



EDUCATIONAL FRAMEWORK ON CLOUD COMPUTING

First part



A-CCT



Co-funded by
the European Union



acctproject.eu



Maja Pucelj, Annmarie Gorenc Zoran, Nadia Molek, Ali Gökdemir, Ioan Ganea,
Christina Irene Karvouna, Petter Grøttheim, Leo Mršić, Maja Brkljačić, Monika
Rohlik Tunjić, Alojz Hudobivnik

EDUCATIONAL FRAMEWORK ON CLOUD COMPUTING

First part

Novo mesto, 2023



Co-funded by
the European Union

EDUCATIONAL FRAMEWORK ON CLOUD COMPUTING – FIRST PART

Maja Pucelj, Annmarie Gorenc Zoran, Nadia Molek, Ali Gökdemir, Ioan Ganea, Christina Irene Karvouna, Petter Grøttheim, Leo Mršič, Maja Brkljačić, Monika Rohlik Tunjić, Alojz Hudobivnik

Funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

Published by: Faculty of Organization Studies in Novo Mesto

Copyright © 2023 in part and in full by the author and the Faculty of Organization Studies in Novo mesto, Novo mesto.

All rights reserved. No part of this material may be copied or reproduced in any form, including (but not limited to) photocopying, scanning, recording, transcribing, without the written permission of the author or another natural or legal person to whom the author has transferred the material copyright.

Published on: <https://www.fos-unm.si/si/dejavnosti/zaloznistvo/>

Content

1	INTRODUCTION.....	11
2	CLOUD COMPUTING TRAINING MATERIALS	11
2.1	Introduction to cloud computing technologies and types of cloud computing	11
2.2	Pricing vs. Market Comparison between AWS, Azure and GCP	19
2.2.1	What Does Cloud Computing Offer?	20
2.2.2	The 3 key players on the Market	24
2.2.3	The Cloud Market Share Comparison	27
2.2.4	Analysis of Pricing Structures	30
2.3	Selecting and setting the infrastructure	32
2.3.1	Deploying Servers and Load Balancers on all Computing Platforms	32
2.3.2	Storage Services in the cloud.....	37
2.3.3	Identity access management	52
2.3.4	Database Services in the cloud	59
2.3.5	Considerations for Domain Set up	68
2.4	Connectivity Types of Network Services and Setting Them	72
2.4.1	About cloud architecture	72
2.4.2	Cloud access connectivity principles	76
2.4.3	Cloud network set-up.....	86
2.5	Cloud System Management (Monitoring and Notification Service).....	93
3	APPLICATIONS.....	100
3.1	Access to a database using a person's fingerprint as a password.....	100
3.2	Active Directory Server	100
3.3	AI Behaviour analysis Systems	101
3.4	Application for managing the activity of renting tools and equipment from a company to natural persons	102
3.5	Application for monitoring autonomous room cleaning equipment (vacuum cleaners) at the headquarters of small and medium-sized companies or in private homes	102
3.6	Asset Tracking.....	103
3.7	Attendance tracker for students.....	104
3.8	Automated Facilities Management	104
3.9	Automation of tasks using cloud-based services: recommendation engine	106
3.10	Back-Up / Disaster relief	106
3.11	Chatbot for indicating free places in public parking lots in a city.....	106
3.12	Chatbot to personalize the learning activity of students in vocational high school education	107
3.13	Chatbot for students in EDU institution	107
3.14	Cloud-based e-learning.....	108
3.15	Communication/ Information Exchange Application/ Channels.....	108
3.16	Continuous monitoring of the operation of some industrial installations using cloud computing and IoT technologies.....	109
3.17	Continuous patient monitoring.....	110
3.18	Create test environments	110



3.19	Creating a didactic application to help students learn a foreign language.....	110
3.20	Data backups and archiving	111
3.21	Data loss prevention cloud-based system.....	112
3.22	Data management system about a company's employees.....	112
3.23	Digital asset certification using distributed ledger/blockchain	113
3.24	Digital identity.....	113
3.25	Digital twinning.....	114
3.26	Disaster prevention platform.....	114
3.27	Distribution of parcels in a geographical region with the help of autonomous.....	115
3.28	Document similarity detection and document information extraction system.....	115
3.29	Document translation	115
3.30	Dynamic website hosting	116
3.31	Dynamic website with data storage in a database	116
3.32	E-commerce Application	117
3.33	Electronic catalogue with students' school results	118
3.34	Facilities Access Control.....	118
3.35	Facilities Management.....	119
3.36	Facilities Occupancy Data.....	120
3.37	File Comparison	121
3.38	File storage system using hybrid cryptography cloud computing.....	123
3.39	Handling traffic spikes.....	123
3.40	Host a static website using AWS (or other clouds)	123
3.41	Instant Messaging applications.....	124
3.42	Manage virtual network	125
3.43	Migrate to cloud	125
3.44	Monitoring the activities carried out by agricultural machinery on a given surface	126
3.45	Monitoring the physiological parameters of athletes during training.....	126
3.46	Operate several projects simultaneously.....	127
3.47	Reconfiguration of public transport routes in a city	127
3.48	Remote-controlled smart devices in smart home/office	127
3.49	Resource and application access management	128
3.50	Rule-based phishing website classification	128
3.51	SAP Build.....	129
3.52	Set up load balancers.....	130
3.53	Smart traffic management.....	130
3.54	Supply real-time sales data	131
3.55	The graphic interface for programming at a car service combined with a website	132
3.56	Video conference system.....	132
3.57	VoD offering.....	133
3.58	Water supply management using distance readers in water supply networks.....	134
3.59	Web application for the online completion of a company's staff timesheet	134
3.60	Website hosting with static content	135
3.61	Webstore.....	135
	REFERENCE	137
	APPENDIX.....	140

CONTENT OF THE TABLES

Table 5.1. The five rules with the largest lift.....	150
Table 5.2. Circuit current without optimization.....	158
Table 5.3. Current through the water sensor.....	158
Table 5.4. Current with reduced microprocessor clock speed.....	158

CONTENT OF THE FIGURES

Figure 2.1. Suggestive image of the term cloud computing.....	12
Figure 2.2. The hierarchy of the three basic levels in cloud computing services.....	15
Figure 2.3. Aspects inside a data centre that provides cloud computing services.....	18
Figure 2.4. Business Benefits of Cloud Implementation.....	21
Figure 2.5. Cloud Provider Market Share Trend.....	25
Figure 2.6. The costs of cloud infrastructure services for Q1 2021 in USA compared to years 2019 and 2020.....	28
Figure 2.7. The costs of cloud infrastructure services for Q1 2021 in China compared to years 2019 and 2020.....	29
Figure 2.8. AWS Vs. Azure Vs. GCP Cloud Cost Comparison.....	32
Figure 2.9. Application Load Balancer for AWS.....	33
Figure 2.10. Network Load Balancer.....	34
Figure 2.11. Load Balancers.....	35
Figure 2.12. Choosing a Cloud Load Balancer.....	35
Figure 2.13. Hybrid deployment with an external global HTTP(S) load balancer.....	36
Figure 2.14. Network Load Balancer in a user case.....	37
Figure 2.15. A price comparison between Hot storage and Cool storage with AWS S3.....	38
Figure 2.16. Infrequent Access prices.....	38
Figure 2.17. S3 Standard prices.....	39
Figure 2.18. S3 Standard prices.....	39
Figure 2.19. S3 Glacier Instant Retrieval, Flexible Retrieval and Deep Archive.....	40
Figure 2.20. S3 console.....	42
Figure 2.21. Create bucket in S3 console.....	42
Figure 2.22. Setting a bucket in S3 console.....	43
Figure 2.23. Bucket versioning in S3 console.....	44
Figure 2.24. Object ownership in S3 console.....	44
Figure 2.25. Finishing configuration in S3 console.....	45
Figure 2.26. Uploading files to newly created bucket in S3 console – first step.....	45
Figure 2.27. Uploading files to newly created bucket in S3 console – second step.....	46
Figure 2.28. Uploading files to newly created bucket in S3 console – third step.....	46
Figure 2.29. Uploading images to newly created bucket in S3 console.....	47
Figure 2.30. Properties in S3 console.....	48
Figure 2.31. Upload in S3 console.....	49
Figure 2.32. Success-message of upload in S3 console.....	50
Figure 2.33. Information about the stored data in S3 console.....	50
Figure 2.34. Retrieving files in cloud in S3 console.....	51
Figure 2.35. Deleting objects from a bucket in S3 console.....	51
Figure 2.36. Deleted status in the bucket in S3 console.....	52

Figure 2.37. Deleting the bucket in S3 console.....	52
Figure 2.38. Authorization.....	54
Figure 2.39. Granting access to specific resources within AWS.....	55
Figure 2.40. Role-based access control	57
Figure 2.41. Attribute-based access control.....	58
Figure 2.42. Relationship between the book and library.....	59
Figure 2.43. Database with Amazon RDS using the Amazon Aurora MySQL.....	61
Figure 2.44. Creation of a new database – first step	61
Figure 2.45. Creation of a new database – second step	62
Figure 2.46. Settings for the database.....	63
Figure 2.47. Creating an Aurora Replica	64
Figure 2.48. Connectivity settings	65
Figure 2.49. Creating database.....	66
Figure 2.50. Created database visible in Amazon RDS console page	67
Figure 2.51. Endpoints of created database	67
Figure 2.52. Using MySQL workbench for connecting to new database.....	68
Figure 2.53. List of different cloud services.....	71
Figure 2.54. Management services, which utilize IoT tools	72
Figure 2.55. Services overview	73
Figure 2.56. Service Model Types	75
Figure 2.57. Public, Private, and Hybrid Cloud Deployment Example	76
Figure 2.58. Web 2.0 Interfaces to the Cloud	77
Figure 2.59. Cloud connectivity.....	78
Figure 2.60. Connect to the cloud – decision tree	79
Figure 2.61. Cloud connectivity using the public internet (advantages and disadvantages)	80
Figure 2.62. Cloud connectivity using public internet and cloud prioritisation (advantages and disadvantages)	81
Figure 2.63. Direct Ethernet cloud connect (advantages and disadvantages).....	82
Figure 2.64. MPLS IP VPN cloud connect (advantages and disadvantages).....	83
Figure 2.65. SD WAN cloud connect (advantages and disadvantages)	85
Figure 2.66. Virtual Networks.....	87
Figure 2.67. Building Blocks of Cloud Network	88
Figure 2.68. Surveying Network Configuration Options.....	89
Figure 2.69. Dynamic or Private Ports.....	91
Figure 2.70. Servicing Your Cloud Network.....	92
Figure 2.71. Determine granting access to the Cloud Network.....	93
Figure 2.72. Cloud System Management.....	93
Figure 2.73. Cloud Management Components.....	96
Figure 5.1. LUIS in Action	143
Figure 5.2. Quick Replies	144
Figure 5.3. Showing the module for entering the diploma.....	145
Figure 5.3. Show the diploma verification module.....	146
Figure 5.5. Transnational data source	147
Figure 5.6. ETL relations.....	148
Figure 5.7. The variables after applying ETL procedure.....	148
Figure 5.8. Bar plot of the support of the 25 most frequent items bought.....	149
Figure 5.9. A scatter plot of the confidence, support and lift metrics.....	150
Figure 5.10. Graph-based visualisation of the top ten rules in terms of lift.....	151

Figure 5.11. Water flow sensor connection diagram..... 155

Figure 5.12. Central transceiver antenna position and measuring range..... 156

Figure 5.13. LoRa LPWAN..... 160

Figure 5.14. Comparison of different methods for feature selection..... 161

Figure 5.15. Pruned tree, using the full set of features 162

Figure 5.16. Classification results for C 4.5 and SVM, experiment 1 uses only selected features.
 Experiment 2 uses selected features plus Country and ASN of client 163

Figure 5.17. Creating an S3 bucket – first step..... 168

Figure 5.18. Creating an S3 bucket – second step..... 169

Figure 5.19. Creating an S3 bucket – third step..... 169

Figure 5.20. Creating an S3 bucket – fourth step 170

Figure 5.21. Creating an S3 bucket – fifth step..... 170

Figure 5.22. Creating an S3 bucket – sixth step 171

Figure 5.23. Upload web files to S3 bucket – first step..... 171

Figure 5.24. Upload web files to S3 bucket – second step 172

Figure 5.25. Upload web files to S3 bucket – third step 172

Figure 5.26. Create IAM Role – first step 173

Figure 5.27. Create IAM Role – second step 173

Figure 5.28. Create IAM Role – third step..... 174

Figure 5.29. Create IAM Role – fourth step..... 174

Figure 5.30. Create IAM Role – fifth step 175

Figure 5.31. Create IAM Role – sixth step..... 175

Figure 5.32. Create an EC2 instance – first step..... 176

Figure 5.33. Create an EC2 instance – second step..... 176

Figure 5.34. Create an EC2 instance – third step..... 177

Figure 5.35. Create an EC2 instance – fourth step..... 177

Figure 5.36. Create an EC2 instance – fifth step..... 178

Figure 5.37. Create an EC2 instance – sixth step..... 178

Figure 5.38. Create an EC2 instance – seventh step..... 179

Figure 5.39. Create an EC2 instance – eight step..... 179

Figure 5.40. Create an EC2 instance – ninth step..... 180

Figure 5.41. Create an EC2 instance – tenth step..... 180

Figure 5.42. Create an EC2 instance – eleventh step 181

Figure 5.43. Create an EC2 instance - eleventh step..... 181

Figure 5.44. Connecting to EC2 by using MobaXterm - first step..... 181

Figure 5.45. Connecting to EC2 by using MobaXterm - second step..... 182

Figure 5.46. Connecting to EC2 by using MobaXterm – third step..... 182

Figure 5.47. Connecting to EC2 by using MobaXterm - fourth step..... 183

Figure 5.48. Installing a LAMP web server on Amazon Linux 2..... 184

Figure 5.49. Successful deployment a dynamic website on EC2..... 185

Figure 5.50. Host a static website using AWS – first step..... 186

Figure 5.51. Host a static website using AWS – second step..... 186

Figure 5.52. Host a static website using AWS – second step..... 187

Figure 5.53. Host a static website using AWS – third step 187

Figure 5.54. Host a static website using AWS – fourth step..... 188

DICTIONARY

TERM	MEANING
<i>Agility</i>	Agility in the context of cloud computing refers to the prompt and effective capacity of cloud resources and services to adjust to evolving business and technological requirements.
<i>Backend</i>	The server-side components of a cloud-based application. It encompasses many functionalities such as data management, business logic implementation, application hosting, and data processing. These backend components work in conjunction with the user-facing frontend, facilitating its operation and functionality.
<i>Backhaul</i>	The network backbone serves as a conduit for transmitting data to the central core of the network.
<i>Back-up data</i>	Data duplication is the process of creating an additional copy of data that is already stored in another location, with the purpose of mitigating the risk of data loss.
<i>Blockchain</i>	A distributed ledger that records all transactions occurring within a network.
<i>Blowfish</i>	A symmetric-key block cipher is employed to ensure secure transmission of data.
<i>Bucket</i>	A logical entity for storing data within object storage systems, such as AWS S3.
<i>Cloud computing</i>	Leveraging computational resources, such as servers, storage, and databases, via internet-based infrastructure commonly referred to as "the cloud."
<i>Cloud computing technologies</i>	The technologies that facilitate the utilization of computing services via the internet.
<i>Cluster</i>	A network of interconnected computers that collaborate closely to do activities.
<i>Durability</i>	Data durability refers to the capacity of a system to prevent the loss of data within a specified timeframe
<i>Elasticity</i>	The capacity to dynamically allocate computing resources based on the prevailing workload.
<i>Firewall</i>	A network security device that performs the functions of monitoring and filtering both inbound and outbound network traffic.
<i>Flexibility</i>	The capacity to effectively and flexibly adjust to alterations and fluctuations in workload.
<i>Frontend</i>	The components pertaining to the user interface and user experience within a given system.
<i>Flux</i>	Flux is the new generation of scalable decentralized cloud infrastructure.
<i>Headers</i>	The preamble, which is commonly employed for the purpose of including routing information, is an extra set of data located at the beginning of a data packet.
<i>Health</i>	The current state or functional state of a system or process.

<i>Health checks</i>	The implementation of monitoring systems is crucial for the purpose of ensuring that services are operating at their highest level of efficiency.
<i>Health probe</i>	A test inquiry is conducted to verify the responsiveness and overall health of a service.
<i>Hub</i>	A widely utilized interface for establishing connectivity between devices within a network.
<i>Industrial revolution</i>	The aforementioned time denotes a significant phase of industrial development, maybe alluding to the concept of Industry 4.0 within a contemporary information technology framework. This paradigm encompasses the integration of the internet of things and cloud computing.
<i>IT technology</i>	The utilization of computer systems and telecommunications technologies for the purpose of storing, retrieving, transmitting, and manipulating data.
<i>Latency</i>	The time delay experienced in a system.
<i>Listener</i>	A network monitoring system or protocol that actively detects and responds to network connections and requests.
<i>Local computer network</i>	A local area network (LAN) refers to a network that encompasses a limited geographical area, such as a residence, workplace, or educational institution.
<i>Main frame computer</i>	A high-performance computing system utilized for the execution of computationally intensive tasks on a wide scale.
<i>Mapping</i>	The process of establishing a relationship between elements belonging to one set and elements belonging to another set.
<i>Patch</i>	A software update intended to rectify or enhance its functionality.
<i>Proxy</i>	An intermediary server that acts as a mediator between end-user clients and the destinations they access for browsing purposes.
<i>Push code</i>	The act of transmitting code to a repository or environment for the purpose of implementing modifications.
<i>Routing</i>	The process of ascertaining the route for data packets to traverse within a network.
<i>Scalability</i>	The ability of a system to expand and adequately handle an increased level of demand.
<i>Virtual machine scale sets</i>	The Azure compute resource referred to is a platform that enables users to deploy and oversee a collection of indistinguishable virtual machines (VMs).
<i>Virtual machines</i>	A computer system simulation is a software-based representation that replicates the capabilities of a physical computer.

1 INTRODUCTION

In the year 2021, project partners from Slovenia, Croatia, the Netherlands, Norway, Romania and Turkey successfully received the European Erasmus+ project entitled: "Digital content development for integration of cloud technologies in formal and distance vocational education". One of the results of the project is also course teaching material content on cloud technologies, supported by sample applications, prepared as a guide for teachers of formal and distance vocational education. Below the teachers can find a first part of the mentioned guidelines.

In this document, teachers will find a number of value propositions that the project partners have identified as the most appropriate to start teaching students about cloud services. The focus has been on the convergence of industries today, so teachers will find a combination of best practices from different industries to provide their students with tailor-made solutions with maximum efficiency for their students. The topics of the Cloud teaching material content are as follows: 1. Introduction to Cloud Computing and Types of Cloud Computing, 2. Pricing vs Market Comparison between AWS, Azure and GCP, 3. Deploying Servers and Load Balancers on All Computing Platforms, 4. Storage Services on AWS, Azure and GCP, 5. Security Services - Identity and Access Management, 6. Types of Network Services and Setting Them, 7. Database Services on AWS, Azure and GCP, 8. Domain Setup and 9. Monitoring and Notification Service.

Below teacher can also find 61 practical examples of applications, suitable for teaching VET students about cloud technology. In appendix 1 a teacher can find one more detail example of applications and in appendix 2, a teacher can find Code Snippets for some of the applications below, which teachers can use as templates that make it easier for them to explain to VET students how to enter repeating code patterns.

2 CLOUD COMPUTING TRAINING MATERIALS

2.1 Introduction to cloud computing technologies and types of cloud computing

Difficulty Level: Easy

Completion Period: hours

Objectives:

After reading the material, the reader will understand the concept of cloud computing as it is perceived in IT technology and the main services it includes. You will also know the main advantages and disadvantages of cloud computing technologies.



Achievements

After completing this application, you will be able to:

- know the history of the term cloud computing,
- understand the meaning of the term cloud computing,
- know the services offered by cloud technologies,
- know the advantages and disadvantages of cloud computing technologies.

Starting with the first industrial revolution, human society as a whole evolved and scientific progress continued. Humanity has gone through three industrial revolutions, each with its own characteristics. The beginning of the third millennium is marked by the emergence of the fourth industrial revolution characterized by the large-scale use of industrial robots, artificial intelligence, and cloud computing technologies.

All these things bring profound transformations in people's activity and life. If the terms robots and artificial intelligence are somewhat suggestive and do not offer much ambiguity, the term cloud computing seems to be more of a jargon than a technical term. And yet this term has an important technical meaning for the IT industry.

The term cloud is actually a metaphor for the term Internet. Moreover, the icon related to the Internet is the representation of a cloud, and it means everything contained in Internet technology unseen by the user. In other words, the icon wants to express the fact that everything that belongs to the Internet is hidden in a nebula for the Internet user.




Figure 2.1. Suggestive image of the term cloud computing

Cloud computing technologies create a special impact on economic activities all over the world. Although the immediate meaning of this term would be the data storage service on a server at a company that has the technical ability to store data safely, the full term of cloud computing is broader. It has its origin in the moment when computer systems appeared.

Thus, according to the opinion of the majority of scientists and researchers, the concept of cloud computing would have been stated in a very simple form in 1955 when a computer scientist came up with the idea that certain computer resources should be shared between different users through rental because IT technologies at that time had exorbitant costs and many users could not afford to purchase them. This idea belongs to the researcher John McCarthy and is regarded as the beginning of the concept of cloud computing.

Fourteen years later another researcher J.C.R. Licklider developed the local computer network in the institution where he worked, which is now considered the ancestor of the Internet. The purpose of Licklider's network was to facilitate the exchange of IT resources (software and data) between researchers from the respective institution. McCarthy's concept of renting IT and network resources realized by J.C. R. Licklider for the exchange of IT resources led to the development of what we call today the Internet. Initially this was called Ethernet.

In 1972 the IBM company created the first mainframe computer VM/370 or Virtual Machine Facility/370. Any researcher or scientist could access the data stored on this system using a Hercules emulation program. If until the 80s of the last century computer technologies were accessible only to scientists, researchers, or large companies, but in the period 1980-1989 home computers appeared and the technologies used to create communication networks between computers were improved. The communication network was called Ethernet and was standardized. Some companies like Ms_Dos and Novel have made an important contribution to the improvement of communication networks between computers. The IT resources were hosted on servers that could be accessed from anywhere and by anyone who had a computing system connected to the computer network.

The Internet grew exponentially between the years 1990-1998. In 1996, a group of researchers from the Compaq Computer company introduced the concept of cloud computing for the first time. The launch of the Salesforce.com application in 1999 made it possible to sell information to collaborating companies or to store it through a web portal. This was the beginning of a period in which other companies began to offer the same services and contributed to the improvement of the Internet. The appearance on the computer products market of Web Services offered by Amazon was an important moment. This service offered data storage, access to programs and virtualization.



Between 2006 and 2012, the Google company consolidated its presence on the Internet services market by launching Google Apps. In 2011, the Apple company announced the launch of its own data storage solution on servers accessed via the Internet under the name Apple iCloud. A year later, the Google Drive application was launched by the Google company, which united all the facilities offered under a single service.

Between 2012-2017, cloud services were expanded and due to the appearance of high-performance mobile devices, cloud services were accessed by more and more users, which stimulated IT companies to improve the services offered. The research in the field of IT has led to the increase of the technical level of the networks for data transfer and thus the speed of the Internet has also increased.

Today, the term cloud is used more and more without knowing its true meaning in IT. The simplest definition of the term cloud computing is to have easy access to IT resources (programs and data) or to other services that are not installed on your own computer. For the home consumer, cloud services can mean access to electronic mail services, storing data in Google Drive or using specialized services for transferring large files that cannot be sent by email (e.g. Drop Box). It can also mean accessing movies, music or games via the Internet.

From the point of view of some small and medium-sized enterprises, cloud computing services can be defined by the safe storage of software applications and own data in locations located outside the company that can be easily accessed from anywhere and by anyone authorized by the company's management. This brings significant financial benefits to the company because it is not necessary for it to purchase its own equipment for data storage or software applications, nor is it necessary for the presence of specialists to manage specific IT activities.

In order to dispel ambiguities in the definition of the term cloud computing, the US National Institute of Standards and Technology (NIST) defined cloud computing services in 2011 as follows:

„Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications and services). It can be rapidly provisioned and released with minimal management effort or service provider interaction”.

NIST also specified five essential characteristics that cloud computing must have:

- on demand self-service;
- broad network access;
- resource pooling;
- rapid elasticity or expansion;
- measured service.



Cloud computing services can be provided by a company that works in the IT field or can be accessed by a company with a different IT profile, by individuals or by communities.

That is why NIST defined four types of cloud computing:

- public;
- private;
- community;
- hybrid;

Each of the four types of cloud computing specified above can offer the following basic services:

1. software, (Software As A Service) – (SAAS);
2. platform (Platform As A Service) – (PAAS);
3. infrastructure (Infrastructure As A Service) – (IAAS);

Within a cloud computing service that provides all three basic services listed above, they are structured as shown in the following image (see Figure 2.2 below):

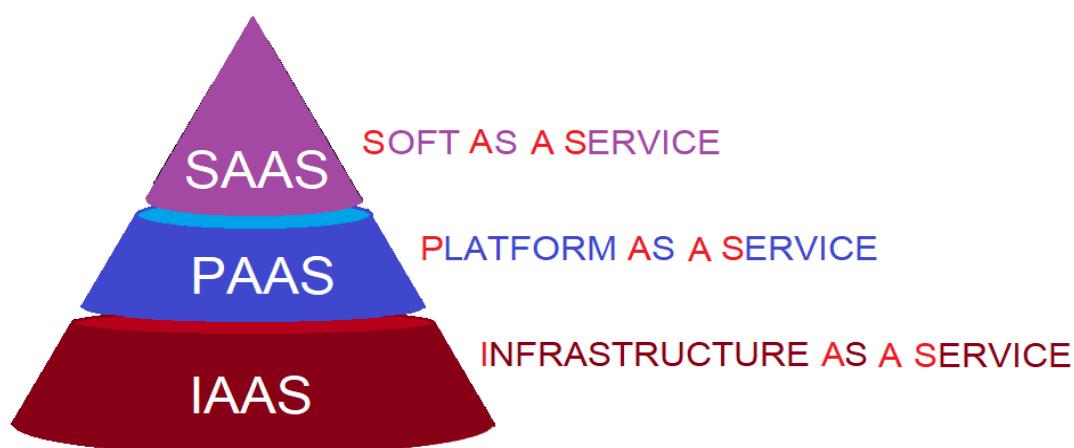


Figure 2.2. The hierarchy of the three basic levels in cloud computing services

In addition to the three services shown so far, IT companies also provide other services that have the following names and acronyms:

- 4 Gaming As A Service (GAAS).
- 5 Communications As A Service (CAAS).
- 6 Database As A Service (DBAAS).
- 7 Desktop As A Service (DAAS).
- 8 Hardware As A Service (HAAS).
- 9 Intity As A Service (IDAAS).

10 Storage As A Service (STAAS).

In the following we will explain the definition of each service:

- Software As A Service (SAAS) consists of providing services such as: Gmail, Youtube or other similar services to the user. By using these services sometimes an access fee is paid or they are free.
- Platform As A Service (PAAS) offers software developers a platform for writing codes for different applications and testing them on that platform.
- Infrastructure As A Service (IAAS) is a service that consists in subletting servers and networks to a company that can in turn offer them as services to other users.
- Gaming As A Service (GAAS) is a service provided by some companies through which users can access software that offers fun games in the virtual environment. This software can run on computers or mobile devices.
- Communications As A Service (CAAS) are messaging services, video conferencing for communities whose members are not in the same place or remote communications by voice or text. This category includes applications such as those offered by Skype, Facebook or Twitter.
- Database As A Service (DBAAS) means the provision of database services that involve the storage of data belonging to companies, communities or individuals on the servers of IT companies specialized in this regard. These data can be easily and safely accessed by the data owner. The data owner pays a rental fee for this service. The service is profitable because the creation and management of a specialized database for a certain field requires a special financial effort for many companies.
- Desktop As A Service (DAAS) is a service through which a certain person can use his computer by accessing it from another device when he is in another location. The process is called virtualization and allows accessing a computer under Windows, Mac or Linux operating systems through cloud technologies, using icons, shortcuts, etc. of the accessed computer.
- Hardware As A Service (HAAS) allows a company to rent hardware from a supplier. All hardware elements: computers, printers, mobile phones, tablets, etc. are the property of the supplier during the time they are used by the company that rented them. This service is considered to be part of cloud technology, although it seems different from other technology-specific services.
- Identity As A Service (IDAAS) ensures secure access to IT resources through software that identifies the fingerprint or iris detection of the person who wants access to the data. In addition to these elements used for detection, there may be other procedures for verifying the identity of the person requesting access to the stored data.
- Storage As A Service (STAAS) Google drive and Dropbox are two examples of this type of service. In principle, this service allows the storage of data belonging to the employees of a company or individuals. These data can be accessed at any time by their owner, guaranteeing data security.

As it is mentioned above, the term cloud is considered a metaphor for the Internet. Over time, in addition to the communication services between computers known as the Internet, (which required



the existence of a network and specialized software), companies created several facilities that later became services. The term cloud can refer to the local internet networks of some companies that provide IT services in a geographical area or to the entire internet network spread across the globe. As a result, we can talk about a local cloud and a general cloud. On the world market there are four giant companies that offer cloud computing services. These are: Microsoft with the Microsoft OneDrive service, Amazon that offers the Cloud Services service, Apple offers the iCloud service and Google that offers Gmail, Drive, etc. services. In addition to these are the Dropbox Cloud services.

An IT company can choose to create its own local cloud system that it can rent out to end users who can be individuals, local communities, or companies with a different activity profile than IT.

A cloud computing system consists of:

- **The local Internet network of one or more users.** All computers, printers and other hardware components of a user are connected to one or more local switches.
- **The router is the device** through which the user's switch is connected to the Internet network of an ISP (Internet Service Provider).
- **A portal or a website** ensures the connection to the company's cloud service that has its own local servers and servers. In the company's cloud servers, data can be stored, or software applications can be run. Communication between the company's servers and the user is done through a frontend portal. All the company's cloud servers are interconnected and form a cluster. Server clusters can be located anywhere in the world and can be located in different places at great distances from each other.

The company that owns the server clusters ensures secure user access, maintains the database and updates the software programs offered to clients. A simpler definition of a cloud service is a data centre in which hundreds of servers are interconnected that offer the possibility of storing and running some software to which companies or individuals have free or paid access. In addition to the possibility of storing data or running application software, cloud services can also offer some of the services listed above.





Figure 2.3. Aspects inside a data centre that provides cloud computing services

The use of cloud computing services offered by an IT company has the following advantages:

- ✓ **Easy access from anywhere in the world.** The data stored on the server can be accessed from anywhere in the world by the person who owns the data. The condition is that the person can have access to the Internet and possess a device through which they can access the Internet.
- ✓ **Reducing the company's costs** - because the company does not have to invest in the purchase of hardware equipment and hire IT specialists to create software and manage databases. Many times, the investments in hardware and software equipment are greater than the benefits that a company with a different activity profile realizes from them.
- ✓ **Flexibility** - characterizes the fact that the features of the software or user interfaces can be easily changed according to the client's wishes. This can lead to business performance.
- ✓ **Permanent updating of IT technologies.** The IT technologies related to the databases but also to the software used for the transfer of the databases are in continuous progress. Cloud computing service providers acquire new technologies in order to keep pace with technical progress. Thus, the user of cloud services can benefit from the latest news in the field of these technologies.
- ✓ **Data protection in the case of natural** disasters that would affect the owner company. Data and software applications used by a company, or a community can be lost if a fire or a natural disaster affects the company that owns the data if it is stored on servers or other local devices. Since the company's data is stored on servers located at a distance from the owner company, its data is safe.
- ✓ **Collaboration between employees of a company or several companies** through access to programs or data. Employees of a company who collaborate on a project can easily access the same data that is stored on the server much more easily.
- ✓ **Data security.** Access to data or programs stored on the servers of an IT company is secured and based on access passwords. If the data stored on the server were kept on a local storage system, CD-ROM, stick or even in a laptop, its loss or theft leads to the irreparable loss of data. Also, the company that offers the cloud services takes strict measures to stop the access of unauthorized persons to the stored data.

