

University of Business and Technology in Kosovo

UBT Knowledge Center

Theses and Dissertations

Student Work

Spring 5-2021

CLOUD COMPUTING AND SECURITY OF DATA

Elbekir Krasniqi

Follow this and additional works at: <https://knowledgecenter.ubt-uni.net/etd>



Part of the [Computer Sciences Commons](#)



Programi për Shkenca Kompjuterike dhe Inxhinierise

CLOUD COMPUTING AND SECURITY OF DATA
Shkalla Bachelor

Elbekir Krasniqi

Maj / 2021
Prishtinë



Programi për Shkenca Kompjuterike dhe Inxhinierisë

Punim Diplome
Viti akademik 2012 – 2013

Elbekir Krasniqi

CLOUD COMPUTING AND SECURITY OF DATA

Mentor: Dr. Sc. Naim Preniqi

Maj / 2021

Ky punim është përpiluar dhe dorëzuar në përmbushjen e kërkesave të
pjeshme për Shkallën Bachelor

ABSTRACT

Cloud computing presents a new model for IT services and delivery and it usually involves over-a-network, on-demand, self-service access, which is dynamically scalable and elastic, utilizing pools of often virtualized resources. Through these features, cloud computing has the potential to improve the way businesses and IT operate by offering fast start-up, flexibility, scalability and cost efficiency. Even though cloud computing provides compelling benefits and cost-effective options for IT hosting and expansion, new risks and opportunities for security exploits are introduced. Security standards, policies and controls are therefore of the essence to assist management in protecting and safeguarding systems and data. Cloud computing risks should be analyzed and understood to be able to protect environments and avoid data being exposed. The focus of this paper is analyzing current and future trends on cloud computing and mitigation for cloud computing security risks as a fundamental step towards ensuring secure cloud computing environments. It is of essence to understand the capabilities and risks of cloud computing before a decision is made to move to a cloud computing provider.

ACKNOWLEDGMENT

I would like to express my sincere gratitude to all those (mentioned, or not) who helped me complete this part of my journey.

A special thanks goes to Professor Dr. Sc. Naim Preniqi who supervised this project. He provided me with countless hours of advice, guidance, and support. I thank him for his commitment during all the easy and difficult moments I had until the completion of this paper. The same thanks go to my colleagues from the University of Business and Technology, for their contribution.

I dedicate this paper to my parents Sejfi and Ditrije Krasniqi, for their endless support.

TABLE OF CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	v
1 INTRODUCTION	1
2 LITERATURE REVIEW	2
2.1 Public Cloud Benefits	3
2.1.1 Reduction of costs	3
2.1.2 Scalability and agility	6
2.1.3 Deployment time	8
2.2 Security Challenges	9
2.3 Architecture and Scalability.....	11
2.3.1 Three tier architecture.....	12
2.3.2 Serverless architecture.....	13
2.3.3 Containerization Technology	15
2.4 Technologies and Networks.....	15
2.4.1 Elastic Compute Cloud – EC2.....	16
2.4.2 Simple Storage Service – S3	16
2.4.3 Identity Access Management – IAM.....	17
2.4.4 Virtual Private Cloud – VPC	17
3 PROBLEM STATEMENT	19
3.1 Public Cloud benefits.....	19
3.1.1 Reduction of costs	20
3.1.2 Scalability and agility	20
3.1.3 Deployment time	20
3.2 Relevant security mechanisms.....	21
3.2.1 EC2 and Security	21
3.2.2 S3 and security mechanisms.....	22
3.2.3 IAM Policies and Roles	23

3.2.4	VPC Configurations	24
3.3	Security challenges	24
3.3.1	Distributed Denial of Service (DDoS).....	25
3.3.2	Access Control.....	25
3.3.3	Multitenancy Security and Privacy.....	26
4	METHODOLOGY	27
5	CASE STUDY	29
5.1	Netflix on AWS in numbers	29
5.2	Netflix Benefits	30
6	CONCLUSION	31
7	REFERENCES	33
8	BIBLIOGRAPHY.....	37
9	APPENDIXES.....	38

LIST OF FIGURES

Figure 1: Cloud computing benefits	6
Figure 2: The standard enterprise cloud	7
Figure 3: AWS Shared Responsibility Model	10
Figure 4: Three tier architecture High Level Overview	13

LIST OF TABLES

Table 1. Time (min:sec) needed to manage the cloud instances.....	8
---	---

1 INTRODUCTION

Since the beginning of the computer era in the early 60s, the industry has tried to make computers as wide spread as possible in terms of exposure to everyday people. Although it is difficult to pinpoint the beginning for Cloud Computing due to its rapid evolution and changes in what it was supposed to offer, we can probably start since the early days when computers started integrating in high profile jobs when the requirement came from DARPA (Defense Advanced Research Agency) to MIT (Massachusetts Institute of Technology) to develop a technology to allow a computer to be used by two people simultaneously. This primitive technology has marked the first step to what we collectively known today as Cloud Computing.

There are many milestones that had to be reached to get to where we are today in Cloud Computing, starting from the Personal Computer era which started in the mid-70s with Altair 8800 being sold as a construction set and known as a Home Computer to the Internet we know today, which can be traced back to a project in ARPA (Advanced Research Projects Agency) also known as ARPAnet.

The term Cloud Computing was first coined in 2007, typically referring to a join hardware and software deployment concept [1].

As always with the rise of new technologies in the IT sector, there is always an aspect that we have to take into consideration, and that being security of data and it's challenges to moving to a Cloud Computing Provider in terms of data privacy.

There are two major concerns when it comes to Cloud Computing, data security and data privacy due to online exposure and different market regulations that are put in place like GDPR. In the Cloud Computing world, the user is able to access a shared virtual environment where customer data is stored. Consequently, security concerns and challenges arise.

This paper will focus on these concerns and challenges and try to address different mitigation methods from industry best practices on key architectural tiers.

2 LITERATURE REVIEW

The decision to use a Cloud Computing Provider whether you are a small company or big company is one of the key decisions that need to be made. For both cases this can affect day to day operations and yearly investments depending on the size of the infrastructure required to keep your services running. One of the biggest advantages of using Cloud is the cost reduction and effort that need to be put in, as there is no need to spend huge amounts of money on purchasing and maintaining equipment. This drastically reduces capex costs as there is no investment in hardware, facilities or building servers which in the end will likely be underutilized. One other benefit is that it reduced downtime significantly weather maintenance downtime or unintentional. In cases where resiliency and uptime is of utmost importance, services which are running on on-premise datacenters, find it very difficult to navigate in such requirements, mainly due to the amount of cost and effort that is needed to keep a Highly Available and Geo Redundant service. This, in most cases includes a secondary site that stays in standby mode to take over in case there is failure or any sort of downtime, thus doubling the original costs for the primary environment. In these scenarios, a Cloud based solution, offers varying degree of scalability and reliability depending on the architecture and cost versus importance metrics as it is easy and considerably faster to put in such architecture in place.

2.1 Public Cloud Benefits

It is no surprise that cloud computing has attracted so much attention and popularity in the last couple of years, mainly due to the allure and promise of flexibility for enterprises, from saving time and money to improving scalability and agility.

