
Academic Report on Virtualization, Virtual Machines, Authentication and Authorization and their use in Small and Medium Enterprise

Cloud Computing Architecture /ACCA5025/

Lecturer: Koushik Modak

Valentin Latinov Miladinov /2024732/

University of Wales Trinity Saint David

Submission Deadline: 06/04/2023

Table of content

Introduction:	3
Virtualization:.....	3
Virtual machines:	5
Cloud Services:	6
Evaluation of Availability:	9
Evaluation of security:	10
Conclusion:.....	11
References:	12

Introduction:

This report aims to describe virtualization and virtual machines, to analyse Cloud Computing Services and their use for a small or medium-size enterprise; to evaluate two budget scenarios and to analyse cybersecurity concerns, where authentication and authorization are defined and Multifactor authentication, applying the least-privilege permissions model, network security groups and Bastions with their separate purpose are explained.

Virtualization:

Virtualization, being a key in Cloud Computing, defines in the process of creation of a virtual version of either an operating system, network resources, desktop, or server. This separates the physical resource and the physical delivery of the resource or application. The great idea behind this is to let the user run a set of desktop operating systems and applications at the same time on the same hardware and software. This process by no means slows down or impacts the system's efficiency whatsoever. It also makes the virtual emulation of products or services in the same machine certain to happen. [1]

While providing access to resources and storage is within Cloud Computing's power, virtualization allows virtual version of software or hardware resources, network, or storage devices to be created. Yes, Cloud vendors could rent an organisation certain amount of storage, physical hardware and network resources and software applications. No, they cannot guarantee no overloading of the storage capacity, if every customer was to ask for more resources. However, virtualization allows cloud vendors to have their infrastructures and servers virtually copied, then having those copies of resources assigned to the customer, would it be a single user or an organisation. Virtualization's role in Cloud Computing is characterized in the creation of Virtual Machines and their functionality with multiple operating systems. It is useful at making virtual ecosystem for server operating systems and multiple storage devices. [2] The main advantage of this technology is that the main processes, ran in the virtual machine, behave almost like the real deal (using computers). [3] Other advantages virtualization provides are in tight collaboration with the cloud computing, which already delivers efficiency, economic value, and scalability; together they deliver a more cost-efficient way of running multiple operating systems, while using one dedicated resource. [2]


There are certain characteristics to virtualization:

- Distribution of resources;
- Accessibility of server resources;
- Resource isolation;
- Security and authenticity;
- Aggregation.

Furthermore, there are different kinds to virtualization:

- Application;
- Network;
- Desktop;
- Storage;
- Server;
- Data.[1]

Presented is an image of the infrastructure of virtualization:



Virtual machines:

The term “virtual machine” made its way through the report few times so far, but what does it stand for? Virtual machines are software which run applications or programs with no need to be tied to a physical machine. Virtual machines are welcoming to more than one guest machines to run on a host computer. No virtual machine is dependent on the other one, as it has its own operating system and functions on its own. Virtual Machines are usually run on servers, but it is no surprise to be run on desktop systems or perhaps on embedded platforms too. [4]

For Virtual Machines to work properly with the hardware of the host machine, Hypervisors exist. Hypervisors are the isolating layer which manages both the physical hardware and the virtual machines. This makes them the perfect target for hacking, since it accesses all the virtual machines’ data in CPU registers, Memory, and I/O. Should a Hypervisor be attacked successfully, no virtual machine on the given physical machine is safe as long as the attack is due. Hypervisors go by two types:

1. Type 1 Hypervisor (Native/Bare Metal Hypervisor):

- Installed and run straight on the physical layer (hardware);
- No host operating system is required;
- Has direct access and control over hardware resources;
- Proxmox, VMware ESXI, Citrix Xenserver, Ovir, Hyper-V.