Although cloud computing services have advantages that recommend them to be used on a large scale, these services also have some disadvantages. These disadvantages are:

- ❖ Updating the software that manages the operation of the servers can lead to the loss of stored data. An example is the one from 2011 when the Amazon company lost its customers' data.
- ❖ The lack of connection to the Internet is a major disadvantage, for the person who is in a place where he does not have access to the Internet, he cannot use the data from the server.
- ❖ In the case of some companies that offer cloud services, the expenses can increase and force the company to suspend the services offered to clients.
- ❖ Inability to access a company's server even if the Internet connection is possible. It happened in the past even at well-known companies where the message HTTP Error 503: The server is unavailable appeared. Fortunately, this rarely happens.
- ❖ Government access to personal or company data. Governments can force cloud computing companies to give them access to data stored on their servers in order to obtain confidential information about citizens or companies that have data stored on those servers. In order to maintain the secrecy of the stored data, some companies have moved their servers to the territories of other states and thus come out of the jurisdiction of the state that requests access to the data stored on their servers.
- ❖ Servers can be attacked by Hackers. In this case, data security is at risk. There have been situations in which famous people complained that their personal data stored on the servers of cloud companies was stolen.

Despite all the disadvantages listed above, cloud computing services are increasingly used throughout the world and many companies in the IT field invest to increase the quality of cloud computing services.

2.2 Pricing vs. Market Comparison between AWS, Azure and GCP

Difficulty Level: **Easy**

Completion Period: **45 minutes per unit, 4 units in Module**

Objectives:

Cloud computing is one of the hottest buzzwords in the IT industry right now, as cloud providers offer the advantages of easy setup, high scalability, and affordability everywhere.

The following units in this module will familiarize you with the top cloud providers available on the market today. Amazon Web Services (AWS), Google (GCP), and Microsoft (Azure) are the best-known public cloud providers and hold billions of dollars of market share in cloud computing.

As we progress through the units, we will be moving from a general overview of these 3 providers to a focused analysis on what they offer for the price. It is a common reality that cutting-edge cloud solutions come with a price, which is no different for these big three providers—AWS, Azure, and



Google—as we will analyse how their prices vary according to their plans, service selection, features, discount options, resource usage, and more.

Achievements

After completing this module, you will be able to:

- understand the demand for cloud computing platforms
- recognize their influence in the business management sector, among others
- recognize some of the similarities and differences between the cloud platforms from a technical perspective
- learn about the market presence of the 3 platforms, in comparison to one another
- learn how pricing options are established and how they are related to market demand after 2019.

2.2.1 What Does Cloud Computing Offer?

Why would you turn to a cloud platform for your needs? This unit looks at what cloud computing provides for someone like yourself, hoping to manage a business, or looking for any type of IT assistance.

Let's talk basics.

These platforms are similar in key factors as to why they are dominating the market, yet each offers different resources when it comes to computing, networking, and storage options.

It is clear that when you are looking for the best cloud computing platform for your business it's important to keep track of your goals, expected growth, and budget.

What does cloud computing offer?

Let's look at some of the main reasons cloud computing is great for what you need when help managing a business:

- **Reduced IT costs:** Cloud implementations enable you to pay only for computing capacity according to your business needs, reducing the ongoing costs of purchasing, deploying, maintaining, and managing on-premises infrastructure.
- **Faster time to market:** The cloud is enabled in minutes. No waiting to get started.
- **High scalability and flexibility :** Cloud implementations can automatically scale workloads in response to changing market demands.
- **Improvement in business reliability:** Implementing data backup and disaster recovery in the cloud is typically much easier, less expensive, and less disruptive than on-premise which is risky and time-consuming.



- Continuous Performance Improvements: Because it is in real time, cloud infrastructure is regularly updated with the latest and the most powerful computing, storage, and networking hardware.
- Ensure security measures : Easily meet basic security and compliance requirements with the most flexible and secure cloud environment available today.

This image below shows how using the cloud reduces the cost of IT overall in managing a business and why it is so appealing to users:



Figure 2.4. Business Benefits of Cloud Implementation

Here are some basic definitions of the three providers:

What is the AWS Cloud Platform?

AWS, or Amazon Web Services, is a cloud services platform from Amazon that provides computing, storage, delivery, and other services to users. Taken together, all of these SaaS (Software-as-a-Service), Infrastructure-as-a-Service (IaaS), and Platform-as-a-Service (PaaS) offerings can be effectively used by you, as they offer the following features:

- 18,000+ services
- Computing

- Storage solutions
- Cloud app integration
- Analytics and machine learning
- Productivity tools
- Developer and management tools

Amazon Web Services is the most popular storage service for object archives, which is a main reason why it dominates the current cloud market. It consists of tools for IoT, security, database, management, analytics, enterprise applications, and more.

From Amazon come three separate tiers of developer support, business support, and enterprise support, offering a combination of tools, cloud technology, and experts.

Many of AWS' strengths relate to its position as a premier provider of modern cloud services and the sheer scale of its global operations. Taken together, these factors fuelled AWS's growth and enabled the company to offer a large list of non-stop services to businesses around the world.

Here are some of the strengths of AWS:

- Supports all major operating systems including macOS (unlike other vendors)
- Offers wide range of services
- Continued growth of service offerings
- Sophisticated and readily available
- Can handle large numbers of end users and resources
- Very easy to access and start

Here are some of the disadvantages:

- A relatively high cost
- Additional charges for essential services
- Additional cost for Customer Technical Support
- Steep learning curve after engaging the platform

Microsoft Azure

As there is also an integrated platform on the cloud that offers storage, the same database opportunities, and computing that Amazon does, it also has various cloud types that cater to specific requirements. It is one of the best options in the cloud for companies that need a high amount of data storage space, with options such as Data Lake Storage and Queue Storage. Bulk storage is ideal for businesses with a large amount of unstructured data, whereas file storage is ideal for businesses with specific file storage requirements. Azure takes the base from current Microsoft office suite software another business tools to offer the following features, in a configured format-



- A development platform on the cloud
- Blockchain tech
- Predictive software
- IoT integration tools

An important feature of Azure, just like Amazon, is the tiered approach to support services, which includes a developer plan offering unlimited support during business hours, and the standard plan, which also includes unlimited access. For more structured support for businesses, the professional plan for the cloud is the best option.

Users enjoy the particular features of Azure due to its:

- Widespread availability
- Service contract vouchers for Microsoft cloud computing users
- Intuitive configuration with the Microsoft family of software
- Built-in apps that support multiple languages (including Java, Python, .NET, and PHP)

Some of the problems that may be encountered include:

- Inadequate data management
- Reports of core network difficulties
- Some people believe it is more difficult to master than other platforms
- The design may appear less professional than on other platforms
- Reported technical support issues

Google Cloud Platform (GCP)

Due to its endless IT expertise and internal research, Google has proven to be a market contender. It features many hosted services such as Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) for computing, storage, and application development.

Google was first made public in 2004, but it has only recently begun to pose a serious threat to both AWS and Azure.

GCP is quickly catching up to the competition thanks to Google's extensive global presence and seemingly limitless capacity for innovation.

At the moment, it provides services like:

- Managing productivity in businesses and other areas
- storing data
- cloud application development studio
- Engines for AI and machine learning, such as cloud speech API, vision API, and others



- Business analytics and other supplementary components

In contrast to the other two services, Google's storage options are fairly simple, with cloud storage and persistent disk storage rounding out the list. In addition to its own internal transfer service, Google also gives users access to a growing number of online transfer services. Sadly, Google's backup options—Nearline backup for frequently accessed data and Coldline backup for rarely accessed data—are also rather basic.

Several outstanding features provided by GCP include:

- high degree of scalability
- straightforward configuration and installation
- use of widely used programming languages like Python and Java
- Reasonable long-term savings
- Data load balancing and quick responses

Disadvantages include the following:

- Inadequate advanced features
- Less variation in features
- Fewer service options
- There are fewer global data centres

Questions to Consider

1. What is a cloud platform, and what advantages does it provide?
2. Name 3 of the business benefits from the cloud in the graph and why they appeal to you.
3. Who would you choose as your provider and why?
4. Check the source, titled Techfunnel (2022) and respond with what you learned.

2.2.2 The 3 key players on the Market

Three major cloud service providers control the majority of the market in 2021, accounting for 64% of the total market share. As seen in the chart below, AWS holds the top spot with a 33% market share, followed closely by Azure with 21% and Google Cloud with 10%, as seen in the figure (chart) below.






Figure 2.5. Cloud Provider Market Share Trend

Due to these cloud providers' extensive global network, these numbers can be explained. With a market that is only getting bigger, Amazon is an interesting case because its market share has stabilized at about 33%. In other words, over the past few years, AWS cloud revenues have been rising steadily. Although the competition is getting stronger, AWS has been selling its cloud products for 11 years and continues to be the market leader. Others imitate what Amazon does when it adopts a new technology or business strategy.

According to Jeff Bezos, CEO of AWS, "AWS had the unusual advantage of a seven-year head start before facing like-minded competition." Because of this, AWS services are by far the most advanced and functional. AWS reported a \$62 billion revenue and \$18.5 billion net profit in 2021. In comparison to last year's turnover, it represents a 38% increase.

But here's the catch: Microsoft is unquestionably AWS's biggest competitor; the Intelligent Cloud Division of Microsoft generated 60 billion dollars in revenue last year, which is very close to AWS's revenue, but here's the catch: this division also includes many other services, including Microsoft Azure, GitHub, Windows Server, Microsoft SQL Server, and other versions of those products. The intelligent cloud division's revenue increased by 24% from 2020 to 2021.

Google Cloud is the third largest cloud provider after AWS and Azure.

Its revenue grew from \$13B in 2020 to \$19B in 2021. Google Cloud operating loss decreased \$2.5 billion from 2020 to 2021. The decrease in operating loss was primarily driven by growth in revenues.

Like Microsoft Azure, the Google cloud division also incorporates feedback from other places, like Google workspace. The previous years saw significant investments made by Google Cloud, which resulted in operating losses, in order to catch up to AWS and Azure. In year 2022, Ruth Porat, the CFO of Google and Alphabet, predicted this in the following way: "Looking forward, we will continue to focus on revenue growth driven by ongoing investment in products and the go-to-market organization... Scale will eventually reduce operating loss and improve operating margin."

Here is a quick look at some key aspects of each:

Azure Virtual Network: Azure is currently accessible in 54 regions worldwide and keeps as much traffic as possible inside the Azure network rather than over the internet. In the end, it is a networking solution that performs better than even AWS is quick and secure. Additionally, because the Azure Virtual Network is so flexible, businesses can use a hybrid networking strategy or bring their own IP addresses and DNS servers.

Amazon Direct Connect: To guarantee consistent service and dependable performance at all times, Amazon has created a comprehensive global framework centered around 114 edge locations, 14 data centers, and 22 different global regions. As a result, AWS is able to provide quick cloud deployment models, quick delivery, and instantaneous response times for its broad range of services. In particular, its 802.1q VLANs, which are industry standards, enable a dedicated connection between private networks and AWS via any of the numerous Direct Connect Locations.

GCP: Despite not having the same scope as the other two providers, Google's renowned innovation capabilities support the Google Cloud Platform. In addition to a huge number of data centers located all over the world, Google currently has 21 regions and is continuously adding more with the addition of undersea cabling. Hybrid connectivity products like Cloud Interconnect and Cloud VPN enable you to establish secure direct connections or IPsec VPN connections.

In order to understand the cloud market share for each of the three major providers, you should also be familiar with each company's current share figures:

AWS: With a 32% market share overall, Amazon rules the global market. It actually outperformed the other two most popular cloud platforms in terms of revenue, bringing in a respectable \$11.6 billion and experiencing a 29% growth rate this quarter.

Azure: With Azure, which holds a 19% market share, Microsoft has a sizeable market share. Microsoft reported a growth rate of 48% over the previous quarter, despite the fact that it does not publicly disclose Azure's revenue figures.



Google Cloud Platform: GCP is still rapidly expanding and is currently in third place with a 7% market share. Its growth is actually 45% year over year, with \$3.44 billion in total revenue this quarter.

Following the pandemic that accelerated the adoption of cloud computing over the last two years, we can see that the numbers continue to rise, and that the crisis was more of a long-term booster for the cloud market than a short-term effect.

It was recently noted that businesses that embraced cloud computing in recent years have increased their usage and are now moving more and more toward multi cloud strategies. The Flexera State of the Cloud 2022 report has also shown that businesses are investing an increasing amount of money in these technologies and that new issues like security, multi cloud management, and Kubernetes adoption are emerging as a result. Since the stakes are always higher, it is crucial for businesses to better understand and utilize resources as efficiently as possible.

Companies are making significant investments on a global scale. Spending on public clouds will increase from \$408 billion in 2021 to \$474 billion by the end of 2022, according to Gartner's forecast.

Questions to Consider

1. What are the current market share percentages among the providers?
2. Account for the differences among providers in cloud market share.
3. Why would market share increase or decrease: name some factors and indicate what some of your future predictions may be for the 3 providers.

2.2.3 The Cloud Market Share Comparison

To better understand how the market fares globally, let's look at global market shares held by the big 3 in the following major markets: in the United States, in Europe, and in China.

The US cloud market

It should not be surprising that the US cloud market, which accounts for 44% of all global spending, is by far the largest. The top three cloud service providers still have the same market share: AWS has 37%, Azure has 23%, and GCP has 9%. AWS, Azure, and Google Cloud opened new data centres in the US in 2021. Microsoft Azure, for instance, began operating in Georgia and Arizona in 2021, and this number will keep increasing as they recently announced plans to build 50 to 100 new data centres each year throughout the world. From the figure below we can see the costs of cloud infrastructure services for Q1 2021 compared to years 2019 and 2020.



Figure 2.6. The costs of cloud infrastructure services for Q1 2021 in USA compared to years 2019 and 2020

The US cloud market is by far the largest and accounts for 44 % of total spending, which is not surprising. On the graph above, you can see significant growth spikes (38 %) during the COVID crisis and more recently, a growth of 29% in Q1 2021 to reach a record of \$18,6B.

Cloud Market in Europe

Although it has increased during the Covid era, the European cloud market is still only the third largest after the US and China.

National cloud service providers like Deutsche Telekom, OVH, Scaleway, Orange, and various national telcos are available in the European market. These providers are competing with the top three cloud service providers in the world, AWS, Azure, and GCP, which now control 66% of the market, up from 50% three years ago.

The European cloud market is anticipated to grow very strongly in the coming years, with new data centres sprouting up all over the continent, despite its tardiness relative to other significant regions. By 2030, according to various projections, the European market will be worth more than \$300 billion, which would be equal to the size of the global market today.

The Cloud Market in China

The Chinese cloud market is still growing at a rate twice as fast as that of the US (60 vs. 30%), outpacing the rest of the world. China accounted for 14% of the global cloud market in Q2 2021, with cloud infrastructure spending that exceeded 6 billion dollars, as seen in the figure (chart) below.



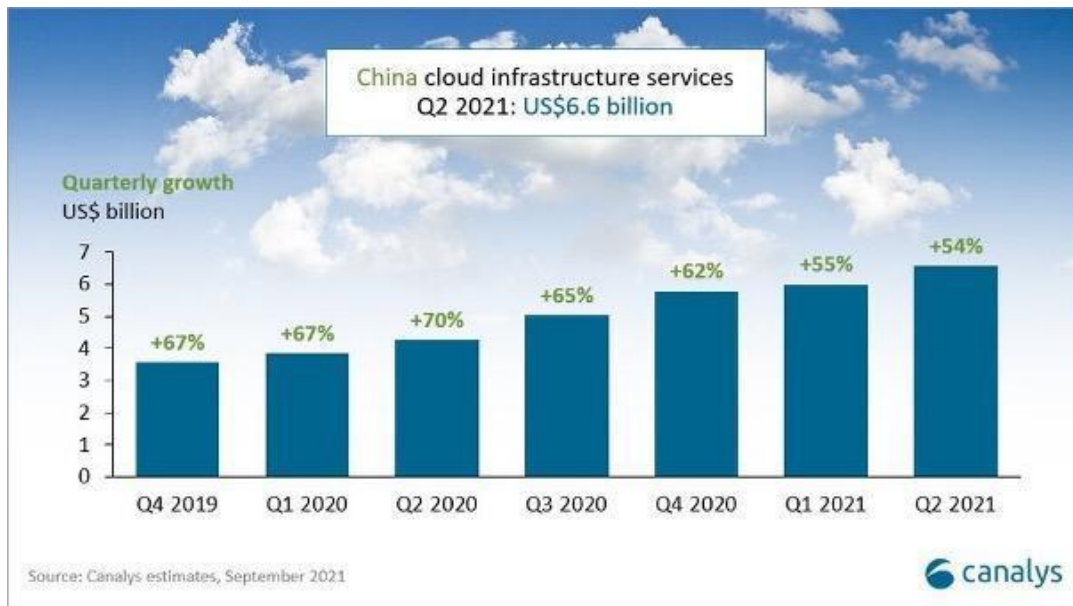


Figure 2.7. The costs of cloud infrastructure services for Q1 2021 in China compared to years 2019 and 2020

The pandemic has accelerated growth in the same way it has in other market regions; in Q2 2020, growth reached a peak of 70%. There are other underlying reasons for this fast growth: China is the only significant economy that reported economic growth for 2020, with a GDP growth of 2.6%. The Chinese government has made cloud computing a top priority through its “internet plus” strategy in 2015.

The government is promoting and subsidizing the cloud industry. Chinese tech giants such as Alibaba, Tencent, Baidu, Huawei offer cloud solutions and can compete against their American rivals as they have an equivalent size. Alibaba, Huawei Cloud, Tencent, and Baidu AI Cloud, which together account for more than 80% of total expenditure, are the primary cloud providers in the Chinese cloud market. American companies struggle because of laws that favour Chinese businesses.

Chinese cloud service providers are now aiming to grow in Europe, Asia, and developing nations. We can anticipate a digital competition between the US and China, similar to the 5G network.

Questions to Consider

1. What are the key statistics of the US Cloud market?
2. How much is the global market predicted to be worth in Europe by 2030?
3. Name two ways that China can hope to challenge the US market, and do you predict that it will be successful in its ambitions?

2.2.4 Analysis of Pricing Structures

To understand pricing structures, it is key to know that each of the three main platforms has two things in common: a free tier with very few options and a per-hour or minute on-demand pricing model for all resources. Comparing prices can be challenging because they can differ significantly depending on resource usage, service preferences, and other factors.

In general, a pricing war is always in play among the top three: by lowering their prices, Microsoft and Google attempt to challenge AWS. Users of AWS services pay only for what they use, with no additional fees or termination charges due after the service has been completed. This is known as a pay-as-you-go model.

Here are key features of the pricing models of all three providers:

Pricing for AWS. It has been stated that the pricing structure provided by Amazon is "so complex, you'll need a third-party app to manage it." Amazon does, however, provide a 12-month period of 750 hours per month of EC2 services as part of its free tier, as well as up to a 75% discount for a 1–3-year commitment.

- high cost in comparison
- Additional charges for necessary services
- Customer technical support charges are additionally charged.

The pricing structure within AWS is so complex that you need a third-party app to manage all of these services. The minimum instance, which has 2 virtual CPUs and 8 GB of RAM, will run you about USD 69 per month, while the maximum instance, which has 128 virtual CPUs and 3.84 TB of RAM, will cost you about USD 3.97 per hour.

Azure Pricing. Azure users frequently use a third-party app to manage costs because it is complex in a way similar to AWS. Similar to AWS, Azure offers a free tier that allows users to use 750 hours of virtual machines per month for 12 months in exchange for a steep discount if they commit to a one- to three-year period.

- Discounts on service agreements for users of Microsoft's cloud computing services
- Affordable on-demand prices
- The use of high redundancies to cut downtime.

A number of variables, including the location, required capacity, and management level, affect the price of Azure. It also offers a free tier, which permits free use of some models for the first 12 months only, as well as free use of some models for ever.



Pay-as-you-go pricing is an option with Azure, just like with AWS. It also offers a different way to pre-pay for its service, which it refers to as "Reserved Instance" (upfront commitment). Additionally, it provides spot instances, allowing customers to buy virtual machines (VMs) from Azure's excess capacity at a discount.

Users can start or stop the service as needed and only pay for the seconds they actually use when using the pay-as-you-go method. The Reserved Instance, on the other hand, is designed for continuous use and is based on the cost for an entire month (730 hours), while the pay-as-you-go model also relies on the 730-hours analysis, according to the pricing calculator. Microsoft Azure allows a wide range of services such as computing, networking, storage, and analytics. Therefore, its pricing model depends on various factors, including the capacity required, the location, the type of service, and the management level.

Google Pricing. It is evident that Google made an effort to learn from the mistakes of its rivals and adopted a cost per second model that is fairly simple. In addition, GCP offers a \$300 credit for a year of service, one free micro-instance per month for the first year of its free tier, and a 30% discount for continued use.

It offers a number of pricing options, such as pay-as-you-go pricing, long-term reservations, and free tier options. The cost of Google Cloud is also influenced by a number of factors, including pricing for compute, SQL, networks, storage, and serverless. You should consider these factors when choosing a cost structure for any business.

Google offers its customers USD 300 credit for free as customers can spend their amount on their Google Cloud products. Users can also make use of a variety of free products, including the most popular cloud services currently available on the market for computing, storage, databases, IoT, and artificial intelligence. Additionally, the US tech giant offers significant discounts for products that are "committed to use," or used at a specific level for one or three years in advance.

Google offers its users a special choice known as "Sustained Use discounts." If you use the services for a certain percentage of the month, this offer will automatically be applied on a sliding scale. Additionally, you are not required to make any upfront payments or sign any commitments in order to combine non-overlapping instances and receive the benefits of a percentage discount up to the maximum level.

Here is a chart showing the price comparisons among the platforms:



AWS Vs. Azure Vs. GCP Cloud Cost Comparison

Detail	Amazon AWS	Microsoft Azure	Google GCP
Minimum Instance	2 virtual CPUs, and 8 GB of Ram will price you around – USD 69/month	2 virtual CPUs, and 8 GB of Ram will price you around – USD 70/month	2 virtual CPUs, and 8 GB of Ram will price you around – USD 52/month
Maximum Instance	3.84 TB Ram, 128 vCPUs will price you around – USD 3.97/hour	3.89 TB Ram, 128 v CPUs will price you around – USD 6.97/hour	3.75 TB Ram, 160 v CPUs will price you around – USD 5.32/hour
Type of Discount	Reserved Instances (RIs)	Reserved Instances (RIs)	Committed Use Discount (CUD) Sustained Use Discount (SUD)
Commitment	1 or 3 years	1 or 3 years	Committed Use Discount (CUD) – 1 or 3 years Sustained Use Discount (SUD) – no commitment
Discount percentage	Up to 75 percent	Up to 72 percent	Committed Use Discount (CUD) – for 1 year up to 37 percent or 3 years up to 55 percent Sustained Use Discount (SUD) – up to 30 percent
Is cancellation available?	Yes, it offers to sell your products on the marketplace	Yes, they will charge a 12% cancellation fee	No cancellation is available
Payment options	3 options are available on AWS – no up-front, partial up-front, all up-front	All up-front	No up-front
High Profile Customers	LinkedIn, Facebook, BBC, Airbnb, Twitch, Netflix, Adobe, ESPN, Lamborghini, etc.	Apple, HP, Coca-Cola, LG Electronics, Verizon, Xbox, Fujifilm, etc.	Twitter, Intel, Yahoo, PayPal, eBay, Target, 20th Century Fox, etc.

Figure 2.8. AWS Vs. Azure Vs. GCP Cloud Cost Comparison

Questions to Consider

1. What is the most appealing price structure to you and why?
2. Why is it difficult to make a direct price comparison among the competitors?
3. Name the two aspects all 3 competitors share and consider how you might sell the platforms based on the differences among them.

2.3 Selecting and setting the infrastructure

2.3.1 Deploying Servers and Load Balancers on all Computing Platforms

In this unit, we will look at the roles of Load Balancing, which is a method to help a network avoid annoying downtime and deliver optimal performance to users by processing tasks and directing

sessions on different servers. This is done differently in different cloud networks. In this unit, we will look at the main 3: AWS, Azure and Google Cloud Services.

What is a Load Balancer?

A load balancer divides user traffic between multiple instances of your applications. Load balancing reduces the likelihood of performance issues in your applications by spreading the load. Cloud Load Balancing is a software-defined, fully distributed managed service. Because it is not hardware-based, you are not required to manage a physical load balancing infrastructure.

Load balancers are classified according to their platform, and here we will compare the platforms with some of their key load balancers and graphs illustrating the cases:

Amazon Web Services (AWS)

Elastic Load Balancing (ELB) distributes incoming application traffic across multiple targets and virtual appliances in one or more Availability Zones automatically (AZs). An Application Load Balancer makes application layer (HTTP/HTTPS) routing decisions, supports path-based routing, and can route requests to one or more ports on each container instance in your cluster. Dynamic host port mapping is supported by Application Load Balancers. Below is a figure (graph) outlining the Application Load Balancer for AWS.

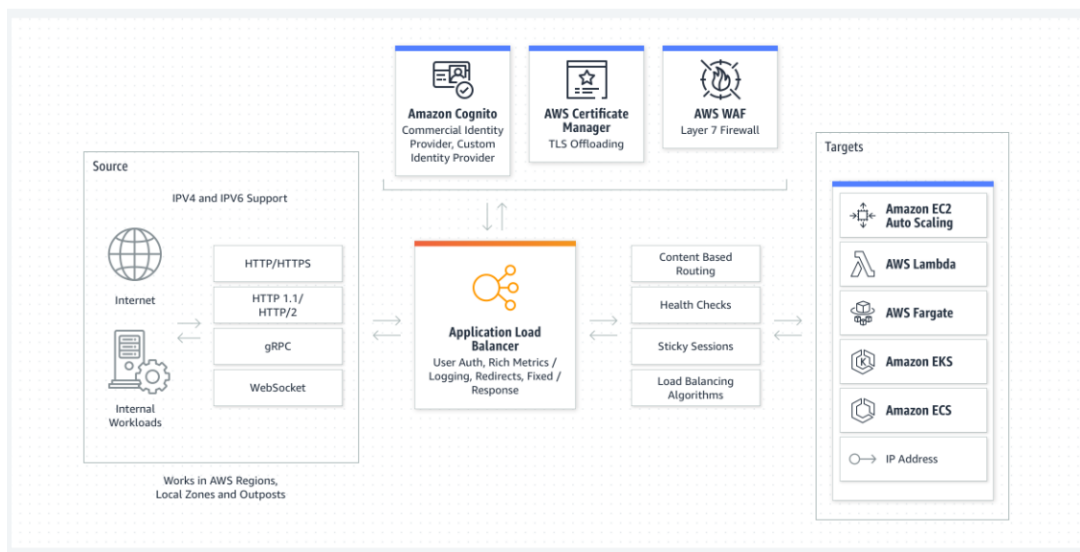


Figure 2.9. Application Load Balancer for AWS

A Network Load Balancer makes routing decisions at the transport layer (TCP/SSL). It can process millions of requests per second. When a connection is received, the load balancer employs a flow hash routing algorithm to select a target from the target group for the default rule. It attempts to establish

a TCP connection to the selected target on the port specified in the listener configuration. It sends the request with the headers unchanged. When the load balancer receives a connection, it uses a flow hash routing algorithm to select a target from the target group for the default rule. Requests are seen as coming from the Network Load Balancer's private IP address when configured with IP addresses as targets. This means that once you allow incoming requests and health checks in the target's security group, services behind a Network Load Balancer are effectively open to the world (as seen from the figure below).

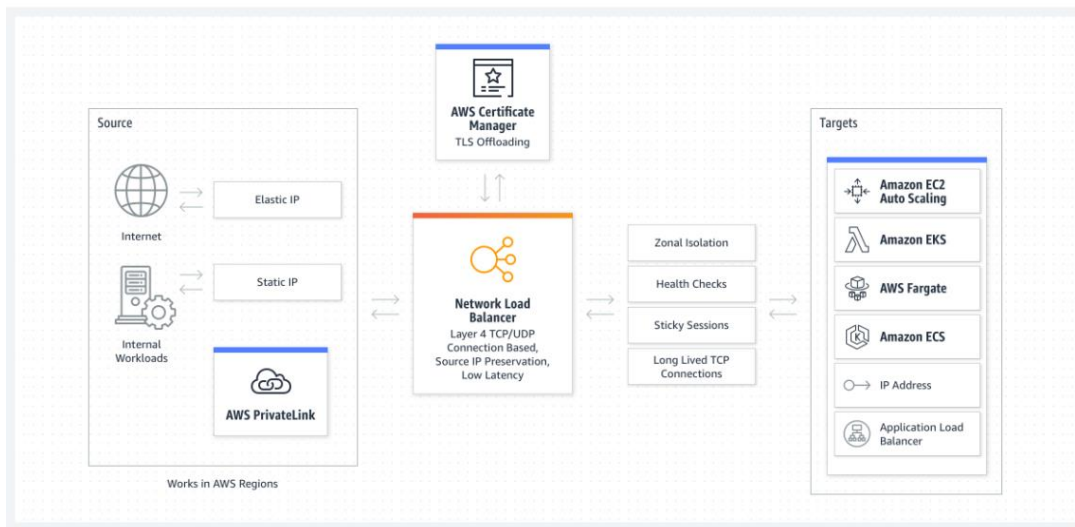


Figure 2.10. Network Load Balancer

Azure

An Azure load balancer is used to distribute traffic loads to backend virtual machines or virtual machine scale sets. You can use a load balancer more flexibly by defining your own load balancing rules. The process of evenly distributing load (incoming network traffic) across a group of backend resources or servers is referred to as load balancing. You can use Azure load balancer to distribute traffic to your backend virtual machines. An Azure load balancer ensures that your application is always available. The Azure load balancer is a self-managed service.

Outbound connections for virtual machines (VMs) within your virtual network can be provided by a public load balancer. These connections are made possible by converting private IP addresses to public IP addresses. Public Load Balancers are used to deliver balanced internet traffic to your virtual machines. When only private IPs are required at the frontend, an internal (or private) load balancer is used. Internal load balancers help to balance traffic within a virtual network. In a hybrid scenario, a load balancer frontend can be accessed via an on-premises network. The Load Balancer are presented below in figure 2.11.

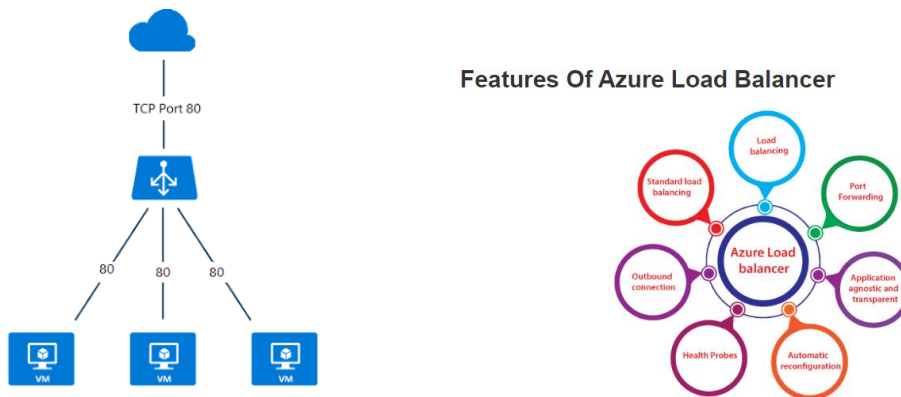


Figure 2.11. Load Balancers


Some of the key scenarios that an Azure does through a Standard Load Balancer include:

- Direct Internal and external traffic to Azure virtual machines
- Distribute resources within and across zones to increase availability.
- Monitor load-balanced resources with health probes.
- Through Azure Monitor, provides multidimensional metrics.

GCS Cloud Load Balancing is built on the same infrastructure that powers Google's frontend. It can handle 1 million or more queries per second while maintaining consistent high performance and low latency. Cloud Load Balancing traffic enters through 80+ distinct global load balancing locations, maximizing the distance travelled on Google's fast private network backbone. You can serve content as close to your users as possible by using Cloud Load Balancing (figure 2.12 below).

Summary of Google Cloud load balancers

The following diagram summarizes the available Cloud Load Balancing products.



Cloud Load Balancing overview (click to enlarge)

Figure 2.12. Choosing a Cloud Load Balancer

To choose a Cloud Load Balancing product, you must first determine what type of traffic your load balancers must handle, as well as whether you require global or regional load balancing, external or

internal load balancing, and proxy or pass-through load balancing. Cloud Load Balancing can load-balance traffic to endpoints other than Google Cloud, such as on-premises data centres and other public clouds accessible via hybrid connectivity.

The figure (diagram) below depicts a hybrid deployment with an external global HTTP(S) load balancer.