Cloud computing differs in one major way from the traditional on-premise deployment. A company hosts everything in-house in an on-premise infrastructure, while in a cloud environment, a third-party vendor like AWS (Amazon Web Services), hosts all that for you. This allows companies to pay on as-needed basis, basically paying only for the time you use that resource. This also allows enterprises to easily scale up or down depending on usage, requirements and growth of the company.

In the following chapters, we will focus on how specifically enterprises can benefit from a cloud platform in comparison to the traditional on-premise solutions [3][4]. The main focus will be on the following:

- Reduction of costs
- Scalability and agility
- Deployment time

2.1.1 Reduction of costs

One of the key indicators that is in favor of Cloud Computing being used by enterprises is the amount of investment that is spent. The overall cost benefit comes down to mainly hardware responsibilities. In traditional on-premise infrastructure, any failure that happens from network layer up to computing layer is done by the enterprise itself, while in Cloud Computing, this is the responsibility of the Cloud Provider. The advantage here is the reduction of hardware costs and how quickly and easily resources can be acquired.

Along with purchase costs, on-premise costs include also power consumption which translates to high operational costs [2][5]. Large datacenters produce a large amount of heat which adds on top of the power consumption because a cooling system needs to be in place

to keep the infrastructure cool and avoid overheating to be able get the maximum performance out of the infrastructure.

A study in 2013 that was conducted by Rackspace explains this best, where 88% of cloud users in United Kingdom and United States point to cost saving. In addition, 56 percent of respondents agree that cloud computing has helped them boost profits. 68 percent also say the use of open source cloud is on the increase [6].

Cloud Computing Providers offer different cost models for their services, but all of them are based on the Pay As You Go (PAYG) model, meaning that enterprises will pay only for the usage of a specific resource type. This is one of the key advantages in cloud computing as this allows enterprises to scale up or down depending on the resources that need to be available for a service to run smoothly without any disruption of the service [7]. This will be covered in more details in the following chapter.

An important factor that has a major impact on costs in the cloud is how well architected and balanced the service itself is. AWS offers guidance with their published whitepapers that specifically target key pillars on how to build secure, high-performing, resilient and efficient infrastructure for applications and their workloads.

The following pillars apply for AWS Well-Architected Framework:

- Operational Excellence Pillar
- Security Pillar
- Reliability Pillar
- Performance Efficiency Pillar
- Cost Optimization Pillar

The operational excellence pillar focuses on running and monitoring systems to deliver business value, and continually improving processes and procedures. Key topics include automating changes, responding to events, and defining standards to manage daily operations.

The security pillar focuses on protecting information and systems. Key topics include confidentiality and integrity of data, identifying and managing who can do what with privilege management, protecting systems, and establishing controls to detect security events.

The reliability pillar focuses on ensuring a workload performs its intended function correctly and consistently when it's expected to. A resilient workload quickly recovers from failures to meet business and customer demand. Key topics include distributed system design, recovery planning, and how to handle change.

The performance efficiency pillar focuses on using IT and computing resources efficiently. Key topics include selecting the right resource types and sizes based on workload requirements, monitoring performance, and making informed decisions to maintain efficiency as business needs evolve.

The cost optimization pillar focuses on avoiding unnecessary costs. Key topics include understanding and controlling where money is being spent, selecting the most appropriate and right number of resource types, analyzing spend over time, and scaling to meet business needs without overspending [7].

Depending on the use case and what services are used in AWS, there is a further breakdown of costs that can be managed on the resource type you choose.

Below on Fig. 1 are listed the benefits of cloud computing arranged from highest occurrence (cited most in literature) to the lowest [8].

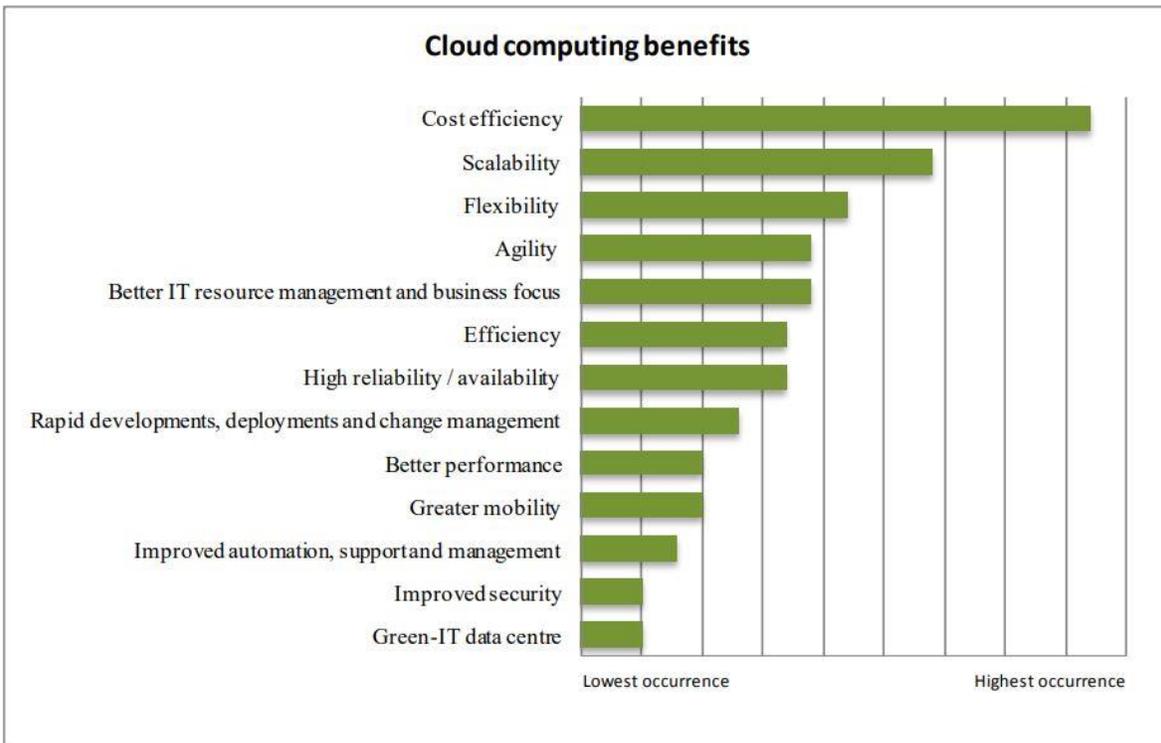


Figure 1: Cloud computing benefits

2.1.2 Scalability and agility

Scalability is an important factor in cloud computing, in Fig 1, it is the second most occurring benefits that has been mentioned in literature. There is a constant need from enterprises to be flexible in their infrastructure to optimize costs and not overspend, but at the same time, as it is usual in different IT services there are peak hours that a service will be used by end users which translates into how scalable the infrastructure is. This depends on how the infrastructure is architected and designed to work but also the amount of flexibility one enterprise has.

A well architected service takes always into account on how scalable the service is to be able to cope with peak usage hours but also at the same time to optimize cost when there is less usage.