2. Type 2 Hypervisor (Hosted Hypervisors):


- Installed on host operating system of the physical machine;
- Coordinates calls for CPU, disk, memory, network and other recourses through the host’s operating system;
- VMware, Workstation, Oracle VM VirtualBox. [5]

Cloud Services:

When it comes to using Cloud Services, there are three options to choose from:

- Infrastructure as a Service (IaaS):
 - Customers can purchase resources such as networking and storage on-demand;
 - Dismisses the need for a costly hardware;
 - Highly scalable;
 - Offers flexibility;
- Platform as a Service (PaaS):
 - Customers receive developers with a framework and tools that would build apps and software, fitting to the customers' needs;
 - Customers have access to data centres and servers;
 - Data centres and servers are managed by a third-party provider;
 - Main use is the creation of a custom SaaS application;
- Software as a Service (SaaS):
 - Most common form of Cloud Computing;
 - Ready to use;
 - Commonly run straight through the customer's web browser;
 - Hosted on remote servers;
 - Fully managed updated and maintained by a third-party vendor. [6]

Presented is a photo of what responsibility each service takes from the customer (blue marks responsibility of the user):



Presented is a table on pros and cons of all the services for a brief explanation[7]:

IaaS		PaaS		SaaS	
Pros	Cons	Pros	Cons	Pros	Cons
Most flexible	Data security	Highly scalable	Data security	No hardware costs	No control
Most dynamic	Inaccessible data when vendor outages	Highly available	Compatibility of existing infrastructure	No initial setup costs	Limited solutions
Cost-effective (pay-as-you-go pricing)	Need for team training	Multi-tenant	Dependency on the vendor	Automated upgrades	Connectivity is required
Easy to use (automated deployment of hardware)		Quickened and simplified development process		Cross-device compatibility	
Less responsibilities (management tasks virtualized)		Creating, testing, and launching apps at a lower cost		Accessible from anywhere	
		Automated company policy		Pay-as-you-go model	
		Less coding required		Scalability	
		Easy migrating to the hybrid cloud		Easy customization	

Infrastructure as a Service shines brightest when there is a small company or a startup that is with no budget for creating own infrastructure, if there is a rapidly growing company with demands unstable and changeable, when there is a large company that aims for effective control over infrastructure but would pay only for the resources used.

Platform as a service does best when developers would prefer to engage more in coding, testing, and deploying own applications, therefore PaaS is most beneficial when multiple developers are onto the same project, other vendors are to be included, there is the need to create own customized app.

Software as a service bests at satisfying needs to have a ready-made software launched quickly, collaborated short-term projects, constant use of applications, especially those that are with both web or mobile access. [7]


Back in the days, retailers used to mainly self-host their software and maintain it with their personal IT team. But since in more recent days there are over two billion people purchasing goods or services only yearly, cloud-based web hosting became the better choice, as it allowed e-commerce retailers build applications that scale efficiently, are secure and cost-effective. Cloud-based hosting solution rewards with 100% uptime, regardless of the numbers of buyers. Presented are top features expected from the top e-commerce web hosts:

- Complexity – possible building cloud-based applications always uptime and customizable;
- Pay as you use;
- Scalability;
- Security;
- Disaster recovery. [8]

Given this information, it's safe to say that the most logical service for a small to medium enterprise would be Software as a service – it's simple, producing results, secure, with a good payment model, not too costly, accessible, automated. That is if the company isn't willing to invest in more in count and more experienced IT specialists or furthermore isn't willing to present a specific training to the staff overall. Noted should be the following diagram of Cloud Computing Market Size in the period of 2018-2024, provided by T4. SaaS' growth is still the largest of the group by a lot, regardless being cheaper than the other two services: [9][10]

Cloud Computing Market Size, 2018-2024

www.T4.ai



Evaluation of Availability:

Location in Azure (or region) is a set of data centres that were deployed within a latency-defined perimeter and connected through Azure's dedicated regional low-latency network. [11]

