. On even a single physical machine, you may have multiple "small" computers, that are hereinafter known as Instances.

Amazon EC2's straightforward web service interface makes it easy to obtain and configure capacity. It offers you absolute control over your computing resources and allows users to access Amazon's tried-and-true computing infrastructure. Amazon EC2 transforms computing fundamentals by allowing you to pay only for the services that you actually use. Instead of acquiring physical servers, you might go for a server with more CPU capacity, RAM capacity, and then you would buy a server over a five-year term, so you would have to plan for five years ahead of time. Amazon EC2 gives developers the tools they really ought to create resilient applications that segregate themselves from typical situations.

B. Why use Amazon EC2

Assuming you are a developer and want to work independently, you purchase some servers, estimate the right capacity, and also have considerable computing power. You must now ensure that the necessary patches are updated on a regular basis, and also resolve any issues that may arise at the server-side. However, if you purchase an EC2 instance, you won't have to worry about any of these issues since Amazon will handle everything; all you have to do is focus on your application.

C. What Exactly is Indeed an AWS Instance

Instances in AWS are generally virtual environments. Such virtual environments are usually segregated from the fundamental base operating system. It is also an on-demand service, which means that users can rent virtual servers (instances) hourly and deploy their applications on them. As Aforementioned, users can easily increase or decrease the number of ec2 instances depending on the dynamism of their website's traffic. AWS EC2 instances also eliminate the need for additional hardware and software expenses and maintenance.

D. Amazon EC2 Main Feature

- 1) An instance is a type of virtual computing environment. EC2 is a platform that enables you to deploy your server in the cloud. In AWS, servers are usually referred to as instances. Since they are virtualized, you might access your instance from almost anywhere and at any time.
- 2) An instance is created from an AMI, or Amazon Machine Image, which is a pre-configured template of operating systems and applications. This eliminates the need to go through all of the application post-OS deployment configuration steps, as well as the overall distribution time for instances.
- 3) As an instance type, EC2 offers a wide range of instances with different CPU, memory, and network configurations. This allows the user to select the best instance based on your CPU, memory, and network architecture.
- 4) Concerning the instance's security, EC2 provides a Key Pair. If you do not have one, you must create one when you create your instance. AWS retains the public key and provides you with the private key. Nobody can ever access your instance unless you provide the private key.
- 5) AWS allocates one optional public IP address and one private IP address to your EC2 instance once you start to create it. The public IP address, on the other hand, is dynamic and might fluctuate if your instance is terminated or rebooted. Elastic IP addresses can be used to acquire a static public IP address.
- 6) The AWS cloud spans many continents and countries. You may choose the location of your instance by utilizing regions and availability zones.

E. Type of Amazon EC2 Instance

In principle, Amazon EC2 instance types offer flexibility and scalability through various compute, memory, and storage parameters within each instance type.

A list of AWS EC2 Instance Types is provided below:

- 1) *General Purpose Instances*: The most widely used; is used for web servers, development environments, and so on. A1, M5, M5a, M4, T3, T3a, and T2 are examples of general-purpose instances.
- 2) *Compute Optimized*: Excellent for compute-intensive scenarios such as scientific modeling or high-performance web servers. It seems to be more expensive than other instance kinds but usually deploys faster. C5, C5n, and C4 are examples of compute-optimized instances.
- 3) *Memory Optimized*: Used as memory-intensive applications including real-time big data analytics or running Hadoop or Spark. R5, R5a, X1, and X1e are examples of memory-optimized instances.
- 4) *Accelerated Computing*: Add additional hardware (GPUs, FPGAs) and provide massive volumes of parallel computing for applications like graphics processing. P3, G3, and F1 are examples of accelerated computing instances.

- 5) *Storage Optimized*: Suitable for activities requiring huge amounts of storage, particularly those requiring sequential read-write, also including log processing. H1, D2, and I3 are examples of storage-optimized instances.

