

Essay1

by Gareth Williams

Submission date: 18-Nov-2022 07:05PM (UTC+0000)

Submission ID: 189492145

File name: 17790_Gareth_Williams_Essay1_460442_369299648.docx (80.89K)

Word count: 3057

Character count: 17332

Assignment 1: Essay

Cloud Computing Architecture ACCA5025

Gareth Williams (1910285)
18th November 2022

Table of Contents

1. Virtualisation & Virtual Machines	3
2. Cloud Services	4
2.1 IaaS – Infrastructure as a Service	5
2.2 PaaS – Platform as a Service	5
2.3 SaaS – Software as a Service	6
3. SME Benefits and Evaluation of moving to IaaS	6
3.1.1 Costs	6
3.1.2 Support	6
3.1.3 Expertise	7
4. Availability.....	7
4.1 High Availability	8
4.2 Acceptable Availability	8
5. Cloud Security Considerations	9
5.1 Authentication & Authorisation.....	9
5.2 Multifactor Authentication (MFA)	9
5.3 Least-privilege permissions.....	9
5.4 Network Security Groups.....	10
5.5 Bastions and Jump-boxes	10
6. Bibliography	11

1. Virtualisation & Virtual Machines

Traditionally, server and PC operating systems and their hosted services are installed directly on to local permanent storage and directly access the physical resources such as CPU, memory, graphics cards and network cards. The hardware the services execute on require maintenance overheads such as firmware patching, tech refresh, cooling and hardware security such as datacentre or a server room on premises. On call staff 24/7 maybe needed to be available and attend, all adding to the cost of running a service for users.

The technical and performance improvements of wider area networks and server hardware allows the ability for services to be hosted by specialist companies rather than your own datacentre. This also has the benefit of economy of scale by having multiple companies hosted in the same datacentre and therefore costs for such items as cooling can be shared. To take the efficiency further, servers and services could be hosted on the same hardware with the aim of sharing costs. One way of achieving this productivity is the use of *virtualisation*.

The IBM web page describes the term *virtualisation* as "... a process that allows for more efficient utilization of physical computer hardware and is the foundation of cloud computing" [1]. An existing instance of a server or pc installed on its own dedicated hardware can be virtualised and run in a virtual environment. Note that the existing operating system will need preparation, key to this is the installation of virtual device drivers and especially the virtual network card that is required for administrative access.

The hardware layer is abstracted away from the guest operating systems by the use of a *hypervisor* layer. The hypervisor controls and shares the physical components with each of the guest operating systems running as virtual machines.

A virtual machine can be created from either clean image, or an existing host can be prepared and virtualised. The ability to virtualise exiting server provides a migration path and flexibility to move services from legacy hardware. As the virtual machine is essentially a file, it can be easily transferred between hardware without the need to rebuild the host or application.

The virtual machine created is now in a file, this can be run locally in organisational datacentres on a virtual platform. The flexibility of the virtual machine ecosystem means we can take the vm file and it can be uploaded and hosted by a cloud provider.

The use of *extensions* can help speed up the process of provisioning of a new virtual machine, the extension is essentially a script that executes to enable and configure services deployed on the server. This means the same scripts can be used to build one to n machines and be confident that each build is configured identically.

2. Cloud Services

The term *cloud computing* is defined by A. Butterfield and G.E. Ngondi in *A Dictionary of Computer Science* as:

“An approach to computing in which the end user connects to a remote network of computers (the cloud) in order to run programs, store data, etc.” [2]

There are many cloud service providers (CSP), with some of the key large multinational companies such as Google, Microsoft, Amazon and IBM.

In the article “Cloud computing adoption and its impact on SMEs’ performance for cloud supported operations”, Abul Khayer et al says of cloud computing:

“It offers an innovative business model for the pervasive, convenient, and on-demand access to the virtualized and distributed resources such as network, servers, storage, applications, and services” [3].

There are three broad groups of cloud services, Infrastructure-as-a-Service, Platform-as-a-Service and Software-as-a-Service. The difference between them is how many of a solutions components are managed by the CSP and how many are managed by the customer. The following diagram taken from the RedHat [2] website, provides a comparison of the main types of cloud services.




Figure 1 – RedHat Service Type Comparison [2]

A customer isn't limited to using one type of category, a solution may be made up of a mix of On-site, IaaS, PaaS or SaaS. Each of the components can be seen as *resources*, and the resources are combined to make up the overall solution.

