

A Comparative Study of Amazon Web Service and Windows Azure

Rajeev BV¹, Vinod Baliga², Arun kumar³

Abstract

This paper compares features of two major cloud vendors Amazon Web Services (AWS) and Microsoft's Windows Azure, and helps one to decide the cloud vendor suitable for their requirements to migrate their application. The paper focuses on three main cloud characteristics- Storage, Platform and Compute Service and looks in depth on offerings of AWS and Azure on each of these characteristics. A summary table briefs the distinct and unique features provided by both the vendors.

Keywords

AWS, Windows Azure, Cloud Storage, Compute, Platform.

1. Introduction

Current generation business and industry needs are growing rapidly and hence require scalability, availability etc. features for their business or products. Most of these features required by today's businesses are inherent characteristics of cloud computing and hence are offered by most cloud vendors. But there are a host of vendors to choose from and features provided by every single one of them differ in many ways. In this paper we have opted to compare the features provided by two major cloud vendors – Amazon Web Services and Microsoft Windows Azure with an aim to help the reader decide which cloud vendor is more suitable to their requirements.

2. Features Comparison of Amazon Web Services (AWS) and Windows Azure

In this section we compare Storage, Platform and Compute services of both Amazon Web Services (AWS) and Windows Azure [1].

Rajeev BV, Technology Evaluation Competence, Siemens Technology and Services Pvt Ltd, Bangalore, India.

Vinod Baliga, Technology Evaluation Competence, Siemens Technology and Services Pvt Ltd, Bangalore, India.

Arunkumar, Technology Evaluation Competence, Siemens Technology and Services Pvt Ltd, Bangalore, India.

A. Storage services

Storage Services are classified as follows:

Structured storage: Structured storage is used to store data that has multiple attributes and allow user to query the storage based on these attributes. Structured storage is further classified into relational and non-relational data stores. Relational database has a fixed predefined schema.

The tables that form the database are related to each other and users can query the data using SQL. Non – relational structured storage are mostly key-value based data store. They do not necessarily have features such as table relationships, schema or indexes. Amazon SimpleDB and DynamoDB are the structured storage services offered by AWS [2]. Windows Azure Table is provided by Microsoft as a part of its structured storage services [3].

B. Comparison of Amazon RDS and SQL Azure

Amazon's Relational database service is called Amazon RDS and Windows Azures relational database is called SQL Azure.

Some of the features like Provisioned IOPS (Input/output per second) are provided only by Amazon Web Services. With Provisioned IOPS, user can choose the I/O rate between 1000 IOPS to 10000 IOPS as per his needs. Amazon RDS can be deployed in two ways of storage namely Standard and Multi-AZ deployment. In a Standard deployment, backups are stored in the same zone whereas in Multi-AZ deployment backups are stored in different zones.

Comparison of Amazon SimpleDB and Windows Azure Table storage

Amazon SimpleDB and Windows Azure tables are the structured storage services provided by AWS and Microsoft respectively [4].

One of the major differences between the two vendors is the consistency model used by them. SimpleDB offers Eventual Consistency wherein the latest updated record may not necessarily be returned to the GET request that follows. Azure Table offers Strong Consistency wherein each and every GET request is guaranteed to fetch the latest updated value. In case of eventual consistency, the time taken by the update request could be lesser, as the response

doesn't wait for all replicas of the data to have the latest updated value.

Comparison of Amazon DynamoDB and Windows Azure Table storage

Amazon DynamoDB supports Throughput Provisioning where users can specify number of I/O's required for the DynamoDB, so that Amazon reserves the required resources to serve the requested I/O's.

Unstructured storage: Unstructured storage consists of data that does not have definite structure. E.g. images, videos and documents etc. Simple Storage Services (S3) and Blob Storage are the unstructured storage services supported by AWS and Windows Azure respectively [6].

Comparison of Amazon Simple Storage Service (S3) and Windows Azure Blob storage

C. Compute services

Compute services describe compute of both the cloud vendor AWS and Windows Azure.

