

MorganCobmstock2014028

by Morgan Combstock

Submission date: 18-Nov-2022 09:21PM (UTC+0000)

Submission ID: 189499669

File name: 24982_Morgan_Combstock_MorganCobmstock2014028_460442_1055712495.docx (195.92K)

Word count: 1455

Character count: 7756



NOVEMBER 1, 2022

CLOUD COMPUTING ARCHITECTURE

ACCA5025 – CW 1

MORGAN COMBSTOCK
2014028

Contents

Extract 2

Introduction & Purpose 2

Virtualisation 2

Virtual Machines 2

As a Service Offering 2

 Infrastructure as a Service (IaaS) 3

 Platform as a Service (PaaS) 3

 Software as a Service (SaaS) 4

Scenario 4

Resilience 4

Security 4

 Authorisation and Authentication 4

 Network Security Groups 5

 Bastions 5

Conclusion 5

References 6

Table of Figures

Figure 1 - IaaS, PaaS & SaaS Comparison 3

Extract

Cloud computing offers consumers the ability to run virtual machines and provision architecture in a few clicks. There are different offerings from service providers, each providing a different level of management or administration on the consumer's part. When operating an infrastructure as a service model, there are more considerations than using software or platform as a service such as, operating systems, OS, data and applications. Security, access and connectivity are also key considerations for using this service.

Introduction & Purpose

The purpose of this document is to explain the terms virtualisation and virtual machine in the context of cloud computing. It will also evaluate the benefits and liabilities associated with migration to the cloud, the different options available and address any security concerns that these may give way to.

Virtualisation

In the context of cloud computing, virtualisation means sharing hardware across multiple systems and users. Where the traditional approach to computer architecture is on-premise, meaning the individual or company has to purchase all the hardware and keep them in a given location, virtualisation allows someone to use hardware that is owned by a provider and located in centralised warehouses.

Virtual Machines


A virtual machine is effectively a computer. Unlike a physical computer though, a virtual machine is a software computer. It's hardware is centralised and components such as it's memory, network interfaces and processor are not physically attached to the computer. The software that it uses to replicate these concepts could also be shared across multiple virtual machines; it is not dedicated resource.

As a Service Offering

The as a service offering is a service model whereby providers offer differing levels of control over virtualisation. These come in three different options:

- Infrastructure as a Service
- Platform as a Service
- Software as a Service

Figure 1 shows the differences in services managed by the consumer and the supplier based on each offering.



[1]

Figure 1 - IaaS, Paas & Saas Comparison

Infrastructure as a Service (IaaS)

Infrastructure as a Service is the first step into virtualisation from on-premise. It delivers fundamental compute, network, and storage resources to consumers on-demand, over the internet, and on a pay-as-you-go basis [2]. This allows the consumer to administer their systems in their own way and have greater control. This however requires more specialism than the other two offerings. The provider has the physical drives and computing power that the virtual machines run on but, the administrator for the consumer is responsible for provisioning virtual machines, configuring the virtual networks, subnet and security therein. This allows the consumer more control over the service but carries an inherent need for more specialist skills.

Platform as a Service (PaaS)

This option is the in between stage of the three. Where Infrastructure as a service requires you to administer and manage the platform from which you develop, platform as a service provides that platform for you to build on. The consumer is only required to build their own applications and manage how their data is stored on the supplied platform.

Software as a Service (SaaS)

Software as a Service is the option that requires the least amount of configuration; none at all other than personal, cosmetic preferences. This type of service is when a product that is ready for use is simply provided to a company or end user. The app itself, the platform it runs on and the infrastructure that it runs out of are all configured by the application provider to minimal to no administration is required by the consumer.

Scenario

For the purpose of this report, the scenario has been given that a company is moving from on-premise to the infrastructure as a service model. The company has two big projects with multiple virtual machines requiring high availability. In the first instance, there is no cost limit or budget. In the second one, there is a reasonable budget and some loss of service is considered acceptable.

Resilience

A key consideration for both legacy systems and more modern cloud systems is resilience. Any data that is saved on a drive, physical or virtual needs to be accessible to a user or an application as required and also needs to be protected against loss.