What is a *managed service* – as a customer you choose which components of a solution stack you wish to be managed by the cloud service provider.

Whether using IaaS, PaaS or SaaS the physical infrastructure costs such as datacentre floor space, physical network, hardware, power and cooling are shared with economies of scale helping to reduce costs. In addition, as the infrastructure is shared, it allows the compute and storage to be pooled and used more efficiently.

2.1 IaaS – Infrastructure as a Service

Infrastructure-as-a-Service is essentially a platform that hosts virtual machines and provides access to them via the internet. The cloud service provider will manage physical networking, storage, servers (physical), and the virtualisation platform (hypervisor). T, Lynn et al in *Measuring the Business Value of Cloud Computing* state:

“...IaaS provides users with a high level of flexibility but requires a high level of IT skills in order to optimise and manage the infrastructure.” [5]

The operating system and any middleware, runtime environments, data and applications are the responsibility of the customer. This responsibility will include any security patching, operating system updates, runtime library management.

Although hardware components are managed, the expertise and skills switch from an understanding of physical networks and hardware resilience concepts to the equivalent virtual component. An understanding of traditional infrastructure concepts of networking, firewalls, domain name services (DNS), hardware resilience etc. are still required to ensure the virtual environment is configured correctly.

2.2 PaaS – Platform as a Service

Platform-as-a-Service (PaaS) is where application code and data are the only components that need managing by an organisation.

The application code is deployed to, and executed in, an *Application Service*. The patching of runtime environments is taken care of by the cloud service provider. In addition, the data may be held in a *managed* database, again the database engine is patched and managed by the cloud provider. Only the data itself and the database schema is the responsibility of the customer.

A key benefit with PaaS is that existing applications would need little modification to be deployed and run in a PaaS environment.

PaaS can be confused with, *serverless* computing, however applications would need to be re-architected for deployment to serverless environments.

2.3 SaaS – Software as a Service

Software as a Service is a fully managed service. A customer is essentially purchasing licenses for access to an instance of an *out-of-the-box* application.

An example of SaaS application is Microsoft and Salesforce who both offer provides Customer Relationship Management software. Microsoft also offer Office services such as Office 365 and Sharepoint.

SaaS will still require cloud infrastructure expertise, for example, the SaaS service tend to rely on Active Directory or the CSPs identity access management (IAM) solutions.

3. SME Benefits and Evaluation of moving to IaaS

This section describes the benefits of moving a small eCommerce application to the cloud platform and specifically IaaS.

3.1.1 Costs

When using IaaS, infrastructure costs are reduced as essentially the server is sharing a datacentre meaning that hardware, power and cooling costs are shared between many companies. As a typical eCommerce site is 24/7 then these costs are high during quiet periods.

At the time an on-premises server hardware reaches end of life and requires a replacement there will need to be a procurement, lead-times and migration of services to new hardware.

Scalability, to expand the ability of the eCommerce application to cope with higher demand, additional compute and storage can be added quickly – and sometimes without disruption to the application. Additional servers can also be added for sideways scaling.

A key benefit to moving the existing application to IaaS is that the application itself doesn't change, therefore, the investment in the eCommerce application code and business processes are protected. This is also a low-risk approach as it doesn't involve the time and cost of re-platforming by redeveloping and testing applications to use PaaS.

The existing developer skill sets do not change.

As the provision and tear-down of VMs is quick, they are useful in application testing where new deployments can be tested easily. Applications can also be tested on older operating systems to test compatibility.

3.1.2 Support

As physical network and server hardware components are taken care of by the IaaS provider, there is no longer need for infrastructure support to be available and attend site.

The virtual machines will be hosted on resilient hardware, any component failures will be taken care of, and the customer may never know there was an issue.

Only local IT desktop support would be required in organisation offices to resolve access to the internet. User access to ecommerce application will be via the internet.

eCommerce application deployment and application support will change little with only connectivity being the main change.

There is no longer a hardware support requirement for someone with hardware skills to attend site or be available in the event of component failure. The responsibility for physical server resources and server reduced server network infrastructure are now the IaaS service providers will offer *Service Level Agreements* with support arrangements.

3.1.3 Expertise

Using IaaS means the existing company investment into the eCommerce application and developer skills and tools don't change – even the DNS name can remain the same. The migration would be transparent to users. Code can the way the website is hosted doesn't change. This means that development costs and application deployment costs are not affected.