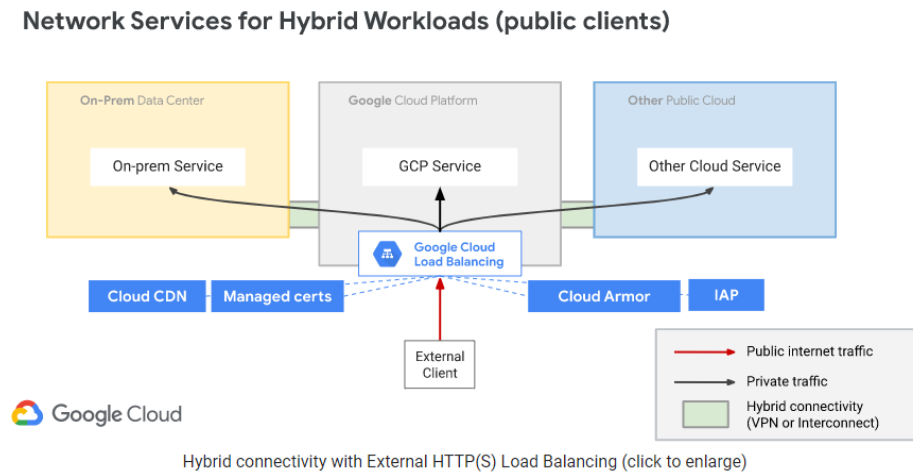


Figure 2.13. Hybrid deployment with an external global HTTP(S) load balancer

A GCS network load balancer can accept traffic from

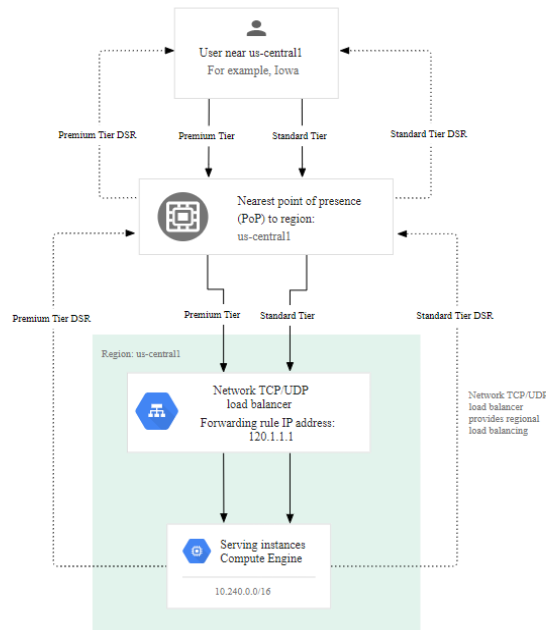
- any internet client.
- Google Cloud VMs with external IPs
- Google Cloud VMs that have internet access through Cloud NAT or instance-based NAT

The following are the characteristics of network load balancing in GCS:

- A managed service is network load balancing.
- Andromeda virtual networking and Google Maglev are used to implement network load balancing.
- Load balancers on networks are not proxies.
- Backend VMs receive load-balanced packets with the source and destination IP addresses, protocol, and, if the protocol is port-based, the source and destination ports unchanged.
- The backend VMs terminate load-balanced connections.

Below, you can find an example of a Network Load Balancer in a user case:

In the following diagram, traffic is routed from a user in Iowa to the network load balancer in us-central1 (forwarding rule IP address 120.1.1.1).



Network Load Balancing example for a user in Iowa (click to enlarge)

Figure 2.14. Network Load Balancer in a user case

Questions to Consider:

1. Why should you use a load balancer?
2. Name one feature that is useful from each cloud platform to consider
3. Fill in the blanks in this statement: Cloud Load Balancing is a _____, _____ managed service. Because it is not _____, you are not required to manage a physical load balancing infrastructure.
4. Name two of the key scenarios that an Azure does through a Standard Load Balancer.

2.3.2 Storage Services in the cloud

Three of the biggest cloud providers Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure (Azure) all offer three main types of storage on their services. Object storage, also known as blob storage in Microsoft Azure, Block storage and File storage, all having their pros and cons and different use cases.

For Object/Blob storage the three main services are AWS' Simple Storage Service (S3), Google's Cloud Storage and Microsoft's Azure Blobs. These three all do mostly the same things, with some variations

in the policies and storage tiers they offer, and the price points they have for storage per GB and for accessing the files.

All three providers have at least three generalised storage tiers categorised into what is called Hot, Cool, and Cold storage. These names indicate how often the data that is being kept in the storage is accessed.

Hot storage is for data that would be accessed frequently and with as low latency as possible. An example of the kind of data that should be stored in Hot storage would be product images in a e-commerce shop. Customers want to be able to see photographs of the items in the store with as low latency as possible, not having to wait for the website to fetch and load the image in their browser.

Cool storage is for data that should be accessed infrequently. An example of Cool storage would be a aggregate sales report. The data in the report is accessed maybe just once a month to update with data for the previous month, otherwise access is minimal. It is much cheaper to store data in a Cool storage tier than in hot data but comes at the expense of a much higher price for accessing the data, and with a minimum storage time.

The S3 standard storage tier with a price of close to twice as much as the Infrequent Access level as seen from figure below.

S3 Standard - General purpose storage for any type of data, typically used for frequently accessed data	
First 50 TB / Month	\$0.024 per GB
Next 450 TB / Month	\$0.023 per GB
Over 500 TB / Month	\$0.022 per GB

Figure 2.15. A price comparison between Hot storage and Cool storage with AWS S3

S3 Infrequent Access offering a very low per GB price.

S3 Standard - Infrequent Access** - For long lived but infrequently accessed data that needs millisecond access	
All Storage / Month	\$0.0131 per GB
S3 One Zone - Infrequent Access** - For re-createable infrequently accessed data that needs millisecond access	
All Storage / Month	\$0.01048 per GB

Figure 2.16. Infrequent Access prices

But the S3 standard tier offers a much lower price for accessing the data stored in the buckets.

	PUT, COPY, POST, LIST requests (per 1,000 requests)	GET, SELECT, and all other requests (per 1,000 requests)	Lifecycle Transition requests into (per 1,000 requests)	Data Retrieval requests (per 1,000 requests)	Data retrievals (per GB)
S3 Standard	\$0.0053	\$0.00042	n/a	n/a	n/a
S3 Standard - Infrequent Access **	\$0.01	\$0.001	\$0.01	n/a	\$0.01
S3 One Zone - Infrequent Access **	\$0.01	\$0.001	\$0.01	n/a	\$0.01

Figure 2.17. S3 Standard prices

Cold storage is used for data that is very infrequently accessed, once or twice per year. The most common use case is archival data that needs to be archived for several years for regulatory reasons, but retrieval speed is less of a factor, having retrieval speeds ranging from several minutes up to 12 hours.

One archive data that differs slightly is certain health data where access is needed very infrequently, but when the need arises it has to be near instant access.

Cold storage is the cheapest of the storage types when it comes to the storage. But the low cost of storing the data comes at the cost of a much higher price for the access and retrieval of the data.

S3 Glacier Instant Retrieval*** - For long-lived archive data accessed once a quarter with instant retrieval in milliseconds	
All Storage / Month	\$0.005 per GB
S3 Glacier Flexible Retrieval (Formerly S3 Glacier)*** - For long-term backups and archives with retrieval option from 1 minute to 12 hours	
All Storage / Month	\$0.00405 per GB
S3 Glacier Deep Archive*** - For long-term data archiving that is accessed once or twice in a year and can be restored within 12 hours	
All Storage / Month	\$0.0018 per GB

Figure 2.18. S3 Standard prices

From figure 2.19 below the S3 Glacier Instant Retrieval, Flexible Retrieval and Deep Archive comparison of the prices is showed.

	PUT, COPY, POST, LIST requests (per 1,000 requests)	GET, SELECT, and all other requests (per 1,000 requests)	Lifecycle Transition requests into (per 1,000 requests)	Data Retrieval requests (per 1,000 requests)	Data retrievals (per GB)
S3 Glacier Instant Retrieval ***	\$0.02	\$0.01	\$0.02	n/a	\$0.03
S3 Glacier Flexible Retrieval ***	\$0.0318	\$0.00042	\$0.0318	See below	See below
Expedited	n/a	n/a	n/a	\$10.50	\$0.0315
Standard	n/a	n/a	n/a	\$0.053	\$0.0105
Bulk ***	n/a	n/a	n/a	n/a	n/a
Provisioned Capacity Unit ****	n/a	n/a	n/a	n/a	\$105.00 per unit
S3 Glacier Deep Archive ***	\$0.06	\$0.00042	\$0.06	See below	See below
Standard	n/a	n/a	n/a	\$0.106	\$0.021
Bulk	n/a	n/a	n/a	\$0.02625	\$0.005

Figure 2.19. S3 Glacier Instant Retrieval, Flexible Retrieval and Deep Archive

Block Storage: Amazon EBS, Azure Disks, Google Persistent Disk or Local SSD

Block storage is a storage type where the storage volumes act like storage drives, much like the disk drives on a physical laptop or desktop computer.

The data is saved to these drives in fixed size blocks of data. These blocks are given unique addresses allowing the block storage software to quickly find the location of the data needed. These block storage drives can also be shared between several different virtual machines and are often used for storing data that is needed by applications being run in many different virtual machines.

One of the benefits of block storage compared to an object storage is for data where large files need to be changed and updated often. With a block storage you only need to update the blocks where there is data that is being updated, while in an object storage you would need to update the entire file every time a change is made.

Another use case for block storage is as persistent storage for applications running on virtual machines. If a VM only used the local storage allocated to that particular VM, all the data that it would write would be lost whenever the VM had to restart, since you would never have any guarantee that the server running one particular instance of a VM would be the same the next time the VM instance is run, thus all data written would have been lost.

File storage: Amazon EFS, Google Filestore, Azure Files



When choosing the class/tier of storage it's important to not only consider the price, but also things like availability of the service, what kind of access patterns will be used (will the data be accessed several times per hour or once a month, hot cool cold?) how long does the data need to be stored?

For example, AWS has the S3 Intelligent Tier that monitors the access patterns of your data and moves between the S3 standard and the S3 Infrequent Access tiers to help decreasing storage costs, which is an excellent solution if the access pattern or your data is not fully known.

Another consideration would be which provider is being used in the rest of the enterprise, and the familiarity of co-workers in the provider ecosystem.

Different providers also have data centres in different parts of the world so considerations should be taken to which regions are available with which services. Having your storage deployed in regions as close as possible to your users will decrease the latency for accessing the files stored.

These are all considerations that needs to be taken when choosing which storage type, and which storage tier, is best suited for the data you have, and which provider offers the best overall solution for your particular business need.

As of now Amazon offers 27 different regions with Microsoft Azure having the most regions with 42. Google comes in at 34 regions.

How to create a bucket using the Amazon AWS Console:

When in the console home page, click on the icon in the top left that says 'Services'. This will create a dropdown menu with a list of AWS services. Scroll down to the bottom and click on 'Storage'.

This will open a side panel listing the different Storage services offered by AWS. Click on 'S3'. This will take you to the Amazon S3 console.

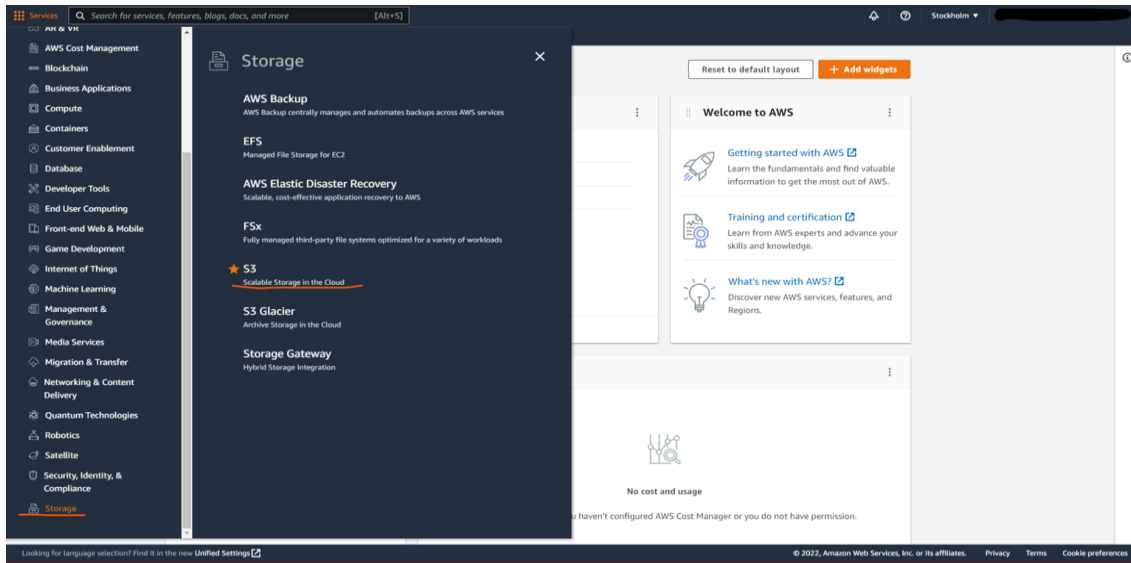


Figure 2.20. S3 console

When in the S3 console you will get a list of all the S3 buckets on your account (as shown in figure 2.20 above).

If this is the first time you have opened the S3 console, no buckets will be listed.

Click on the orange button on the right that says, 'Create bucket' (as shown in figure below).

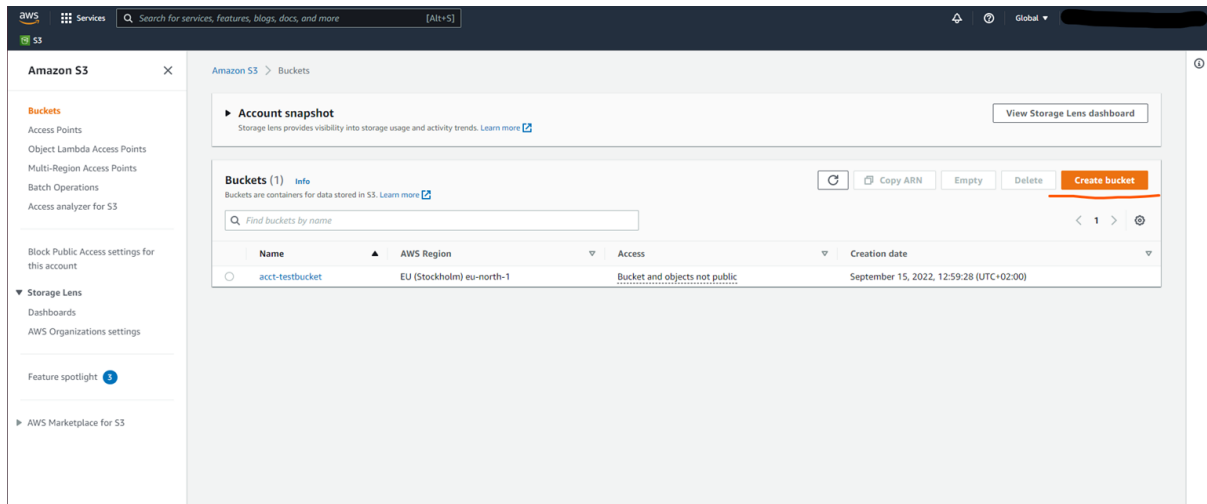
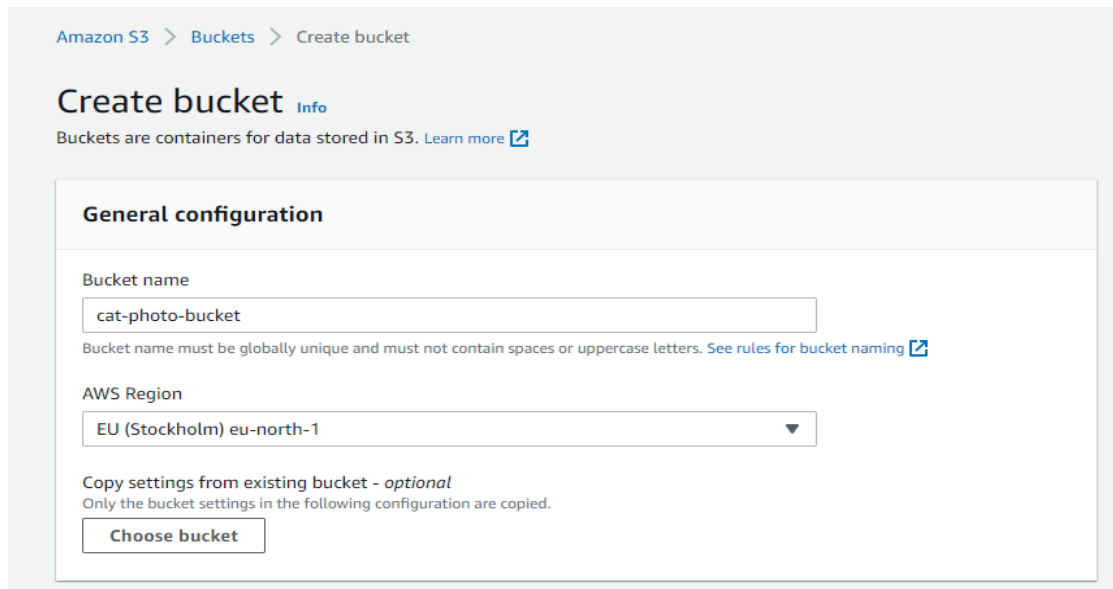


Figure 2.21. Create bucket in S3 console

When you have clicked the button, you will be presented with the 'create bucket wizard'.

Here you will set the configuration for the bucket (see figure below). These include the globally unique name for the bucket. And the AWS region the bucket will be stored.



Amazon S3 > Buckets > Create bucket

Create bucket [Info](#)

Buckets are containers for data stored in S3. [Learn more](#)

General configuration

Bucket name

Bucket name must be globally unique and must not contain spaces or uppercase letters. [See rules for bucket naming](#)

AWS Region

Copy settings from existing bucket - *optional*
Only the bucket settings in the following configuration are copied.

Figure 2.22. Setting a bucket in S3 console

Setting the correct region is important as setting a bucket in a region that is far away from your userbase can introduce latency in accessing the files stored in the bucket.

Next you set the ownership of the objects that are being stored in the bucket. We will choose the recommended setting leaving the Access Control List (ACL) disabled. This means that the ownership of the objects stored will be left with the account that the bucket belongs to.

The second setting in this image is the Public Access. This setting allows you to decide whether or not the objects in the bucket are accessible from other accounts based on the different criteria described in the wizard.

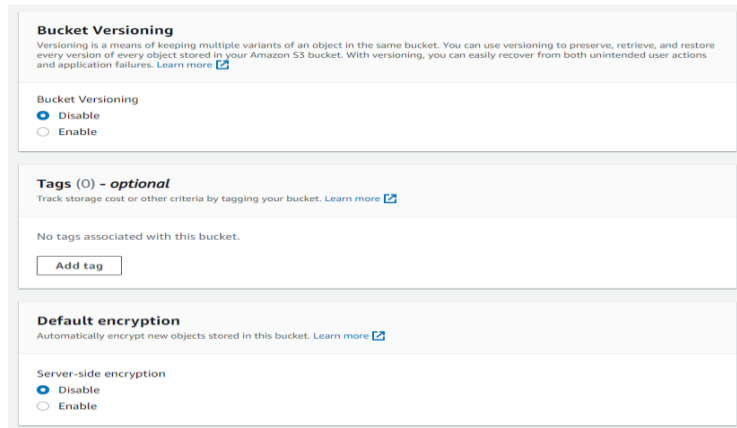


Figure 2.23. Bucket versioning in S3 console

Bucket versioning (see figure above) is used to keep an archive of all the different iterations of the objects in the bucket. Using versioning lets you keep a log of changes and edits to the bucket, and also lets you reroll or retrieve objects in the case of an error, such as an unintended deletion.

Tags can be used to give your buckets an easy way to group buckets together so that they can be used for e.g. cost allocation, ensuring that the costs associated with a particular project are tracked properly.

Default encryption lets you decide if you want the objects in your bucket to be encrypted before AWS saves it to the bucket, leaving it encrypted while at rest, and only decrypts it when it gets downloaded again. Enabling encryption requires you to set up a key for encrypting and decrypting the objects using either Amazon S3-managed keys (SSE-S3) or the AWS Key Management Service.

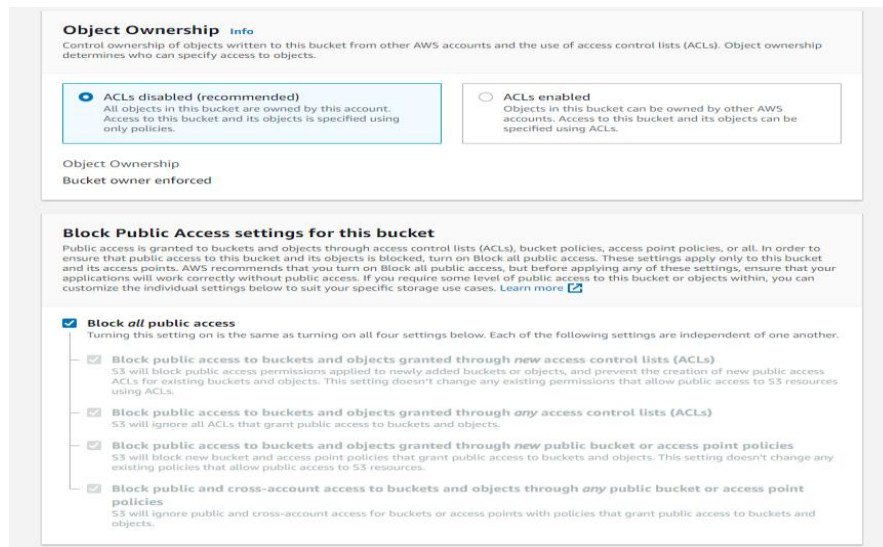


Figure 2.24 Object ownership in S3 console

Under the advanced settings we can set the bucket to have an object lock (see figure above). Enabling Object lock means the objects being stored cannot be deleted or changed while the lock is in effect. This is called a Write-Once-Read-Many model, or WORM model.

When all the configuration has been done, click on the ‘Create bucket’ button.

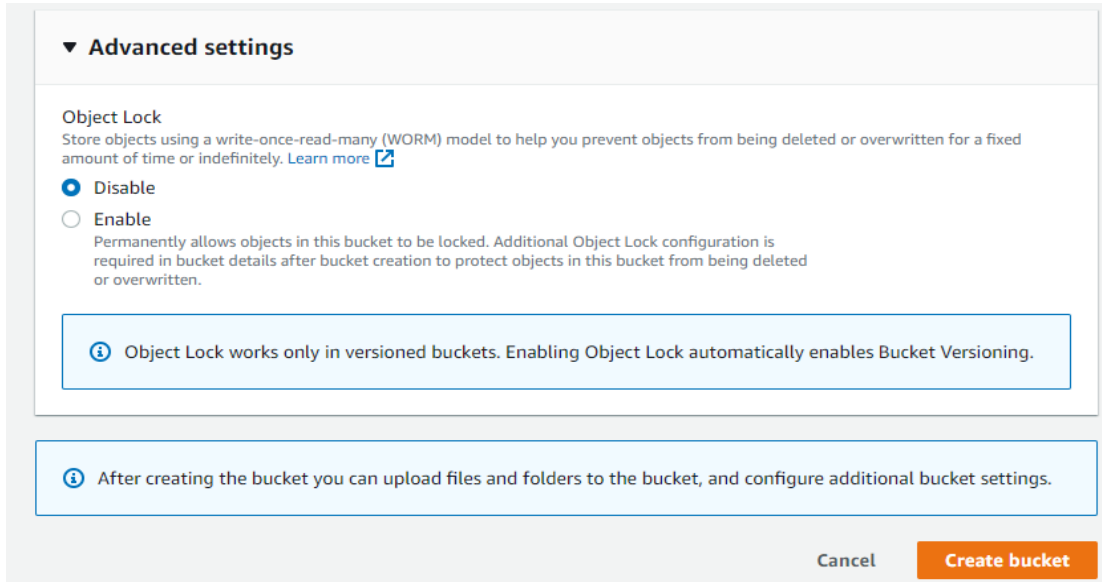


Figure 2.25. Finishing configuration in S3 console

When you have created the bucket, you will be taken back to the S3 console page and your new bucket will be listed in the table of buckets and is now ready to store your files.

To start uploading files to this newly created bucket, you click on the name. This will open the bucket (see figure below).

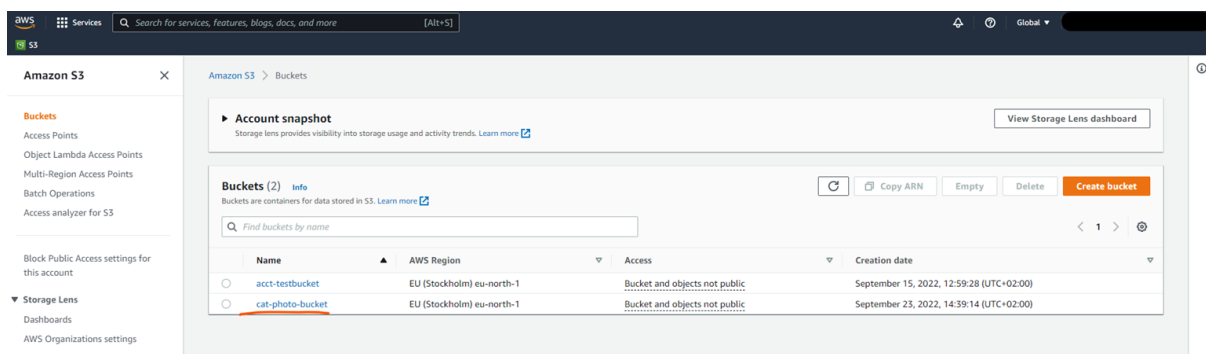


Figure 2.26. Uploading files to newly created bucket in S3 console – first step

Here you can see a lot of information about the bucket, such as the objects that are stored in the bucket, and in the properties tab you can see and edit some of the configuration of the bucket that was set during the creation.

To upload a file to the bucket you can click on either of the two Upload buttons, or you can drag and drop the files from your file explorer (see figure below).

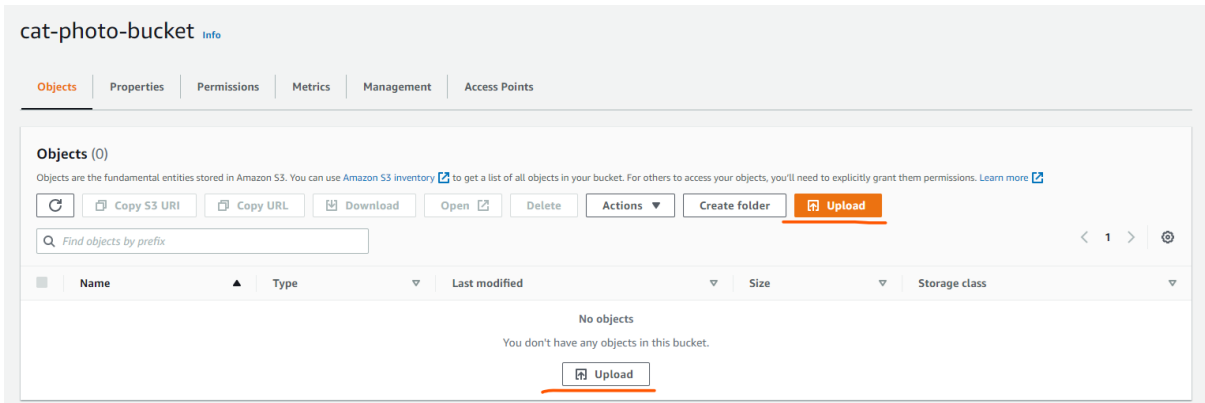


Figure 2.27. Uploading files to newly created bucket in S3 console – second step

Clicking on one of the 'Upload' buttons takes you to the next screen, here you are given a choice between uploading individual files, or an entire folder. You can either click on one of the 'Add' buttons, this will open up a new file explorer and you can choose the files or folders you want to upload, depending on which of the two buttons you clicked.

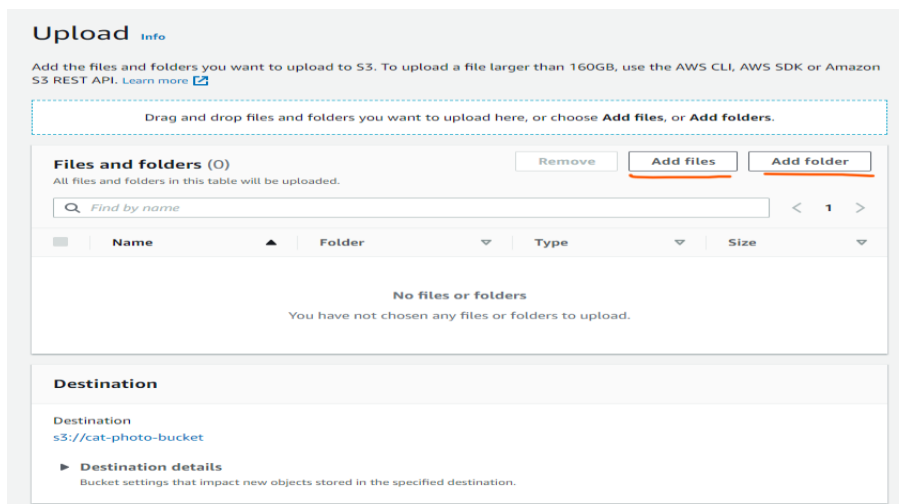


Figure 2.28. Uploading files to newly created bucket in S3 console – third step

In our example, we have uploaded three images. Note that the destination is the bucket we created (see figure below). Opening the Destination details will show some of the bucket settings that was specified. Versioning, Default Encryption and Object Locking.

Upload Info

Add the files and folders you want to upload to S3. To upload a file larger than 160GB, use the AWS CLI, AWS SDK or Amazon S3 REST API. [Learn more](#)

Drag and drop files and folders you want to upload here, or choose **Add files**, or **Add folders**.

Files and folders (3 Total, 109.4 KB) Remove Add files Add folder

All files and folders in this table will be uploaded.

Find by name < 1 >

<input type="checkbox"/>	Name ▲	Folder ▼	Type ▼	Size ▼
<input type="checkbox"/>	cat-1045782__340.jpg	-	image/jpeg	63.3 KB
<input type="checkbox"/>	cat.jfif	-	image/jpeg	5.4 KB
<input type="checkbox"/>	cute-cat-photos-1593441022.jpg	-	image/jpeg	40.7 KB

Destination

Destination
s3://cat-photo-bucket

► **Destination details**
Bucket settings that impact new objects stored in the specified destination.

Figure 2.29. Uploading images to newly created bucket in S3 console

Then we have the properties (see figure below). This is where you can set which storage class you wish to use for the files or folders being uploaded.

▼ Properties
Specify storage class, encryption settings, tags, and more.

Storage class
Amazon S3 offers a range of storage classes designed for different use cases. [Learn more](#) or see [Amazon S3 pricing](#)

	Storage class	Designed for	Availability Zones	Min storage duration	Integrity
<input checked="" type="radio"/>	Standard	Frequently accessed data (more than once a month) with milliseconds access	≥ 3	-	-
<input type="radio"/>	Intelligent-Tiering	Data with changing or unknown access patterns	≥ 3	-	-
<input type="radio"/>	Standard-IA	Infrequently accessed data (once a month) with milliseconds access	≥ 3	30 days	1
<input type="radio"/>	One Zone-IA	Recreatable, infrequently accessed data (once a month) stored in a single Availability Zone with milliseconds access	1	30 days	1
<input type="radio"/>	Glacier Instant Retrieval	Long-lived archive data accessed once a quarter with instant retrieval in milliseconds	≥ 3	90 days	1
<input type="radio"/>	Glacier Flexible Retrieval (formerly Glacier)	Long-lived archive data accessed once a year with retrieval of minutes to hours	≥ 3	90 days	-
<input type="radio"/>	Glacier Deep Archive	Long-lived archive data accessed less than once a year with retrieval of hours	≥ 3	180 days	-
<input type="radio"/>	Reduced redundancy	Noncritical, frequently accessed data with milliseconds access (not recommended as S3 Standard is more cost effective)	≥ 3	-	-

Figure 2.30. Properties in S3 console

You can also turn on additional checksums, this lets you set your own checksum function to make sure the integrity of the objects is valid.

Tags are similar to the ones mentioned earlier during the creation of the bucket, and Metadata is data that in some way describes the data itself, such as the content type or the username of the person creating the original file.