"Azure availability zones are connected by a high-performance network with a round-trip latency of less than 2ms. "is what the official site of Microsoft tells. There can be many different reasons for things to go wrong – software or hardware failures, events like earthquakes, flood, fires – which means countermeasures should be taken. To be resilient, there are at least three separate availability zones in any availability zone-enabled region. These zones help the data remain accessible and synchronized when things go downhill. Availability zones are designed in a way that if one of them is temporarily down, the other two would compensate and support the regional services, capacity, and high availability. [12]

When it comes to resource groups, they are containers that hold related resources for an Azure solution. They could contain all the resources needed for the solution or just those, asked for by the given user. Allocating the resources to resource groups is all in the freedom of the user to do. Metadata about the resources is all stored in the resource groups. This means by choosing a location for them the metadata they carry would be stored there. Due compliance, particular regions would need to be the choice for storage of these resource groups. [13]

Truth be told, so far there isn't anything like 100% availability. But there is the rule of the Nines. As availability solutions generally deliver 99%, it seems it still is a good deal to take, however this means 87.6 hours of unplanned downtime each year, which is way over three days. In accordance with a research Aberdeen group provides, an hour of downtime and the data loss it causes costs over £130 000. The rule of the Nines mentioned before, it helps to separate the good business practice from the bad one. The more nines there are after the decimal point of 99%, the more profitable a company is. Availability of 99.9% allows up to 8.76 hours of unplanned downtime every year. It feels like a cigarette break next to over 87 hours. It is considered as a good business practice. To achieve it, x86 is to be used. [14]

Availability needs to go even higher, that's why further solutions are needed; solutions better and more enduring. Cluster technology brings this idea to life. Essentially, clusters are simply two or more physical servers that are connected in one network. Should one server fail, application support is resumed on the other server. Depends on the cluster build-up and the speed of the failovers, the range of availability can range from 99.95% to 99.99%. [14]

To reach the fifth Nine, fault-tolerant systems are the tool needed. It works through faults without disrupting the applications in the least, therefore preventing downtime due system failure. They deliver 99.999% availability or even better, which results at less than five minutes unplanned downtime per year. Virtualization brings the fifth Nine, as software fault tolerance uses industry-standard servers that run in parallel, thus enabling one application to live on two virtual machines at the same time. Should one Virtual Machine go down, the other would keep on running the application with no interruptions or any data loss whatsoever. Something to be aware of is that some emulate fault tolerance, however they create lots of overhead, which weighs on the performance. [14]

Based on the prices of the tools from the previous examples, enterprise should pick whichever its budget suffices.

Evaluation of security:

“Authentication is the act of validating that users are whom they claim to be. This is the first step in any security process.”, according to okta.com. By the same source: “Authorization in system security is the process of giving the user permission to access a specific resource or function. “. Authentication makes users prove to be who they say they are. Some examples of authentications would be usernames and passwords, authentication apps, one-time pins, biometrics. Examples of authorization would be giving permission to someone to download a given file on a server or providing them with administrative access to an application. [15]

Despite of the similar sounding of they have with one another, they are two separate steps in the logging process. Authentication is like a key, if it fits, it will unlock, and the person will be able to get in the room. Once the system receives the right credentials, the user is allowed to be in. Authorization is a permission; it is like what the person is allowed to do in the room – use the kettle, use the dishwasher, sleep on the couch – in the system it might mean that the user could be allowed to go through specific documentation, to manipulate it, to create new documentation and so on. [15]

The table presented shows what capabilities both processes have:

	Authentication	Authorization
Purpose	Verifying credentials	Granting or denying permissions
Ways to operate	Through passwords, biometrics, apps, or one-time pins	Thorough settings
Visibility to the user	Full	None
User control	Partial	None
Data movement	Through ID tokens	Through access tokens

To put it simple, authentication allows users to join in, whereas authorization allows users to take specific actions. [15]


Multifactor authentication is a colossal must, due to how beneficial it is to have the users identify themselves with more than just their username and password. These credentials are comparably easier to attack and obtain by third parties. It would be rather hard to go around MFA factors like a physical hardware key or a thumbprint, which increases the confidence of the organisation. [16]