The Cloud model is composed of five essential characteristics, three service models, and four deployment models [9] as illustrated in Fig. 2.

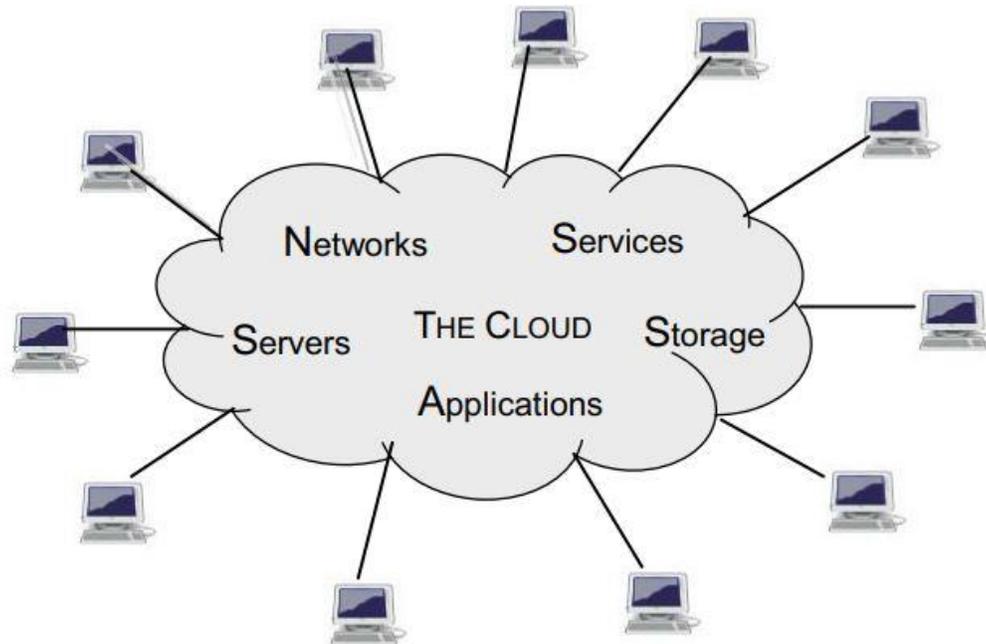


Figure 2: The standard enterprise cloud

In all the layers illustrated in Fig 2. there needs to be some level of scalability, Servers and Storage being the most important as they are the building blocks where scalability is often focused.

Cloud solutions should be ideally build with the concept “build today with tomorrow in mind,” meaning solutions need to cater to current scale requirements as well as the anticipated growth of the solution. This growth can be either the organic growth of a solution or it could be related to a merger and acquisition type of scenario, where its size is increased dramatically within a short period of time.

Still, when a solution scales, many enterprises experience added complexity to the overall architecture in terms of its manageability, performance, security, etc. By architecting solutions or application to scale reliably, you can avoid the introduction of additional complexity, degraded performance, or reduced security as a result of scaling [10].

2.1.3 Deployment time

With the constant need of faster deployments time, cloud computing providers like AWS offer highly scalable and dynamic solutions to cope with different business requirements. This is achieved with the power of virtualization. Resources like EC2 instances on AWS can be deployed in a matter of minutes. On traditional physical infrastructure this could lead to time delays of weeks, with processes like infrastructure procurement, installation, licensing etc.

In Table 1 [11] below, we illustrate the deployment time of compute instances for different cloud providers.

Provider	Boot 1 instance	Boot 4 instances	Delete 1 instance	Delete 4 instances
EC2	2:34	2:34	0:33	0:33
Azure	3:06	12:43	3:07	12:26
Rackspace	12:37	12:57	0:23	0:23

Table 1. Time (min:sec) needed to manage the cloud instances.

Besides the short deployment time that cloud computing providers offer, there also different built in services that are offered in cloud providers to shorten even further the deployment time of services that require bigger infrastructure like AWS Elastic Beanstalk. The added bonus in these cases is that these type of built in services are free, excluding the costs of the resources that get created by them [12].

Amazon Web Services (AWS) comprises over one hundred services, each of which exposes an area of functionality. While the variety of services offers flexibility for how you want to manage your AWS infrastructure, it can be challenging to figure out which services to use and how to provision them.

With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without having to learn about the infrastructure that runs those applications. Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload

your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring [12].

2.2 Security Challenges

When multiple organizations share resources there is a risk of data misuse. So, to avoid risk it is necessary to secure data repositories and also the data that involves storage, transit or process. Protection of data is the most important challenges in cloud computing. To enhance the security in cloud computing, it is important to provide authentication, authorization and access control for data stored in cloud. The three main areas in data security [13] are:

- Confidentiality: Top vulnerabilities are to be checked to ensure that data is protected from any attacks. So security test has to be done to protect data from malicious user such as Cross-site Scripting, Access Control mechanisms etc..,
- Integrity: To provide security to the client data, thin clients are used where only few resources are available. Users should not store their personal data such as passwords so that integrity can be assured.
- Availability: - Availability is the most important issue in several organizations facing downtime as a major issue. It depends on the agreement between vendor and the client.

Due to the complication that virtualization introduces to the security aspect of having solutions reside in the cloud, there needs to be a clear line of who takes care of which layer. On a high level overview AWS does this separation with a clear distinction of these layers. Security and Compliance is a shared responsibility between AWS and the customer. This shared model can help relieve the customer's operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall. Customers should carefully consider the services they

choose as their responsibilities vary depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations. The nature of this shared responsibility also provides the flexibility and customer control that permits the deployment. As shown in the chart below, this differentiation of responsibility is commonly referred to as Security “of” the Cloud versus Security “in” the Cloud [14].