When all of these are set you click the Upload button, and your files will be stored in the cloud (see figure below)!

Additional checksums

Checksum functions are used for additional data integrity verification of new objects. [Learn more](#)

Additional checksums

Off
Amazon S3 will use a combination of MD5 checksums and Etags to verify data integrity.

On
Specify a checksum function for additional data integrity validation.

Tags - optional

Track storage cost or other criteria by tagging your objects. [Learn more](#)

No tags associated with this resource.

Metadata - optional

Metadata is optional information provided as a name-value (key-value) pair. [Learn more](#)

No metadata associated with this resource.

Figure 2.31. Upload in S3 console

After the upload is completed, we see that we have a success-message at the top and we can see the list of our three images in the files and folders-table with some additional data about the type and size of the files and the status message.

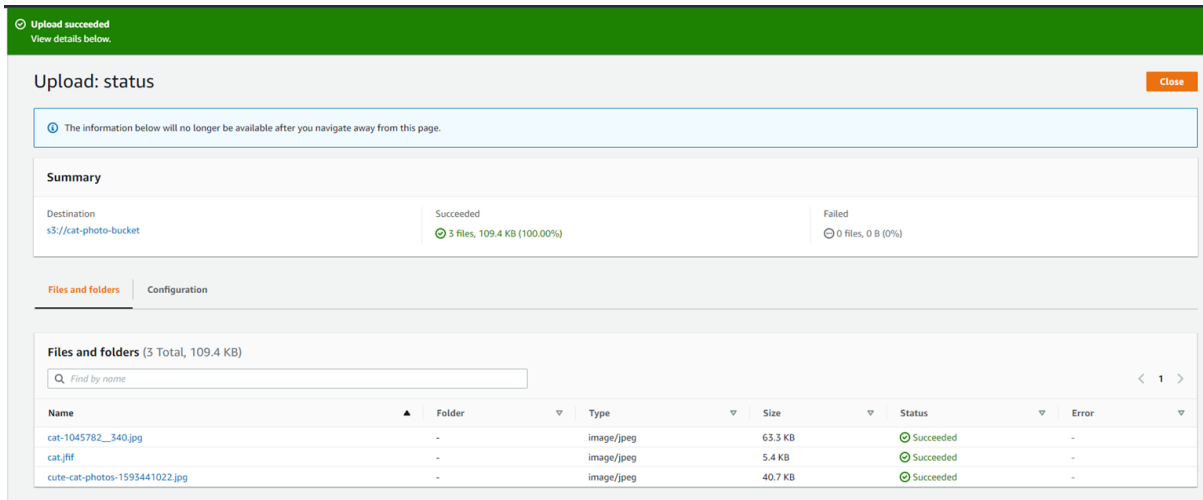


Figure 2.32. Success-message of upload in S3 console

Clicking the close button, we are taken back to our bucket, and we can now see that we have three files in our objects-table, and some information about the files such as the type, size and which storage class is used to store it (see figure below).

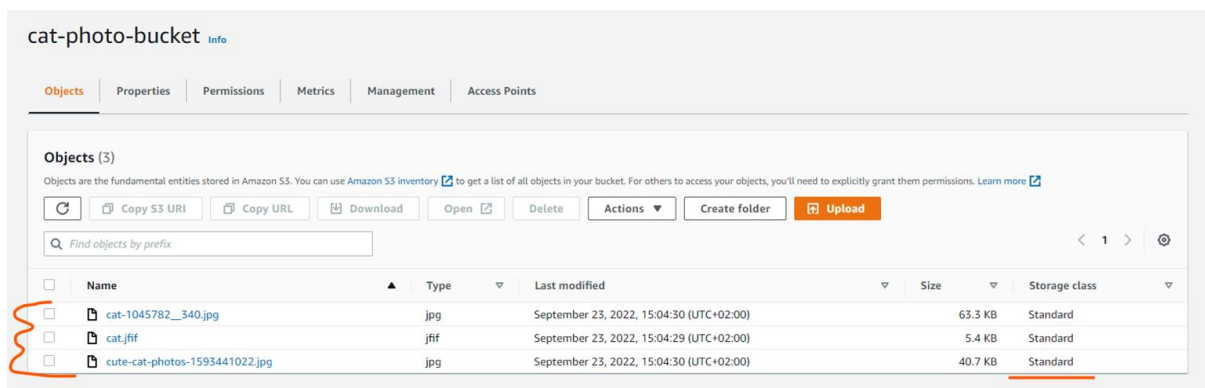


Figure 2.33. Information about the stored data in S3 console

Now that we have our files in the cloud, we can retrieve these files from the cloud. Mark the checkbox next to the name of the file you wish to retrieve, and you will notice the previously greyed out buttons on the row above the table are now available. Click on the Download button and you will start the download of the file onto your computer. The copyable URL and S3 URI can also be used to access the objects, but in our case pasting the URL into the browser will only give you an error message stating we do not have access.

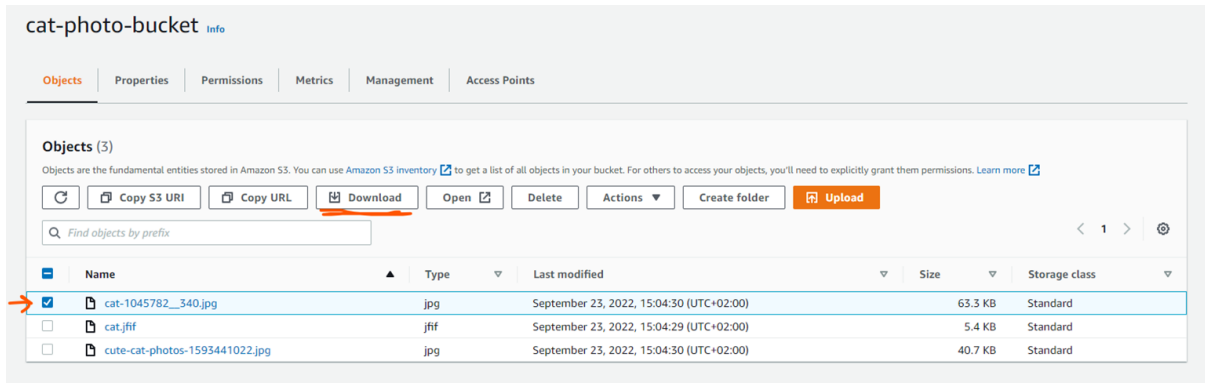


Figure 2.34. Retrieving files in cloud in S3 console

Sometimes you need to delete objects from a bucket (see figure below). Do to so you just have to select the files or folders you wish to delete, checking the box next to here it says 'Name', selecting all objects in the bucket, or selecting each individual file like we did when we retrieved the file.

After all the objects you wish to delete have been selected, click the delete button.

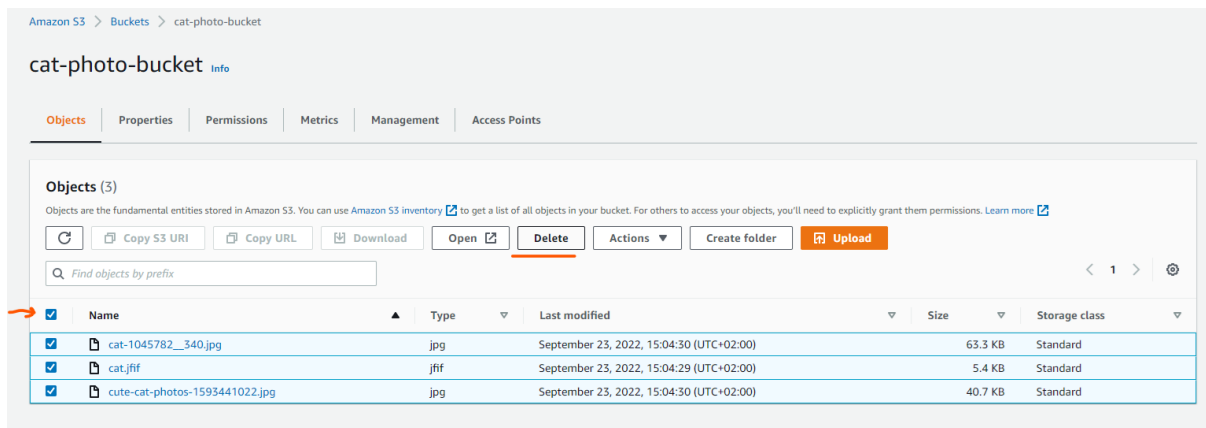


Figure 2.35. Deleting objects from a bucket in S3 console

After clicking the delete button, you will be asked to confirm your deletion with a warning about the consequences of the action. Confirm the deletion by entering the prompted text 'permanently delete' in the text field and click the button to proceed. After the button has been clicked you will be redirected to a summary of the action showing if it was successful or if any errors occurred.

Click close and you will be returned back to the bucket page where all the objects in the bucket have now disappeared.

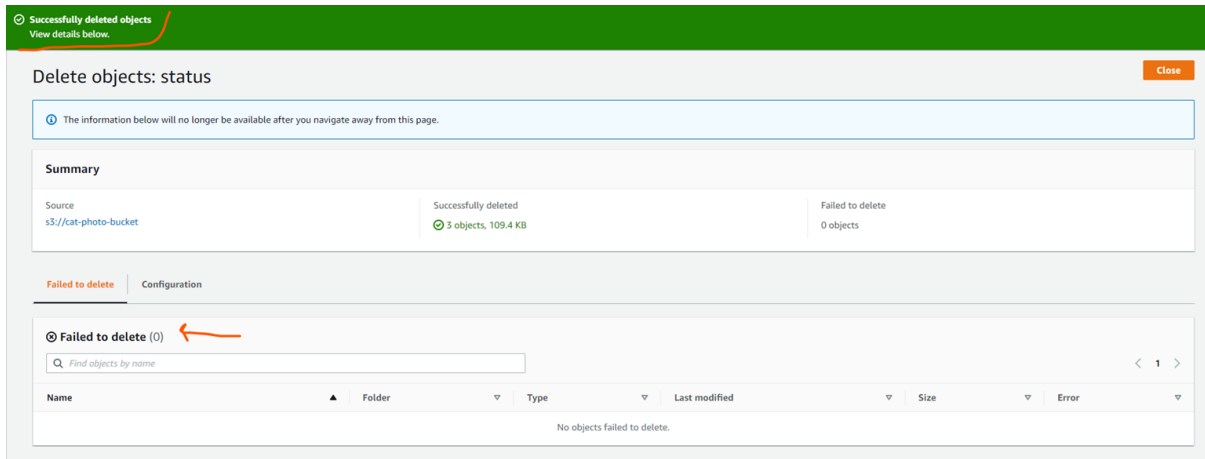


Figure 2.36. Deleted status in the bucket in S3 console

Now that the bucket is empty, we can safely remove the bucket from our account. To delete the bucket itself you just have to select the bucket you want to delete by checking the radio button for the correct bucket and click the delete button next to the create bucket button we used earlier (see figure below).




Figure 2.37. Deleting the bucket in S3 console

As when we deleted the objects in the bucket you will be prompted to confirm your deletion by entering the name of the bucket and clicking the Delete bucket button. After the deletion has completed you will be redirected to the main S3 page, where your bucket will no longer be listed in the bucket table.

2.3.3 Identity access management

Identity Access Management(IAM) is a way to handle both the authentication of a principal, be it a human user or a machine accessing through an API, and the authorization of that same principal, and allows access to the members of an account or organisation to the cloud infrastructure based on the permissions they are granted from the IAM service.

The IAM service can set policies at several levels such as individual users or for groups.

So what is authentication and authorization, and what is the difference?

Authentication is the act of validating the whoever is attempting to access your cloud resources, is who they are claiming to be. This can be done using:

- Username and password
 - The most common way to authenticate users. This requires whoever is trying to login to supply a combination of username and password that is then checked by a system, and if it matches what is registered in that system, the user has verified that they are who they claim to be.
- One-time pins
 - This is a way of validation where the user requests access to the system through an automatically generated PIN that typically only lasts for the duration of the user's session or for a single transaction.
- Authentication apps
 - A trusted third party system generates a password for a user to use.
- Biometrics
 - Biometrics requires the user to verify their identity through a fingerprint , eye scan or face recognition

More and more we see Multi-Factor Authentication (MFA) being used. This requires whoever is trying to authenticate to successfully verify themselves through two or more of the aforementioned methods. These methods are often placed in three main categories; Something you know, something you have and something you are.

These will very often be your password, a phone application, and a biometric part respectively. After the user has been successfully verified, they also need to be authorised before they can start accessing the resources in the cloud system.

Authorisation is, in this context, a process where the system checks to see if the user, who was previously authenticated, has the permissions required to perform the action they are trying to do. One example of this could be an image repository where regular users are allowed to view and download the images on the site, but only users with administrator-privileges are allowed to upload images to the repository (see figure below).



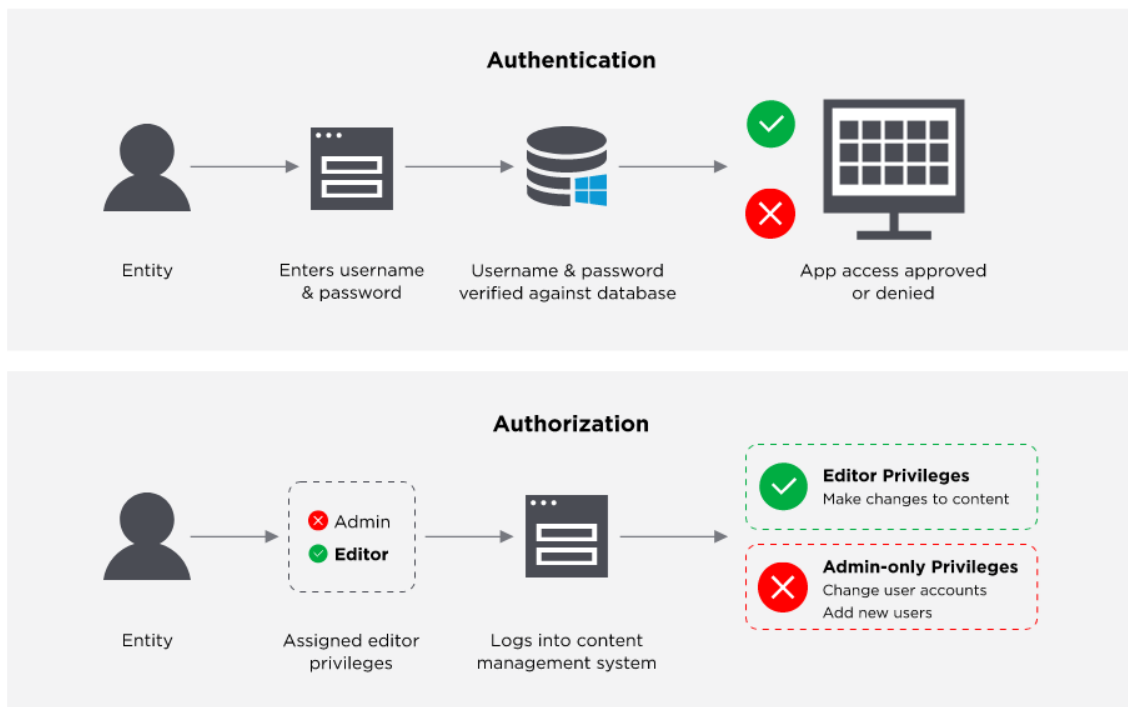


Figure 2.38. Authorization

All of the big three cloud service providers offer an Identity and Access Management service. Microsoft Azure calls it Azure Active Directory. Amazon has named theirs AWS IAM and in Google Cloud it's just called IAM.

Identity and Access Management helps you securely control the access to the Cloud services you use, such as an AWS S3 bucket or a Microsoft Azure CosmosDB instance. Letting you create several IAM users under the umbrella of your main account that handles all the resources.

Without using, for example, AWS IAM to manage the access to your cloud resources, you would have had to create multiple AWS accounts, each of which would have their own separate billing and subscriptions to the various AWS products. Or all employees within your organisation needing to use AWS would have to share the credentials for a single AWS account with no way to restrict the employees from accessing resources they do not need access to.

With IAM however it is possible to set up several users within a single AWS account, starting with the root level user that AWS automatically creates when creating the account. Each subsequent user added to the account will have their own credentials. These users, be they human or machine, can also be granted access to specific resources within AWS through the use of policies which in AWS are defined

in JSON format (see figure below).

```
{  
  'Version': '2012-10-17',  
  'Statement': [{  
    'Effect': 'Allow',  
    'Action': 'iam:ListUsers',  
    'Resource': '*' } ]  
}
```

Figure 2.39. Granting access to specific resources within AWS

These policies are attached to users either directly, or via a user group.

A user group is an IAM resource that you can use to add several IAM users to, so that you can easily attach several policies to any user by adding the user to a user group. For example, if you have a role within your organisation that requires the users to be able to create and delete S3 buckets, whenever a new person is given that role the IAM administrator can simply add that user's IAM user account to the user group, rather than attaching all the policies needed to the user manually.

All three provide the same basic functionality of authenticating the users associated with their account or organisation and authorising these users to access the resources they need access to through policies that are attached to the users in some way.

IAM Resources

The user, group, role, policy, and identity provider objects that are stored in IAM. As with other AWS services, you can add, edit, and remove resources from IAM.

IAM Identities

The IAM resource objects that are used to identify and group. You can attach a policy to an IAM identity. These include users, groups, and roles.

IAM Entities

The IAM resource objects that AWS uses for authentication. These include IAM users and roles.

Principals

A person or application that uses the AWS account root user, an IAM user, or an IAM role to sign in and make requests to AWS. Principals include federated users and assumed roles.

Models and Principles

Principle of Least Privilege:

The principle of least privilege, or Just-Enough-Access, is one of the corner stones of access management and states that a user or application should only be granted the least amount of access necessary to perform the task it is doing. For example, if an application is being used to display images that are being stored in an object storage i.e. Azure blob storage, that application will only ever need to read from that storage, as such it should not be granted anything beyond read-access.

Zero-trust model:

The zero-trust model is a security model where it is assumed that the integrity of the network has been compromised, and that there are no inherently safe access-points.

This is contrary to the old traditional security models where a network was closed off from the rest of the internet, and only trusted and managed computers would be allowed to join. The network would then grant access to these computers and devices based on their location and them having been given access to the network.

With the Zero-trust model, all devices are treated as if coming from an unsafe location, and requires everyone to authenticate to prove their identity before gaining access to the assets and resources they require.

Just-In-Time:

Just in time access is a security model where a firewall will restrict any and all inbound traffic to a resource until a user requests access. The user would then have their authorization checked, and if the request is approved the rules for inbound traffic for the requested resource are temporarily changed to allow that user access and then changing them back to disallowing any traffic.

RBAC & ABAC

Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) are two of the most common methods for securing access to the resources in the cloud.



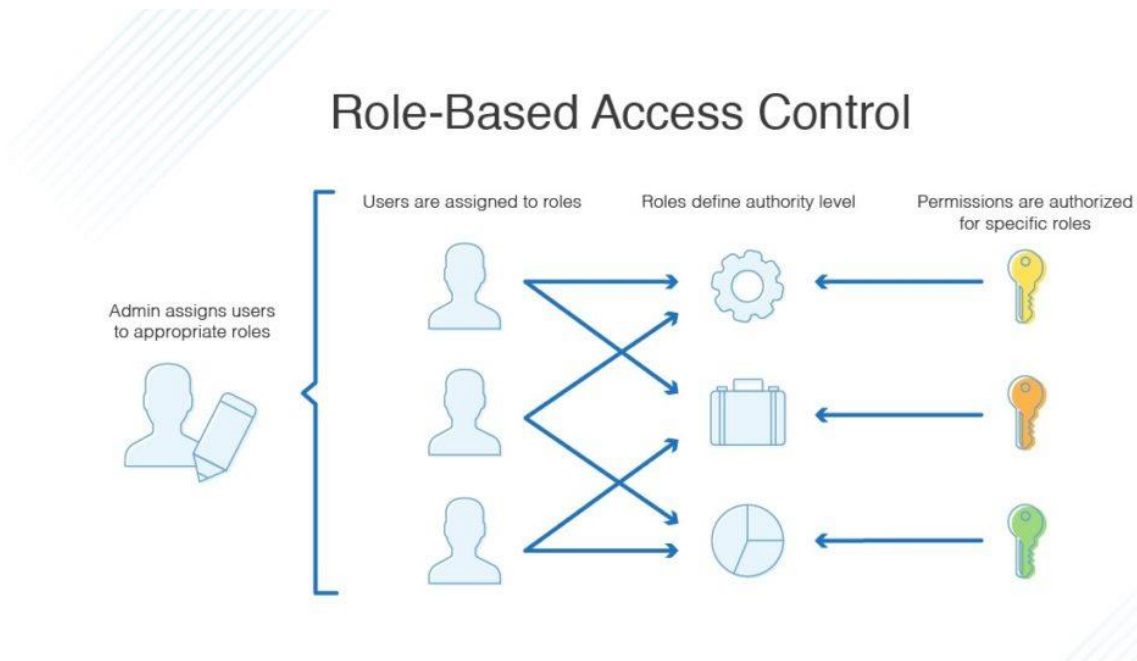


Figure 2.40. Role-based access control

In role-based access control, the access is given based on the roles granted to the user by the administrator of the cloud ecosystem. The rules for access are defined in policies assigned to the roles. Example roles could be one role for a developer who needs read and write access to a database, and one role for an accountant who needs access to the billing information for the same database. Whenever a person need access to the resources in the roles, the user can have that role assigned to them. One user can also have several roles assigned to them, and one role can have several users associated with it (see figure above).

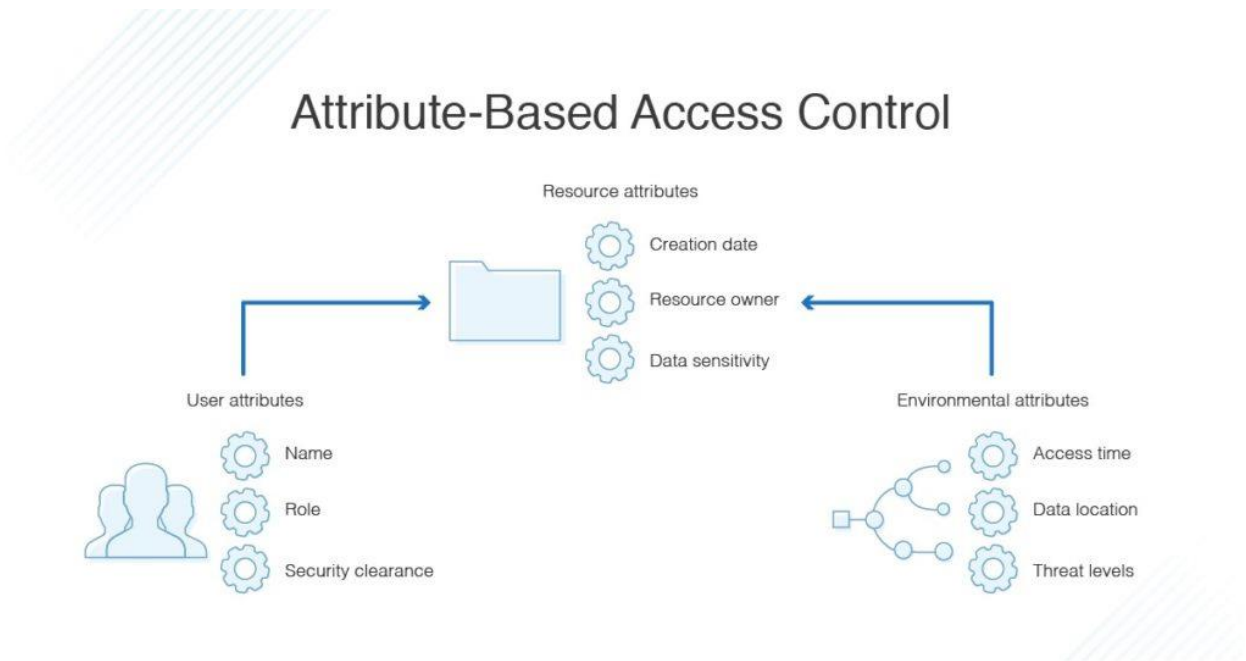


Figure 2.41. Attribute-based access control

In an attribute-based access control system, the access to resources is given to users based on some attributes defined in the policy of the resource. This lets the users create new resources that authorized users have instant access to because they have been given access via an attribute, such as a tag. This means the administrator does not have to create or update policies to accommodate access to new resources being created (see figure above).

RBAC vs. ABAC – pros and cons

ABAC pros

- High level of control and granularity
- Can avoid time consuming work when managing an overwhelming amount of roles

ABAC cons

- Can be time-consuming to set up
- Must be implemented from the very beginning

RBAC pros

- Easy-to-use and straightforward, less complex rules

RBAC cons

- Can lead to role explosion where one would have to manage an excessive number of different roles

When to choose an RBAC model?

- Small companies who manages few cloud resources and with small teams where there is little risk of 'role-explosion'

- If the organisation structure is simple and with well-defined roles.

When to choose an ABAC model?

- If you are working with temporary or distributed teams where you may need to grant access based on the location they are accessing from and the time-zones they are in.
- If there is a lot of collaboration on files and documents where access needs to be based on the type of document/file rather than the role who wants to access it.

In many cases you will want to have a combination of both models where RBAC provides access at a higher level but use ABAC to achieve a finer more granular control.

2.3.4 Database Services in the cloud

When choosing a database and database provider there are, like with choosing the storage type/provider, many different considerations that needs to be taken. There are several different kinds of databases, and they all have their strengths and weaknesses, depending on what kind of data is being stored.

The traditional Relational databases using SQL(Structured Query Language), such as mySQL or PostgreSQL is great when working with datasets that are well-defined from the start and where there will be no changes to the format of the data over time, and there are strong and clear relationships between the different parts of your dataset.

For example, you can have the relationship between a book and a library. In this case a library will have many books, but one book can only belong to a single library (see figure below).




Figure 2.42. Relationship between the book and library

But sometimes the data you have isn't well structured and you have less control over how the data is going to change over time. For these cases it might be best to use a NoSQL(Not Only SQL) database....

Amazon RDS (Relational Database Service) supports a range of the most popular relational databases such as MySQL, MariaDB, and OracleSQL but also provides their own relational database called Amazon Aurora.

Microsoft Azure offer several relational database services, such as Azure SQL Database, Azure Database for PostgreSQL/MariaDB/MySQL.

Google Cloud Platform has Cloud SQL, AlloyDB, and Cloud Spanner. They also have a solution optimized for datawarehouses named BigQuery.

For the NoSQL solutions Google offers their own Document database called Firestore and a key-value database named Cloud Bigtable.

Microsoft Azure has a NoSQL solution called Cosmos DB that supports a broad variety of other NoSQL API's, such as Apaches Cassandra and MongoDB, but also has support for SQL.

Amazon's NoSQL services include DocumentDB and DynamoDB which as both managed services.

When choosing the type of database, and the provider some consideration has to be made. Key things to consider are the type of data that is being stored. Is the data highly structured with strong relationships? The best choice might be a relational database.

In addition to deciding what kind of database is best suited, deciding what instance class is needed is important. The DB Instance classes determine the amount of memory, CPU, the I/O throughput is available to the database server.

DB instance class determines Memory, CPU, I/O network storage throughput - can be managed in AWS management console, AWS CLI, RDS API.

DB security: Proximity to internet. - Virtual Private Cloud. Network gateway, access control uses IAM, users and roles can be used to determine access to DB actions (Getting, Posting)

AWS uses AES-256 while at rest.

Amazon Availability Zones to increase durability in case of Infrastructure failure.

A practical walk-through for how to instantiate a database with Amazon RDS using the Amazon Aurora MySQL (see figure below):

When in the AWS console home page, click the Services button on the top left and then locate 'Database' in the drop-down menu. Click it and a new pane opens. Locate 'RDS' and click it.

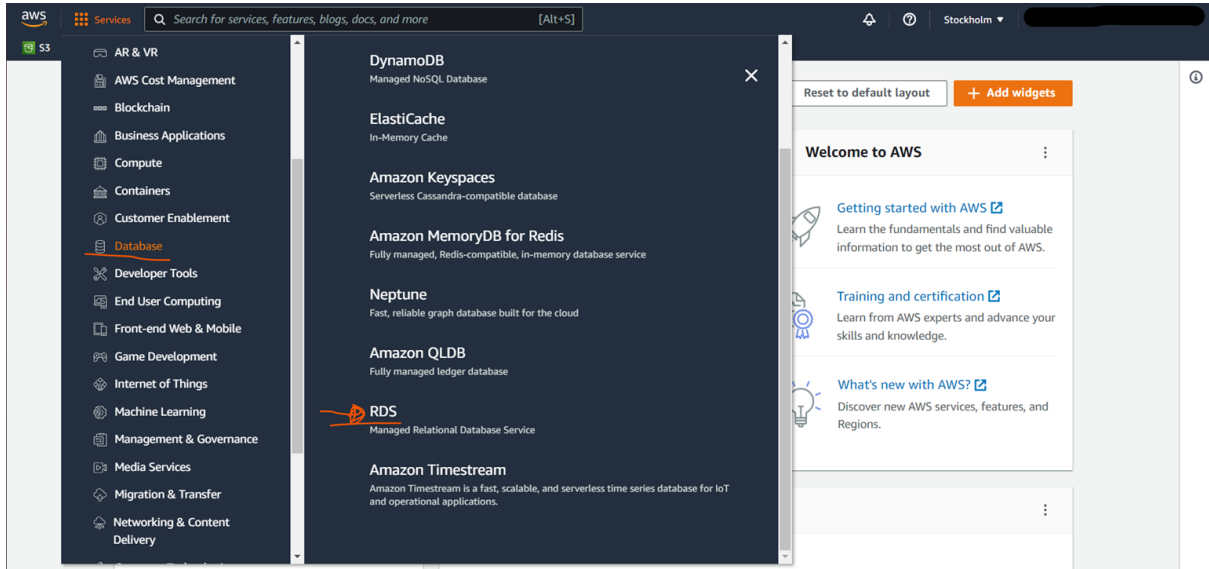


Figure 2.43. Database with Amazon RDS using the Amazon Aurora MySQL

In the Amazon RDS console home page, click on Databases on the menu on the left, this will take you to the panel with the overview of all databases connected to your AWS account.

To create a new database, you click on the button that says 'Create database' which will open the database creation wizard (see figure below).

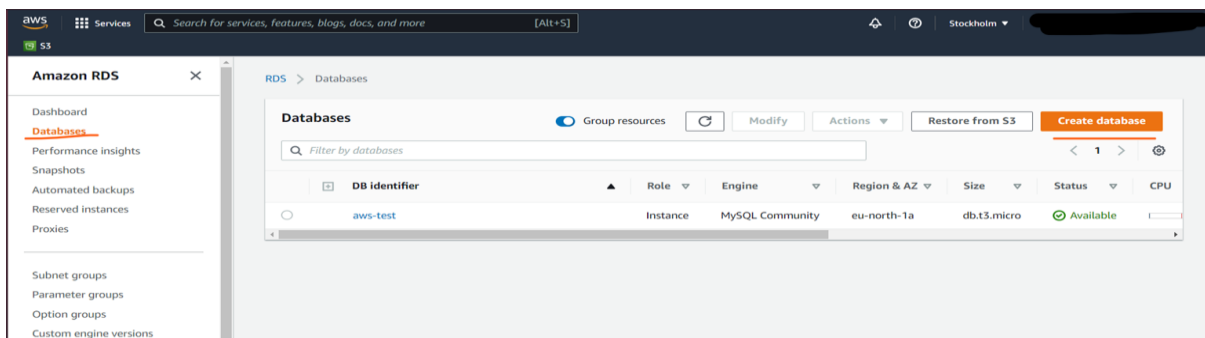


Figure 2.44. Creation of a new database – first step

In the wizard you will get several options for which relational database you wish to use, and also which versions of the database engines you want to use. We will leave the default choose the Amazon Aurora MySQL-compatible edition and have it run on the MySQL 5.7 version (see figure below).