AWS Compute

AWS Compute is provided in terms of VM's. Amazon provides a service called EC2 (Elastic cloud computing), by which we can create VM's and computation can be carried out. Following are the various components of EC2:

- **Amazon Machine Images (AMI) and Instances:** An EC2 instance is created by selecting an AMI. An AMI is a collection of software, E.g. an operating system, application server and applications etc. The selected AMI is copied to virtual server and starts running. Amazon provides multiple AMI's but still user can customize AMI according his needs and start running the instance. Select an instance type based on the amount of memory and computing power that you need for the application or software that you plan to run on the instance.
- **Regions and Availability Zones:** Amazon data centers are present in many regions, the user can select any of the regions to deploy his instance and pricing is based on the selected region. A region consists of multiple Availability Zones, where each Availability Zone is isolated from other to avoid failures. All zones are connected through low latency network.
- **Storage:** Amazon EC2 offers the following storage options

- Amazon Elastic block storage

- Amazon EC2 instance store

Amazon Simple storage services (S3)

- **Database:** If the application running on Amazon EC2 needs a database, then use Amazon RDS. This enables the application to access the relational database in the cloud.
- **Networking and Security:**
 - Elastic IP: IP is assigned to the AWS account instead of EC2 instance and can be assigned to any EC2 instance that belongs to IP assigned AWS account.
 - Network security group: Acts like a firewall, where we can define security rule for inbound traffic.
- **Monitoring , Auto Scaling and Load Balancing**
 - **Monitoring:** User can get monitoring data like CPU utilization, disk read of the instance, by using the Management console or Cloud watch API. Management console generates the graphs and Cloud watch API gets the data.
 - **Auto Scaling:** Auto Scaling is a web service designed to launch or terminate EC2 instances automatically based on user-defined policies, schedules, and health checks. Auto Scaling is useful for maintaining a fleet of Amazon EC2 instances that can handle the presented load.
 - **Load Balancing:** Amazon Web Services (AWS) provides Elastic Load Balancing to automatically distribute the incoming application traffic across multiple Amazon Elastic Compute Cloud (Amazon EC2) instances in an AWS account. It detects unhealthy instances and reroutes traffic to healthy instances until the unhealthy instances have been restored. Elastic Load Balancing automatically scales its request handling capacity in response to incoming traffic [11].

Windows Azure Compute

Applications in Windows Azure are implemented as one or more roles and azure can run multiple instance of each role, by using built-in load balancer.

Different roles of implementation are:

- **Web Role:** - Mainly used for hosting front facing application such as ASP.NET,

Silverlight etc. Web role have preconfigured IIS7, so deploying application become easy. Developers can also create applications in native code—using the .NET Framework isn't required. This means that they can install and run non-Microsoft technologies as well, including PHP and Java.

- **Worker Role:** - Mainly used for running tasks comparable to background Windows Services. Difference between worker role and web role is worker role does not have preconfigured IIS. Applications running in worker role aren't meant to be hosted; they are long running applications meant for heavy computation. Developer is free to use .Net Framework or other software that runs in windows or other software runs on windows.
- **VM Role:** The VM Role gives users the freedom to control the operating system image. With this control comes the loss of automated consistent OS servicing. The developer is responsible for making sure that the OS image is up to date through the creation of difference disks. OS servicing increases the cost of development, testing and maintenance.

3. Example Use Cases

Scenario – 1

- ASP Application on cloud (1 instance)
- DB on premise – SQL Size: 200 GB

Outbound data Transactions: 10 GB / month



Fig 1:- Application and Database are on premise on windows platform (Without cloud).




Fig 2:- Application is in cloud and Database is on premise on windows platform (With cloud).

Recommendation:

- **Pricing:** AWS: As shown in Fig 2 the Database is on premise, there is no database instance on cloud, so only data transfer charges are incurred here. Data transfers charges depend on type of transaction. Refer the transaction in AWS S3 table. Medium Instance: AWS: 0.230 \$ / Hr. Data transfer charges=\$0.120 per GB (1st GB Free) = 9*0.120=\$ 1.08. Azure: Data transfer charges in Azure are fixed. Medium Instance: Azure: 0.16 \$ / Hr. Data transfer charges up to 10 GB charges are \$0.1 per GB (free for 1st 5 GB). = 5*0.1=\$0.5. Recommended: Azure because its transaction costs is lesser than AWS.
- **Maintainability:** AWS: - As the application is deployed in a VM, which have an existing OS image or a customized image so we need to take care of updating the OS.