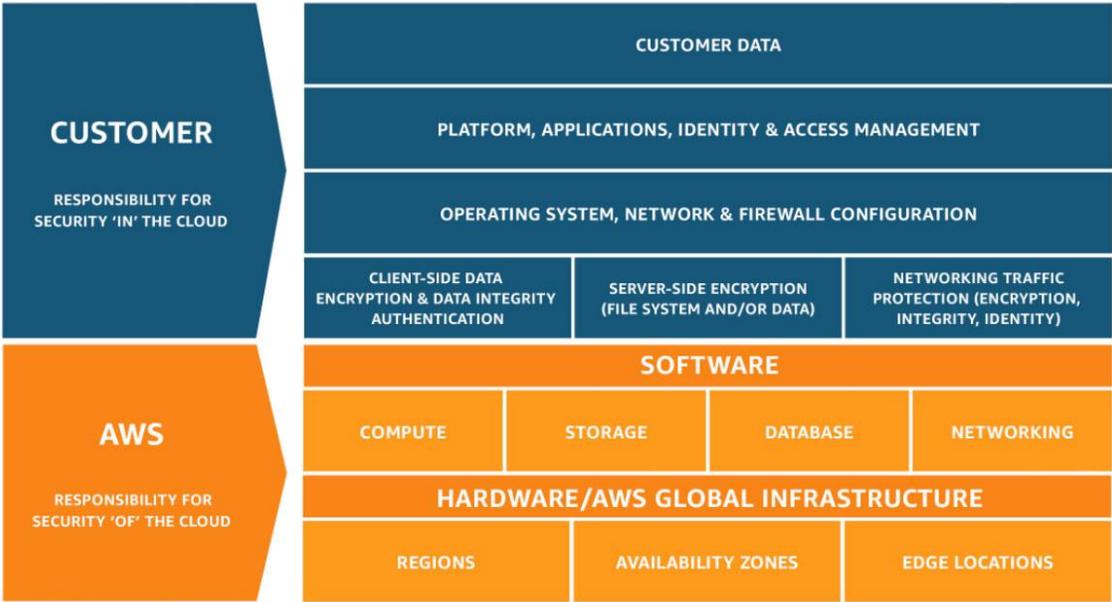


Figure 3: AWS Shared Responsibility Model

AWS responsibility “Security of the Cloud” - AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. Customer responsibility “Security in the Cloud” – Customer responsibility will be determined by the AWS Cloud services that a customer selects. This determines the amount of configuration work the customer must perform as part of their security responsibilities. For example, a service such as Amazon Elastic Compute Cloud (Amazon EC2) is categorized as Infrastructure as a Service (IaaS) and, as such, requires the customer to perform all of the

necessary security configuration and management tasks. Customers that deploy an Amazon EC2 instance are responsible for management of the guest operating system (including updates and security patches), any application software or utilities installed by the customer on the instances, and the configuration of the AWS-provided firewall (called a security group) on each instance. For abstracted services, such as Amazon S3 and Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and customers access the endpoints to store and retrieve data. Customers are responsible for managing their data (including encryption options), classifying their assets, and using IAM tools to apply the appropriate permissions.

2.3 Architecture and Scalability

Building a solid cloud computing environment is not an easy task, it needs a good understanding of what components are involved, what the high level and detailed requirements are and how everything fits into the high level architecture, starting from user requirements and their type of access (end users or daily users which make use of the service or admin level access including also operational users) up to what type of resources are needed.

A well architected solution should consider with heavy focus the consumers of the service. It is obvious that they play an important role in business while their demand and expectation are valuable in IT industries. Therefore, by understanding the cloud consumer expectation of cloud computing, it will support continuously the development of cloud services [15]. On the following chapters we will focus on the mainstream and recent architectures which are commonly used by industry giants and also suggest by industry standards and best practices.

2.3.1 Three tier architecture

The three-tier is one of the most common architecture models that is widely used in the IT industry. The three-tier client/server architecture is an evolution of the traditional two-tier model, and is receiving increased interest, particularly for large business applications. The main difference is that in a three-tier architecture, most of the functionality is separated out in a middle layer, called application servers [16]. It offers a robust architecture in terms of failure isolation making it easier to troubleshoot, scalability at the highest level from a component point of view while also separating the layers of security and those are:

- Frontend Tier
- Backend Tier
- Secure Backend Tier

The FrontEnd layer is the upper-most security layer that is considered to be the less trusted layer. In this layer are only systems in operation that need to have direct access from external resources.

The application layer or backend tier is the middle security layer that has a medium trusted level. The separation from business logic and databases takes place inside the application layer. The application layer provides the business logic to control and monitor the state of the service. The business logic itself is implemented on the different server types within this layer. Sensitive data shall be processed but not stored permanently inside this layer.

The Secure Backend Tier or data base layer as it is often referred is the security layer that has the highest trusted level. This layer is designed to store sensitive data that needs to be protected. The data base layer is divided from the application layer by using a Firewall.

With this design it is ensured that the database layer is not accessible from insecure networks and this layer is not able to access to insecure networks. Figure 4 below is an example of the high level three tier architecture.

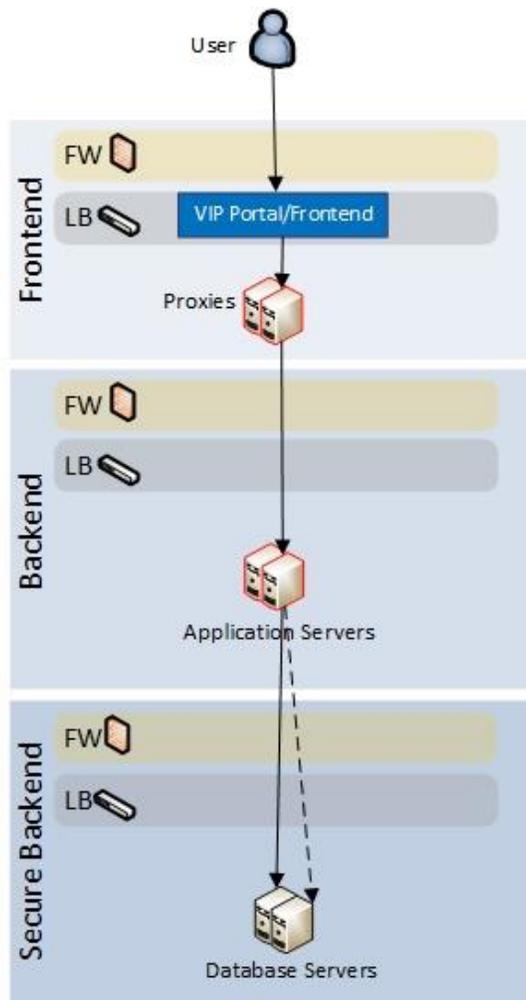


Figure 4: Three tier architecture High Level Overview

2.3.2 Serverless architecture

Serverless architecture describes a way for companies to build and run applications but not have to manage infrastructure. It provides a way to remove architecture responsibilities from your workload, including provisioning, scaling, and maintenance. Scaling can be automatic, and you only pay for what you use. This architecture is on the as best found out by a survey done by O’Rilley where 40 percent of organizations adopted serverless architecture [18]. Both serverless computing and containers enable developers to build applications with far less overhead and more flexibility than applications hosted on traditional servers or virtual

machines. Which style of architecture an enterprise should use depends on the needs of the application, but serverless applications are more scalable and usually more cost-effective [17].

AWS offers the built in service called Lambda. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers, creating workload-aware cluster scaling logic, maintaining event integrations, or managing runtimes. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Serverless architecture has primarily five principles [19]:

- **No Servers:** In the platform, the user functions are stored in permanent media storage until they are invoked. Only then, it is executed (in RAM). Essentially, there is only one running and listening service called API Gateway which maintains all the clients, other function calls for all functions in the system, contrarily with the traditional Server architecture that obliges each developer to configure and maintain his applications and their servers by himself (all servers are in continuous listening).
- **Single-Purpose Stateless functions:** A function should only take one task. That makes it easy to test/ reuse and minimizes its side effects.
- **Event-Driven invocation:** Functions are loosely coupled that are triggered by event occurrences.
- **Powerful front-ends:** Serverless functions are running in back-end that is invoked from front-end client applications. Therefore, reducing the task of a back-end to the very principle and making the front-end communicate with other services (third-party services) results in less execution time for Serverless functions and, hence, cheaper costs.
- **Third-Party Service usage:** This allows the reuse of already existing APIs and services which reduces the task of building from scratch.at any scale

Lambda functions can be written in most high level programming languages like Python, Go, Java etc which helps enterprises focus only on their code rather the whole management and operations that need to be taken care of in an environment.