Expertise in monitoring spend on resources is required. Finance operations are required to monitor costs and projecting spend on cloud resources. In addition, and understanding of how to configure and tag resources to ensure money is not wasted e.g. ensure unused services are shutdown.

Understand such aspects as automatic scaling is key, where not only are more resources added during peak times, but resources are also removed during quiet times.

4. Availability

The white paper *Resilience in Microsoft Azure* describes *availability* as “The proportion of time that a system is functional and working” and *resilience* as “A system's ability to recover from failures and continue to function” [6].

Cloud Service Providers will take care of hardware component failures as the physical host servers will have resilience built in, for example multiple CPUs, Memory banks, Disk Arrays with RAID configurations. However, there is still possibility of failure of a rack, CSP datacentre network, power outage or disaster. To guard against such failures CSPs provide the ability for customers to choose how resilient they wish to make a service with additional costs becoming a factor.

In Azure, *regions* are abstract geographical areas such as “uk south” and “uk west”. In each region there maybe one or many datacentres, the datacentres within a region will have

dedicated low latency “diverse” [6] network connectivity between them. When deploying resources, a customer can choose which “location” to deploy resources.

Resource Groups are used to containerise and control a collection of different types of resources, the resource group itself can be configured so all resources contained will failover to a chosen *availability zone*. *Azure Availability Zones* is the Microsoft solution to allow resources to be configured to failover between zones without interruption. A mixture of locations, resource groups and availability zones allow customers to build an availability strategy for the services they have deployed.

In the article *On Resilience in Cloud Computing: A Survey of Techniques across the Cloud Domain*, Welsh and Benkhelifa state:

“... the security and resilience of cloud computing is greater than that of traditional paradigms, largely because cloud infrastructures are typically hosted in DCs with greater facilities than those possible to finance or manage by a single organisation. Highly redundant resources, excellent power resilience, excellent physical security, and strong network links to the internet backbone. Therefore, these underlying characteristics ensure these environments are resilient by nature.” [7]

However, there are additional options to further enhance availability and resilience by adding redundancy and automatic failover of services.

4.1 High Availability

To achieve high availability, additional tools and techniques are offered by CSPs and will come with additional costs. The balance for a customer is deciding how much downtime a service can afford versus the additional service costs.

Within Azure - High availability can be achieved by the configuration of *availability zones*. By providing active-active replication of services and data between *zones*. CSPs can provide *Zone-redundant services* with automatic failover.

High cost comes as essentially the resources and storage are duplicated between datacentres. CSP services such as Azure Traffic Manager and load-balancing will detect failures and route network traffic to the available zone.

4.2 Acceptable Availability

Where a company can tolerate some service loss or can't afford high availability approaches, an active-passive strategy maybe more cost effective. This is known as *Zonal services*. In this scenario, a resource is only active in one datacentre and changes to servers and data replicated to a *cold standby*. In the event of failure, there may be a manual process to spin up the cold standby and reconfigure Domain Name Service to switch to the standby endpoint.

If a level of data loss is acceptable, a cheap option maybe to simply take snapshots of virtual machines during quiet times.

5. Cloud Security Considerations

In most cases, cloud services are accessed via the internet and therefore it is key that policies and procedures are in place to protect the cloud resources. This section describes some of the common security techniques.

5.1 Authentication & Authorisation

Authentication – the ability to prove an individual's identity to gain access to services. multifactor authentication is now commonplace to prove someone's identity.

Once granted access to a service or resource, the process of *authorisation* determines what permissions have been granted to use within that service, create new resource, configure existing, delete resource.

CSPs implement Identity and Access Management (IAM), which controls access to resources and permissions to provision, edit or remove resources and resource hierarchies.

5.2 Multifactor Authentication (MFA)

As cloud resources are administered via the internet it means access is available from anywhere in the world (backend servers and not just application front ends), a more robust way of authenticating is required as just username and password pairs can be discovered and used with malicious intent and could provide unlimited access to data and services.

In the article *Multifactor Authentication: Its Time Has Come*, J Reno states:

“Multifactor authentication requires the presentation of two or more of the three authentication factor types: “What you know”, “What you have”, and “What you are”. After presentation, each factor must be validated by the other party for authentication to occur.” [8]

What you know maybe a password, *what you have* could be a phone or fob that produces a code, and *what you are* maybe biometrics such as fingerprint or face-id.