Azure: Azure being PaaS, platform updates like OS updates, patches are automatically taken care of, while one has to incur upgrade cost on AWS.

Recommended: Azure.

- **Portability:** Both Platforms offer equal capability to port the existing app to cloud.

Recommended: Equals (Either one)

Scenario – 2

- ASP Application on cloud (1 instance)
- DB on cloud – SQL Size: 200 GB, as shown in Fig 3.
- Outbound data Transactions: 10 GB / month.



Fig 3:- Application and Database both are on cloud on windows platform

Recommendation:

- Pricing: AWS: AWS database instance considers: storage charges + data transfer charges + DB instance charges.

DB storage charges= \$0.10 per GB-month
i.e. $200 \times 0.10 = \$ 20$ / month.

Data transfer charges=\$0.120 per GB (1st GB Free) =
 $9 \times 0.120 = \$ 1.08$.

DB instance charges=\$0.17 per hour for small instance. (For web edition of SQL Server)

DB instance charges=\$0.65 per hour for small instance. (For standard edition of SQL Server)

Total cost (considering web edition of SQL Server):
143.48\$ / month.

Total cost (considering standard edition of SQL Server): 489\$ / month.

Azure: Azure does not charge for DB instance because we don't know exactly whether database is running in separate VM or it's shared. So consider storage charges + data transfer charges.

Storage Charges=\$ 225.99/month for 150 GB + Data transfer charges.

Data transfer charges=up to 10 GB charges are \$0.1 per GB (free for 1st 5 GB). = $5 \times 0.1 = \$0.5$.

Total cost: 226.49\$/month.

Recommended: For standard and higher versions of SQL server, Azure seems to be more economical than AWS. However, for basic versions of SQL Server, AWS seems to be good.

- Portability: AWS: If SQL database size is within 1TB then, we can directly export DB to Amazon RDS.

Azure: In our example our DB size is 200 GB which is more than SQL azure limit; in this case we need to shard the database and store it in two different

servers.

Recommended: Either one (However note, In SQL Azure we'll have to shard DB size after 200 GB).

- Maintainability: AWS: Amazon RDS automatically updates the under-lying database and their some extra feature like Automatic Host replacement in case VM failure on which DB was running.

Azure: Azure also updates any new patches of database if available.

AWS: As the application is deployed in a VM, which have an existing OS image or a customized image so we need to take care of updating the OS.

Azure: Azure being PaaS, platform updates like OS updates, patches are automatically taken care off, while one has to incur upgrade cost on AWS.

Recommended: - Azure.

Scenario – 3

- Java based MVC Application with MySQL DB are in Cloud.

Recommendation:

- Platform: LINUX
- Pricing for my MySQL;

AWS: Small Instance: 0.09\$/Hr. Medium instance: 0.18\$/Hr.

Azure: Small Instance: 0.08\$ - 0.12\$/Hr. Medium instance: 0.16\$ - 0.24\$/Hr.

4. Conclusion

The paper is helpful for users to choose one of the cloud vendor based on the features like storage, compute and platform. All 3 features are compared in depth. AWS provides some specific features like Bit Torrent, Server Side Encryption and many more features. While Azure provides features like Silverlight tool to work with database etc. Some example use cases are shown to get clarity and the paper suggests the best suitable vendor for the described scenarios.

References

[1] Shivani Goyal,IA Comparative study of cloud computing servc providers,International Journal of Advanced Research in Computer Science and Software Engeering, Volume 2,Issue 2, February 2012.
[2] James Murty, Programing Amazon Web

Services, o'reilly media, Inc., 1005 Gravenstein North, Sebastopol, United States of America.