2.3.3 Containerization Technology

Applications developed to fulfil distributed systems needs have been growing rapidly. Major evolutions have happened beginning with basic architecture relying on initiated request by a client to a processing side referred to as the server. Such architectures were not enough to cope up with the fast ever-increasing number of requests and need to utilize network bandwidth. Mobile agents attempted to overcome such drawbacks but did cope up for so long with the growing technology platforms. Service Oriented Architecture (SOA) then evolved to be one of the most successful representations of the client-server architecture with an added business value that provides reusable and loosely coupled services. SOA did not meet customers and business expectations as it was still relying on monolithic systems. Resilience, scalability, fast software delivery and the use of fewer resources are highly desirable features. Microservices architecture came to fulfil those expectations of system development [20].

Although serverless architecture and containers are similar how exactly are they similar? Both serverless computing and containers are cloud-based, and both greatly reduce infrastructure overhead – serverless computing more so than containers. In both kinds of architecture, applications are broken down and deployed as smaller components. In a container-based architecture, each container will run one microservice.

Microservices are segments of an application. Each microservice performs one service, and multiple integrated microservices combine to make up the application. Although the name seems to imply that microservices are tiny, they do not have to be.

One of the advantages of building an application as a collection of microservices is that developers can update one microservice at a time instead of updating the entire application when they need to make changes. Building an application as a collection of functions, as in a serverless architecture, offers the same benefit but at a more granular level.

2.4 Technologies and Networks

One of the key benefits for small, medium or large enterprises to start or move to the cloud is the abundance of services that are offered in AWS. Whether it is a small web service or big data solution, chances are customers will be able to find themselves their tailored service.

Amazon Web Services offers a broad set of global cloud-based products including compute, storage, databases, analytics, networking, mobile, developer tools, management tools, IoT, security and enterprise applications. These services help organizations move faster, lower IT costs, and scale. In the following chapter we will focus on the most basic services available on AWS.

2.4.1 Elastic Compute Cloud – EC2

EC2 is one of the most basic services that allows customers to order compute resources and a variety of choices in compute power, RAM, Graphics and interfaces. Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, 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 computing environment.

EC2 offers the broadest and deepest compute platform with choice of processor, storage, networking, operating system, and purchase model. There are also powerful GPU instances that can be used for machine learning training and graphics workloads, as well as one of the lowest cost-per-inference instances in the cloud.

2.4.2 Simple Storage Service – S3

Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. This means customers of all sizes and industries can use it to store and protect any amount of data for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics. Amazon S3 provides easy-to-use management features so you can organize your data and configure finely-tuned access controls to meet your specific business, organizational, and compliance requirements. Amazon S3 is designed for 99.999999999% (11 9's) of durability [21].

2.4.3 Identity Access Management – IAM

AWS Identity and Access Management (IAM) is a web service that helps securely control access to AWS resources. IAM is used to control who is authenticated (signed in) and authorized (has permissions) to use resources.

When first an AWS account is created, you begin with a single sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account root user and is accessed by signing in with the email address and password that you used to create the account. It is strongly recommended that the root user is not used for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create the first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks [22].

IAM is used for:

- Shared access to your AWS account
- Granular permissions
- Secure access to AWS resources
- Multi-factor authentication (MFA)
- Identity federation etc.

2.4.4 Virtual Private Cloud – VPC

Amazon Virtual Private Cloud (Amazon VPC) is a service that lets you launch AWS resources in a logically isolated virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can use both IPv4 and IPv6 for most resources in your virtual private cloud, helping to ensure secure and easy access to resources and applications.

As one of AWS's foundational services, Amazon VPC makes it easy to customize VPC's network configuration. It allows you to create a public-facing subnet for your web servers that have access to the internet. It also lets you place your backend systems, such as databases

or application servers, in a private-facing subnet with no internet access. Amazon VPC lets you to use multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Amazon VPC provides advanced security features that allow you to perform inbound and outbound filtering at the instance and subnet level. Additionally, you can store data in Amazon S3 and restrict access so that it's only accessible from instances inside your VPC. Amazon VPC also has monitoring features that let you perform functions like out-of-band monitoring and inline traffic inspection, which help you screen and secure traffic [23].

3 PROBLEM STATEMENT

Cloud computing has become a trend in the last couple of years offering secure and highly scalable environments for customers, it comes also with its own benefits for small, medium and large enterprises for different kinds of services. Due to its nature of virtualizing all the necessary layers it enables customers to deploy their environments in a quick and easy way without lacking or sacrificing any security controls.

This chapter and the previous chapter will present the necessary answers to questions listed below:

- What are the public cloud benefits?
- What are the relevant security mechanisms?
- What are the security challenges?

3.1 Public Cloud benefits

Benefits of cloud computing vary depending on the chosen architecture a customer wants to build. For example, in cases where web services are planned to be deployed the built in services offered in AWS make it possible for customers to further choose how scalable and efficient the service will be. For these type of services AWS offers CloudFront which is a Content Delivery Network (CDN) that is used primarily as caching service to securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment [24].

In chapter two the benefits of cloud computing have been presented which show different areas where consumers have seen significant efficiency in costs and deployment time.

3.1.1 Reduction of costs

Cost efficiency is the main driver for cloud computing adoption. Other primary benefits include scalability, flexibility, agility, better IT resource management and business focus, efficiency, higher reliability and availability, rapid development, deployment and change management, better performance and greater mobility. Improved automation, support and management, improved security, and green-IT data centers were also cited as valuable drivers for moving to the cloud [8].

3.1.2 Scalability and agility

With the increasing number of resources in data centers, scalability and agility in the cloud is one of the key drivers that consumers seek in their environments. This is due to the increasing demand of IT services by end users hence driving forward these requirements [25].

3.1.3 Deployment time

In today's industry where time is of crucial importance, consumers constantly seek shorter times to deploy their services to make it available for their consumers. Whether deploying a new environment or enhancing their existing ones. Table 1 shows the time needed to deploy instances in the cloud environment.

There are also different services in AWS that can be used to automatically scale an environment based on the load the service is currently facing. One such functionalities is Auto Scaling. AWS Auto Scaling lets you set target utilization levels for multiple resources in a single, intuitive interface. You can quickly see the average utilization of all of your scalable resources without having to navigate to other consoles. For example, if your application uses Amazon EC2 and Amazon DynamoDB, you can use AWS Auto Scaling to manage resource provisioning for all of the EC2 Auto Scaling groups and database tables in your application [26].

3.2 Relevant security mechanisms

Information security is of paramount importance to Cloud Computing customers. Security is a core functional requirement that protects mission-critical information from accidental or deliberate theft, leakage, integrity compromise, and deletion. Under the AWS shared responsibility model, AWS provides a global secure infrastructure and foundation compute, storage, networking and database services, as well as higher level services. AWS provides a range of security services and features that customers can use to secure their assets. Customers are responsible for protecting the confidentiality, integrity, and availability of their data in the cloud, and for meeting specific business requirements for information protection.