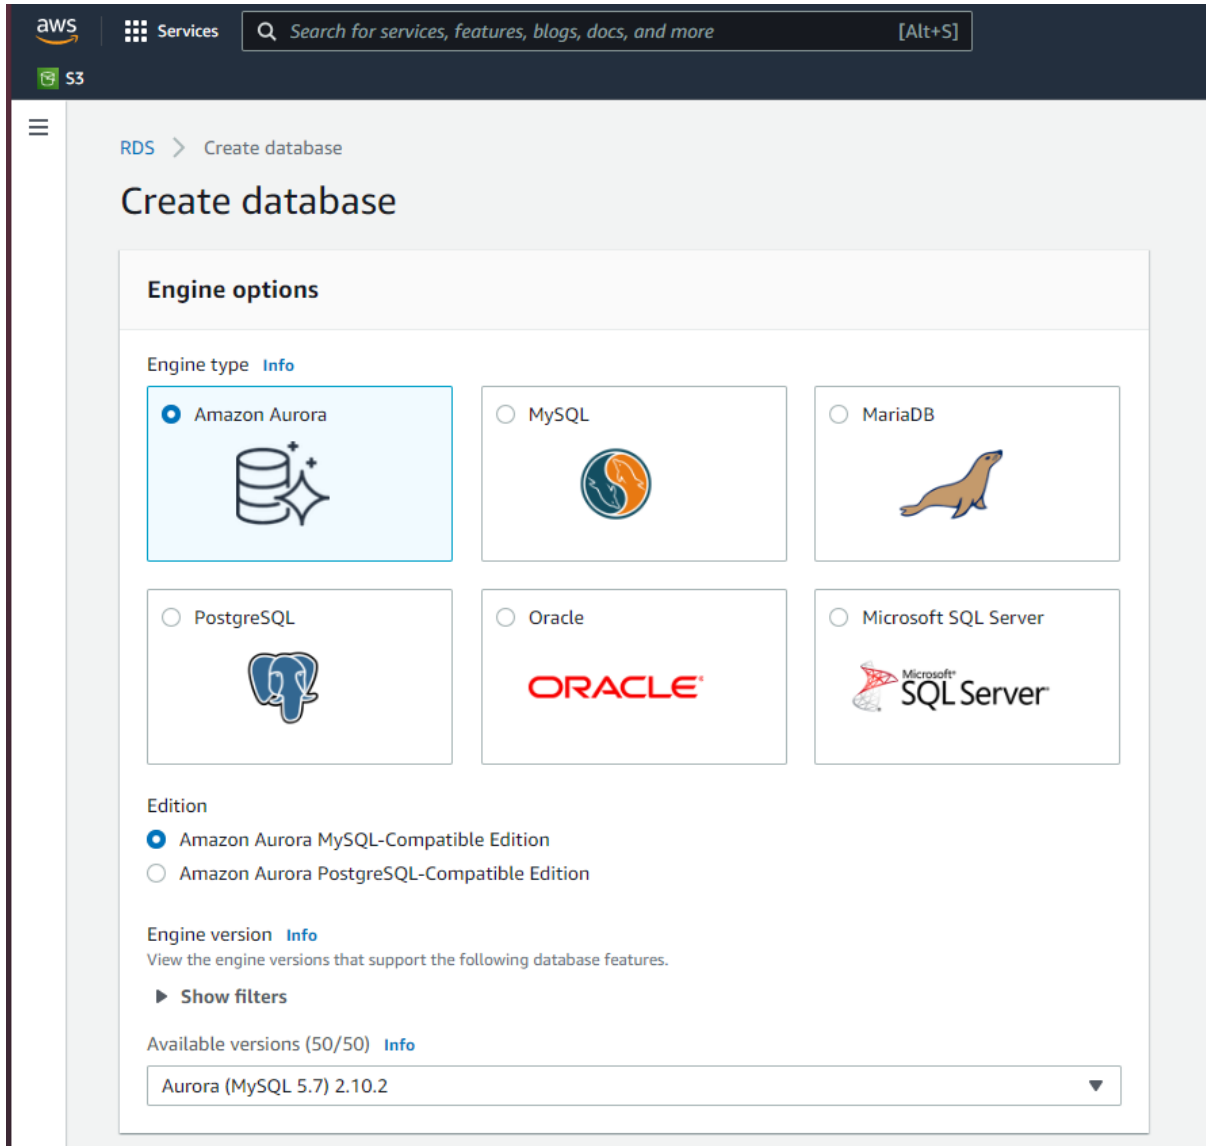


Figure 2.45. Creation of a new database – second step

Scrolling down we can set the settings for the database, defining the cluster name (or just the database name if using mysql) and the credentials such as the username and the password (see figure below).

Settings

DB cluster identifier [Info](#)
Type a name for your DB cluster. The name must be unique across all DB clusters owned by your AWS account in the current AWS Region.

The DB cluster identifier is case-insensitive, but is stored as all lowercase (as in "mydbcluster"). Constraints: 1 to 60 alphanumeric characters or hyphens. First character must be a letter. Can't contain two consecutive hyphens. Can't end with a hyphen.

▼ **Credentials Settings**

Master username [Info](#)
Type a login ID for the master user of your DB instance.

1 to 32 alphanumeric characters. First character must be a letter.

Auto generate a password
Amazon RDS can generate a password for you, or you can specify your own password.

Master password [Info](#)

Constraints: At least 8 printable ASCII characters. Can't contain any of the following: / (slash), '(single quote), "(double quote) and @ (at sign).

Confirm password [Info](#)

Figure 2.46. Settings for the database

In the instance configuration setting you decide which instance class to use for the database. In our case we will choose the smallest burstable class for cost reasons, but in a real application one would have to consider what kind of datasets it should handle, and what kind of access patterns and what kind of throughput it needs to be able to handle.

Creating an Aurora Replica can be chosen to create replicas in different Availability Zones, so that if one AZ goes down or experiences any issues, you can quickly swap over to another AZ with minimal downtime (see figure below).



Instance configuration

The DB instance configuration options below are limited to those supported by the engine that you selected above.

DB instance class [Info](#)

Memory optimized classes (includes r classes)
 Burstable classes (includes t classes)

db.t3.small
2 vCPUs 2 GiB RAM Network: 2 085 Mbps

Include previous generation classes

Availability & durability

Multi-AZ deployment [Info](#)

Create an Aurora Replica or Reader node in a different AZ (recommended for scaled availability)
Creates an Aurora Replica for fast failover and high availability.

Don't create an Aurora Replica

Figure 2.47. Creating an Aurora Replica

In the connectivity settings (see figure below) you will set up whether or not you would like to connect your database to an Amazon Elastic Cloud Compute, or EC2 resource.

It is also necessary to create a Virtual Private Cloud (VPC). In this VPC you can create special rules for who is allowed to access the resource within it.

The subnet group is used to define which IPs the Database is allowed to use within the VPC. We will leave both these as default.

Public access defines if anything or anyone that is not within the VPC can access the database through a public IP address created by the wizard. Usually, you want to turn this off, so that only resources that are inside the VPC can access the database, minimising the risk of unauthorized access.

VPC security groups are like access lists for which IP addresses are allowed to access the database.

Connectivity Info

Compute resource
Choose whether to set up a connection to a compute resource for this database. Setting up a connection will automatically change connectivity settings so that the compute resource can connect to this database.

Don't connect to an EC2 compute resource
Don't set up a connection to a compute resource for this database. You can manually set up a connection to a compute resource later.

Connect to an EC2 compute resource
Set up a connection to an EC2 compute resource for this database.

Virtual private cloud (VPC) Info
Choose the VPC. The VPC defines the virtual networking environment for this DB cluster.

Default VPC (vpc-03544722a135fae8b)

Only VPCs with a corresponding DB subnet group are listed.

After a database is created, you can't change its VPC.

DB Subnet group Info
Choose the DB subnet group. The DB subnet group defines which subnets and IP ranges the DB cluster can use in the VPC that you selected.

default-vpc-03544722a135fae8b

Public access Info

Yes
RDS assigns a public IP address to the cluster. Amazon EC2 instances and other resources outside of the VPC can connect to your cluster. Resources inside the VPC can also connect to the cluster. Choose one or more VPC security groups that specify which resources can connect to the cluster.

No
RDS doesn't assign a public IP address to the cluster. Only Amazon EC2 instances and other resources inside the VPC can connect to your cluster. Choose one or more VPC security groups that specify which resources can connect to the cluster.

VPC security group (firewall) Info
Choose one or more VPC security groups to allow access to your database. Make sure that the security group rules allow the appropriate incoming traffic.

Choose existing
Choose existing VPC security groups

Create new
Create new VPC security group

Existing VPC security groups
Choose one or more options

Availability Zone Info

No preference

► Additional configuration

Figure 2.48. Connectivity settings

You can choose which availability zone within your region you would prefer the database to be located in.

The authentication lets you decide if only the database password is enough, or if any authentication also has to include an AWS IAM user/role.

Monitoring watches the resource usage of your database. Now that we have configured all our settings, we can create our database. Click the create database button (see figure below).



Database authentication

Database authentication options [Info](#)

- Password authentication
Authenticates using database passwords.
- Password and IAM database authentication
Authenticates using the database password and user credentials through AWS IAM users and roles.

Monitoring

Monitoring

- Enable Enhanced monitoring
Enabling Enhanced monitoring metrics are useful when you want to see how different processes or threads use the CPU.

Granularity

60 seconds ▼

Monitoring Role

default ▼

Clicking "Create database" will authorize RDS to create the IAM role rds-monitoring-role

▶ **Additional configuration**

Database options, encryption turned on, failover, backup turned on, backtrack turned off, maintenance, CloudWatch Logs, delete protection turned off.

i You are responsible for ensuring that you have all of the necessary rights for any third-party products or services that you use with AWS services.

Cancel Create database

Figure 2.49. Creating database

Our database has been created and we can now see it in the list in our Amazon RDS console page (see figure below).

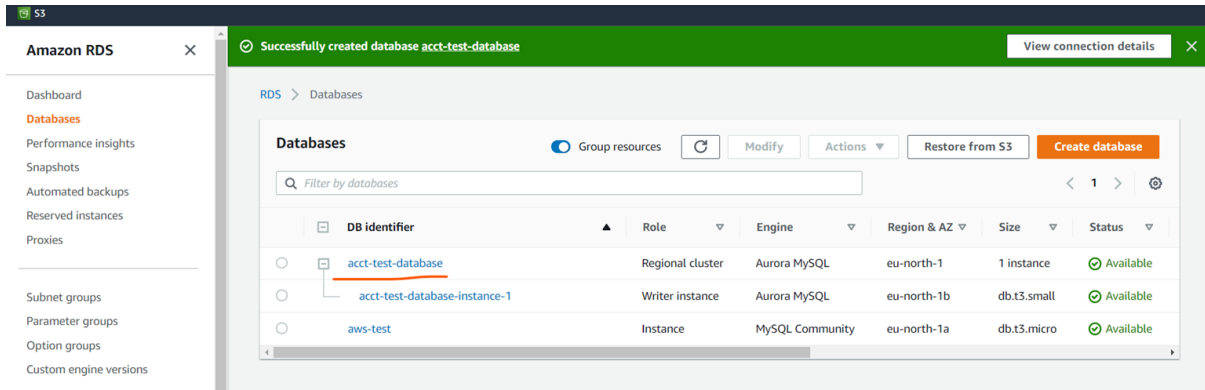


Figure 2.50. Created database visible in Amazon RDS console page

Now that we have created our database, we want to connect to it and create some tables. Clicking on the name of the database we just created we can see the endpoint of the database (see figure below). This is the address we have to connect to.




Figure 2.51. Endpoints of created database

Next, we are going to use MySQL workbench to connect to this database. Clicking on the circled in plus next to where it says “MySQL connections” will open up the window in the image. The hostname is where you will paste the endpoint from the server, and the username will be the master DB username you chose when creating the database. After you have supplied these two, you can click ‘test connection’. You will then be prompted to supply the password you set, and if this all works you will get a message saying a connection was made. You can then give the connection a name and click OK.

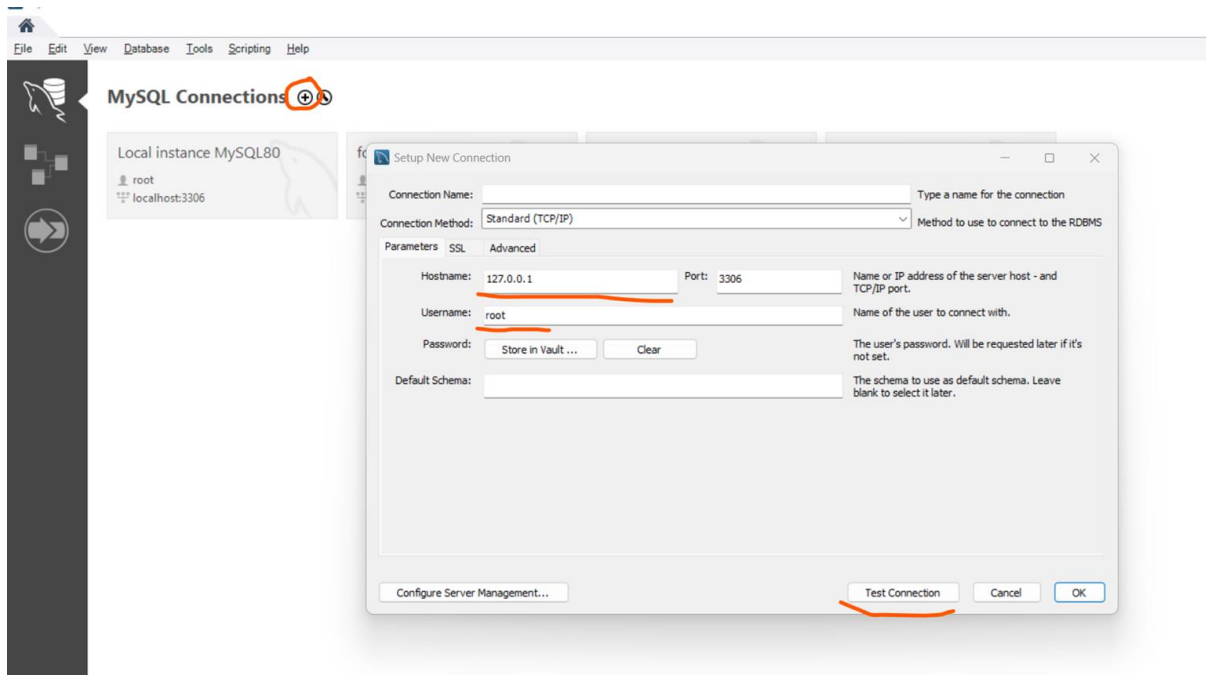


Figure 2.52. Using MySQL workbench for connecting to new database

After successfully connecting to the database and logged in you will be able to start creating databases with tables and information using SQL statements.

2.3.5 Considerations for Domain Set up

In this unit you will learn that when choosing the right cloud service provider, businesses should consider more than just product suites. Almost all companies use an Infrastructure as a Service (IaaS) or Platform as a Service (PaaS) provider. Organizations are likely to start with one of three cloud providers, Amazon Web Services (AWS), Azure, or Google Cloud Platform (GCP) and may decide to mitigate risk by offering a diverse services base through multiple providers, optimize by deploying the right workloads in the right cloud, and minimize vendor lock-in.

Key Differences among Cloud Platforms

1. A look at Amazon Web Services (AWS)

When AWS first launched in 2006, it primarily provided compute, storage, and database services used by developers. As the first cloud provider, AWS remains innovative because it had an earlier foundation to build on.

Most companies use the following services on AWS:

- AWS Elastic Compute Cloud (EC2): scalable, scalable compute power for software hosting or machine learning
- AWS Relational Database Service (RDS): A customizable database engine for hosting database servers and working with NoSQL databases.
- AWS Lambda Functions as a Service (FaaS): Event-driven serverless computing for background processes such as image transformation, real-time data processing, and streaming data validation.
- AWS Simple Storage Service (S3): initially for developers with persistent storage, but also for archiving and cost-effective data migration
- AWS Elastic Container Service (ECS): Container management for starting, stopping, and managing cluster containers.
- AWS CloudFront Content Delivery Network (CDN): Stores data at the edge to deliver data, video, images, applications, and APIs.

2. What does Azure offer?

Azure tends to promote enterprise organizations that have already invested in Microsoft products and services.

Most Azure companies use the following services:

- Azure Hybrid: A service for workloads that combines on-premises Windows Server and SQL Server licenses
- Azure Virtual Desktop (AVD): Virtual Desktop Interface (VDI) for remote access to Windows 10 and applications
- Azure Sentinel: Security Information Event Management (SIEM) and Security Orchestration Automated Response (SOAR) for threat detection, detection, visibility and response
- Azure Cosmos DB: NoSQL database with open API for mobile/web, gaming and e-commerce/retail applications
- Azure Active Directory (AD): An identity service that syncs across on-premises and cloud Microsoft environments with single sign-on and multi-factor authentication.

3. Google Cloud Platform (GCP) and how it compares

Not to be outdone, Google launched a beta version of GCP in 2008. While AWS offers IaaS services, GCP initially focused on PaaS services. Developers can develop and run their web applications in data centers managed by Google. Over time, GCP has expanded its offerings to include Google suites, big data technologies, and management tools.

GCP generally focuses on developers who want to build and run applications. It tends to focus on organizations that want to build applications but lack the local data centers to support them. Most companies use the following services in GCP:

- Google Compute Engine: A preconfigured or customizable kernel-based virtual machine (KVM) for Linux and Microsoft servers



- Google Cloud Storage (GCS): block, file and object storage with lifecycle management rules for different data types
- Google Kubernetes Engine (GKE): A managed, managed staging environment for deploying microservices
- BigQuery Machine Learning (ML): Machine Learning Models for Business Insights

How many availability regions does each provider have?

This is important when determining a company's compliance requirements, which according to the General Data Protection Regulation (GDPR), companies should store and process data in one of the EU countries.

Here is how the competitor's rank:

- AWS: 26 geographic regions
- Azure: 60+ regions
- GCP: 29 regions

Typically, each region has multiple availability zones. This means that you should consider the following:

- AWS: 84 Availability Zones in total
- Azure: 3 availability zones per region, at least 180 in total
- GCP: 88 Availability Zones

Additional considerations may include specialized service options that each provider has that differs:

- AI/Machine Learning
- Internet of Things (IoT)
- Augmented reality/virtual reality
- business analysis
- robot technology

Pricing structure

Each of the three major providers offers different pricing models based on an organization's cloud usage. All three providers find pricing and invoicing difficult which means you need to be aware of the following when considering which server to use:

- Governance
- Billing format
- Monitor consumption and budget
- Changes in the pricing model
- Long term vs. pay-as-you-go pricing value

Management Tools

As has already been mentioned, you may use different cloud services to combine resources and tools to streamline and centralize business needs. It is important to note however that AWS and Azure are

more business oriented than GCP and AWS offers the most range in outsourced services. This may be a major consideration for those businesses that need the most robust options (see figure below).

Here is a figure (chart) detailing the differences:

Management and governance			
	AWS	Azure	Google Cloud
Automation	AWS CloudFormation, AWS Proton, AWS OpsWorks	Azure Resource Manager, Azure Automation	Cloud Deployment Manager, Cloud Foundation Toolkit, Cloud Scheduler
Anomaly detection	CloudWatch Anomaly Detection	Anomaly Detector	Anomaly Detection
Application portfolio and data governance	AWS Service Catalog	Azure Managed Applications, Azure Blueprints (preview), Azure Purview (preview)	Dataplex, Private Catalog, Service Directory
Automated Windows Server management	N/A	Azure Automanage (preview)	N/A
Configuration management	AWS Config	Azure App Configuration	Cloud Asset Inventory
Health dashboard	Personal Health Dashboard	Resource Health, Azure Service Health	Cloud Monitoring
Hybrid and multi-cloud management	Amazon EKS Anywhere (preview), Amazon ECS Anywhere	Azure Arc	Google Anthos, Network Connectivity Center (preview)
License management	AWS License Manager	N/A	N/A

Figure 2.53. List of different cloud services

Also key to management services are those which utilize IoT tools and how they difer (see figure below).

IoT

	AWS	Azure	Google Cloud
Cloud-device connections, data collection and management	AWS IoT Analytics, AWS IoT Core, AWS IoT Device Defender, AWS IoT Device Management, AWS IoT Events, AWS IoT SiteWise	Azure IoT Central, Azure IoT Hub, Azure Defender for IoT, Azure Sphere	Cloud IoT Core
IoT edge compute	AWS Greengrass	Azure IoT Edge, Azure Percept (preview)	Edge TPU
Microcontroller OS	FreeRTOS	Azure RTOS	N/A
Virtual modeling	AWS IoT Things Graph	Azure Digital Twins	N/A

Figure 2.54. Management services, which utilize IoT tools

Questions to consider:

1. What are the main differences among the three servers, and what platform would be more appealing to a start of business versus one with very strong management needs?
2. What kind of pricing structure would appeal to you most from your current perspective?
3. In the same vein, how would regional coverage affect your choice?
4. Read this article and consider what are your main the considerations as a current or potential business owner when deciding on a platform? <https://www.netsolutions.com/insights/how-to-choose-cloud-service-provider/>

Additional Resources: Samoshkin (n. d.), Cloud Industry Forum (2022), Rathore (2022), CloudSigma (2023).

2.4 Connectivity Types of Network Services and Setting Them

2.4.1 About cloud architecture

The term “cloud” appears to have its origins in network diagrams that represented the internet, or various parts of it, as schematic clouds. “Cloud computing” was coined for what happens when applications and services are moved into the internet “cloud.” Cloud computing is not something that suddenly appeared overnight; in some form, it may trace back to a time when computer systems remotely time-shared computing resources and applications. More currently though, cloud computing refers to the many different types of services and applications being delivered in the internet cloud,



and the fact that, in many cases, the devices used to access these services and applications do not require any special applications.

Businesses often seek to find the best cloud solution to fit their unique organizational needs. A large part of this decision is selecting a cloud service provider. There are four primary cloud service providers that control the majority of global cloud resources. However, there are other lesser-known cloud solutions that offer specific services to niche markets.

The four most widely used cloud service providers all offer SaaS, PaaS, IaaS and many other cloud services on a global scale. The major cloud service providers include:

- Google Cloud Services
- Microsoft Azure
- Amazon Web Services (AWS)
- IBM Cloud

GOOGLE CLOUD SERVICES	MICROSOFT AZURE	AMAZON WEB SERVICES (AWS)	WHAT IT DOES
Google Compute Engine	Azure Virtual Machines	Elastic Compute Cloud (EC2)	Infrastructure as a Service (IaaS)
Google App Engine	Azure Cloud Services	AWS Elastic Beanstalk	Platform as a Service (PaaS)
Google Cloud SQL	Azure SQL Database	Amazon Relational Database Service	Database as a Service (DaaS)
Google Cloud Bigtable	Azure Table Storage	Amazon Dynamo DB	Scalable SQL database services
Google BigQuery	Azure SQL Database	Amazon Redshift	Relational Databases
Google Cloud Functions	Azure Functions	AWS Lambda	Serverless Applications
Google Cloud Datastore	Azure Cosmos DB	Amazon Simple DB	Highly Scalable NoSQL Database Services
Google Storage	Azure Storage	Amazon Simple Storage Service (S3)	Storage of object, blocks and files. Also for cool and cold storage of data.

Figure 2.55. Services overview

Some other cloud solutions offering specific services include the following:

- **Heroku:** Large provider of PaaS cloud services, including app development, deployment, management and scaling.
- **GitHub:** A large version-control repository service used for collaborative app development. Developers and managers can review code, manage projects and build software as joint effort.
- **QuickBooks Online:** SaaS version of accounting software offered by QuickBooks.
- **BackBlaze:** Provides a cloud service of data backup and recovery for personal and business use.

- **ClearDATA:** Provides cloud solutions specific to the healthcare industry. Designed to help institutions comply with industry regulations.
 - **Salesforce.com:** Runs its application set for its customers in a cloud, and its Force.com and Vmforce.com products provide developers with platforms to build customized cloud services.
- This is merely scratching the surface of the various cloud solutions that are available. However, these cloud service providers offer a solid base for understanding what kind of services are available.

Characteristics

Cloud computing has a variety of characteristics, with the main ones being:

- **Shared Infrastructure** — Uses a virtualized software model, enabling the sharing of physical services, storage, and networking capabilities. The cloud infrastructure, regardless of deployment model, seeks to make the most of the available infrastructure across a number of users.
- **Dynamic Provisioning** — Allows for the provision of services based on current demand requirements. This is done automatically using software automation, enabling the expansion and contraction of service capability, as needed. This dynamic scaling needs to be done while maintaining high levels of reliability and security.
- **Network Access** — Needs to be accessed across the internet from a broad range of devices such as PCs, laptops, and mobile devices, using standards-based APIs (for example, ones based on HTTP). Deployments of services in the cloud include everything from using business applications to the latest application on the newest smartphones.
- **Managed Metering** — Uses metering for managing and optimizing the service and to provide reporting and billing information. In this way, consumers are billed for services according to how much they have actually used during the billing period.

In short, cloud computing allows for the sharing and scalable deployment of services, as needed, from almost any location, and for which the customer can be billed based on actual usage.

Service Models

Once a cloud is established, how its cloud computing services are deployed in terms of business models can differ depending on requirements. The primary service models being deployed (see Figure 2) are commonly known as:

- **Software as a Service (SaaS)** — Consumers purchase the ability to access and use an application or service that is hosted in the cloud. A benchmark example of this is Salesforce.com, as discussed previously, where necessary information for the interaction between the consumer and the service is hosted as part of the service in the cloud. Also, Microsoft has made a significant investment in this area, and as part of the cloud computing option for Microsoft® Office 365, its Office suite is available as a subscription through its cloud-based Online Services.
- **Platform as a Service (PaaS)** — Consumers purchase access to the platforms, enabling them to deploy their own software and applications in the cloud. The operating systems and network access are

not managed by the consumer, and there might be constraints as to which applications can be deployed. Examples include Amazon Web Services (AWS), Rackspace and Microsoft Azure.

- **Infrastructure as a Service (IaaS)** — Consumers control and manage the systems in terms of the operating systems, applications, storage, and network connectivity, but do not themselves control the cloud infrastructure.

End user application is delivered as a service. Platform and infrastructure is abstracted, and can be deployed and managed with less effort. Application platform onto which custom applications and services can be deployed. Can be built and deployed more inexpensively, although services need to be supported and managed. Physical infrastructure is abstracted to provide computing, storage, and networking as a service, avoiding the expense and need for dedicated systems.




Figure 2.56. Service Model Types

Deployment Models

Deploying cloud computing can differ depending on requirements, and the following four deployment models have been identified, each with specific characteristics that support the needs of the services and users of the clouds in particular ways (see Figure 3).

- **Private Cloud** — The cloud infrastructure has been deployed and is maintained and operated for a specific organization. The operation may be in-house or with a third party on the premises.
- **Community Cloud** — The cloud infrastructure is shared among a number of organizations with similar interests and requirements. This may help limit the capital expenditure costs for its establishment as the costs are shared among the organization. The operation may be in-house or with a third party on the premises.
- **Public Cloud** — The cloud infrastructure is available to the public on a commercial basis by a cloud service provider. This enables a consumer to develop and deploy a service in the cloud with very little financial outlay compared to the capital expenditure requirements normally associated with other deployment options.
- **Hybrid Cloud** — The cloud infrastructure consists of a number of clouds of any type, but the clouds have the ability through their interfaces to allow data and/or applications to be moved from one cloud to another. This can be a combination of private and public clouds that support the requirement to retain some data in an organization, and also the need to offer services in the cloud.




Figure 2.57. Public, Private, and Hybrid Cloud Deployment Example

Challenges

The following are some of the notable challenges associated with cloud computing, and although some of these may cause a slowdown when delivering more services in the cloud, most also can provide opportunities, if resolved with due care and attention in the planning stages.

- **Security and Privacy** — Perhaps two of the more “hot button” issues surrounding cloud computing relate to storing and securing data and monitoring the use of the cloud by the service providers. These issues are generally attributed to slowing the deployment of cloud services. These challenges can be addressed, for example, by storing the information internal to the organization, but allowing it to be used in the cloud. For this to occur, though, the security mechanisms between organization and the cloud need to be robust and a Hybrid cloud could support such a deployment.
- **Lack of Standards** — Clouds have documented interfaces; however, no standards are associated with these, and thus it is unlikely that most clouds will be interoperable. The Open Grid Forum is developing an Open Cloud Computing Interface to resolve this issue and the Open Cloud Consortium is working on cloud computing standards and practices. However, keeping up to date on the latest standards as they evolve will allow them to be leveraged, if applicable.
- **Continuously Evolving** — User requirements are continuously evolving, as are the requirements for interfaces, networking, and storage. This means that a “cloud,” especially a public one, does not remain static and is also continuously evolving.
- **Compliance Concerns** — The EU has a legislative backing for data protection across all member states, but in the US data protection is different and can vary from state to state. As with security and privacy mentioned previously, these typically result in Hybrid cloud deployment with one cloud storing the data internal to the organization.

2.4.2 Cloud access connectivity principles

For service developers, making services available in the cloud depends on the type of service and the device(s) being used to access it. The process may be as simple as a user clicking on the required web page or could involve an application using an API accessing the services in the cloud.

Accessing through Web APIs

Accessing communications capabilities in a cloud-based environment is achieved through APIs, primarily Web 2.0 RESTful APIs, allowing application development outside the cloud to take advantage of the communication infrastructure within it (see Figure 4).

These APIs open up a range of communications possibilities for cloud-based services, only limited by the media and signalling capabilities within the cloud. Today's media services allow for communications and management of voice and video across a complex range of codecs and transport types.

By using the Web APIs, these complexities can be simplified, and the media can be delivered to the remote device more easily. APIs also enable communication of other services, providing new opportunities and helping to drive Average Revenue per User (ARPU) and attachment rates, especially for Telcos.




Figure 2.58. Web 2.0 Interfaces to the Cloud

Communications Scalability

To deliver on the scalability requirements for cloud-based deployments, the communications software should be capable of running in virtual environments. This allows for easily increasing and decreasing session densities based on the needs at the time, while keeping the physical resource requirement on servers to a minimum.

Cloud connectivity option selection

Many network service providers (NSP's) have a range of options when it comes to cloud connectivity, though a lack of industry standards and confusing terminology can make things difficult to understand. Not so long ago, the only option available to connect to a Cloud Service Provider (CSP) was over the public Internet. However, with the rapid shift to cloud computing, customers quickly began to demand more - better security, lower latency, higher throughputs and increased reliability.

CSPs soon realised better end-to-end cloud performance wasn't going to be possible using the public Internet. They also understood that they didn't have the expertise or the infrastructure to manage interconnectivity between dozens of network service providers and colocation racks in their own data centres.

CSPs also quickly realised the answer was in the hundreds of carrier neutral data centres spread all over the world, also known as Internet Exchange Points (or IXPs). All network service providers were already present at these locations, so CSPs could extend their backbone connectivity to meet them there. This provided the potential for a direct physical link between the network service provider network and the cloud service provider network (known as a cross-connect), bypassing the regular Internet and providing a pseudo-private network. This interconnectivity, known as private peering, enabled direct, end-to-end connectivity and brought with it a whole range of security, latency and performance improvements (in addition to cost efficiencies for customers moving high volumes of data from cloud environments to their locations).

Today, cloud connectivity falls into two buckets, one that relies on the public Internet, and another that uses private, dedicated connectivity. Within these 2 buckets are typically 5 different connectivity options available (see figure below).

Internet connectivity	Dedicated connectivity
Public Internet	Ethernet
Public Internet with cloud prioritisation	MPLS IP VPN
	SD WAN

Figure 2.59. Cloud connectivity

We'll walk you through 5 cloud connectivity options and explain the pros and cons of each, so that you can choose the most suitable cloud access solution for your needs (see below).