The principle of least privilege gives access control to a user as much access control that user would need to do their job fully, nothing over it. This reduces the risk of virtual or personal data compromise drastically. [17]

A network security group is based on security rules that would either allow or deny inbound network traffic to or outbound network traffic from some types of Azure resources. Any rule has a source and destination, port, and protocol to be specified by the user. This filters network traffic from and to Azure resources. [18]

Bastion (Azure Bastion) allows customers to connect and access Virtual machines more securely. It connects local computers, Virtual machines, and loud resources with no exposure to public network connections. [19]

Something to be considered on the security point: the hypervisor security. There is a privilege level to any processor, divided in levels – from zero to three. Level 0 is the Operating System Kernel, most privileged of them all. Level 1 and 2 are Operating System Services, which are less privileged. Then, there is level 3 where all the applications programmes run on. The Hypervisor works on Level 0, whereas the Guest Operating System works on Level 1 and 2, user applications are on Level 3, having no privilege over the hardware overall. The Hypervisor is responsible for the hardware resources which need to be accessed by the guest operating system. This model more or less reminds of the principle of least privilege and it makes sense, since Hypervisors are the most attacked and valuable target, having so many permissions. Presented visually [5]:



Conclusion:

As this project reaches its finalization, concluded shall be this outcome: cloud computing has progressed by leaps and bounds in past years, allowing virtualization of what else would be a physical product, thus evading the need to waste physical resources, polluting the world even further in the making of the physical product, company's costs and so many more. This is profitable both for the world and the business. Sure, there are new possible breaches in security, but it looks like a gamble worth taking. Investing in Cloud becomes the better option by the minute, knowing how to do it becomes more and more important. The world spins way too fast for enterprises not to aim for the highest availability possible, while maintaining a good security culture.

References:

- [1] R. Peterson, "Virtualization in cloud computing: Types, architecture, advantages," *Guru99*, 25-Feb-2023. [Online]. Available: <https://www.guru99.com/virtualization-cloud-computing.html>. [Accessed: 01-Apr-2023].
- [2] S. U. R. B. H. I. RASTOGI, "Virtualization," in *Cloud computing simplified: Explore application of cloud, cloud deployment models, service models... and mobile cloud computing*, S.I.: BPB PUBLICATIONS, 2021, pp. 112–113.
- [3] P. K. Das and G. C. Deka, "Applications of Virtualization Technology in Grid Systems and Cloud Servers," in *Design and use of virtualization technology in cloud computing*, Hershey, Pennsylvania: IGI Global, 2018, pp. 2–2.
- [4] K. Shaw, "What is a virtual machine, and why are they so useful?," *Network World*, 17-Jul-2022. [Online]. Available: <https://www.networkworld.com/article/3583508/what-is-a-virtual-machine-and-why-are-they-so-useful.html>. [Accessed: 01-Apr-2023].
- [5] Z. Aalam, V. Kumar, and S. Gour, "A review paper on hypervisor and Virtual Machine Security," *Journal of Physics: Conference Series*, vol. 1950, no. 1, pp. 1–2, 2021.
- [6] "IaaS vs. PaaS vs. SaaS - differences, examples and diagram," *LeanIX*. [Online]. Available: https://www.leanix.net/en/wiki/saas/iaas-vs-paas-vs-saas?utm_term=&utm_source=adwords&utm_medium=ppc&utm_campaign=EMEA-UKI_UK%2B%7C%2BSMP%2B%7C%2BSaaS%2BManagement%2B%7C%2BDSA%2B%7C%2BENG&hsa_ver=3&hsa_cam=15150382310&hsa_grp=133953171750&hsa_acc=9751618594&hsa_kw=&hsa_mt=&hsa_net=adwords&hsa_src=g&hsa_tgt=dsa-1465784135448&hsa_ad=563932871663&gclid=EAIaIQobChMIxOWf--mS_gIVh9LtCh1fqg3YEAAYAiAAEgJfnPD_BwE#differences-between-iaas-paas-saas. [Accessed: 01-Apr-2023].
- [7] N. Sakovich, "IaaS vs. paas vs. SAAS. advantages and disadvantages," *Sam Solutions*, 03-Feb-2023. [Online]. Available: <https://www.sam-solutions.com/blog/iaas-vs-paas-vs-saas-whats-the-difference/>. [Accessed: 01-Apr-2023].
- [8] D. Lim, "The top cloud-based e-commerce hosting solutions in 2023," *fabric blog*, 08-Mar-2022. [Online]. Available: <https://fabric.inc/blog/ecommerce-hosting/>. [Accessed: 01-Apr-2023].
- [9] "Cloud computing market share," *T4*, 15-Feb-2021. [Online]. Available: <https://www.t4.ai/industry/cloud-computing-market-share>. [Accessed: 01-Apr-2023].
- [10] "Cloud computing: The difference between IAAS, paas and SAAS," *CMI*. [Online]. Available: <https://www.newcmi.com/blog/uk-london-it-provider-cloud-computing-the-difference-between-iaas-paas-and-saas#:~:text=SaaS%20vs.,the%20cost%20is%20relatively%20lower>. [Accessed: 01-Apr-2023].