3.2.1 EC2 and Security

One of the most used services in AWS is EC2 and as such there needs to be a tight control on security policies and security configurations. These configurations go hand in hand with VPC security. EC2 can be secured using firewalls known as Security Groups. AWS Security Groups help you secure your cloud environment by controlling how traffic will be allowed into your EC2 machines. With Security Groups, you can ensure that all the traffic that flows at the instance level is only through your established ports and protocols.

When launching an instance on Amazon EC2, you need to assign it to a particular security group. You can add rules to each security group that allow traffic to or from designated services including associated instances.

Like whitelists, security group rules are always permissive. It's not possible to create rules that deny access. For example, you may have traffic coming from an Elastic Load Balancer (ELB) to a subnet with web servers. Your AWS Security Group can list that ELB as their sole permitted source. Security groups are stateful, which means that if an inbound request passes, then the outbound request will pass as well [27].

3.2.2 S3 and security mechanisms

Amazon S3 provides a number of security features to consider for customers. The following best practices are general guidelines and don't represent a complete security solution.

Unless explicitly required for anyone on the internet to be able to read or write to S3 buckets, customers should ensure that your S3 bucket is not public. The following are some of the steps they can take [28]:

- Use Amazon S3 block public access. With Amazon S3 block public access, account administrators and bucket owners can easily set up centralized controls to limit public access to their Amazon S3 resources that are enforced regardless of how the resources are created. For more information, see [Blocking public access to your Amazon S3 storage](#).
- Identify Amazon S3 bucket policies that allow a wildcard identity such as Principal "*" (which effectively means "anyone") or allows a wildcard action "*" (which effectively allows the user to perform any action in the Amazon S3 bucket).
- Similarly, note Amazon S3 bucket access control lists (ACLs) that provide read, write, or full-access to "Everyone" or "Any authenticated AWS user."
- Use the ListBuckets API to scan all of Amazon S3 buckets. Then use GetBucketAcl, GetBucketWebsite, and GetBucketPolicy to determine whether the bucket has compliant access controls and configuration.
- Consider implementing on-going detective controls using the s3-bucket-public-read-prohibited and s3-bucket-public-write-prohibited managed AWS Config Rules

3.2.3 IAM Policies and Roles

IAM is one of the key security controls that can help customers secure their accounts and environments. There are three pillars in IAM listed below which are also known as Identities [29]:

- IAM users
- IAM user groups
- IAM Roles

An IAM user is an entity that you create in AWS. The IAM user represents the person or service who uses the IAM user to interact with AWS. A primary use for IAM users is to give people the ability to sign in to the AWS Management Console for interactive tasks and to make programmatic requests to AWS services using the API or CLI. A user in AWS consists of a name, a password to sign into the AWS Management Console, and up to two access keys that can be used with the API or CLI. When you create an IAM user, you grant it permissions by making it a member of a user group that has appropriate permission policies attached (recommended), or by directly attaching policies to the user. You can also clone the permissions of an existing IAM user, which automatically makes the new user a member of the same user groups and attaches all the same policies.

An IAM user group is a collection of IAM users. You can use user groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users. For example, you could have a user group called Admins and give that user group the types of permissions that administrators typically need. Any user in that user group automatically has the permissions that are assigned to the user group. If a new user joins your organization and should have administrator privileges, you can assign the appropriate permissions by adding the user to that user group.

An IAM role is very similar to a user, in that it is an identity with permission policies that determine what the identity can and cannot do in AWS. However, a role does not have any credentials (password or access keys) associated with it. Instead of being uniquely associated

with one person, a role is intended to be assumable by anyone who needs it. An IAM user can assume a role to temporarily take on different permissions for a specific task. A role can be assigned to a federated user who signs in by using an external identity provider instead of IAM. AWS uses details passed by the identity provider to determine which role is mapped to the federated user.

3.2.4 VPC Configurations

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS Cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC.

To secure resources on the VPC, it is a best practice to separate public and private subnets. This also falls on the three-tier architecture best practices where resources are divided into layers from lower trusted level (Frontend) where proxies and loadbalancers are hosted to highly trusted level (Secure Backend) where databases are hosted [30].

One key functionality that VPCs offer are Network ACL (Access Control List). A network access control list (ACL) is an optional layer of security for VPCs that acts as a firewall for controlling traffic in and out of one or more subnets. It is an added benefit set up network ACLs with rules similar to security groups in order to add an additional layer of security to VPCs.

3.3 Security challenges

Ever since the mainstream adoption of cloud computing that happened over the course of the last decade, there has been a major debate over the security of the cloud. Many businesses have yet to make the move to the cloud for fear of having their personal information moved outside of the organization and into the cloud. Some argue that on-premise is the most secure as your data stays within an organization, behind the company firewall. While some argue that the cloud is just as secure if not more. One of the main pillars of separating responsibilities is the Shared Responsibility Model by AWS.

In the following sub-chapters, we will focus on some of these challenges.

3.3.1 Distributed Denial of Service (DDoS)

A good example of a common security challenge is Distributed Denial of Service (DDoS) attacks are attempts by a malicious actor to flood a network, system, or application with more traffic, connections, or requests than it is able to handle. To protect web applications against DDoS attacks, customers can use AWS Shield, a DDoS protection service that AWS provides automatically to its customers at no additional charge. Customers can use AWS Shield in conjunction with DDoS-resilient web services such as Amazon CloudFront and Amazon Route 53 to improve the ability to defend against DDoS attacks [32].

3.3.2 Access Control

In many application scenarios, such as those in enterprises or organizations, users' access to data is usually selective and highly differentiated. Different users enjoy different access privileges with regard to the data. When data are outsourced to the cloud, enforcing secure, efficient, and reliable data access among a large number of users is thus critical. Traditionally, to control the dissemination of privacy-sensitive data, users establish a trusted server to store data locally in clear, and then control that server to check whether requesting users present proper certification before letting them access the data. From a security standpoint, this access control architecture is no longer applicable when we outsource data to the cloud. Because data users and cloud servers aren't in the same trusted domain, the server might no longer be fully trusted as an omniscient reference monitor⁹ for defining and enforcing access control policies and managing user details. In the event of either server compromise or potential insider attacks, users' private data might even be exposed. One possible approach to enforce data access without relying on cloud servers could be to encrypt data in a differentiated manner and disclose the corresponding decryption keys only to authorized users. This approach usually suffers from severe performance issues, however, and doesn't scale, especially when a potentially large number of on-demand users desire fine-grained data access control. Researchers have been working on how to realize a fine-grained access control design that fully leverages the cloud's computation resource richness. Via this approach, data users would be able to securely delegate to the cloud most cumbersome user/

data management workloads — such as handling frequent user access privilege updates in large dynamic systems — while still preserving the underlying data confidentiality against any unauthorized Access [8] [31].