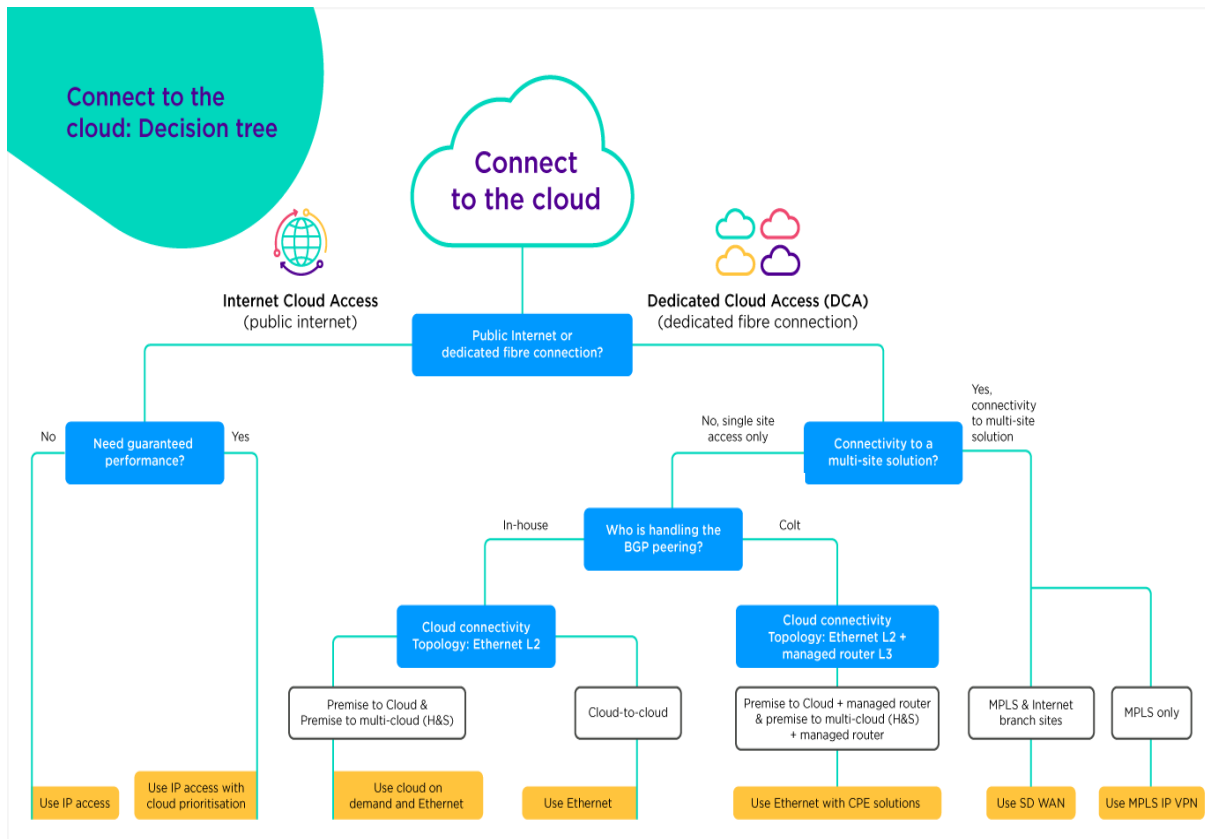


Figure 2.60. Connect to the cloud – decision tree

Cloud connectivity using the public internet

Arguably the cheapest and easiest way to connect to the cloud is through your standard Internet connection over the public Internet, sometimes referred to as IP access or IP transit.

Using your public Internet access is easy to set up and versatile, as accessing the cloud is just one of the many use cases for a standard Internet access connection. It provides a cost-efficient access method where you don't have specific performance needs.

However, accessing cloud applications via the public Internet can also result in performance inconsistencies and increased security risks. Historically, the term IP transit was used to reflect situations where providers had no direct access to the destination network and needed to 'transit' over other networks and network providers.

You can think of public Internet routes like a highway - they're dynamic and shared which can result in congestion at times, and when the most direct link is not available, data is routed through the next best option, which you have no control over resulting in packet loss and increased latency (delays). Additionally, multiple hand-offs between ISPs creates instability in the connection and increased risk.

Essentially the more pops and routers involved in delivering your data to its final destination, the more points of potential failure and a wider surface area for security attacks. Despite this, the growth of cloud connectivity via public Internet has shown no sign of slowing down. The public Internet remains by far the most common way to access the cloud (see figure below).

Advantages	Disadvantages
Best for single locations	A best-effort service not suited for critical applications
Cost-effective for low and medium data transfer volumes	Shared and dynamic routes mean no performance optimisation or guaranteed performance
Suitable for most topologies (premise/wan to single cloud, premise/wan to multi-cloud)	Not suitable for cloud-to-cloud connectivity
Use your existing business-as-usual internet connection	Becomes expensive for higher data transfer rates due to per Gigabyte out billing (egress)
Easy to get up and running, no need for a dedicated circuit	Exposed to security risks, such as DoS and DDoS attacks against routers and links
On demand delivery and scaling typically available	The least secure connectivity option

Figure 2.61. Cloud connectivity using the public internet (advantages and disadvantages)

Cloud connectivity using public internet and cloud prioritisation

Internet connectivity with cloud prioritisation enables you to dynamically reserve a portion of your normal Internet bandwidth for select cloud applications. Traffic prioritisation is effective for both incoming and outgoing traffic enabling a consistent, SLA-backed user experience specifically for your traffic to the cloud.

Cloud prioritisation is offered by network service providers that have direct peering services with cloud providers, such as Microsoft. For example, Microsoft Azure Peering Services (MAPS for short) enables end-users' direct access to Microsoft cloud services through certified network providers.

Once in place, your cloud traffic stays completely on your providers network, bypassing the public Internet and avoiding any other intermediary ISPs.

Cloud prioritisation combines the benefits of optimised routing and direct peering infrastructure with traffic prioritisation over the last mile, between the customer router and provider edge.

Advantages	Disadvantages
An add-on to standard Internet access services	Offerings are dependent on your connectivity and cloud providers
Consistent and guaranteed SLA-backed performance to the closest peering point	Layer 3 access only
Dynamically reserved bandwidth for cloud applications	No dedicated connection
Works for both incoming and outgoing bandwidth	
Optimised routing selects the shortest path to the cloud network edge	
Avoids network contention and unpredictable routing changes	
30 milli second Round Trip Delay (RTD)	
Traffic congestion control *	

* only available from some MAPS providers

Figure 2.62. Cloud connectivity using public internet and cloud prioritisation (advantages and disadvantages)

Direct Ethernet cloud connect

Dedicated connectivity through Ethernet connectivity services is the fastest and safest route for cloud connectivity, and the first of the Internet-bypass solutions. It is the result of service providers, like Amazon, Microsoft, Google, Oracle and IBM working together with network service providers to



enhance end-to-end cloud connectivity and automation capabilities - without touching the Internet. End-users are probably already familiar with the names of these CSP's direct interconnect programmes - like AWS Direct Connect, Microsoft ExpressRoute and Google Cloud Interconnect - that enable end-to-end secure connectivity through a Network Service Provider towards the customer location.

Direct Ethernet connectivity to the cloud renders performance, quality of service and security problems obsolete. It's provided by cloud on-ramps at data centres where the cloud service provider is present. This connects your premises or facilities through a NSP to the cloud provider via a dedicated layer 2 link.

Direct cloud connectivity provides the secure, high performance, end-to-end connectivity needed to run critical applications that can't be rivalled when only using the Internet. Cloud Service Providers typically charge data transfer fees - which are different when connecting to the Cloud through direct Ethernet connectivity vs. through the Internet, so direct connectivity can be particularly cost-effective if you are likely to be transporting large amounts of data out from your cloud environment (known as 'egress') towards your location.

Advantages	Disadvantages
Supports all topologies (Premise to cloud, premise to multi-cloud and cloud to cloud)	Only suitable for a single site (not multisite/WAN connectivity)
Bandwidth services up to 100Gbps available	Requires a dedicated circuit
Bandwidth is fully dedicated and guaranteed end-to-end	Customer to handle BGP peering
On demand delivery and scaling typically available	By <u>default</u> a layer 2 service, some NSP's provide managed router (L3)
End to end connectivity SLA with deterministic latency and performance	
Very suited and cost efficient for higher data transfer - due to lower price per Gigabyte (egress) out billing vs through the Internet	
Not subject to DDOS attacks as traffic bypasses the public Internet	

Figure 2.63. Direct Ethernet cloud connect (advantages and disadvantages)



MPLS IP VPN cloud connect

Integrating cloud connectivity into an IP-VPN (also known as IP-VPN cloud connect or MPLS-WAN technology) is a scalable and cost-effective way to access cloud services.

MPLS IP-VPN provides direct, high bandwidth and secure cloud connectivity to Cloud Service Providers. It's suited to customers that require secure access to the cloud across multiple sites and has traditionally been a common way for businesses connect to cloud providers.

The cloud connection is directly integrated into the IP VPN, so that it is completely private, with no reliance on the Internet. The cloud locations are integrated into the private WAN and effectively seen as another site (or sites) on the IP-VPN, meaning there is no need to redesign large corporate networks. Different customer locations in the IP-VPN then share the connectivity to access their resources in the cloud.

Advantages	Disadvantages
Very suitable for integration in existing and new MPLS IP-VPN networks	MPLS only, no Internet Branch sites
Highly secure, part of private IP-VPN	Layer 3 connectivity
No need to redesign large corporate networks	Dedicated connection required
Fully integrated in IP-VPN (any-to-any), avoids the need to backhaul traffic	Can increase latency – depends on where branch sites are located
Cost-effective as multiple locations on the IP-VPN share the connectivity toward the cloud	
Support different topologies: Single Cloud, Multi-Cloud and Cloud-to-Cloud	

Figure 2.64. MPLS IP VPN cloud connect (advantages and disadvantages)



SD WAN cloud connect

SD WAN (sometimes called SDWAN, SD WAN Cloud Access or SD WAN Multi-Cloud) can connect your software-defined WAN infrastructure to multiple cloud service providers (such as AWS, Microsoft Azure and Google Cloud) to enable direct, high performance and secure multi-cloud connectivity. Each branch office benefits from seamless end-to-end connectivity to your public cloud providers.

For cost-effective, direct connectivity into multiple cloud environments, SD WAN is likely the optimal solution.

SD WAN offers sophisticated and comprehensive connectivity capabilities, with features including prioritisation, optimisation, security, analytics, automated provisioning, and deployment. It brings together a single cohesive view of the enterprise network, tying together WAN sites, IaaS/SaaS cloud, and branch site connectivity, typically all within a single online portal. Coupled with on-demand capabilities such as zero touch site provisioning and real-time bandwidth upgrades, SD WAN is an extremely powerful solution.

Prior to SD WAN, traffic was typically backhauled to a central site or regional hub where a physical hardware stack provided functionality that was cost prohibitive to deploy at satellite sites (such as security and analytics). SD WAN now enables this functionality to be deployed in software on a common hardware platform. These software stacks comprise of various software functions that can be dynamically loaded and deployed in a modular fashion with a range of functionality, including:

- Networking & routing
- Analytics
- Security
- Traffic optimisation
- Remote access
- and more.

By tying together WAN sites and cloud infrastructure SD WAN can deliver end-to-end security, performance and visibility.

Building on MPLS IP VPN above, SD WAN offers private connectivity into multiple cloud providers in a single solution, combined with end-to-end performance backed by a SLA, end-to-end security, and end-to-end analytics.



Advantages	Disadvantages
The best way to manage multi-cloud infrastructures (MPLS and Internet branch sites)	Can require significant network changes and redesign to leverage all the benefits
Completely avoids the need to backhaul traffic from a branch site to a CSP or data centre	Newer services such as on demand capabilities may be limited
Bandwidth is fully dedicated and guaranteed end-to-end	Check support for your specific cloud provider (CSP) requirements
Automatic provisioning and deployment	Check support and roadmap for features and functionality such as application optimisation, analytics, SASE and more
Dynamic path selection - intelligent and dynamic routing to the best available path	Can increase latency – depends on where branch sites are located
Additional security features like FW/NAT to support the CSP public domain	
End-to-end visibility and management of the entire enterprise network	
Supports all topologies - WAN to cloud, WAN to multi-cloud and cloud to Cloud	
Also supports Internet only branch sites connecting directly to CSP through SD-WAN	

Figure 2.65. SD WAN cloud connect (advantages and disadvantages)

There is no ‘one-size-fits-all’ solution for enterprises as they connect to the cloud. Here are top ten questions and considerations to ensure you remain future-proofed by a new provider:



1. What level of partnership do you have with the major cloud providers?
2. How many public cloud points of presence do you have?
3. How many data centres are currently connected to your network?
4. How many offices are currently connected to your network?
5. Do you provide on demand capabilities via a self-serve software portal?
6. Are your data centre and cloud service provider neutral?
7. Who owns your fibre network - is it privately owned or leased from a 3rd party?
8. Do you provide end-to-end connectivity, including the last mile?
9. Do you provide guaranteed SLAs including for latency, packet loss and throughput?
10. What bandwidths are supported for cloud connectivity?

2.4.3 Cloud network set-up

While it often goes unnoticed by the average user, networks are implemented to isolate data from the outside world. Organizations rely on networking to connect their devices and integrate their systems across geographical barriers, while ensuring safe passage for information. This quick-start guide walks you through the basics of setting up your cloud network.

Virtual Network

Virtual networks can be thought of as separate networks within a larger network. Administrators can create a separate network segment consisting of a range of subnets (or a single subnet) and control traffic that flows through the cloud network. Depending on your business needs, you can implement your network using cloud technology from a cloud service provider (CSP).

The key difference for cloud administrators and architects when it comes to designing cloud networking solutions is the amount of control needed to have over the hardware. When you implement cloud networking with a CSP, you have little control over — and likely little knowledge about — the design of the CSP's network. Because of this limitation, virtual networks are often the go-to choice when you want to provide secure network isolation.

With a cloud solution, these virtual networks are known as VNets or Virtual Private Clouds (VPC). These act as a representation of a network in the cloud, giving you a cloud network.

Virtual networks provide the following benefits:





Figure 2.66. Virtual Networks

- **Isolation**

You can keep networks isolated from one another to ensure security and for purposes of development, quality assurance and deployment of cloud networks.

- **Internet Connectivity**

Each virtual network can be configured to access or deny access to the internet, or to limit access to specific destinations on the internet if needed.

- **Connection to Other Cloud Services**

Virtual networks often need a connection to CSP services. This allows the network to utilize services offered by the CSP. Providers typically allow for configuration of routing tables, domain name resolution, firewall and related items to manage the connections to your virtual networks.

- **Connection to Other Virtual Networks**

This allows you to interconnect your virtual networks when necessary while maintaining control over connections.

- **Connection to On-Premises Infrastructure**

Part of the flexibility of a virtual network is the ability to control connections. You can connect your virtual network to on-premises systems. Often this type of configuration is for end users to access a secure private cloud network or done as part of a hybrid cloud implementation.

- **Traffic Filtering**



Most secure connections involve filtering. Normally, this involves filtering items by source IP address and port, destination IP address and port, and particular protocol. This gives cloud computing engineers increased control over the communications occurring on your network.

Building Blocks of Cloud Network

As a cloud administrator or cloud computing engineer, your ability to create a virtual network will typically be reliant on virtual machine software or a cloud network provided by a CSP. Virtual machine software allows cloud administrators to designate and configure virtual network parameters associated with a host's physical Network Interface Card (NIC). When you configure multiple hosts to operate using the same parameters, you are adding those hosts to the virtual network. Virtual networks must have the following components:



Figure 2.67. Building Blocks of Cloud Network

- **Virtual Switch**

Virtual switches give you the capability to create segments on your network and connect those components together. You can connect one or more virtual machines to a virtual switch.

- **Virtual Bridge**

This component allows you to connect virtual machines to the LAN used by the host computer. The virtual bridge connects the network adapter on the virtual machine to the physical NIC on the host computer. Multiple virtual bridges can be configured to connect to multiple physical NICs.

- **Virtual Host Adapter**

The adapter makes it possible for your virtual machines to communicate with the host. Virtual host adapters are common in host-only and Network Address Translation (NAT) configurations. These cannot connect to an external network without a proxy server.

- **NAT Service**

NAT services allow multiple devices within your cloud network to connect to the internet.

- **DHCP Server**

The DHCP server allocates IP addresses to virtual machines and hosts. This applies to host-only and NAT configurations.

- **Ethernet Adapter**

This is a physical network adapter installed on hosts that connect to the network.



Many CSPs provide cloud services that make it easier to configure virtual networks and cloud networks. With cloud networks, you configure your virtual network and add your resources to them, rather than configuring them at the virtual machine level. Cloud networks also typically offer capabilities to simplify monitoring, management, connections and security.

Surveying Network Configuration Options

If you want to make use of a virtual network, you must also configure the following components:



Figure 2.68. Surveying Network Configuration Options

- **Subnets**

Subnets are a required part of a virtual network. You need TCP/IP subnets, which will designate addresses that are used on that network. Public and private address ranges are often used. When that's not possible, addresses are often assigned by CSPs. Virtual networks can be segmented into one or more subnets.

- **Routers or Routing Tables**

For any network, you must configure routers or routing tables on any virtual machine connected to the network so that packets can be routed appropriately.

- **DNS**

DNS server addresses must be provided, either assigned by you or your CSP.

- **CSP Region or Zones**

Virtual networks operating in different CSP regions must be specified. Doing so will also allow you to connect virtual networks in different regions. If needed, you can configure isolation between regions as well.

- **Traffic Filters**

Configuring your traffic filters to the specifications of your security protocols will only allow approved traffic to pass through your network. Filters can be applied at NIC in virtual machines, to a subnet or to a cloud service. When necessary, you will do this with a network virtual appliance.

Cloud Network Design Tips

When designing cloud networks, consider the following:

- As you design your cloud network, take the time to compare virtual network services offered by cloud providers. A hosted cloud network may be the only way you can create virtual networks the way you want them. Often, these cloud networks are easier to configure and manage.
- If you plan to filter traffic (and most companies should!), plan testing of the filter into your deployment to avoid future user complaints due to blocked traffic.
- If you choose to go with a CSP, work with their personnel to configure your cloud network components, such as routing tables, network virtual appliances and subnets. Save yourself some hassle up front.

Cloud Network's Ports and Protocol

One of the key steps you need to take to secure your cloud network is drilling down into the nitty gritty to uncover what people, services and technologies need access to the network. Ports are an essential part of your cloud network. The port is the endpoint of your connection.

Users connect to the cloud network through a designation port. All ports are assigned a number ranging from 0 to 65,535. The Internet Assigned Numbers Authority (IANA) separates port numbers into three ports, based on their numbers. TCP and UDP ports are assigned based on these ranges. Hackers commonly go after well-known ports but have been known to target open registered or dynamic ports, as well.

The three ports are:

- **Well-known Ports**

Preassigned to system processes by IANA, these include 0 to 1,023 and are most prone to attacks.

- **Registered Ports**

Available to user processes and listed by IANA, these registered ports go from 1,024 to 49,151 and are known to be too system-specific for direct target by hackers. However, hackers sometimes scan for open ports in this range. Don't turn your back, but you can avert your gaze occasionally.

- **Dynamic or Private Ports**

Assigned by a client operating system as needed, these are the ports numbered from 49,152 to 65,535. Dynamic ports are constantly changing (hence, the name dynamic), so it is difficult to directly target numbers. But again, hackers have been known to scan for open ports. As far as watching for hackers is concerned, maybe you can turn your back on dynamic or private ports, but not for too long!



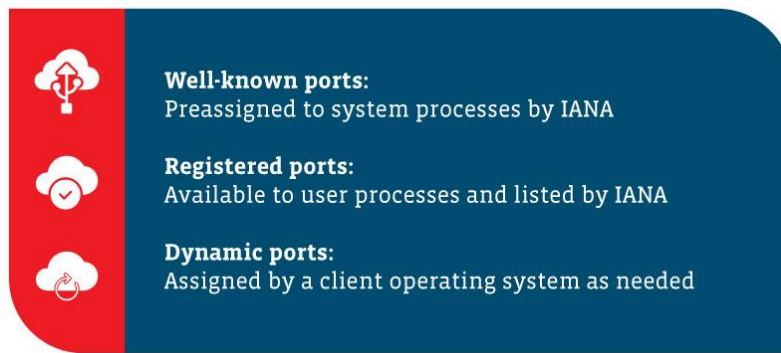


Figure 2.69. Dynamic or Private Ports

So, what are these ports used for? Here is a list of some of the most common default network ports used in the tech world:

- 21 FTP (File Transfer Protocol)
- 22 SSH (Secure Shell)
- 25 SMTP (Simple Mail Transfer Protocol)
- 53 DNS (Domain Name System)
- 80 HTTP (Hypertext Transfer Protocol)
- 110 POP3 (Post Office Protocol)
- 139 NetBIOS Session Service
- 143 IMAP (Internet Message Access Protocol)
- 443 HTTPS (Hypertext Transfer Protocol Secure)
- 3389 RDP (Remote Desktop Protocol)

Servicing Your Cloud Network

Services and apps that float among the cloud are similar in many ways to the services and apps that remain grounded in your on-premises infrastructure. Take cloud-based web apps and directory services, for example. Many will use the same ports and protocols that are used by their on-premise counterparts. Management tools, whether CSP-based, third-party or those built by your IT team, will also utilize port and protocol requirements.

If you decide to make the jump from the ground to the cloud, you will need to review your ports to determine what needs to be based in the cloud and what needs to remain housed on your own infrastructure. Take a close look at what needs internet access, in order to communicate with outside services or apps, and what type of access is required from inside the cloud.

Once you narrow it down, you can configure firewalls and set the necessary filters to ensure your cloud network will remain secure. As you work to deploy your cloud network, make sure you consult the following resources:



Figure 2.70. Servicing Your Cloud Network

- App and service configuration guides to identify the necessary ports and protocols each one uses.
- CSP security and deployment guides or white papers to locate the ports and protocols you need to access cloud services such as websites, databases, directory services and so on.
- Third-party deployment guides that are similar to the cloud network you are implementing.
- Your own (yes, your own) documentation to reference your firewall, routing, and other related information that could help you understand your own port and protocol usage. It will be tough to implement a successful cloud deployment if you have no idea from where you are jumping.
- If the fates forbid you from uncovering what ports and protocols are used by a legacy application that you want moved to the cloud, you might want to gather some helpful tools such as a port scanner or protocol analyser to unlock the guarded secrets of your predecessors.

Before launching any cloud network, take a fine-tooth comb through all your apps and services to ensure all ports and protocols are toeing the line.

Determine granting access to the Cloud Network

Before you go giving those magic entry passes away and granting access to your cloud network, consider these guidelines in addition to the information already provided:

- Don't assume you know all ports related to an app service. You know what assuming does, right? Don't be on the receiving end of that.
- Pay close attention to the direction of traffic flow when you are creating inbound and outbound rules for network access.

Cloud networks are still an emerging technology, showing lots of possibility for the future of IT.

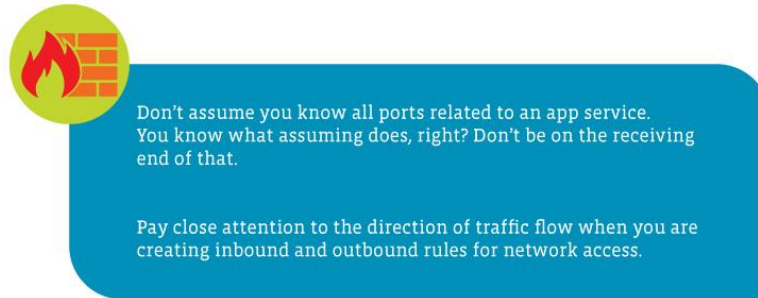


Figure 2.71. Determine granting access to the Cloud Network

2.5 Cloud System Management (Monitoring and Notification Service)

Difficulty Level: Easy

Completion Period:

Objectives

After reading the material, the reader will understand the concept of Cloud Management, Cloud Management systems and Monitoring Tools. You will also know the main goals and characteristic of Cloud Management, platforms, tools and vendors.



Figure 2.72. Cloud System Management

Achievements

After completing this application, you will be able to:

- know to what Cloud Management refers to

- how does Cloud Management work
- the importance of Cloud Management
- Cloud Management goals and characteristics
- the 4 types of Cloud Management
- know to what Cloud Monitoring refers to
- the challenges of Cloud monitoring
- analysis of Cloud management platforms, tools and vendors

Introduction of Cloud Management and Cloud Management Systems

What is Cloud Management?

Cloud management refers to the exercise of control over public, private or hybrid cloud infrastructure resources and services. A well-designed cloud management strategy can help IT experts control dynamic and scalable computing environments.

Cloud management is **the process of monitoring and maximizing efficiency in the use of one or more private or public clouds**. Organizations typically use a cloud management platform to manage cloud usage. Furthermore, Cloud Management is **a method of reviewing, observing, and managing the operational workflow in a cloud-based IT infrastructure**. Manual or automated management techniques confirm the availability and performance of websites, servers, applications, and other cloud infrastructure.

Why is Cloud Management being used?

Organizations are increasingly deploying enterprise applications to the cloud in order to reduce the high upfront investments they would otherwise have to make for on-site infrastructure. Public cloud environments provide on-demand computing power and data storage that is consistent with the growing, fluctuating demand for data and services. Through cloud service management, administrators oversee cloud activities ranging from resource deployment and utilization to lifecycle management of resources, data integration and disaster recovery.

How Does Cloud Management Work?

Summarizing all the above, Cloud management is a discipline which is facilitated by tools and software. To realize the control and visibility required for efficient cloud management, enterprises or any other interested party should see their hybrid IT infrastructure through a consolidated platform that pulls relevant data from all the organization's cloud-based and traditional on-premises systems.



Cloud management platforms help IT teams secure and optimize cloud infrastructure, including all applications and data residing on it. Administrators can manage compliance, set up real-time monitoring, and pre-empt cyberattacks and data breaches.

So how does it work? Typically, a cloud management system is being installed on a mentioned targeted cloud. After capturing information on activity and performance an analysis to a web-based dashboard is being send. There administrators can observe and react accordingly. If any issue may occur administrators can share comments back to the cloud through the cloud management platform.

Importance of Cloud Management

Business's/ Organizations are more likely to improve cloud computing performance, reliability, cost containment and environmental sustainability. The management of applications contains repetitive tasks which through Cloud Management servers and push code can be provisioned automatically via APIs instead of manual management.

Cloud Management can play an important role in managing the security status and vulnerability of IT assets.

Cloud Management goals and Characteristics

Without a doubt, the biggest challenge to cloud management is cloud sprawl (Cloud sprawl is the uncontrolled proliferation of an organization's cloud instances, services or providers) - IT staff loses track of cloud resources, which then multiply unchecked throughout the organization. Cloud sprawl can increase costs and create security and management problems, so IT shops need governance policies and role-based access controls in place.





Essential areas of cloud management include the automated and orchestrated instances and configurations, secure access and policy adherence, and monitoring at all levels -- all done as cost-efficiently as possible.

Figure 2.73. Cloud Management Components

Cloud management platforms provide a common view across all cloud resources to help monitor both internal and external cloud services. Management platform tools can help guide all individuals that touch an application's lifecycle. Regular audits can keep resources in check. Finally, consider third-party tools to help fine-tune enterprise usage, performance, cost and business benefits.

Metrics needs to be settled in order to help identify trends and provide guidance on what the user want to measure and track over time. There are plenty of potential data points, but every enterprise/ interested party should choose the ones that matter most to their business/ organization/ project.

More analytically the following needs to be considered:

- Data about the utilization of a compute instance's volume and performance (processor, memory, disk, etc.) provides insight about the application's overall health.
- Storage consumption refers to storage tied to the compute instances.
- Load-balancing services distribute incoming network traffic.
- Database instances help pool and analyse data.
- Cache instances use memory to hold frequently accessed data and thus avoid the need to use slower media, such as disk storage.



- Functions, also called serverless computing services, are used to provision workloads and avoid the need to supply and pay for compute instances. The cloud provider operates the service that loads, executes and unloads the function when it meets trigger parameters.

Types of Cloud Management

There are four (4) main types of computing which are categorized to **private clouds, public clouds, hybrid clouds, and multi-clouds.**

More analytically:

- **Private clouds:** are defined as computing services offered either over the Internet or a private internal network and only to select users instead of the general public. Also called an internal or corporate cloud, private cloud computing gives businesses/ organizations many of the benefits of a public cloud - including self-service, scalability, and elasticity - with the additional control and customization available from dedicated resources over a computing infrastructure hosted on-premises. Private clouds deliver a higher level of security and privacy through both company firewalls and internal hosting to ensure operations and sensitive data are not accessible to third-party providers.
- **Public clouds:** IT models where public cloud service providers make computing services—including compute and storage, develop-and-deploy environments, and applications—available on-demand to organizations and individuals over the public internet.
- **Hybrid clouds:** sometimes called a cloud hybrid—is a computing environment that combines an on-premises datacentre (also called a private cloud) with a public cloud, allowing data and applications to be shared between them.
- **Multi-clouds:** a company's/ organization's use of multiple cloud computing and storage services from different vendors in a single heterogeneous architecture to improve cloud infrastructure capabilities and cost. It also refers to the distribution of cloud assets, software, applications, etc. across several cloud-hosting environments.

Cloud Management and Monitoring Tools

Cloud monitoring is **a method of reviewing, observing, and managing the operational workflow in a cloud-based IT infrastructure.** Manual or automated management techniques confirm the availability and performance of websites, servers, applications, and other cloud infrastructure.

Cloud monitoring measures the conditions of a workload and the various quantifiable parameters that relate to overall cloud operations. Results are monitored in specific, granular data, but that data often lacks context.

Cloud observability is a process similar to cloud monitoring in that it helps assess cloud health. Observability is less about metrics than what can be gleaned from a workload based on its externally visible properties. There are two aspects of cloud observability: methodology and operating state.



Methodology focuses on specifics, such as metrics, tracing and log analysis. Operating state relies on tracking and addresses state identification and event relationships, the latter of which is a part of DevOps.

Cloud monitoring challenges

One of the biggest challenges to Cloud monitoring is for IT teams is to keep up with modern and distributed application designs. As applications evolve, IT teams needs always to adjust their monitoring strategies.

Effective cloud monitoring is a complex task. The tools that an organization currently uses may no longer be the ones they need, as different types of applications will need to be monitored in different ways.

Where does the success depend on?

The success of any cloud management strategy depends not just on the proper use of tools and automation but also on having a competent IT staff in place. IT and business teams must collaborate naturally in order to assimilate to a cloud culture and understand the business's/ organization's goals.

IT teams must also test cloud application performance, monitor cloud computing metrics, make critical infrastructure decisions, address patch and security vulnerabilities, and update the business rules that drive cloud management.

Business's/ Organizations that lack a skilled IT staff can always seek support from third parties. Third-party apps support budget threshold alerts that can notify finance and line-of-business stakeholders so they can monitor their cloud spending. Cloud brokerages often have a service catalog and some financial management tools. The time to scrutinize cloud spending is early on when apps go into production.

Cloud management platforms, tools and vendors

As cloud computing expands across the enterprise, a general cloud management platform can help deploy, manage and monitor all cloud resources. Enterprise IT must form a clear idea on what it wants to monitor before evaluating cloud management platforms to fit those needs -- whether it's individual tools that solve a single problem, such as network performance or traffic analysis, or a comprehensive suite that looks at everything. Some of these decisions will weigh tools from cloud providers, such as security tools from cloud platform vendors or from third-party providers.

The most comprehensive cloud management products offer features that cover these five categories:

- automation and orchestration for applications and individual VMs;
- security, including identity management and data protection and encryption;



- policy governance and compliance, including audits and service-level agreements;
- performance monitoring;
- cost management.

Many multi-cloud management vendors offer a range of tools, each with strengths and weaknesses. Some of the more prominent ones are VMware (a virtualization and cloud computing software provider based in Palo Alto, Calif. Founded in 1998, VMware is a subsidiary of Dell Technologies), CloudBolt Software (a hybrid cloud management platform developed by CloudBolt Software for deploying and managing virtual machines (VMs), applications, and other IT resources, both in public clouds (e.g., AWS, MS Azure, GCP) and in private data centers (e.g., VMware, OpenStack)), Snow Software (which acquired Embotics, is a market-tested developer of software asset management tools), Morpheus Data (a novel approach in providing remote access to micro data of official statistics), Scalr (an information technology (IT) vendor that offers a management platform for cloud computing) and Flexera (specializes in IT management software, optimization & solutions). Also in this mix are traditional IT service management vendors, such as BMC Software (baseboard management controller (BMC) is a specialized service processor that monitors the physical state of a computer, network server or other hardware device using sensors and communicating with the system administrator through an independent connection), CA Technologies (one of the largest independent software companies in the world. The company, which was formerly known as Computer Associates International, is an American multinational publicly held corporation), Micro Focus (a British multinational software and information technology business) and ServiceNow (a cloud-based workflow automation platform that enables enterprise organizations to improve operational efficiencies by streamlining and automating routine work tasks), which typically serve big companies with ITSM governance processes (ITSM: IT Service Management Software).

IT shops that use a single public cloud might want to stick with tools offered by that service provider because such tools are designed to enhance those native management platforms. For cloud monitoring, Google Cloud Operations (formerly Stackdriver) monitors Google Cloud as well as applications and VMs that run on AWS Elastic Compute Cloud. Microsoft Azure Monitor collects and analyses data and resources from the Azure cloud. There are also many open source cloud monitoring options for enterprises comfortable working with open source tools.