F. Pricing Option for Amazon EC2 Instance

Amazon offers a huge free tier for experimentation to many of its products, such as 750 free hours for EC2 instances over a period of 12 months. Time to purchase an instance, users get three options:

- 1) *On-Demand Instances*: Only pay for the services you utilize. There is no long-term commitment or additional fees.
- 2) *Reserved Instances*: You must be committed to a one-year or three-year term. That dedication allows you to save money. The machine is never turned off.
- 3) *Spot Instances*: Purchase unused EC2 instances, which can drastically lower your Amazon EC2 prices.

To sum up, in pricing options, a reserved instance is suitable if you have a long-term demand for an always-on instance. On the other hand, you can save money by using an On-Demand or Spot instance. However, if it is required that the machine be ready when you need it, On-Demand is your best alternative.

G. Key components of AWS EC2

Key components of AWS EC2 service that you can leverage to use those features. Since there are so many components, we will keep it short. We will cover most of the components in-depth in its separate blog post

1) Instances

- a) *Instances*: Instances are essentially servers that are hosted in the AWS cloud through the EC2 service.
- b) *Launch Templates*: Launch templates allow you to save launch settings so you don't have to enter them every time you start an instance.
- c) *On-demand Instance*: Provides a non-predefined use duration instance where you pay as you go. When you launch a new instance, this is the default instance type.
- d) *Spot Instances*: Spot Instances provide excess computation power in the AWS cloud at significant savings over On-Demand instances.
- e) *Reserved Instances*: Reserved instances offer savings over on-demand instances. A reserved instance is reserved for a predetermined period with various payment methods.
- f) *Dedicated Instance*: Runs on a virtual private server (VPS) with hardware dedicated to you.
- g) *Dedicated Host*: The physical server is entirely devoted to you.

2) Images

- a) *AMI*: AMI offers you a template with the OS and application pre-configured to help you decrease the number of instances you need to deploy.
- b) *Bundle Tasks*: A Windows instance store-backed AMI may be created by bundling a Windows instance.

3) Elastic Block Store

- a) *Volumes*: EBS volumes are long-lasting, block-level storage devices that might be attached to the single EC2 instance.
- b) *Snapshots*: Point-in-time snapshots may be used to back up the data on your Amazon EBS volumes to Amazon S3. Snapshots are incremental backups, which means they store just the blocks on the device that have changed since the last snapshot.
- c) *Lifecycle Manager*: Amazon Data Lifecycle Manager (DLM) for EBS Snapshots gives a convenient, automated method of backing up data saved on Amazon EBS volumes. By defining lifecycle rules based on tags, you may specify backup and retention schedules for EBS snapshots. You no longer need to rely on bespoke scripts to make and maintain backups with this functionality.

4) Network & Security

- a) *Security Groups*: A security group serves as a virtual firewall for one or more instances, controlling traffic.
- b) *Elastic IP*: The Elastic IP address is a static IPv4 address intended for use in dynamic cloud computing. Elastic IP allows you to connect to the internet.
- c) *Placement Groups*: Instances can be launched or started in a placement group, which governs how instances are put on the physical machine.

