Report: Cloud Platform Amazon

Aigerim Zhalgasbekova
PERCCOM Cohort 3

Course: Computing clusters, grids, clouds
By Professor Andrey Y. Shevel

ITMO UNIVERSITY

Introduction

Today cloud computing plays an important role in technology development. It is a paradigm where computing resources are available when needed, and you pay for their use in much the same way as for household utilities. Just as water is piped to your home and you pay for as much or as little as you use, cloud computing resources are available whenever needed and charges are based on how much you use them. When you turn it off, the water that you would have used is available for use by others and, in the same way, shared cloud resources can be used by others when not used by you.

Widespread cloud computing is made possible by the Internet, and this is the most common way of accessing cloud resources. There is a wide range of platforms which offer different cloud services provided by Amazon, Microsoft, Google, IBM and others. In this report, Amazon Web Services Platform will be described.

1 What is AWS? Benefits.

Amazon Web Services is cloud computing platform that "provides a highly reliable and scalable infrastructure for deploying web-scale solutions, with minimal support and administration costs and flexibility" [1].

There are some benefits which AWS provides to their costumers:

- *Flexible.* AWS enables organizations to use the programming models, operating systems, databases, and architectures with which they are already familiar. In addition, this flexibility helps organizations mix and match architectures in order to serve their diverse business needs.
- *Cost-effective.* With AWS, organizations pay only for what they use, without up-front or long-term
- commitments.
- *Scalable and elastic.* Organizations can quickly add and subtract AWS resources to their applications in order to meet customer demand and manage costs.
- *Secure.* In order to provide end-to-end security and end-to-end privacy, AWS builds services in accordance with security best practices, provides the appropriate security features in those services, and documents how to use those features.
- *Experienced.* When using AWS, organizations can leverage Amazon's more than fifteen years of experience delivering large-scale, global infrastructure in a reliable, secure fashion [5].

Amazon takes the first place in the market with 57% respondents currently running applications in this cloud.

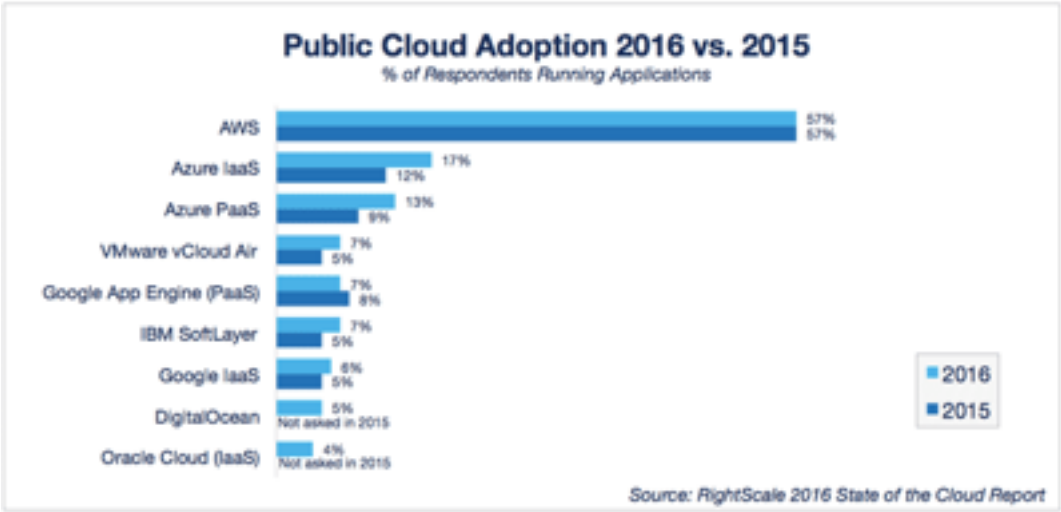

Figure 1: Rating of Cloud Platforms according a popularity [2]

2 AWS Platform Architecture

AWS Platform structure is illustrated on the figure below. In this report, first, the whole platform will be described generally, then the second tier from the bottom will be described in more detail.