In terms of data loss, there are a few things that can be considered. Firstly, availability zones. Each provider has data warehouses where the physical infrastructure that support virtual machines are stored. These warehouses are split into separate availability zones, eg. UK-South, UK-West. Data stored in one zone can be replicated in other zones. This means that, if one zone were to go down for any reason, not only is the data not lost, but it is still accessible as it is backed up in another zone which should still be operational. The level of resilience provided will depend on the cost of the service. The more resilience that is required, the more storage across locations is required and this leads to an increase in cost.

Additionally, in terms of access and availability, using availability zones in other countries or areas must be done with great caution. One reason for this is latency. Data housed in the USA or Asia will have a longer ping time for UK customers than data house in the UK. Also, consideration needs to be given to GDPR as laws differ depending on the country you're using.

Security

Authorisation and Authentication

Two key terms regarding security are authorisation and authentication. Authorisation is the process of granting a user or a group permission to access a system or network. Authentication is a means to identify somebody's identity, be it by means of a password, biometrics or an approval code [3]. A

common method for authentication is multi-factor authentication (MFA). This allows a user to authenticate using a second source of validation before being granted access [4].

When authorising a user or a group, the principle of least privilege is a good way to provision access. This principle means giving the lowest amount of access or privilege that user requires to perform the tasks they need to [5].

Network Security Groups

One way to provision a user, group or machine access to resources is using network security groups (NSG). An NSG can be applied to the network interface card (NIC) of a virtual machine or it can be applied to subnet. However, NSGs need to have rules applied to them such that define the input and output traffic that should be allowed or denied.

Bastions

A bastion allows you to connect to an individual virtual machine instance on a virtual network. While the same could be achieved by using a jumpbox or connecting by connecting using a NAT, a bastion allows more secure access. A bastion is, in itself, a virtual machine but it has the sole purpose of allowing access to the other VMs on the network. No public IPs are required as the individual VM private IPs are utilised and there is also no need for an NSG. Transport layer security is used rather than message level.

Conclusion

For the first project, where there is no budget constraints and no room for error with up time, the best option would be to have a higher service level agreement which promises greater uptime. When configuring the virtual machines, the storage and compute capability should be spread across multiple availability zones to prevent data being inaccessible in the event of an issue at a given data centre. This would incur a higher cost, but this could be mitigated to an extent. This could be done with automated scaling of VMs and deallocation. The VMs should be scripted to provision when needed and deallocate once no longer required.

For the second project, as some downtime is allowable, a lower-level service agreement could be opted for. This would allow some server drop, but this would be minimal. For this, a different licence agreement could also be used, rather than a year long subscription.

References

- [1] K. Sookocheff, "IaaS, PaaS, SaaS and Infrastructure," 23 10 2018. [Online]. Available: <https://sookocheff.com/post/opinion/iaas-paas-saas-infrastructure/>. [Accessed 17 11 2022].
- [2] IBM Cloud Education, "IaaS (Infrastructure-as-a-Service)," IBM, 12 07 2019. [Online]. Available: <https://www.ibm.com/cloud/learn/iaas>. [Accessed 17 11 2022].
- [3] Microsoft Learn, "Authentication vs. authorization," 11 02 2022. [Online]. Available: <https://learn.microsoft.com/en-us/azure/active-directory/develop/authentication-vs-authorization>. [Accessed 18 11 2022].
- [4] CISCO, "What Is Multi-Factor Authentication (MFA)?," 2022. [Online]. Available: <https://www.cisco.com/c/en/us/products/security/what-is-multi-factor-authentication.html>. [Accessed 18 11 2022].

FINAL GRADE

65/100

GENERAL COMMENTS

Instructor

Virtualization/VM: 15

Cloud Services: 17

Availability: 11

Security: 22

Feedback:

++++++

I was hoping for a little more discussion of the nature of virtualisation and the components of VMs.

Generally the report is a little brief.

The discussion of security considerations was quite good.

PAGE 1

PAGE 2

PAGE 3

PAGE 4

PAGE 5

PAGE 6

PAGE 7