- d) *Key Pairs*: To encrypt and decode login information, Amazon EC2 uses public-key cryptography. Public-key cryptography encrypts a piece of information, such as a password, using a public key, and then the receiver decrypts the data with a private key. A key pair is made up of public and private keys.
 - e) *Network Interface*: A network interface (This documentation refers to a network interface) that is elastic Is a virtual network card-like logical networking component in a VPC.
 - f) *IAM Roles*: For AWS access management, the Identity & Access Management (IAM) role is utilized.
 - g) *VPC*: The Virtual Private Cloud, or VPC, allows you to deploy AWS resources within a predefined virtual network. This virtual network is quite similar to a typical network that you would run with your own data center, but with the added benefit of utilizing AWS's scalable infrastructure.
- 5) *Load Balancing*
- a) *Load Balancers*: Load balancing spreads incoming application or network traffic over different destinations in several Availability Zones, such as Amazon EC2 instances, containers, and IP addresses. Load Balancing automatically scales your load balancer as traffic to your application varies over time and can handle the overwhelming majority of workloads.
 - b) *Target Groups*: Target groups route traffic from a load balancer to a specified destination.
- 6) *Auto Scaling*
- a) *Launch Configurations*: An Auto Scaling group can utilize a launch configuration to launch EC2 instances. When you build a launch configuration, you provide data for such instances, including the Amazon Machine Image ID or AMI, instance type, key pair, one or more security groups, and mapping of block devices
 - b) *Auto Scaling Groups*: An Auto Scaling group is indeed a collection of EC2 instances with comparable properties and for the purposes of instance scaling and management, are considered as a logical grouping.
- 7) *Monitoring*
- a) *Cloud Watch*: Amazon CloudWatch provides real-time monitoring of your AWS resources as well as the apps you host on AWS. CloudWatch can be used to gather and track metrics, which are variables that can be measured for both resources and applications. CloudWatch alarms provide messages or make adjustments to the resources you're monitoring depending on the criteria you define.
 - b) *Status Check*: Amazon EC2 does automatic checks on all running EC2 instances to detect hardware and software problems. The results of these status checks can be viewed to discover specific and detectable issues.
 - c) *Events*: For any announcement, update, maintenance, or disaster, AWS publishes events for each region.
- 8) *System Manager*: AWS Systems Manager is a suite of tools for setting and administering Amazon EC2 instances, on-premises servers and virtual machines, and other AWS services at scale. Systems Manager offers a uniform interface for quickly centralizing operational data and automating activities across your AWS services.
- 9) *Region & Availability Zone*: Amazon EC2 is available in a variety of locations across the world. These areas are divided into regions and Availability Zones. Each region is a geographical area in its own right. Availability Zones are isolated places found in each region. Amazon EC2 allows you to deploy resources, including such instances and data, in various locations. Unless expressly requested, resources are not duplicated between regions.

H. How to Enhance your Amazon EC2 performance

While you could begin conveying an application on a General-Purpose example and randomly pick an instance size, it is improbable that you will have picked the most ideal choice for ideal application execution. The mark of Amazon's particular example types families, case size choices, and added instance highlights is to make it simpler to coordinate with your application's presentation needs to the specific abilities given by a particular case type. Coordinating with your application to a proper EC2 instance type and size is a decent initial phase in upgrading application execution. To go to a higher level, you will need, nay need, to screen your application's real presentation. Regardless of whether you pick what is by all accounts the 'right' instance, it is conceivable that impediments in the product, framework, or by and large framework engineering cause execution issues. At the point when execution issues are recognized, it is then conceivable to re-evaluate the example decision and go through the testing interaction again to approve the new decision. There are many burdens testing and application profiling devices that can be utilized to acquire the essential bits of knowledge.

I. Recent Amazon EC2 Challenges

EC2 would require frequent redesign with the security patches which cause weight to oversee. Amazon delivers new AMI (Amazon Machine Images) which has most recent fixes oftentimes. The client needs to redesign the EC2 with the most recent AMI. This isn't straightforward work yet could be sensible to utilize programmed scripts. Be that as it may, it is a monotonous undertaking.

III. CONCLUSION

This review paper focuses on Amazon EC2 which is a key component of the AWS nature that is most sought after in the domain of cloud computing environments compared to other web services providers, and it presents an overview of Amazon EC2 uses, features that include the multiple types of EC2 instances, as well as pricing. We can deploy an EC2 instance using provided AMI (Amazon Images, a pre-combined package of OS + Software groups as per requirements), and can specify security protocols / open ports/access limitations using security groups. To prevent data loss, and instance recovery we have a great option of Snapshots, to distribute static files globally, CloudFront can be integrated, and recent challenges with EC2 instances are explained, a practitioner to consider that contributes significantly to improving the efficiency of EC2 instances.

IV. ACKNOWLEDGEMENT

I Anurag Chaudhary, student of Master of Technology, Grateful to IET., Dr. Rammanohar Lohia, Avadh University, Ayodhya. I have completed my Research Paper under the supervision of my guide Er Pradeep Verma and Co-Guide Er. Piyush Rai.