[3] Roberto Brunetti , —Windows Azure step by step, o'reilly media ,United States Of America.

[4] Tejaswari Redkar and Tony Guidici, —Windows Azure Platform , APress Second Edition.

[5] Amazon SimpleDB Pricing [Online] Available : <http://aws.amazon.com/simpledb/pricing/>

[6] AWS Simple Storage Services [Online] Available: <http://awsdocs.s3.amazonaws.com/S3/latest/s3-gsg.pdf>

[7] AWS Simple Storage Services [Online] Available at: <http://docs.aws.amazon.com/AmazonS3/latest/dev/S3TorrentCharge.html>

[8] Amazon DevPay [Online] Available: <http://aws.amazon.com/devpay/>

[9] Windows Azure [Online] Available: http://www.microsoft.com/industry/government/products/azure_platform/default.aspx

[10] Amazon DynamoDB [Online] Available: <http://aws.amazon.com/dynamodb/>

Arun Kumar is holding Bachelors Degree in Computer science from Appa Insititute of Engineering and technology and MTECH in Computer Science and Engineering at Siddaganga Insititute of technology Tumkur, Karnataka, INDIA and is currently working at Siemens Technology and Services Pvt. Ltd, Bangalore - 560100. He has total experience of 11 months spread across designing and developing Multi-Dimensional data structures and web application and interested in cloud computing.

Vinod Baliga is holding a Bachelor's degree in Information Science and Engineering and has 8 years experience in working with Microsoft Technologies like ASP.NET, WCF and Windows Azure. He currently works at Siemens Technology and Services focusing mainly on Cloud, BigData and NoSQL.

Rajeev B V is holding a Bachelor's Degree in Mechanical Engineering from Siddaganga Institute of Engineering and technology. Currently working as a Technical Lead at Siemens Technology and Services Pvt. Ltd, Bangalore. INDIA-560100 His areas of interests are Cloud, UI technologies and patterns

Table 1: Comparison of Amazon RDS and SQL Azure

	Features	Cloud Vendors	
		Amazon RDS	SQL Azure
1	Database Supported	MySQL, SQL Server, Oracle.	SQL Server.
2	Provisioned IOPS Storage	Amazon RDS supports provisioned IOPS storage for I/O sensitive applications.	Does not provide Provisioned IOPS Storage.
3	Database sizes	AWS RDS offers, from 100 GB to 1TB	SQL Azure offers from 1GB to 150GB
4	Backup/Restore	Turned on by default and we can take backup within the retention period, maximum retention period is 35 days.	Taking backup data in SQL Azure is done by the user in one of the following ways: Creating copy of database using Database copy or using BACPAC which is a technique to export and import files to or from SQL database.
5	Pricing	RDS charges based on the type of deployment selected. Deployment can be Standard and Multi-AZ deployment. For example: -In Standard deployment, if we select medium instance then charges are \$0.180 per hour. For Multi-AZ deployment if we select medium instance then charges are \$0.360 per hour.	SQL Azure charges are based on storage space utilized. Up to 5 GB charges are \$25.99/month, per GB \$5.20. Any further usage up to 150 GB charges are \$225.99/month, per GB \$1.51.
6	Disaster recovery	Turned on by default.	Disaster recovery is done by using some tooling setup.
7	Tooling support	There is no such tool available.	SQL Azure provides Silverlight Database tool, using this tool we can create, delete table and add

			data, create stored procedures.
8	Feature set	Amazon RDS supports every feature that MySQL provides.	SQL Azure has subset of features that SQL Server provides.
9	DB Snapshot	Supports DB Snapshot	
10	Rename Database Instances	Supports renaming existing DB instances	Renaming can be done using queries
11	Automatic Software Patching	Deployed instance stays up-to-date with the latest patches	Software patching is done automatically
12	Automatic host replacement	Compute instance gets replaced automatically in case of hardware failure	

Table 2: Comparison of Amazon SimpleDB and Windows Azure Table

	Features	Cloud Vendors	
		Amazon SimpleDB	Table
1	Consistency Provided	Eventually consistent	Strongly consistent
2	Storage Limi	10 GB per storage account.	100 TB per storage account.
3	Pricing [5]	Charges are based on machine utilization. For example first 25 Amazon SimpleDB Machine Hours consumed per month are free. Any further usage will be charged \$0.140 per machine hours consumed. Data transfer charges upto 10 TB/ month is \$0.120 per GB.	Azure table storage charges up to 1TB are \$0.140 per GB. Data transfer charges are applicable.
4	Indexing	Simple DB indexes the data.	Azure table indexes the data based on the partition key and row key.