3.3.3 Multitenancy Security and Privacy

Multitenancy is an essential attribute of cloud computing. To optimize resource utilization, Cloud Computing Providers often use hardware virtualization to hide a computing platform's physical characteristics. This lets multiple users run their distinct application instances simultaneously on the same physical infrastructure without seeing each other's data. Multitenancy increases use of the underlying hardware resources and, with virtualization, eases the management burden for Cloud Computing Providers, allowing for efficient and effective resource provisioning and re-allocation without the need for any upfront hardware purchase or setup. Despite its benefits, this multitenant cloud environment also presents severe security threats and privacy vulnerabilities to both the cloud infrastructure and cloud users. Virtualized environments share similar functionalities with existing operating systems and applications in the physical environment, so software bugs and newly identified security vulnerabilities in these systems remain the primary threat to any virtualized multitenant environment. Considering the scale of cloud systems, the potential threat from these security risks can be even bigger compared to that for a nonvirtualized computing environment. Furthermore, for resource management in the cloud, different virtualized application instances must be constantly provisioned, allocated, or even migrated between multiple physical machines. Consequently, such dynamic features in the multitenant environment further exacerbate the problem's complexity and make achieving and maintaining consistent security difficult [31].

4 METHODOLOGY

The focus of this paper is to analyze the benefits in the cloud computing field and what security challenges arise from its usage. The best answers to this study are provided by secondary research and is believed to be the most cost-effective method which gives the ability to get an easy and accurate insight. Data made public by different universities, research groups, governments and organizations and institutes will provide general knowledge of the past decades, current and future trends in the cloud field. Due to the nature of the study, secondary data sources are referred to, which are obtained from various sources to assist with the conclusion for this study. The selection of the secondary method was also based on that principle, besides it is a method which is faster and less costly also makes it possible to see the trends worldwide and compare different data and opinions done by trustworthy organizations and entities. Secondary sources are considered data that already exist, therefore extensive research was performed by acquiring data at various sources such as [33]:

- Data from the web: a cost-effective way of utilizing existing available data makes it the most convenient method. A downside is that these sources have to be trusted and must have a certain degree of credibility with the data they provide and users must be careful as that may compromise the study.
- Digital libraries: while the cloud computing field is relatively new there is always a need for updates in trends which often presents a challenge in this area since while this paper is being written there are changes in trends!
- ICT Tech Institutes: trusted institutes around the world keep monitoring trends changes and often have yearly reports on trends in ICT trends updates.
- Commercial sources: such as magazines, and market research provide valuable information and are excellent resources for secondary data collection.
- Universities: scholarly paper or study will provide the most answers for this paper.

The steps in conducting this research have been listed below in the following format:

- The research topic has been identified.

- Literature review selection to be used for research.
- Methodology
- Data analysis and review
- Conclusion.

5 CASE STUDY

Although AWS has helped tens of thousands of customers worldwide with their services like Verizon, Amdocs, Volkswagen and Coca Cola. This study will take as a use case the streaming platform Netflix to analyze the benefits and how it makes use of the cloud virtualized environments on AWS.

Netflix is the world's leading internet television network, with more than 200 million members in more than 190 countries enjoying 125 million hours of TV shows and movies each day. Netflix uses AWS for nearly all its computing and storage needs, including databases, analytics, recommendation engines, video transcoding, and more—hundreds of functions that in total use more than 100,000 server instances on AWS [34].

5.1 Netflix on AWS in numbers

Netflix uses a large amount of resources to deliver content to its users in a high quality and distributed manner. Due to the nature of streaming services and its sensitivity to quality and latency of the content they deliver, Netflix has to be uncompromising when it comes to these two areas, on top of these it needs to provide the highest level of security to assure its users that their personal information and credit cards are safe and not easy to compromise. To understand exactly the scale of resources that Netflix uses, we need to first understand the numbers that Netflix delivers [34].

- Provides services to over 200 million users.
- Operates in over 190 countries.
- Provides over 125 million hours of TV shows and Movies each day.
- Sees over 10,000+ stream starts every second at peak times.
- Operates in 3 AWS regions.
- It is deployed in 12 Availability Zones.
- Has over 100,000+ instances.

5.2 Netflix Benefits

Due to the incredible growth of Netflix in the recent years it needed the right infrastructure and technologies to power their services to be able to cope with the high demand and peak hours. Below are listed the key challenges which Netflix has solved [35]:

- Redundancy and Failover - Redundancy and Failover are the most critical of factors that enterprises face when dealing with media content delivery as it deals with sensitive metrics such as latency (lag) and content availability. Netflix has benefited from AWS due to deploying their services in AWS Regions which are located on key geographic continents to cover all of the globe.
- Time consumption on infrastructure – Instead of worrying about taking care of such massive infrastructure, Netflix has solved this by giving this responsibility to AWS. Instead, Netflix engineers are focused on building and improving the business
- Horizontal scaling – Due to unpredictability of customer growth or device engagement, Netflix will not spend time predicting customer growth, instead features like Auto Scaling takes care of it by scaling the infrastructure needed based on the customer demand, effectively reducing cost of unused instances while at the same time providing necessary resources to deal with customer growth.
- Keeping up with the trend – AWS is one of the key players in the industry on setting trends and implementing them first. This allows customers like Netflix to have highest level of exposure to new technologies and trends to gain competitive advantage with its competitors.

6 CONCLUSION

The purpose of this chapter is to highlight the key aspects and research that has been done on this topic. This paper was initiated based on the latest trend of cloud computing and how specifically enterprises can benefit. Due to the immense pace of technologies being developed and IT industry changing and adapting to new paradigms, it was seen as a worthy topic for a thesis. This chapter focuses on the findings on this paper by trying to address the following objectives:

1. What are the benefits in cloud computing and is it worth it?
2. Identifying security challenges and proper security controls.

The first objective of this research was assisted by available data made online by different organizations, universities and different use cases which were made public. When a move to cloud computing is considered, a number of metrics need to be taken into account, like architecture, data privacy and costs.

The second objective of this research has been elaborated due to the need to identify security challenges which come with the public cloud. An extensive research has been done to analyze these challenges and understand the impact they have on various environments. Fortunately, quite a number of resources are available for data gathering and literature review to be able to focus on key concepts and challenges we elaborate in this paper. The benefits, along with the challenges of cloud computing are classified within one framework. The second chapter presents this classification while the third chapter aims to answer the questions and relevant security challenges.

The material presented in Chapter 2 was the basis for the research activities that followed. A good number of papers and data have been analyzed to serve as a building block to be able to sum up and structure key concepts and best practices in the cloud computing industry.

On Chapter 3 we aim to present the most important challenges that need to be addressed while considering a move to a cloud computing provider. As stated many times during in this paper, different use cases require different technologies which need to be investigated

separately and thoroughly while maintaining a degree of skepticism whether cloud computing is the most beneficial for particular use cases.

The case study (Chapter 5) presents how companies can benefit if they utilize the full power of cloud computing and how it adds value in their business processes while freeing up time for employees to focus on business problems, process efficiency and how to make better products. The case study also shows how in cases where customer growth is unpredictable, this can be managed by using the elasticity that cloud computing provides while saving operational time and saving costs.