REFERENCES

- [1] G. A. Kulkarni and J. Gambhir, "CLOUD COMPUTING-INFRASTRUCTURE AS SERVICEAMAZON EC2," 2011. [Online]. Available: <https://www.researchgate.net/publication/234166267>.
- [2] J. Lancon et al., "AWS EC2 Instance Spot Price Forecasting Using LSTM Networks." [Online]. Available: <https://scholar.smu.edu/datasciencereviewhttp://digitalrepository.smu.edu>. Available at: <https://scholar.smu.edu/datasciencereview/vol2/iss2/8>.
- [3] R. Saini and R. Behl, "An Introduction to AWS—EC2 (Elastic Compute Cloud)," in Proceedings of the International Conference on Research in Management & Technovation 2020, Jan. 2020, vol. 24, pp. 99–102, doi: 10.15439/2020km4.
- [4] K. Ravi Chythanya, G. Sunil, K. Sudheer Kumar, S. N. Korra, and A. Harshavardhan, "Security and safety in amazon EC2 service – A research on EC2 service AMIs," Int. J. Innov. Technol. Explor. Eng., vol. 8, no. 6 Special Issue 4, pp. 736–738, Apr. 2019, doi: 10.35940/ijitee.F1149.0486S419.
- [5] Association for Computing Machinery. and SIGAPP., Proceedings of the 27th annual ACM symposium on applied computing 2012 : Symposium on Applied Computing : Riva del Garda, Trento, Italy, March 26-30, 2012. ACM Press, 2012.
- [6] Institute of Electrical and Electronics Engineers. and IEEE Computer Society., E-Science 2009 : proceedings of the 2009 5th IEEE International Conference on e-Science Workshops : 9-11 December 2009 : Oxford, UK. IEEE Computer Society, 2009
- [7] "What is AWS ec2?" <https://www.jigsawacademy.com/blogs/cloud-computing/what-is-aws-ec2/> (accessed Jul. 17, 2021).
- [8] "Everything You Need To Know About Instances In AWS | Edureka." <https://www.edureka.co/blog/ec2-instances-in-aws/> (accessed Jul. 17, 2021).
- [9] "AWS EC2 Tutorial for AWS Solution Architects | Edureka Blog." <https://www.edureka.co/blog/ec2-aws-tutorial-elastic-compute-cloud/> (accessed Jul. 17, 2021).
- [10] "Cloud computing - Wikipedia." https://en.wikipedia.org/wiki/Cloud_computing (accessed Jul. 17, 2021).
- [11] "Introduction to Amazon Web Services (AWS)." <https://www.simplilearn.com/introduction-to-amazon-web-services-aws-article> (accessed Jul. 17, 2021).
- [12] "What is Cloud Computing? - Salesforce IN." <https://www.salesforce.com/in/learning-centre/tech/cloudcomputing/> (accessed Jul. 17, 2021).
- [13] "Key Features Of AWS EC2 Service | CloudAffaire." <https://cloudaffaire.com/key-features-of-aws-ec2-service/> (accessed Jul. 17, 2021).
- [14] "AWS EC2 Instance Types Explained - Cloud Academy Blog." <https://cloudacademy.com/blog/aws-ec2-instance-types-explained/> (accessed Jul. 17, 2021).
- [15] "AWS EC2 Instance Types | Overview of EC2 Instance Types in AWS." <https://mindmajix.com/aws-ec2-instance-types> (accessed Jul. 17, 2021).
- [16] "What is AWS: Introduction to Amazon Web Services." <https://www.simplilearn.com/tutorials/aws-tutorial/what-is-aws> (accessed Jul. 17, 2021).
- [17] "Amazon Elastic Compute Cloud - Wikipedia." https://en.wikipedia.org/wiki/Amazon_Elastic_Compute_Cloud (accessed Jul. 17, 2021).
- [18] "Amazon EC2." <https://aws.amazon.com/ec2/?ec2-whats-new.sort-by=item.additionalFields.post> (accessed Jul. 17, 2021).
- [19] "AWS EC2 - Javatpoint." <https://www.javatpoint.com/aws-ec2> (accessed Jul. 17, 2021).



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)