- [11] A. Hawkes, "A Guide To Microsoft Azure Regions," *Console Connect*, 14-Mar-2022. [Online]. Available: <https://blog.consoleconnect.com/a-guide-to-microsoft-azure-regions#:~:text=Overview%20of%20Azure%20locations&text=An%20Azure%20region%20is%20a,data%20residency%20and%20compliance%20boundaries>. [Accessed: 01-Apr-2023].
- [12] A. Harris, "What are Azure Regions and availability zones?," *Microsoft*, 15-Mar-2023. [Online]. Available: <https://learn.microsoft.com/en-us/azure/reliability/availability-zones-overview>. [Accessed: 01-Apr-2023].
- [13] M. Sabhi, "Manage resource groups - azure portal - azure resource manager," *Microsoft*, 28-Feb-2023. [Online]. Available: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/manage-resource-groups-portal>. [Accessed: 01-Apr-2023].
- [14] D. Todd, "The rule of nines: Why you need 99.999% availability in an always-on world," *Stratus Blog*, 18-May-2022. [Online]. Available: <https://blog.stratus.com/rule-nines-availability-always-on-world/#:~:text=What%20does%2099.99%20availability%20correspond,resumes%20on%20a%20second%20server>. [Accessed: 01-Apr-2023].
- [15] Updated: 02/14/2023 - 10:40 Time to read: 4 minutes, "Authentication vs. authorization," *Okta*, 14-Feb-2023. [Online]. Available: <https://www.okta.com/identity-101/authentication-vs-authorization/>. [Accessed: 01-Apr-2023].
- [16] "What is multi-factor authentication (MFA)," *What is Multi-Factor Authentication (MFA) and How Does it Work?* [Online]. Available: <https://www.onelogin.com/learn/what-is-mfa#:~:text=Why%20is%20MFA%20Important%3F,be%20stolen%20by%20third%20parties>. [Accessed: 06-Apr-2023].
- [17] D. Walkowski and R. Pompon, "What is the principle of least privilege and why is it important," *F5 Labs*, 21-Dec-2020. [Online]. Available: <https://www.f5.com/labs/learning-center/what-is-the-principle-of-least-privilege-and-why-is-it-important#:~:text=Information%20security%20is%20a%20complex,of%20any%20information%20security%20program>. [Accessed: 01-Apr-2023].
- [18] Asudbring, "Network Security Group - How It Works," *Microsoft*, 10-Feb-2023. [Online]. Available: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-group-how-it-works>. [Accessed: 01-Apr-2023].
- [19] J. Kozłowicz, "What's a Jumpbox or bastion host, anyway?," *Lunavi*, 10-Jan-2020. [Online]. Available: <https://www.lunavi.com/blog/whats-a-jumpbox-or-bastion-host-anyway>. [Accessed: 01-Apr-2023].