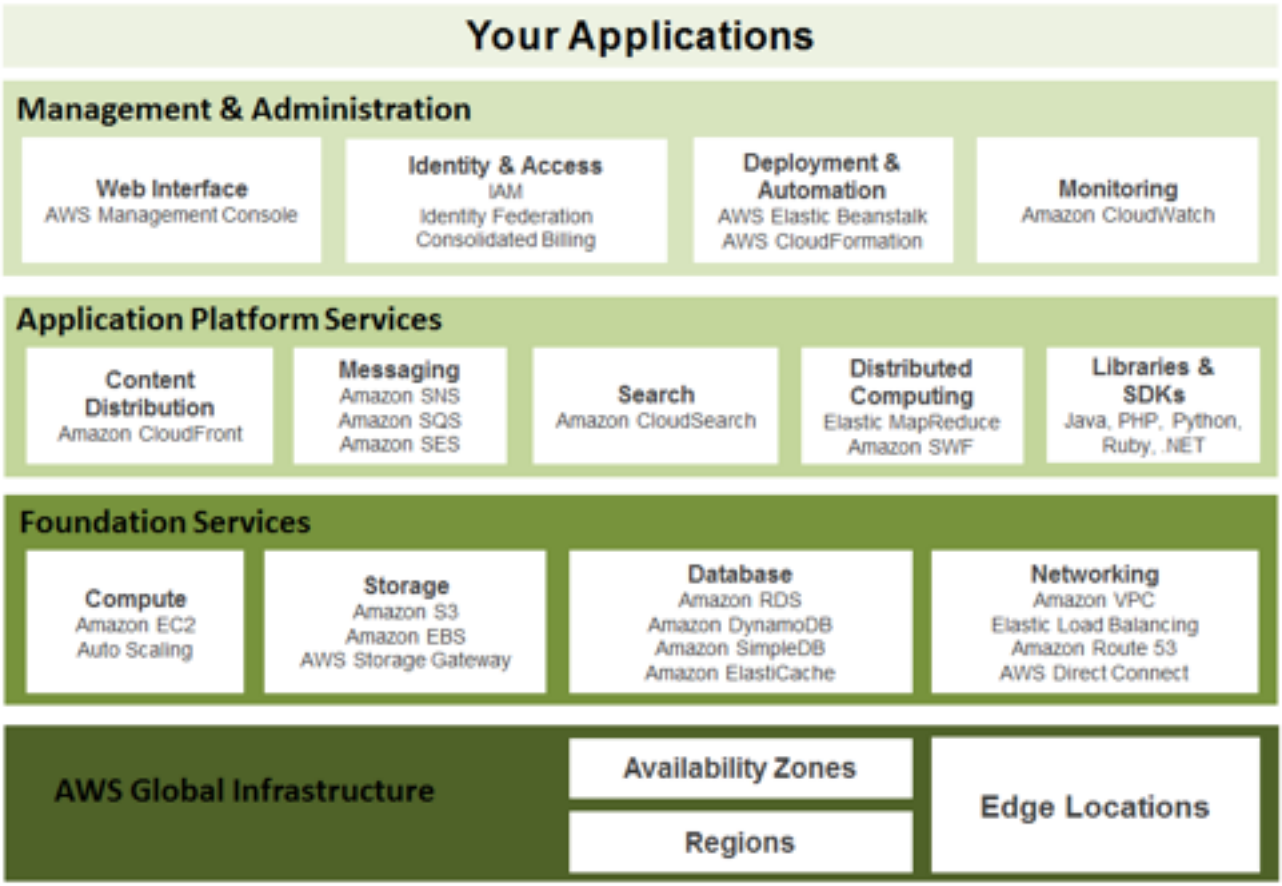

Figure 2: AWS Platform structure

The bottom level is AWS Global Infrastructure which consists of Regions, Availability Zones and Edge Locations. The former are regional datacenters. Availability Zones are clusters of datacenters within the same region interconnected by the low latency networking, that allow customers to run instances in several isolated locations to avoid a single point of failure [3]. So, using Regions and Availability Zones, developers can deploy their services, applications, store the data in the regions they choose, moreover, they can distribute the applications and data across multiple datacenters. Edge Locations provides a content delivery network Amazon Cloud Front and Route 53 DNS services.