5.3 Least-privilege permissions

Microsoft Azure Learn website describes the term least-privilege as:

“The information security principle of least privilege asserts that users and applications should be granted access only to the data and operations they require to perform their jobs.” [9]

This requires cloud resource administrators to have strategies in place to understand, configure, monitor and review the privileges. Excessive privileges can lead to security being compromised and allow the spread of malicious code between cloud resources. Restricting accounts ensure that if the account is compromised then damaged is limited.

5.4 Network Security Groups

The use of Network Security Groups allows the segregation or grouping of resources within the virtual network to provide a level of security. This provides control over what network traffic is allowed inbound and outbound to resources in each NSG. The NSG can also be configured to the port and network protocol.

In its simplest form, a virtual machine's private IP address will be added to an NSG rule and filters will be applied to network traffic to and from the machine.

Note that as a virtual machine can have more than one virtual network card, it is possible to build complex rules.

5.5 Bastions and Jump-boxes

AWS Security Blog describes *bastion* as:

“A bastion host is a server whose purpose is to provide access to a private network from an external network, such as the Internet. Because of its exposure to potential attack, a bastion host must minimize the chances of penetration.” [10]

The bastion server is hardened against attacks and acts as a jump-box to allow access to cloud resources for administrative purposes.

6. Bibliography

- [1] IBM, “<https://www.ibm.com/cloud/learn/virtualization-a-complete-guide>,” IBM Cloud Education, 19 June 2019. [Online]. Available: <https://www.ibm.com/cloud/learn/virtualization-a-complete-guide>. [Accessed 13 November 2022].
- [2] A. B. a. G. E. Ngondi, A dictionary of computer science, vol. Seventh Edition, Oxford: Oxford University Press, 2016.
- [3] M. S. T. Y. B. M. N. H. Abul Khayer, “Cloud computing adoption and its impact on SMEs’ performance for cloud supported operations: A dual-stage analytical approach,” *Technology in Society*, vol. 60, 2020.
- [4] RedHat, “iaas-paas-saas,” RedHat, 16 August 2022. [Online]. Available: https://www.redhat.com/en/topics/cloud-computing/iaas-vs-paas-vs-saas?sc_cid=7013a000002pgRIAAI&gclid=Cj0KCQiAyMKbBhD1ARIsANs7rEE8YQxw9jJzLS7nmccCeEjkPe_YoGwAHqmMVCYdZVrSBSe1PdXSjscaAoRZEALw_wcB&gclidsrc=a.w.ds. [Accessed 13 November 2022].
- [5] J. G. M. P. R. a. G. F. . Lynn, Measuring the business value of cloud computing, Cham: Springer International Publishing, 2020.
- [6] Microsoft, “<https://azure.microsoft.com/mediahandler/files/resourcefiles/resilience-in-azure-whitepaper/Resilience%20in%20Azure.pdf>,” Aug 2019. [Online]. Available: <https://azure.microsoft.com/mediahandler/files/resourcefiles/resilience-in-azure-whitepaper/Resilience%20in%20Azure.pdf>. [Accessed Tue Nov 2022].
- [7] T. W. a. E. Benkhelifa, “On Resilience in Cloud Computing: A Survey of Techniques across the Cloud Domain,” *ACM computing surveys*, vol. 53, no. 3, pp. 1-36, 2020.
- [8] J. Reno, “Multifactor Authentication: Its Time Has Come,” *Technology Innovation Management Review*, pp. 51-58, 2013.
- [9] M. Learn, “Enhance security with the principle of least privilege,” Microsoft, 11 01 2022. [Online]. Available: <https://learn.microsoft.com/en-us/azure/active-directory/develop/secure-least-privileged-access>. [Accessed 17 11 2022].
- [10] AWS, “How to Record SSH Sessions Established Through a Bastion Host,” AWS, 14 Jun 2016. [Online]. Available: <https://aws.amazon.com/blogs/security/how-to-record-ssh-sessions-established-through-a-bastion-host/>. [Accessed 17 Nov 2022].

Essay1

GRADEMARK REPORT

FINAL GRADE

90/100

GENERAL COMMENTS

Instructor

Virtualization/VM: 21

Cloud Services: 23

Availability: 23

Security: 23

Feedback:

+++++++

I would have liked a little more on VM components but this is good.

Could have added a little more definition of the services but again v. good.

Security discussion is very good.

PAGE 1

PAGE 2

PAGE 3

PAGE 4

PAGE 5

PAGE 6

PAGE 7

PAGE 8

PAGE 9

PAGE 10