3 APPLICATIONS

3.1 Access to a database using a person's fingerprint as a password

Goal

Databases sometimes contain very important data for some companies or organizations. Accessing this data is allowed for a small number of people. In order to increase the level of security, access to these data must be based on specific indexes for the persons who have the right of access. The application allows access to the database based on the fingerprint of the people who have the right to access the database.

Expected timeframe to create value

3 weeks – 2 months

3.2 Active Directory Server

Goal

The goal of an Active Directory (AD) Server is to provide a centralized location for managing network resources, such as user accounts, computers, and printers. It is a database repository that stores information about all users and devices connected to a network and allows authorized users to access resources on the network. The AD server acts as a directory service and is responsible for managing the authentication and authorization process for users attempting to access network resources. This enables system administrators to enforce security policies across their organization, ensuring that only authorized users can access specific resources. The AD server also allows for the delegation of administrative tasks to different individuals or groups, which can improve manageability and efficiency within an organization. Overall, the primary goal of an AD server is to simplify network administration, improve security, and provide a centralized management location for all network resources.

Expected Timeframe to Create Value

The expected timeframe to create value with an Active Directory (AD) Server depends on the organization's specific needs and requirements. However, some benefits can be realized soon after deployment, while others may take longer to achieve.

In terms of immediate benefits, an AD server can simplify network administration by centralizing user management. This can improve efficiency and reduce the time and effort required for common IT tasks, such as resetting passwords or creating new user accounts. Additionally, AD can greatly enhance network security by providing a centralized location for enforcing security policies and managing access to network resources. This can help reduce the risk of security breaches and unauthorized access to company data.



Other benefits, such as improved scalability and flexibility, can take longer to realize. For example, the AD infrastructure can support the organization's growth over time by providing a scalable and reliable foundation for user management and authentication. This can help reduce costs and increase efficiency as the organization expands.

Overall, the expected timeframe to create value with an AD server depends on various factors, such as the size and complexity of the organization and the specific technical requirements of the deployment. Some benefits can be realized immediately, while others may take longer to achieve. Nonetheless, an AD server can be a valuable investment in terms of reducing administrative overhead, increasing security, and providing a scalable infrastructure for the organization's long-term growth.

3.3 AI Behaviour analysis Systems

Goal

The goal of AI behaviour analysis systems is to analyse and interpret human behaviour patterns and predict future behaviour based on data-driven insights. It seeks to provide a deeper understanding of human behaviour and decision-making, and identify potential risks, threats or opportunities in various domains such as policing, healthcare, security, and marketing. By leveraging machine learning algorithms and data mining techniques, these systems aim to identify patterns and anomalies in behaviour that could indicate potential threats or issues. The ultimate goal is to leverage insights from behaviour analysis to improve decision-making, reduce risks, and enhance outcomes in numerous fields.

Expected Timeframe to Create Value

The timeframe to create value from AI behaviour analysis systems depends on several factors such as the complexity of the problem being solved, the quality and accessibility of data, and the technology used.

In simpler scenarios, value can be created relatively quickly, such as within a few months. For example, if a company is using behaviour analysis systems to optimize its marketing strategies, it may see results in as little as a few months. On the other hand, more complex scenarios, such as using behaviour analysis systems to detect fraud or prevent security breaches, may require more Expected timeframe to create value and may take several years to fully realize.

Overall, a well-implemented AI behaviour analysis system can bring immediate benefits of improved decision-making and risk mitigation, but the full potential of such systems may take longer to materialize. As the algorithms become more advanced and the data sets become more comprehensive, the value created by these systems will likely continue to increase over time.



3.4 Application for managing the activity of renting tools and equipment from a company to natural persons

Goal

In many situations, individuals who carry out repair activities at their own homes need specific tools for these activities. Some repair or construction activities are rarely carried out and it is not justified to buy the tools or equipment that are necessary for that activity.

One solution is to rent this equipment from companies that have this object of activity. The application manages the activity of renting tools and equipment of a company to people or other companies that use this equipment.

Expected timeframe to create value

1 week – 1 month

3.5 Application for monitoring autonomous room cleaning equipment (vacuum cleaners) at the headquarters of small and medium-sized companies or in private homes

Goal

The application allows monitoring the activity of a vacuum cleaner or several vacuum cleaners that operate autonomously in a closed space. Vacuum cleaners are used for cleaning living rooms or offices. Vacuum cleaners that can be operated by remote control and that can move autonomously without being carried by a person, make cleaning a room easier.

These vacuum cleaners are equipped with different types of sensors that detect the proximity of an obstacle and change the direction of movement of the vacuum cleaner. The direction of movement of the vacuum cleaner depends on the way in which the vacuum cleaner's operating algorithm was written by the manufacturer.

The application creates an algorithm for moving the vacuum cleaner in the space of the room so that the cleaning operation is efficient.

Expected timeframe to create value

4 weeks – 3 months



3.6 Asset Tracking

Goal

The goal of asset tracking is to monitor and manage the physical location and condition of assets such as equipment, materials, and products as they move through the supply chain. Asset tracking systems use advanced technologies such as radio frequency identification (RFID), global positioning system (GPS), and barcoding to provide real-time information about the location, status, and movements of assets.

Some of the key goals of asset tracking include:

1. **Visibility:** Asset tracking systems provide visibility into the location and status of assets, allowing organizations to know where their assets are at all times.
2. **Compliance:** Asset tracking systems help organizations comply with regulations by providing reliable data on the movement and handling of regulated assets such as pharmaceuticals and hazardous materials.
3. **Efficiency:** Asset tracking systems minimize the need for manual inventory checks and improve supply chain efficiency by providing real-time information on asset movements.
4. **Cost reduction:** Asset tracking systems can reduce the costs associated with lost, stolen, or misplaced assets, and can reduce the time and labour required to manage inventory.
5. **Improved decision-making:** Asset tracking systems provide data that can be used to support better decision-making, such as optimizing supply chain operations, forecasting future demand, and identifying inefficiencies.
- 6.

Overall, the goal of asset tracking is to provide organizations with the real-time data they need to effectively manage their assets, improve supply chain performance, reduce costs, and make informed decisions about their operations. By leveraging these insights, organizations can enhance their operations, improve their customer experience, and gain a competitive advantage in their industry.

Expected Timeframe to Create Value

The expected timeframe to create value from asset tracking solutions will depend on the specific needs of the organization and the complexity of the asset tracking solution being deployed. However, in many cases, organizations can expect to see the benefits of asset tracking within a few months to a year of implementation.

In the short term, asset tracking can provide immediate benefits such as reducing the risk of lost or stolen assets, improving inventory accuracy, and optimizing asset utilization. These benefits can be achieved relatively quickly, often within a few weeks or months of implementation.



In the longer term, the value created by asset tracking can increase as the organization gains better visibility into their supply chain operations and identifies opportunities for optimization and improvement. This can lead to further cost savings, higher customer satisfaction, and improved efficiency.

As technology continues to evolve and asset tracking solutions become more advanced, the potential for value creation will continue to grow. Machine learning and predictive analytics, for example, can be used to identify patterns and trends in asset movements, enabling organizations to anticipate disruptions in the supply chain and take preventative action.

Overall, the expected timeframe to create value from asset tracking solutions will vary depending on the specific needs of the organization. However, by implementing an asset tracking solution, organizations can expect to see a positive impact on their operations, efficiency, and bottom line within a relatively short timeframe.

3.7 Attendance tracker for students

Goal

Attendance system is a system that is used to track the attendance of a particular person and is applied in the industries, schools, universities or working places. The traditional way for taking attendance has drawback, which is the data of the attendance list cannot be reuse and tracking and tracing student's attendance is harder. The technology-based attendance system such as sensors and biometrics-based attendance system reduced human involvement and errors. Thus, in this paper, a NFC-based attendance system is presented. A comparative study between this both NFC and RFID is also discussed thoroughly, especially in terms of their architectures, functionality features, benefits, and weakness. Overall, even both NFC and RFID attendance system increases the efficiency in recording attendance, NFC system is providing more conveniences and cheaper infrastructure in both operational and setup cost.

Expected Timeframe to Create Value

3 – 6 months

3.8 Automated Facilities Management

Goal

The goal of automated facilities management is to use technology to streamline and automate building and facility management processes to improve operational efficiency, reduce costs, and enhance the building occupant experience. This includes the use of smart building technologies that enable the



monitoring, control, and optimization of various building systems, including HVAC, lighting, security, and energy usage.

Some of the key goals of automated facilities management include:

- **Improved operational efficiency:** By automating facility management processes, organizations can reduce the time and resources needed to manage their facilities, enabling them to focus more on core business activities.
- **Reduced costs:** Automated facilities management can help organizations reduce energy consumption, minimize maintenance costs and optimize resource allocation.
- **Enhanced building performance:** By leveraging data analysis and real-time monitoring, automated facilities management systems can detect and solve building performance issues quicker, resulting in better building performance and lower operating costs.
- **Improved occupant experience:** Automated facilities management can improve the occupant experience by providing more comfortable and secure environments through real-time monitoring and optimization of various building systems.
- **Compliance:** By automating and standardizing processes, automated facilities management can help organizations comply with regulations and guidelines, reducing the risks of fines, penalties, and litigation.

Overall, the goal of automated facilities management is to leverage technology to enable organizations to achieve better overall management of their buildings and facilities. By enhancing efficiency, reducing costs, and improving the occupant experience, organizations can become more competitive and better serve their customers.

Expected Timeframe to Create Value

The expected timeframe to create value from automated facilities management will depend on several factors, such as the size and complexity of the building or facility, the type of technology used, and the specific objectives of the organization.

In many cases, organizations can expect to see measurable benefits from their automated facilities management systems within a few months to a year of implementation. These benefits may include:

- **Reduced energy consumption:** Automated facilities management can optimize various building systems, reducing energy consumption and resulting in lower energy bills.
- **Streamlined maintenance processes:** By automating maintenance processes, organizations can reduce the need for manual intervention, save time, and reduce costs.
- **Improved occupant experience:** Automated facilities management can improve building comfort levels, resulting in increased occupant satisfaction.



- Better operational efficiency: Automated facilities management can streamline various building management processes, resulting in improved efficiency and reduced organizational costs.
- Predictive maintenance: By adopting predictive maintenance, organizations can improve the lifespan of their building systems and reduce repair costs.

Overall, the expected timeframe to create value from automated facilities management will depend on the specific needs of the organization and the complexity of the systems being implemented. However, by leveraging the benefits.

3.9 Automation of tasks using cloud-based services: recommendation engine

Goal

Market Basket Analysis is a modelling technique based upon the theory that if you buy a certain group of items, you are more (or less) likely to buy another group of items. In retailing, most purchases are bought on impulse. Market basket analysis gives information as to what a customer might have bought if the idea had occurred to them.

Expected Timeframe to Create Value

1 – 6 months

3.10 Back-Up / Disaster relief

Goal

Having an automatic system for backing up critical data across several different regions to minimise the risk of catastrophic failures, so that if there is a failure at an entire region, the back-ups would be unaffected, in contrast to having back-ups set up on different servers within the same region, where a total region failure would result in data loss even with the back-ups.

Expected timeframe to create value

N/A

3.11 Chatbot for indicating free places in public parking lots in a city

Goal

A problem faced by all car drivers is the need to find free places in public parking lots in the city as close as possible to the place where we want to go. This is quite difficult because the driver is in traffic and must orient himself according to the situation in the area.

A solution that solves the problem and eases the driver's task is a chatbot type application on the



driver's mobile phone. The driver communicates with the application by voice and finds out in advance the situation with free parking spaces in the parking lots located near the area where the driver has problems to solve.

Expected timeframe to create value

1 month and half – 10 months

3.12 Chatbot to personalize the learning activity of students in vocational high school education

Goal

Classically, students learn by reading the lesson written on paper or by reading the lesson written in electronic format in a word file or similar files (pdf, etc.). The lesson is usually followed by a set of questions through which the student can check how he learned the lesson.

The application that is proposed now helps the student to learn in an interactive way with more efficiency.

Expected timeframe to create value

1 month – 6 months

3.13 Chatbot for students in EDU institution

Goal

Many software companies try to build at least simple FAQ/Q&A based chatbot recently. Recent works shows that it is really easy to build a bot while to build intelligent one could be an extremely hard (and expensive). Domain specific bots like AI driven Support Centre Automation Bots should consider to be interoperable on many levels and with every new level, level of complexity grows exponentially. In recent years, messaging apps overtaken social networks and become the dominant platforms on smart phones. That enormous potential should be considered to solve one of the issues that any organization larger than 10 participants has. Combining various existing and external data sources company already have access to, most of the first-and second line helpdesk questions could be resolved before they came to support service staff. Robotic Process Automation (RPA) is one of hottest topics among business process experts while one of the fastest growing fields of RPA is Knowledge Mining which is especially applicable in Educational (EDU) environment like any kind of EDU Support System.

Expected Timeframe to Create Value

3 – 9 months



3.14 Cloud-based e-learning

Goal

The increasing research in the areas of information technology have a positive impact in the world of education. The implementation of e-learning is one of contribution from information technology to the world of education. The implementation of e-learning has been implemented by several educational institutions in Indonesia. E-Learning provides many benefits such as flexibility, diversity, measurement, and so on. The current e-learning applications required large investments in infrastructure systems regardless of commercial or open-source e-learning application. If the institution tended to use open-source e-learning application, it would need more cost to hire professional staff to maintain and upgrade the e-learning application. It can be challenging to implement e-learning in educational institutions. Another problem that can arise in the use of e-learning trend today is more likely to institution building their own e-learning system itself. If two or more institutions are willing to build and use an e-learning so they can minimize the expenditure to develop the system and share learning materials more likely happened. This paper discuss the current state and challenges in e-learning and then explained the basic concept and previous proposed architectures of cloud computing. In this paper authors also proposed a model of cloud-based e-learning that consists of five layers, namely: (1) infrastructure layer; (2) platform layer; (3) application layer; (4) access layer; and (5) user layer. In addition to this paper, we also illustrated the shift paradigm from conventional e-learning to cloud-based e-learning and described the expected benefits by using cloud-based e-learning.

Expected Timeframe to Create Value

6 – 12 months

3.15 Communication/ Information Exchange Application/ Channels

Goal

The goal of communication/information exchange applications is to enable seamless and efficient communication and sharing of information between individuals or groups. These applications provide users with a platform to connect with others, collaborate, and access information in real-time, regardless of their location.

With communication/information exchange applications, users can share documents, files, and other forms of data, conduct audio and video conferencing, instant message, and share screens. The ultimate aim is to improve productivity, enhance collaboration, and streamline workflows.

In addition, these applications often provide security features such as end-to-end encryption to protect sensitive information. Some communication/information exchange applications also include AI-



powered features such as document translation, sentiment analysis, and automatic transcription to make communication more efficient and effective.

Overall, the goal of communication/information exchange applications is to facilitate effective communication and collaboration, leading to improved performance, enhanced customer satisfaction, and increased profitability for businesses and organizations.

Expected Timeframe to Create Value

The expected timeframe to create value for a communication/information exchange application may vary based on several factors, including the application's complexity, scope, and the technology stack used to develop it.

For smaller applications with limited functionalities, the value can be created within a few weeks or months. Such applications can be a simple messaging or file-sharing platform that aims to connect remote workers or teammates.

For larger applications with complex functionalities such as group video calls, interactive whiteboards, document collaboration, and other advanced features, it may take several months or even years to create value.

The development time will also depend on the team's resources, experience, and the methodology used in building the application. The agile methodology that involves iterative development and regular feedback from users can help to shorten the development lifecycle and quickly create value.

Overall, a communication/information exchange application can create value as soon as it becomes operational and starts to facilitate efficient collaboration and improve productivity for the users. The key is to focus on creating an application that meets the users' needs, is easy to use, and provides a satisfying experience that will keep them using the application for the long term.

3.16 Continuous monitoring of the operation of some industrial installations using cloud computing and IoT technologies

Goal

The industrial installations that belong to some companies can present a danger if the values of some parameters that characterize these installations are outside the range of normal operation.

An example is the tank storage that contains the liquefied propane-butane mixture that is used to propel the liquid inside the spray containers.



This gas is bottled in smaller containers together with the liquid to be sprayed by pressing a button. The application monitors certain quantities (gas pressure, tank temperature, etc.) of the installation. When the monitored parameters approach dangerous values, measures are taken to bring the operation of the installation to normal parameters.

Expected timeframe to create value

3 weeks 4 months

3.17 Continuous patient monitoring

Goal

A system that uses sensors together with the IoT-hub that can remotely monitor a patient's vitals and raise warnings if certain levels go outside certain thresholds.

Expected timeframe to create value

1 year – 1.5 years

3.18 Create test environments

Goal

Provision and create the resources needed to run test versions of existing deployments so that new functions or possible bug fixes can be written and run without disturbing the current running deployment.

Expected timeframe to create value

1 week

3.19 Creating a didactic application to help students learn a foreign language

Goal

The application is developed for educational purposes. It is designed to facilitate learning a foreign language. There are foreign languages in which for each sound (phoneme) that makes up a word when it is pronounced, the same graphic element (grapheme) is used to record the pronounced word in writing. In other languages, for the same sound, two or three combinations of graphic elements are used to fix the word in writing.

To facilitate the learning of a foreign language, the international phonemic alphabet was introduced, which always uses the same graphic symbol to record the same spoken sound in writing. This makes learning a foreign language easier.



By using the application, the student can learn the correct pronunciation of the words of the foreign language to be learned.

Expected timeframe to create value

3 weeks – 3 months

3.20 Data backups and archiving

Goal

Backup. Archive. Data Storage Method. **The original data remains in place, while a backup copy is stored in another location.** Archived data is moved from its original location to an archive storage location.

Living in a world where cybercrime is the order of the day. No day will pass without cases of major data breaches, which at times become fatal for quite a number of businesses.

Traditional methods of data backup have proven effective in backing up data for a long time. Nonetheless, they are prone to viruses, and due to their portable nature, they can get lost and pose a threat to modern businesses.

Cloud-based backup and archiving is a solution to these challenges. It is easy to implement and provides maximum data security. With this approach, you can backup or archive your sensitive files to cloud-based storage systems. This provides the assurance that your data is still intact even if your live data becomes somehow compromised.

Some cloud computing services allow you to schedule backups to meet your unique needs. Additionally, you can encrypt your cloud backups and make it impossible for hackers and snoopers to access.

With cloud storage, you can get as much space as you require and store as much data as you need and will only pay for what you actually use.

Expected Timeframe to Create Value

The timeframe for creating value may vary depending on the organization's specific needs and the level of backup and archive implementation, but the benefits can start to accrue from the very first implementation.



3.21 Data loss prevention cloud-based system

Goal

DLP (Data Loss Prevention) is a security tool for data protection, and its complexity and technological advancement contribute to a very small understanding of the functionality and capabilities of the tool. With different names and technological approaches, it can be difficult to grasp the ultimate value of the tool and the one that best suits the environments. There is a diverse understanding of what the DLP solution is. Some people consider it to be encrypting or controlling the USB input of DLP while others are looking wider.

DLP is defined as: products that, based on central policies, identify, monitor, and protect sleep, movement, and usage data through deep content analysis. Key DLP features are:

1. content analysis
2. central policy management
3. cover content on multiple platforms and locations

DLP solutions protect sensitive data and help organizations better understand their data and improve their ability to manage content.

Expected Timeframe to Create Value

6 – 12 months

3.22 Data management system about a company's employees

Goal

In general, small companies with a small number of employees manage employee data using excel or word files in which they create tables. The data of the employees are written in these tables. These tables do not have a unitary format and changing some data about employees sometimes requires making changes in several files.

When changes are made, it is possible that some files are not modified and in other files, important data about employees may be accidentally destroyed. In addition, accessing information about an employee requires viewing all excel or word documents.

The proposed application offers the possibility to store employee data in a relational database stored either on one's own server or on another company's server. In addition, the graphic interface used is friendly and suggestive.

Expected timeframe to create value

2 weeks - 2 months



3.23 Digital asset certification using distributed ledger/blockchain

Goal

The main features of blockchain are transparency and decentralization, which today's systems cannot boast. Digital identity combined with blockchain technology will enable people to perform tasks that are faster, simpler and safer, including proof of identity, facts, status and data. Incredibly, the fact is that searching for new employees, checking candidate data, and job application itself could be a process that would take just a couple of mouse clicks on the computer, with the utmost certainty of the data being obtained. But blockchain is just offering it. By placing all the information on our identity on it, with cryptography that makes the whole thing safe and transparent and always accessible through the internet, we spend all the time spent on proving identity, data, facts, and state of affairs on the most important things. Imagine that we can also enclose 3 cryptographic keys with the application for business, so that the employer can easily check with the absolute certainty that we have actually completed the college we have stated in his CV, whether we are unhappy and whether we are at all a person who claims to be. This process would take about a few minutes, while the same process lasts for several days, if not weeks, as the data verification is done by writing queries in each of these systems from which data comes.

Expected Timeframe to Create Value

3 – 9 months

3.24 Digital identity

Goal

Identity is very valuable to us, and not to institutions, we are not behaving accordingly. Because of the lack of awareness and education about identity itself, because of the digital and physical centralization of databases and data about our identities, which creates unavoidable weaknesses that undermine the systematic value of our personal data. Centralized systems are a good booty for attackers with bad intentions because, if they break into the system, they can easily steal (copy) large amounts of data stored in that system. We have witnessed a lot of attacks on centralized systems, not small business systems, but big and globally influential companies such as Yahoo, eBay, Adobe, JP Morgan Chase, Sony and many others.

Blockchain technology offers the solution to this problem that is becoming more and more constant due to constant needs, increased demand and the use of digital identity. But, as we mentioned earlier, this is a new technology and is just in the early stages of the project and we are still investigating all the possibilities and the application of this technology. With the need to prove our identity, we meet each day and in different place. At work, in a bank, in a shop, in travel, in state institutions and in many different places. Currently there are many new and prospective projects and young companies dealing



with this problem and are trying to find their place in the market. In this part we will mention some of them and more specifically explain their business models.

Expected Timeframe to Create Value

1 – 3 months

3.25 Digital twinning

Goal

Create a virtual environment based on a real-world system, using sensors and IoT capabilities and explore possibilities and consequences of changing the environment, monitoring the health of systems so that maintenance and repairs can be done as the need occurs rather than having scheduled inspections.

By observing the data gathered from the sensors, you can simulate changes in the environment to see how the system would respond and gain insights into how to improve the performance of the system. For example, it could be used to improve the performance for a ventilation system by more dynamic usage where it increases workflow at peak times and areas and saving energy when there is less need creating both a more pleasant environment and decreasing energy costs.

Expected timeframe to create value

6 months – 1 year

3.26 Disaster prevention platform

Goal

Using internet enabled environment sensor devices which send data to a cloud-based analysis server which generate alarm and reports based on the analysis of the data.

A comprehensive monitoring solution for collecting, analysing, and responding to telemetry from your cloud and on-premises environments. Maximizing the availability and performance of your applications and services.

Collecting and integrating the data from every layer and component of your system into a common data platform. It correlates data across multiple subscriptions and tenants, in addition to hosting data for other services. Because this data is stored together, it can be correlated and analysed using a common set of tools. The data can then be used for analysis and visualizations to help you understand how your applications are performing and respond automatically to system events.



Expected Timeframe to Create Value

The timeframe to create value for such a platform would depend on various factors such as the complexity of the platform, the resources available, and the level of expertise of the team developing the platform. In some cases, the value of a disaster prevention platform may be immediately apparent, while in other cases, it may take some time to analyse and measure its effectiveness. Ultimately, the success and value of a disaster prevention platform would depend on its ability to prevent or mitigate the impact of disasters.

3.27 Distribution of parcels in a geographical region with the help of autonomous drones

Goal

A few centuries ago, trained pigeons (also called passenger pigeons) were used to transmit messages between the sender and the recipient. At least that's what the stories that have reached us say.

There were advantages and disadvantages of this way of transmitting messages. One way to transmit parcels between a distribution point and various recipients can be using autonomous drones.

Expected timeframe to create value

3 weeks – 6 months

3.28 Document similarity detection and document information extraction system

Goal

It is a human tendency to formulate assumptions while analysing the difficulty of information extraction in documents. We automatically assume it is easier to extract information in the form of named entities from a set of similar documents. Nonetheless, similar-looking documents have a distinct set of problems. The named entities in these document types vary in size, akin to the number of characters, words, height, width, and location. These variations cannot be handled using heuristics or pre-trained language models.

Expected Timeframe to Create Value

6 – 12 months

3.29 Document translation

Goal

Translating documents that contain descriptions of the products sold on a website, so that the business can cater to a broader demographic and can result in increased sales numbers. Making sure a website



is available in different languages is particularly important when trying to appeal to an international audience, or the business is supplying an area or industry that consists of a lot of non-majority speakers. Being in a multilingual space it can also be useful to automatically translate any user-generated content, such as product reviews, or to create and maintain a database of frequently asked questions in several different languages.

Expected timeframe to create value

3 months – 6 months

3.30 Dynamic website hosting

Goal

A web hosting environment contains details specific to the application, such as where the Application is stored at and functions and services essential to managing the entire application. Most common types of web hosting are: static hosting, dynamic hosting and local hosting.

Expected Timeframe to Create Value

1 – 3 months

3.31 Dynamic website with data storage in a database

Goal

Their websites are very popular these days and allow information to be displayed in an attractive and friendly way. The information contained in these sites is in the form of text or images. Some sites can have many pages depending on their purpose. Often the information they transmit has to be changed relatively often due to certain conditions.

For example, a pizzeria that has a web page changes the menu daily. The web page must be updated daily. In this case, the owner of the pizzeria (I chose a pizzeria as an example, but there may be other examples) must contact the person who made the website daily to update the information. A dynamic site that displays information using a database is welcome.

Expected timeframe to create value

2 weeks - 1 month

3.32 E-commerce Application

Goal

The primary goal of an e-commerce application is to enable electronic commercial transactions between businesses and consumers over the internet. E-commerce applications enable businesses to sell their products and services online, and consumers to buy those products and services over the internet.

The goal of an e-commerce application is to provide consumers with a seamless shopping experience while also providing businesses with a cost-effective way to sell their products and services. The application should be designed to be intuitive, easy-to-use, and provide convenient payment options, enabling customers to shop with ease.

Additionally, an e-commerce application should be designed to provide businesses with robust reporting and analytics capabilities, enabling them to track sales data, inventory levels, customer buying patterns, and other key metrics. This helps businesses to identify trends and make data-driven decisions that promote business growth and success.

Overall, the goal of an e-commerce application is to facilitate secure and convenient online purchasing while making it easy for businesses to manage their online transactions. By providing customers with a streamlined, user-friendly shopping experience and businesses with effective management tools, an e-commerce application can significantly increase sales, revenue, and market share for businesses.

Expected Timeframe to Create Value

The timeframe to create value from an e-commerce application depends on several factors, including the size and complexity of the application, the level of customization required, and the resources available for development.

Typically, it can take several months to a year to design, develop, test, and launch an e-commerce application. However, businesses can start generating value from an e-commerce application even before it's fully completed if they follow an agile development approach, which enables them to deliver small increments of value to customers more quickly.

In the early stages of e-commerce application development, businesses should focus on creating a minimum viable product (MVP) that provides a basic set of features and functionality for customers to shop online. This allows businesses to validate their assumptions and test the market before investing more resources into developing additional features.

Once the MVP is launched, businesses can start generating value by measuring key performance metrics such as website traffic, conversion rates, and customer satisfaction levels. They can use this



data to iterate and improve the application continuously, adding new features and functionality to drive customer engagement, sales, and revenue.

Overall, while the expected timeframe to create value from an e-commerce application may depend on various factors, businesses can start realizing its benefits from the early stages of development and continually improve and enhance its features over time to drive customer engagement and growth.

3.33 Electronic catalogue with students' school results

Goal

The application records in a database the results obtained by high school students in the subjects studied at school. The application analyses the results of each student and when the results are below the passing limit or are close to the limit, it notifies this by sending an email or a warning message on the mobile phone to their parents.

Expected timeframe to create value

1 week - 1 month and a half

3.34 Facilities Access Control

Goal

The goal of Facilities Access Control is to ensure that only authorized individuals have access to a particular physical location or facility. Access control helps to prevent unauthorized access, theft, and vandalism, and can also help maintain employee safety and security. By implementing access control measures, an organization can protect sensitive areas of the facility from unauthorized entry, safeguard assets and information, and reduce the risk of harm to employees.

Facilities Access Control typically involves utilizing an electronic system which will require authorized persons to present credentials or identification to gain entry to restricted areas. The system checks the credentials presented against a database of authorized individuals and grants access only if the presented credentials match an authorized entry within the database. Electronic access control systems can be configured to provide access to different levels of security. For example, employees can be granted access to areas relevant to their work, while highly sensitive areas can require additional security measures, such as biometric data or dual authentication.

Overall, the goal of Facilities Access Control is to provide a secure environment for individuals and assets within an organization. Access control measures can help reduce the risk of damage, theft, and unauthorized entry, as well as improve employee confidence and safety.

Expected Timeframe to Create Value



The expected timeframe to create value with Facilities Access Control depends on the specific implementation and requirements of the organization. However, some benefits of access control may be experienced immediately, while others may take longer to realize.

Immediate benefits may include improved security and reduced risk of theft, vandalism, or unauthorized access. This can help protect valuable assets and information, maintain employee safety, and increase overall confidence and peace of mind.

Long-term benefits, such as improved efficiency and cost savings, may take longer to realize. For example, an automated access control system can streamline the process of checking access rights and permissions, reducing administrative overhead and errors. It can also help to avoid the need to hire additional personnel to secure restricted areas. These benefits can add up over time, contributing to ongoing savings and increased efficiency.

Overall, the expected timeframe to create value with Facilities Access Control depends on the specific implementation, the size and complexity of the facility, and the security and access control goals of the organization. Nonetheless, access control is a valuable investment in safeguarding valuable assets, information, and employees, and providing a safe and secure environment.

3.35 Facilities Management

Goal

The goal of Facilities Management (FM) is to ensure that the built environment supports the efficient functioning of an organization's core activities by providing safe, functional, and comfortable facilities. Specifically, the goals of Facilities Management may include:

- **Maintenance and upkeep:** Facilities Management aims to ensure that the built environment is properly maintained, updated and renewed as and when required.
- **Cost optimization:** FM is concerned with optimizing the delivery of facilities services and achieving value-for-money while ensuring high quality service standards are maintained.
- **Asset management:** FM often involves managing the physical assets of an organization, including building structures, equipment, and machinery, ensuring that they are optimally utilized, and driving return on investment.
- **Building performance:** Facilities Management focuses on improving building performance standards such as safety, energy efficiency, environmental performance, and maintenance effectiveness.
- **Occupant satisfaction and productivity:** FM aims to provide a safe and comfortable environment for building occupants, promoting a sense of wellbeing and engagement with the indoor and outdoor working spaces.



In summary, the goal of Facilities Management is to manage and optimize the built environment, supporting organizational activities, enhancing the value of physical assets, optimizing resources and ensuring user comfort and satisfaction.

Expected Timeframe to Create Value

The expected timeframe to create value from Facilities Management depends on several factors, including the state of the current facilities infrastructure, the goals of the organization, and the availability of resources. Here are a few examples:

- **Maintenance and upkeep:** Regular maintenance and upkeep of facilities infrastructure can help to extend its useful life, reduce downtime, and avoid costly repairs. The value can be realized in the short to medium term, depending on the extent of the required maintenance, the complexity of the systems, and the availability of resources.
- **Energy efficiency improvements:** Facilities Management often includes initiatives aimed at reducing energy consumption and promoting sustainable practices. These initiatives can help to reduce energy costs, improve environmental performance, and meet compliance requirements. The value of energy efficiency improvements can be realized over the medium to long term, as they often require more substantial investment and implementation of complex solutions.
- **Building performance:** Facilities Management also involves improving building performance standards such as safety, environmental performance, and maintenance effectiveness. The value of building performance improvements can be realized over the long-term, as they often involve long-term planning, investment and implementation of solutions.

Overall, the expected timeframe for Facilities Management to create value varies depending on the specific goals and context of the organization. However, a well-executed Facilities Management program can provide immediate benefits such as reducing operational costs, improving safety and enhancing the user experience, which can result in long-term cost-saving and productivity improvements.

3.36 Facilities Occupancy Data

Goal

Occupancy data helps to address the needs of the occupants of any working space/ area/ facility space by supplying information to the facilities management team that affects on demand cleaning, hot desking etc., replenishing supplies in frequently used areas like coffee docks.



Devices that report to a cloud based monitoring centre which count the inbound and outbound people traffic in a specific building or locations in order for management to make informed decisions for the employ management, sales turnover, marketing campaign success, etc.