The Global Infrastructure provides services such as Foundation Services, those include Compute, Storage, Database and Networking. AWS is built on AWS, it can be seen as a LEGO blocks, so, Foundation Services are foundation primitives upon which everything else is built. On the next tier there are Applications Services such as mentioned above on the figure 2. On top of those, there is an Application Management and Administration Services layer demonstrated on the figure above. All services provided by AWS are accessible through Application Programming Interfaces (APIs). The users can interact with the platform using web console, command line or directly from programming language applications (libraries and SDKs).

2.1 Compute Services

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides 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 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate themselves from common failure scenarios.

Auto Scaling helps you maintain application availability and allows you to scale your Amazon EC2 capacity up or down automatically according to conditions you define. You can use Auto Scaling to help ensure that you are running your desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs.

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. It enables you to achieve greater levels of fault tolerance in your applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic [1].

2.2 Storage Services

Amazon Simple Storage Service (Amazon S3), provides developers and IT teams with secure, durable, highly-scalable cloud storage. Amazon S3 is easy to use object storage, with a simple web service interface to store and retrieve any amount of data from anywhere on the web. Amazon S3 offers a range of storage classes designed for different use cases including Amazon S3 Standard for general-purpose storage of frequently accessed data, Amazon S3 Standard - Infrequent Access (Standard - IA) for long-lived, but less frequently accessed data, and Amazon Glacier for long-term archive. Amazon S3 also offers configurable lifecycle policies for managing your data throughout its lifecycle. Once a policy is set, your data will automatically migrate to the most appropriate storage class without any changes to your applications.

Amazon S3 can be used alone or together with other AWS services such as Amazon Elastic Compute Cloud (Amazon EC2) and AWS Identity and Access Management (IAM), as well as data migration services and gateways for initial or ongoing data ingestion. Amazon S3 provides cost-effective object storage for a wide variety of use cases including backup and recovery, nearline archive, big data analytics, disaster recovery, cloud applications, and content distribution.

Amazon Glacier is a secure, durable, and extremely low-cost cloud storage service for data archiving and long-term backup. Customers can reliably store large or small amounts of data for as little as $0.007 per gigabyte per month, a significant savings compared to on-premises solutions. To keep costs low, Amazon Glacier is optimized for infrequently accessed data where a retrieval time of several hours is suitable.

Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with EC2 instances. EBS volumes are highly available and reliable storage volumes that can be attached to any running instance that is in the same Availability Zone. EBS volumes that are attached to an EC2 instance are exposed as storage volumes that persist independently from the life of the instance. Amazon EBS is recommended when data changes frequently and requires long-term persistence. EBS volumes are particularly well-suited for use as the primary storage for file systems, databases, or for any applications that require fine granular updates and access to raw, unformatted, block-level storage. Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set.

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. The service allows you to securely store data in the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of your data encrypted in Amazon Simple Storage Service (Amazon S3) or Amazon Glacier [1].

2.3 Database

Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.

Amazon DynamoDB is a fast, fully managed NoSQL database service that makes it simple and cost-effective to store and retrieve any amount of data, and serve any level of request traffic. All data items are stored on Solid State Drives (SSDs), and are replicated across 3 Availability Zones for high availability and durability. With DynamoDB, you can offload the administrative burden of operating and scaling a highly available distributed database cluster, while paying a low price for only what you use Amazon DynamoDB is designed to address the core problems of database management, performance, scalability, and reliability. Developers can create a database table that can store and retrieve any amount of data, and serve any level of request traffic. DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent, fast performance. All data items are stored on solid state drives (SSDs) and are automatically replicated across multiple Availability Zones in a Region to provide built-in high availability and data durability.