Table 3: Comparison of Amazon DynamoDB and Windows Azure Table

	Features	Cloud Vendors	
		Amazon DynamoDB	Table
1	Throughput Provisioning.	Supports Throughput Provisioning.	Throughput is controlled by the system.
2	Storage Limit	There is no storage limit.	100 TB per storage account

3	Pricing	Provision throughput's charges are based on number of writes and reads per hour. Data Transfer charges are applicable. E.g. first 1 GB is free. Any further usage up to 10 TB will be charged with \$0.120 per GB.	Azure table storage charges up to 1TB are \$0.140 per GB. Data transfer charges are applicable.
---	---------	--	---

Table 4: Comparison of Amazon S3 and Windows Azure Blob

	Features	Cloud Vendors	
		Amazon S3	Windows Azure Blob
1	HTTP Protocol Support	S3 support HTTP Via REST and SOAP protocol.	WABS HTTP via REST only.
2	Bit Torrent protocol [7]	S3 support Bit Torrent protocol for downloading data from storage.	WABS does not support Bit Torrent protocol.
3	Requester Pay	S3 support Requester Pay.	WABS does not support Requester Pay.
4	Billing and Account Management	S3 support charging Customers for usage of storage present in your account by enabling feature called DevPay [8].	Windows Azure does not support any online billing services .
5	Redundancy level offered	Standard and Reduced Redundancy level are offered.	Geo redundancy and Local redundancy storage are offered.
6	Auto Deletion of Object/Blob	S3 can auto delete the object after specified time period.	Auto deletion feature is not available in WABS.
7	Server Side Encryption	S3 supports Serer Side Encryption.	WABS does not support Server Side Encryption.
8	Ability to ship disks for importing/exporting huge amount of data	Yes	No
9	Transaction charges	S3 transaction charges are based on the type of transaction for example PUT,COPY,POST and LIST are charged \$0.01 for 1000 transactions and GET is charged\$0.01 for 10,000	Azure transaction charges are fixed (currently 0.01\$ for 10,000 transactions).
10	Storage Charges	Storage charges depend on the level of redundancy for ex: Standard 1 TB - \$0.095/GB. Reduce redundancy Storage - 1TB - \$0.076/GB	Azure provides single level redundancy and charges are: Geographically redundant – 1TB - \$0.095/GB and Locally redundant – 1TB - \$0.070/GB

Platform services

Table 5: Comparison of Amazon Web Services and Windows Azure Platforms

	Features	Cloud Vendors	
		Amazon Web Services	Windows Azure
1	Types of cloud computing provided	Provides all 3 types of cloud computing PaaS, SaaS and IaaS, but major player in IaaS.	Provides all 3 Types of cloud computing PaaS, SaaS and IaaS, but major player in PaaS.
2	Programming Language Supported	AWS supports Java, Ruby, PHP and .NET.	Windows Azure supports .NET, PHP, Ruby, Java, and Python [9].
3	Relational Database Support	SQL, MySql Oracle database are supported.	SQL Server relational database is supported in windows azure called SQL Azure, SQL Azure is a relational database hosted by Microsoft data centre.
4	Virtualization Technology	AWS uses Virtualization technology.	Microsoft system which uses Hypervisor based technique for Hardware virtualization called Hyper-V.
5	Storage Consistency	Storage Consistency varies with the different storage types, for example SimpleDB	Windows Azure provides strong Consistency and High availability storage. is eventually consistent storage and Amazon Dynamo DB is strong consistent storage [10].