7 REFERENCES

- [1] MARKUS BÖHM, STEFANIE LEIMEISTER, CHRISTOPH RIEDL, HELMUT KRCMAR, 2010, Cloud Computing and Computing Evolution
- [2] Bhanu Priya; Emmanuel S. Pilli; Ramesh C. Joshi, 2013, A survey on energy and power consumption models for Greener Cloud
- [3] Daniele Catteddu, 2009, Cloud Computing: Benefits, Risks and Recommendations for Information Security
- [4] Mariana Carroll; Alta van der Merwe; Paula Kotzé, 2011, Secure cloud computing: Benefits, risks and controls
- [5] Hesham M. Elmasry; Ayman E. Khedr; Mona M. Nasr, 2019, An adaptive technique for cost reduction in cloud data centre environment
- [6] Rackspace, “Rackspace survey on cloud”, date of publication: 21.02.2013
[<https://www.rackspace.com/newsroom/88-per-cent-of-cloud-users-point-to-cost-savings-according-to-rackspace-survey#:~:text=Buy%20Now-.88%20per%20cent%20of%20cloud%20users%20point,savings%20according%20to%20Rackspace%20Survey&text=The%20study%2C%20conducted%20by%20Rackspace,the%20cloud%20have%20saved%20money>] date accessed: 03.04.2021.
- [7] AWS, “Architecting for reliable scalability” date of publication: 03.11.2020
[<https://aws.amazon.com/blogs/architecture/architecting-for-reliable-scalability/>], date accessed: 05.04.2021
- [8] Mariana Carroll; Alta van der Merwe; Paula Kotzé, 2011, Secure cloud computing: Benefits, risks and controls
- [9] Irena Bojanova; Augustine Samba, 2011, Analysis of Cloud Computing Delivery Architecture Models
- [10] AWS, “Architecting for reliable scalability” date of publication: 03.11.2020
[<https://aws.amazon.com/blogs/architecture/architecting-for-reliable-scalability/>], date accessed: 10.04.2021.

[11] Eduardo Roloff; Matthias Diener; Alexandre Carissimi; Philippe O. A. Navaux, 2012, High Performance Computing in the Cloud: Deployment, Performance and Cost Efficiency

[12] AWS, “What is AWS Elastic Beanstalk?”, date of publication: February, 2021 [<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>], date accessed: 12.04.2021

[13] R. Velumadhava; K. Selvamani, 2015, Data Security Challenges and Its Solutions in Cloud Computing

[14] AWS, “Shared Responsibility Model”, date of publication: February, 2021 [<https://aws.amazon.com/compliance/shared-responsibility-model/>], date accessed: 10.04.2021

[15] Yanuarizki Amanatullah; Charles Lim; Heru Purnomo Ipung; Arkav Juliandri, 2013, Toward Cloud Computing Reference Architecture: Cloud Service Management Perspective

[16] Amund Aarsten; Davide Brugali; Giuseppe Menga, 2003, Patterns for Three-Tier Client/Server Applications

[17] Cloudflare, “Serverless vs Containers”, date of publication: June 2018, [<https://www.cloudflare.com/learning/serverless/serverless-vs-containers/#:~:text=Both%20are%20cloud%2Dbased%2C%20and,container%20will%20run%20one%20microservice>], date accessed: 13.04.2021.

[18] Newrelic, “What is serverless architecture?”, date of publication: 24.05.2017 [<https://newrelic.com/blog/best-practices/what-is-serverless-architecture#:~:text=Serverless%20architecture%20describes%20a%20way,pay%20for%20what%20you%20use>], date accessed: 20.04.2021.

[19] Boubaker Soltania; Afifa Ghenaia; Nadia Zeghib, 2018, Towards Distributed Containerized Serverless Architecture in Multi Cloud Environment

[20] Tasneem Salah; M. Jamal Zemerly; Chan Yeob Yeun; Mahmoud Al-Qutayri; Yousof Al-Hammadi, 2016, The evolution of distributed systems towards microservices architecture

- [21] AWS, “Amazon S3”, date of publication: August 2016
[<https://aws.amazon.com/s3/>], date accessed: 21.04.2021.
- [22] AWS, “Introduction to IAM”, date of publication: January 2019
[<https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html>], date accessed: 22.04.2021
- [23] AWS, “Virtual Private Cloud”, date of publication: September 2020
[<https://aws.amazon.com/vpc/?vpc-blogs.sort-by=item.additionalFields.createdDate&vpc-blogs.sort-order=desc>], date accessed: 25.04.2021.
- [24] <https://aws.amazon.com/cloudfront/>
- [25] Charlie Baker; Ashiq Anjum; Richard Hill; Nik Bessis, 2012, Improving Cloud Datacentre Scalability, Agility and Performance using OpenFlow
- [26] AWS “AWS Autoscaling”, date of publication: 2021
[<https://aws.amazon.com/autoscaling/#:~:text=AWS%20Auto%20Scaling%20lets%20you%20build%20scaling%20plans%20that%20automate,you%20based%20on%20your%20preference>], date accessed: 26.04.2021.
- [27] Yanpei Chen; Vern Paxson; Randy H. Katz, 2010, What’s New About Cloud Computing Security?
- [28] AWS, “S3 security best practices”, date of publication: June 2020
[<https://docs.aws.amazon.com/AmazonS3/latest/userguide/security-best-practices.html>], date accessed: 02.05.2021
- [29] AWS, “IAM best practices”, date of publication: November 2020
[<https://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>], date accessed: 02.05.2021.
- [30] K Lewis, 2017, Computer and Information Security Handbook (Third Edition)
- [31] Kui Ren; Cong Wang; Qian Wang, 2012, Security Challenges for the Public Cloud
- [32] AWS, “AWS CDN and how to protect against DDoS”, date of publication: January 2021
[<https://aws.amazon.com/blogs/security/how-to-protect-your-web-application-against-ddos-attacks-by-using-amazon-route-53-and-a-content-delivery-network/#:~:text=To%20protect%20your%20web%20application%20against%20DDoS%20attacks%2C%20you%20can,customers%20at%20no%20additional%20charge>], date accessed: 05.05.2021.

[33] Naim Preniqi; Muhamet Gërvalla; Kaltrina Sylaj, 2020, Digital Transformation in e-Learning Education

[34] AWS, “Netflix case study”, date of publication: May 2017
[<https://aws.amazon.com/solutions/case-studies/netflix/>], date accessed: 06.05.2021.

[35] Netflixtechblog, “Four reasons we choose AWS”, date of publication: 26.07.2018 [<https://netflixtechblog.com/four-reasons-we-choose-amazons-cloud-as-our-computing-platform-4aceb692afec>], date accessed: 12.05.2021.

8 BIBLIOGRAPHY

- [1] Keith Lewis, *Computer and Information Security Handbook (Third Edition)*, 2017
- [2] Andreas Wittig; Michael Wittig, *Amazon Web Services in Action (2nd Edition)*, 2021
- [3] Moe Abdula; Ingo Averdunk; Roland Barcia; Kyle Brown; Ndu Emuchay, *The Cloud Adoption Playbook*, 2018
- [4] Teri Radichel, *Cybersecurity for Executives in the Age of Cloud*, 2020

9 APPENDIXES