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases.

Amazon ElastiCache automatically detects and replaces failed nodes, reducing the overhead associated with self managed infrastructures and provides a resilient system that mitigates the risk of overloaded databases, which slow website and application load times. Through integration with Amazon CloudWatch, Amazon ElastiCache provides enhanced visibility into key performance metrics associated with your Memcached or Redis nodes [1].

2.4 Networking

Amazon Virtual Private Cloud (Amazon VPC) allows you to extend your corporate network into a private cloud contained within AWS. Amazon VPC uses IPSec tunnel mode that enables you to create a secure connection between a gateway in your data center and a gateway in AWS.

AWS DirectConnect provide you private connectivity between AWS and your datacenter, in practice, the customer can provision his own cables that are run to peering points connected to the datacenter.

Amazon Route53 is a highly scalable DNS (Domain Name System) service that allows you manage your DNS records by creating a HostedZone for every domain you would like to manage [4].

Conclusion

AWS offers various payment and billing services15 that leverages Amazon's payment infrastructure. All AWS infrastructure services offer utility-style pricing that require no long-term commitments or contracts. For example, you pay by the hour for Amazon EC2 instance usage and pay by the gigabyte for storage and data transfer in the case of Amazon S3. More information about each of these services and their pay-as-you-go pricing is available on the AWS Website.

To sum up, AWS Platform provides:

- Freedom in use of the programming model, language, or operating system (Windows, OpenSolaris or any flavor of Linux) of your choice.
- Freedom in picking and choosing the AWS products that best satisfy your requirements—you can use any of the services individually or in any combination.
- Resizable (storage, bandwidth and computing) resources, therefore, you are free to consume as much or as little and only pay for what you consume.
- Freedom in use of the system management tools you've used in the past and extend your datacenter into the cloud.

References

1. AWS Overview: https://d36cz9buwru1tt.cloudfront.net/AWS_Overview.pdf
2. Cloud Computing Trends: 2016 State of the Cloud Survey: http://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2016-state-cloud-survey
3. A blog about AWS - Inside Amazon's Cloud Computing Infrastructure: http://datacenterfrontier.com/inside-amazon-cloud-computing-infrastructure/
4. AWS Architecting for the Cloud: Best Practices: https://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf
5. Outline of requirements for a modern, robust, industry-leading technology infrastructure platform: https://aws.amazon.com/choosing-a-cloud-platform/

Appendix

Table Appendix - Comparison between Amazon EBS and S3

EBS	S3
A block storage (so you need to format it). This means you are able to choose which type of file system you want.	An object store (not a file system).
As it's a block storage, you can use Raid 1 (or 0 or 10) with multiple block storages	You can store files and "folders" but can't have locks, permissions etc like you would with a traditional file system. This means, by default you can't just mount S3 and use it as your webserver. But it's perfect for storing your images and videos for your website
It is really fast	I/O rate are slower than EBS. However, highly available, redundant. Basically, data loss is not possible (99.999999999% durability, 99.9 uptime SLA).
Storing is relatively cheap. Access to EBS (I/O request) is cheaper than to S3 bucket	Storing megabyte/month is much cheaper than EBS

EBS	S3
With the new announcements from Amazon, you can store up to 16TB data per storage on SSD-s.	Great for short term archiving (e.g. a few weeks). It's good for long term archiving too, but Glacier is more cost efficient.
You can snapshot an EBS (while it's still running) for backup reasons	Great for storing logs
But it only exists in a particular region. Although you can migrate it to another region, you cannot just access it across regions (only if you share it via the EC2; but that means you have a file server)	You can access the data from every region (extra costs may apply)
You need an EC2 instance to attach it to	You can serve the content directly to the internet, you can even have a full (static) website working direct from S3, without an EC2 instance