Expected Timeframe to Create Value

Facilities occupancy data can create value immediately from the moment of implementation. However, the value created can vary depending on the nature of the occupancy data, the facility's type, and how the occupancy data is analysed and used.

Some benefits that can be realized immediately after implementation include:

1. **Efficiency:** Facilities occupancy data can help organizations identify underutilized spaces and optimize their use. This can reduce energy waste and maintenance costs.
2. **Productivity:** Use of occupancy data to understand space utilization can provide insight into the effectiveness of collaborative spaces, providing employees with spaces that aid productivity and create an atmosphere to get-in-the-zone.
3. **Cost reduction:** Accurate occupancy data improves decision making, allowing businesses to reduce the size and expense of facilities that are underutilized.
4. **Environmental benefits:** Effective use of facilities occupancy data can reduce carbon emissions and promote environmental sustainability.

The value of facilities occupancy data continues to evolve over time. With continuous data collection and analysis, occupancy data can be used to optimize space utilization, infer demand patterns, and reduce costs. Moreover, as data from multiple sites is aggregated, broader insights can be generated about utilization patterns across various facilities.

Overall, the expected timeframe to create value from facilities occupancy data depends on different factors, including the size and complexity of facilities, the analytical tools employed, and the organization's internal culture towards data-driven decision-making.

3.37 File Comparison

Goal

The goal of file comparison is to find and highlight the differences between the content of two or more files. The files may be in different formats, such as text documents, spreadsheets, or programs. File comparison is usually done to:



1. **Verify accuracy:** Comparing files can help to validate that data has been imported or exported correctly. For example, comparing a source file to a target file after data migration can help to confirm that all data has been transferred accurately.
2. **Ensure consistency:** Comparing multiple versions of a file can help to ensure consistency across the different versions. For example, comparing two versions of a software program can help to identify any differences or bugs in the code.
3. **Identify changes:** Comparing two versions of a document can help to identify changes that have been made between them. This can be useful for tracking revisions, collaborating on documents, or for identifying plagiarism.
4. **Resolve conflict:** Comparing two different versions of a file can help to detect any conflicts between them, such as when merging code changes made by different developers in a version control system.

Overall, the goal of file comparison is to ensure that files are correct, consistent, and up-to-date, and to identify any changes or errors that may exist between multiple versions of a file.

Expected Timeframe to Create Value

The expected timeframe to create value from file comparison depends on the specific goals and context. Here are a few examples:

1. **Comparing software code:** In this case, file comparison can help to identify issues and inconsistencies in code, aiding in debugging and testing efforts. The value can be realized relatively quickly, depending on the complexity of the code and the number of files that need to be compared.
2. **Comparing data files:** Comparing data files can help to ensure data accuracy and control data quality. The Expected timeframe to create value depends on the size of the data files, the complexity of the comparison process, and the level of validation needed.
3. **Comparing document versions:** Comparing document versions can help to identify changes made by different authors and ensure consistency across versions. The Expected timeframe to create value depends on the complexity of the document and the number of versions that need to be compared.

Overall, the expected timeframe to create value from file comparison can vary widely depending on the specific use case and the complexity of the files being compared. However, file comparison can provide immediate benefits, such as identifying errors or inconsistencies, that can result in time and cost savings over the longer term.

3.38 File storage system using hybrid cryptography cloud computing

Goal

Cloud technology has been used in several fields, manufacturing and defence academies, to supply massive amounts of information. Information extracted from the cloud at the request of the customer. A number of challenges should be resolved in order to keep information in the system. To save data in the cloud, several challenges need to be addressed. A number of techniques could be used in conflict resolution. In this article, we proposed a hybrid Steganography and encryption method for data security. In Internet applications, the use of an optimal solution was not suitable for high-level information protection. We introduced a new security technique on symmetric key cryptography and Steganography. Rivest cipher 6 (RC6), Advanced Encryption Standard (AES), Byte Rotation Algorithm (BRA) and blowfish techniques to provide block safety information and the length of the technical key was 128 bits. A critical data security, Least Signification Bit (LSB)Steganography algorithm was applied.

Expected Timeframe to Create Value

1 – 3 months

3.39 Handling traffic spikes

Goal

The goal of handling traffic spikes is to ensure that your website or application can handle sudden increases in traffic without slowing down or crashing.

Expected timeframe to create value

The time it takes to create value from handling traffic spikes depends on the complexity of the infrastructure, the number of users, and the organization's specific goals. However, organizations can see immediate benefits in terms of improved performance and reliability.

3.40 Host a static website using AWS (or other clouds)

Goal

The trend of hosting static websites on Amazon S3 is a becoming very popular. This approach has been adopted by many organizations due to its advantages over traditional server-based hosting. Static websites are websites that do not require any runtime environment like JRE, .NET, etc. and are mostly based on HTML, CSS, JS, and other static resources (audio/video files, documents, etc.). AWS provides all the necessary services and tools that enable you to build and manage static websites on the AWS cloud very easily. Like other cloud-based hostings, there is no CAPEX investment. However, there is a negligible operational cost for hosting the static website.



Expected Timeframe to Create Value

1 – 3 months

3.41 Instant Messaging applications**Goal**

Instant messaging (IM) technology is a type of online chat allowing real-time text transmission over the Internet or another computer network. Messages are typically transmitted between two or more parties, when each user inputs text and triggers a transmission to the recipient(s), who are all connected on a common network. It differs from email in that conversations over instant messaging happen in real-time (hence "instant"). Most modern IM applications (sometimes called "social messengers", "messaging apps" or "chat apps") use push technology and also add other features such as emojis (or graphical smileys), file transfer, chatbots, voice over IP, or video chat capabilities.

The goal of instant messaging applications is to enable users to send and receive messages instantly in real-time. These applications allow users to communicate with each other regardless of their location, making it convenient and efficient for people to stay connected. Some of the key goals of instant messaging applications include:

1. **Communication:** The primary goal of instant messaging applications is to provide users with a platform to communicate with one another in real-time, whether through text messages, voice calls or video calls.
2. **Convenience:** Instant messaging applications are designed to provide a more convenient and accessible means of communication than traditional methods such as email or phone calls.
3. **Connectivity:** Instant messaging applications allow people to stay connected with one another despite being in different locations and time zones.
4. **Speed:** Instant messaging applications are designed to work in real-time, allowing users to send and receive messages instantly, making communication faster and more efficient.
5. **Privacy:** Instant messaging applications offer various privacy features, such as end-to-end encryption, to protect user data and conversations from unauthorized access.

Overall, the goal of instant messaging applications is to enable people to stay connected and communicate with one another quickly, conveniently, and securely, no matter where they are located.

Expected Timeframe to Create Value

Instant messaging applications can create value from the moment they are deployed and become widely adopted by users. The value that instant messaging applications offer is their ability to connect



people and enable them to communicate effectively in real-time. Thus, the more people that use these applications, the more value they provide.

In many cases, instant messaging applications can create value within a few minutes, as soon as users start using the platform to connect with each other. For example, a group of friends could download an instant messaging application, create a group chat, and start using it to stay connected. In this scenario, the value is created almost immediately.

In a business setting, instant messaging applications may take a little longer to create value, as they may require integration with other business systems, security verification, and adoption among employees. However, once the application is fully implemented, it can provide significant value by enabling employees to communicate more effectively, collaborate on projects, and respond more quickly to customers.

Overall, the expected timeframe for an instant messaging application to create value will vary depending on the situation and context in which it is deployed. However, in general, instant messaging applications can create value relatively quickly by enabling people to stay connected and communicate effectively in real-time.

3.42 Manage virtual network

Goal

The goal of managing virtual networks is to create, configure, and maintain a virtualized network infrastructure that connects virtual machines (VMs) and other resources in the cloud.

Expected timeframe to create value

The time it takes to create value from managing virtual networks depends on the complexity of the network, the number of resources connected, and the organization's specific goals. However, virtual networks can provide immediate benefits in terms of improved scalability, flexibility, and security.

3.43 Migrate to cloud

Goal

The goal of migrating to the cloud is to move your organization's applications, data, and infrastructure from on-premise servers to a cloud-based infrastructure. The objective is to improve agility, reduce operational costs, enhance scalability, and improve security.

Expected timeframe to create value



The time it takes to create value from migrating to the cloud depends on several factors, including the complexity of the existing IT infrastructure, the scope of the migration, and the amount of resources allocated to the project. However, many organizations see significant benefits in terms of improved agility, scalability, and reduced operational costs soon after migrating to the cloud.

3.44 Monitoring the activities carried out by agricultural machinery on a given surface

Goal

The agricultural work performed by the machines driven by the driver is stressful and often demands a lot from the driver. These things are due to the fact that repetitive operations are performed often and sometimes in difficult weather conditions (extreme temperatures, high humidity, etc.) The driver of these machines sometimes works in difficult conditions due to the causes shown above. There are too few possibilities to improve the working conditions of the driver.

Agricultural machines that work without being driven directly by humans seem to be a modern and viable solution. In this case, artificial intelligence and robotic technologies are used to increase the performance of these machines. The proposed application monitors the activity of one or more machines working on a ground surface.

Expected timeframe to create value

Months – 12 Months

3.45 Monitoring the physiological parameters of athletes during training

Goal

Athletes' training is always accompanied by changes in the values of some physiological parameters. The measurement of these parameters and their subsequent processing provide data on how the athlete's body responds to the demands during training.

Exceeding a certain level of demand can lead to accidents. The application proposes monitoring some physiological parameters of athletes during training.

Expected timeframe to create value

3 weeks – 4 months



3.46 Operate several projects simultaneously

Goal

The goal of operating several projects simultaneously using Google Cloud Platform is to manage multiple projects efficiently and effectively, leveraging the scalability, security, and cost-effectiveness of Google Cloud Platform.

Expected timeframe to create value

The time it takes to create value from operating several projects simultaneously using Google Cloud Platform depends on the complexity of the projects, the number of resources involved, and the organization's specific goals. However, organizations can see immediate benefits in terms of improved efficiency, resource utilization, and project outcomes.

3.47 Reconfiguration of public transport routes in a city

Goal

The means of public transport in a locality have a well-defined route that they cover at certain time intervals according to an established schedule. During peak hours, the means of public transport are more congested, and during the rest of the day, the means of transportation are loaded far below their nominal capacity (number of people transported).

The application reconfigures the circulation route of the means of public transport so that they circulate loaded close to their normal capacity and meet the demands of travellers.

Expected timeframe to create value

3 months - 2 months

3.48 Remote-controlled smart devices in smart home/office

Goal

This study discusses the impact of ambient conditions measures on customer behaviour and its application in retail industry. There are basic data sets structure presented that consists of respective data sources, namely IoT sensors, smart meters and internal transactional and analytical databases, and business indicators used to optimize air-quality sub-environments in the stores. Machine learning is proposed to automatize knowledge discovery and pattern discovery from data and as foundation of an interface to interoperable air-conditioning system.

Expected Timeframe to Create Value



6 - 9 months

3.49 Resource and application access management

Goal

Using the IAM services to keep an overview over who has access to which resources and apps.

Using Azure Active Directory's (AD) role-based access control (RBAC) you can set up the permissions for users within an organisation, defining who has access to what by authenticating them with Azure AD credentials, and then authorising them by comparing the roles the user has, and the permissions set up on a particular application or resource. This allows to set up a policy of least privilege.

To increase security implementing different ways of authenticating the user, such as Multi-Factor Authentication (MFA), where you receive a one-time password on your device, or biometrics, where you use face recognition or fingerprints to authenticate.

After the user has been authenticated, the IAM service will look up the user's roles, be it the permanent role, or the Just-In-Time (JIT) role granted through the Privileged Identity Manager (PIM) and compares it to the access policy configured on the resource or application that the user is trying to access.

Setting up the Device Identity you can also make sure that the device the user is currently using is considered safe by using conditional access. This can also be configured to only allow access from certain locations, such as within a certain IP location, within certain timeframes or use risk detection to determine if the user behaviour is considered unusual.

you can monitor the usage and grant authorization to certain users to provision or deprovision extra resources. And is also used to see which users have accessed a specific resource or app at a specific time, which can help troubleshoot where a data breach has happened.

Expected timeframe to create value

3 weeks – 1 month

3.50 Rule-based phishing website classification

Goal

These days various robots are crawling the Internet, they are also called: bots, harvesters, or spiders. Popular search engines use a similar technique to index web pages - they have an autonomous agent



(called robot or bot) that is in charge of crawling various attributes of web sites. Lately, this crawler technique is exploited by malicious users, for example harvesters, which are used for scraping e-mail addresses from websites in order to build a spam list for spambots. Recently, robots are also misused to buy flight tickets or do fast bids in on-line auction system. In this paper we present an intelligent system called Lino which tries to solve the mentioned problem. Lino is a system that simulates a vulnerable web page and traps web crawlers. We collect various features and perform a feature selection procedure to learn which features mostly contribute to the classification of visitor behaviour. For the classification purpose we use state of the art machine learning methods like Support Vector Machine and decision tree C 4.5.

Expected Timeframe to Create Value

1 – 3 months

3.51 SAP Build**Goal**

The primary goal of SAP Build is to enable business users and other stakeholders to easily and quickly create user interfaces and other applications without requiring technical skills. SAP Build is a cloud-based platform that provides a drag-and-drop interface where users can easily create and design web-based applications.

SAP Build's primary objective is to reduce the time and effort required to design and develop user interfaces, which is typically a complex and time-consuming process. It provides a user-friendly, collaborative environment where business users can easily create, visualize, and test their application ideas without requiring assistance from technical developers.

SAP Build is designed to enhance the overall user experience and user interface design of SAP applications. It provides a range of templates, design elements, and patterns that allow users to quickly create intuitive, easy-to-use interfaces that are consistent with SAP's design principles.

Overall, the goal of SAP Build is to empower business users and other stakeholders to take an active part in the design and development of user interfaces, ensuring that the applications meet their needs and requirements while adhering to the best practices in UI design and development.

Expected Timeframe to Create Value

The timeframe to create value using SAP Build depends on various factors such as the complexity of the user interface, the resources available, and the level of expertise of the team developing the



application. However, the use of SAP Build can significantly reduce the time and effort required to design and develop user interfaces, allowing for a faster time-to-market for applications.

Typically, with SAP Build, users can create interactive prototypes and conduct user testing in a matter of weeks, which helps identify any design issues early on and ensures the final application meets user needs.

The use of SAP Build can also improve user satisfaction and productivity by creating more intuitive and user-friendly interfaces, resulting in a more streamlined workflow and better user experience.

Overall, the value of SAP Build can be immediately apparent, particularly in reducing the time and effort needed to design and develop user interfaces, enhancing user satisfaction and productivity, and allowing for a faster time-to-market for applications. The expected timeframe to create value will be determined by the organization's specific goals and design needs.

3.52 Set up load balancers

Goal

During peak load times, the servers might get more traffic than they are able to properly handle, leading to dropped packages, data loss and unresponsive applications which again can lead to loss of users. Setting up an automatic load balancer will solve this problem by distributing the incoming traffic across multiple servers so that no single server becomes overloaded and turning into a bottleneck. This improves the overall performance, availability and scalability being run in the cloud infrastructure.

Expected timeframe to create value

Immediate

3.53 Smart traffic management

Goal

The project solves the problem of the growing need for security (especially in public spaces) and traffic regulation today, the direction of how the area is developing and what the needs will be in the near future. The solution to the problem will be achieved by developing a platform that, using advanced machine learning technologies, transforms monitoring and control systems into tools that open up application possibilities in the field of smart traffic and security. The challenges of other systems on the market is viewed through: (i) on the one side of the market, there are solution vendors which most often send the message that their solutions support a comprehensive approach to surveillance/security and traffic/transport in the best practice way, however those solutions include only a basic or reduced



number of functionalities, and are hard/impossible to support interoperability of that system with others that the user has, such solutions have an emphasized orientation towards a specific (one) manufacturer, demonstrate difficult transition to other solutions once the specific solution is introduced and usually are costly in terms of development and the introduction of the basic system as well as any integration or superstructure ("over-pricing"/"over-promise" problem); (ii) on the other side, small challengers on the market show potential by using advances in technologies (circuitry and software support, i.e. mathematical models), but they usually fail to scale solutions or secure a higher degree of market share due to the high cost of developing basic functionalities, i.e. the fact that the basic investment in development in order to be able to offer even the lowest level of service, it must be large ("laboratory approach" problem); (iii) even though they are presented as such, competition solutions are rarely optimized in the field of surveillance/security and transport/logistics with the absence of clear patterns or studies in which interoperability between different systems is performed and as many users as possible in the field of surveillance/security and transport/logistics considers it necessary because over time they have invested considerable resources in various technologies; (iv) privacy control is also a logical requirement, which to the greatest extent implies control over the models that lead to certain actions or are the basis for understanding behaviour in the field of surveillance/security and transport/logistics; (v) finally, solutions in the field of surveillance/security and transport/logistics are often under special legal regulations and are subject to changes in them, which increases the need for their adaptation through model correction and the construction of such systems on open technologies with a high degree of control over the models that lead until insight.

Expected Timeframe to Create Value

12 - 24 months

3.54 Supply real-time sales data

Goal

To be able to use the information gathered to make on-the-fly changes to promotions, trial new promotions and alter them based on continuous feedback or distributing staff over several locations to plan for peaks in work needing to be done.

Expected timeframe to create value

6 months – 12 months

3.55 The graphic interface for programming at a car service combined with a website

Goal

Companies that offer services to the population such as car services, private medical offices, etc. he plans his activity daily, taking into account the time needed to perform an activity.

For example, if a person's car has a problem, the owner of the car will have to go to a workshop to diagnose the car, propose methods to remedy the situation and fix the defect.

This application offers the customer the possibility of online scheduling at a car service to diagnose the fault of a car.

Expected timeframe to create value

2 weeks – one month

3.56 Video conference system

Goal

The goal of a video conference system is to enable remote communication and collaboration between people or teams, regardless of their physical location. Specifically, the goals of a video conference system may include:

1. Real-time communication: A video conference system aims to provide a platform for real-time, face-to-face interaction between participants, enabling remote teams or individuals to communicate naturally and effectively.
2. Collaboration: A video conference system can facilitate collaboration by allowing participants to share files, documents, and screens, co-author documents, and even brainstorming over virtual whiteboards.
3. Convenience: A video conference system aims to provide convenience and flexibility by eliminating the need for participants to be physically present at the same location, allowing them to participate in meetings from anywhere in the world.
4. Time savings: A video conference system can help to save time by avoiding the need for travel and reducing downtime between meetings, allowing participants to stay productive and engaged.
5. Cost savings: A video conference system can help to save costs associated with travel and accommodation, especially for organizations with multiple offices in different locations or for remote teams that would otherwise require office space to work.



Overall, a video conference system aims to provide a seamless and effective way for remote teams or individuals to communicate and collaborate, enhancing productivity, convenience, and cost savings.

Expected Timeframe to Create Value

The expected timeframe to create value from a video conference system depends on various factors, such as the size of the organization, its operational structure, the frequency of meetings, and the technology ecosystem. Here are some general examples:

1. **Improved collaboration:** Video conference systems can improve collaboration by providing real-time video and audio capabilities, making it easier for teams to collaborate remotely. The value of this feature can be realized in the short term, even during the first few video conferences.
2. **Reduced travel expenses:** Video conference systems can save travel costs by replacing in-person meetings with virtual ones, leading to reduced travel expenses, such as flights, lodging, and transportation. The value from reduced travel expenses can be realized immediately, during the first few video conferences or meetings in which travel is avoided.
3. **Faster decision-making:** Video conference systems can help facilitate faster decision-making by providing instant video and audio access, supporting real-time decision-making. The value of faster decision-making can be realized immediately and throughout the long-term use of the system.

Overall, the expected timeframe to create value from a video conference system can be immediate, especially in terms of cost savings from reduced travel expenses and improved collaboration. Additional value may materialize over the medium to long term as the organization develops a stable ecosystem with well-built processes and technology to support meetings and collaborations.

3.57 VoD offering

Goal

With Video on Demand (VoD), you can create a library of videos that your users can access at any time. You can also control access to your videos by specifying who can view them and when. AMS also provides tools to help you manage your video content, including indexing, search, and analytics.

To use Azure Media Service (AMS) VoD, you first need to upload your videos to the platform. You can do this through the AMS portal, REST APIs, or through a variety of third-party tools and services. Once your videos are uploaded, you can use AMS to transcode them into various formats, create multiple bitrates, and encrypt them for secure delivery.



After your videos are processed, you can use the AMS player to embed them on your website or app. The player supports a variety of features, including adaptive streaming, closed captioning, and multiple audio tracks. You can also customize the player's look and feel to match your brand.

Expected timeframe to create value

2 months – 4 months

3.58 Water supply management using distance readers in water supply networks

Goal

Digital transformation enables significant savings through resource management and business process improvement. It changes the way we use the information we have, the type and amount of data we can collect. To make this data more usable, we use modern analytical and visualization tools whose task is to get useful and timely information from a large amount of different data in a simple and flexible way.

The issues that arise in this area of interest range from how to visualize data, which methods to use to find knowledge hidden in data, and how to develop forecasting models using data.

Researchers and industry are giving special focus to weather data that may have a significant impact on prediction in times of unpredictable climatic changes and weather influences. In technical matter for companies that want to make the first step in this area, they are encountering questions from how to store data into a "cloud"/"big data" container, is it possible to develop data project which "grows together with a company" and more and more acquired data, whether it can all work in real-time and is this "package" available to them in terms of cost and knowledge needed.

Expected Timeframe to Create Value

6 – 9 months

3.59 Web application for the online completion of a company's staff timesheet

Goal

Construction companies carry out work in different work points located in a geographical space. Each work is attended by detached teams for the duration of the work.



At the company headquarters, records of the hours worked by each member of the work team must be kept. The application allows the attendance to be completed up to date for each member of the teams who carry out their activity in different work points.

Expected timeframe to create value

3 weeks - 1 month

3.60 Website hosting with static content

Goal

Having a website is crucial for any business today to stay competitive. It gives the business the opportunity to maintain an online presence for their prospective customers and users, giving the business 24/7 availability, visibility, and accessibility for your current and potential new users. This gives them the opportunity to discover your business without being restricted by things such as opening hours, waiting times on the telephone, and having to visit a physical location.

Even the simplest of websites will allow the business to provide information to visitors about things such as the opening hours of a particular location, contact information or information about the products or services the business offers. It can also be used to show videos and photos promoting the business and its products/services.

This means having a website could potentially give a business a global reach and presence, while reducing time and costs spent on customer service/support by having a lot of commonly asked questions answered on the website. And offers a convenient platform for engaging with customers/users by displaying products and services via promotional material such as videos hosted on the website or sending newsletters with exclusive offers or discounts sent directly to interested customers/users to a global audience.

Expected timeframe to create value

1 week – 6 months

3.61 Webstore

Goal

Selling products, whether it be online, in a brick-and-mortar location or both, having access to accurate and concurrent information about the current state of your product inventory is important to be able

to provide the best possible experience for a customer and minimising the risk of running out of stock which can lead to backlogs of orders and unsatisfied customers.

Keeping records of your customers and the orders they have made is also important as it is both important for making sure you can provide the proper level of customer support to your customer, and it can be used to gain important insight into your customer's behaviour such as, the kind of products they are interested in which can be used to create tailored content for your customers.

Expected timeframe to create value

2 weeks – 2 months



REFERENCE

1. Cloud Industry Forum. (2022). 8 criteria to ensure you select the right cloud service provider. Retrieved from <https://cloudindustryforum.org/8-criteria-to-ensure-you-select-the-right-cloud-service-provider/>
2. CloudSigma. (2023). 10 Steps to Choose the Best Cloud Provider. Retrieved from <https://www.cloudsigma.com/10-steps-to-choose-the-best-cloud-provider/>
3. Colt. (2023). Cloud connect explained. Retrieved from <https://www.colt.net/resources/cloud-connect-explained/>
4. CompTIA. (2022). A cloud networking quick-start guide: around the network in 8 steps. Retrieved from <https://www.comptia.org/content/guides/cloud-network-setup-guide>
5. CompTIA. (2023). Partly Cloudy with a Chance of Computing: A Beginner's Guide to Cloud Types, Solutions and Vendors. Retrieved from <https://www.comptia.org/content/articles/cloud-types-solutions-and-vendors>
6. CompTIA. (n. a.). Partly Cloudy with a Chance of Computing: A Beginner's Guide to Cloud Types, Solutions and Vendors. Retrieved from <https://www.comptia.org/content/articles/cloud-types-solutions-and-vendors>
7. Delta. (2020). Powering Competitiveness in Datacentres. Retrieved from <https://www.deltapowersolutions.com/en/mcis/technical-article-powering-competitiveness-in-datacenters.php>
8. Dialogic. (2017). Introduction to Cloud Computing, White paper. Retrieved from <https://www.dialogic.com/~media/products/docs/whitepapers/12023-cloud-computing-wp.pdf>
9. Dialogic. (2017). Introduction to Cloud Computing, White paper. Retrieved from <https://www.dialogic.com/~media/products/docs/whitepapers/12023-cloud-computing-wp.pdf>
10. Eldh, E. (2013). Cloud connectivity for embedded systems (Master of Science Thesis). KTH Royal Institute of Technology, Stockholm, Sweden.
11. Faddom. (2021). Cloud Computing Costs & Pricing Comparisons for 2023. Retrieved from <https://faddom.com/cloud-computing-costs-and-pricing-comparison/>
12. FERI. (2022). Cloud Calculation. Retrieved from: <https://moja.um.si/studijski-programi/Strani/ucnaenota.aspx?jezik=S&fakulteta=FERI&sifraue=61M252>
13. FRI. (2022). Second Level Master's Study Program Computing and Informatics Presentation Proceedings for Students First Enrolled in The 1st Year In The Academic Year 2022/2023 Ljubljana. Retrieved from: chrome-extension://efaidnbmnnnibpcajpglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Fwww.fri.uni-lj.si%2Fupload%2FZborniki%2F1000471_Ra%25C4%258Dunalni%25C5%25A1tvo_in_informa%2520-%2520Copy%252011.pdf&clen=765163&chunk=true.
14. Google Cloud. (2022a). Cloud Interconnect documentation. Retrieved from <https://cloud.google.com/network-connectivity/docs/interconnect>
15. Google Cloud. (2022b). Google Cloud terms. Retrieved from <https://cloud.google.com/network-connectivity/docs/concepts/key-terms>
16. ITPro Today. (2022a). 2022 Cloud Computing Trends. Retrieved from <https://www.youtube.com/watch?v=PiaouNqFNwA>
17. ITRPro Today. (2022b). Providers Continue to Dominate, Led by AWS. Retrieved from <https://www.itprotoday.com/iaas-and-paas/big-3-public-cloud-providers-continue-dominate-led-aws#close-modal>

18. ITU. (2012). Technical Report: Part 1: Introduction to the cloud ecosystem: definitions, taxonomies, use cases and high-level requirements. Retrieved from <https://www.itu.int/pub/T-FG-CLOUD-2012-P1>
19. ITU. (2022). Focus Group Cloud, Technical Report, Part 1: Introduction to the cloud ecosystem: definitions, taxonomies, use cases and high-level requirements, Version 1.0. Retrieved from <https://www.itu.int/pub/T-FG-CLOUD-2012-P1>
20. Jones, E. (2022). Cloud Market Share: A Look at the Cloud Ecosystem in 2023. Retrieved from <https://kinsta.com/blog/cloud-market-share/>
21. Letica, J. & Buić, N. (2014). Innovation in VET. Retrieved from http://www.refernet.hr/media/1236/innovation-in-vet_croatia.pdf
22. Marinescu, D. (2017). Cloud computing Theory and Practice. USA: Elsevier, Morgan Kaufmann publishing.
23. Markets And Markets. (2019). Retrieved from <https://www.marketsandmarkets.com/>
24. Marko, K. (2021). Cloud providers jockey for 2021 market share. Retrieved from <https://www.techtarget.com/searchcloudcomputing/opinion/Cloud-providers-jockey-for-market-share>
25. Opinion of the European Economic and Social Committee on 'Industry 4.0 and digital transformation: where to go'. (2016). Retrieved from: <chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Feur-lex.europa.eu%2Flegal-content%2FEN%2FTXT%2FPDF%2F%3Furi%3DCELEX%3A52016AE1017%26from%3DEN&pdfilename=CELEX%3A52016AE1017%3AEN%3ATXT.pdf>.
26. Oracle (2023). What is cloud computing? Retrieved from <https://www.oracle.com/cloud/what-is-cloud-computing/top-10-benefits-cloud-computing/>
27. Peterson, R. (2023). Cloud Computing Tutorial for Beginners: What is & Architecture. Retrieved from <https://www.guru99.com/cloud-computing-for-beginners.html>
28. Rathore, A. (2022). How To Find The Best Cloud Server For Small Businesses? Retrieved from <https://kanakinfosystems.com/blog/best-cloud-server-for-small-business>
29. Resonate. (2020). What Are the Different Types of Load Balancers? Retrieved from <https://www.resonatenetworks.com/2020/05/25/what-are-the-different-types-of-load-balancers/>
30. Richter, F. (2023). Big Three Dominate the Global Cloud Market. Retrieved from <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/>
31. Rosencrance, L. (2021). Breaking Down the Cost of Cloud Computing in 2023. Retrieved from <https://www.techtarget.com/whatis/Breaking-Down-the-Cost-of-Cloud-Computing>
32. Samoshki, D. (n. d.). The cloud report. Retrieved from <https://the-report.cloud/how-to-choose-a-cloud-for-your-business/>
33. Sharma, M. (2023). Load balancing in Cloud Computing. Retrieved from <https://www.geeksforgeeks.org/load-balancing-in-cloud-computing/>
34. Sharwood, S. (2022). Cloud a three-player market dominated by AWS, Google, Microsoft. Retrieved from https://www.theregister.com/2022/05/02/cloud_market_share_q1_2022/
35. Slovenska strategija pametne specializacije S4. (2017). Retrieved from: <chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Fwww.gov.si%2Fassets%2Fvladne-sluzbe%2FSVRK%2FS4-Slovenska-strategija-pametne-specializacije%2FSlovenska-strategija-pametne-specializacije.pdf&clen=1536948>.
36. Spaanenburg, L. & Spaanenburg, H. (2010). Cloud Connectivity and Embedded Sensory Systems.



- New York: Springer.
37. Spaanenburg, L., Spaanenburg, H. (2010). *Cloud Connectivity and Embedded Sensory Systems*. Switzerland: Springer.
 38. Spilka, S. (2021). *Cloud Pricing Models - Shedding light upon pricing options*. Retrieved from <https://www.exoscale.com/syslog/cloud-pricing-models/>
 39. Strategija dolgožive družbe. (2017). Retrieved from: chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Fwww.umar.gov.si%2Ffileadmin%2Fuser_upload%2Fpublikacije%2Fkratke_analize%2Fstrategija_dolgozive_druzbe%2Fstrategija_dolgozive_druzbe.pdf&clen=2707481&chunk=true.
 40. Strategija razvoja Slovenije 2030. (2017). Retrieved from: chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Fwww.gov.si%2Fassets%2Fvladne-sluzbe%2FSVRK%2Fstrategija-razvoja-Slovenije-2030%2Fstrategija_razvoja_Slovenije_2030.pdf&clen=4124906.
 41. Strategija višjega strokovnega izobraževanja v Republiki Sloveniji za obdobje 2020-2030. (2017). Retrieved from: chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/viewer.html?pdfurl=https%3A%2F%2Fwww.gov.si%2Fassets%2Fministrstva%2FMIZS%2FDokumenti%2Fvisje-strokovno-izobrazevanje%2Fstrategija-visjega-strokovnega-izobrazevanje-RS-2020-2030%2Fstrategija-visjega-strokovnega-izobrazevanja-v-Republiki-Sloveniji-za-obdobje-20202030.pdf&clen=1259647.
 42. Suhag, A. (2020). *What Are The Different Types Of Cloud Load Balancing?* Retrieved from <https://www.cloudmanagementinsider.com/different-types-of-cloud-load-balancing/>
 43. Techfunnel. (2021). *14 Incredible Benefits of Cloud Computing for Businesses*. Retrieved from <https://www.techfunnel.com/information-technology/benefits-of-cloud-computing/>
 44. *The Complete Cloud Computing Manual*. (2022). Retrieved from <https://online.fliphtml5.com/dslwu/jeti/>
 45. Tripney S. & Hombrados J. (2013). *Technical and vocational education and training (TVET) for young people in low- and middle-income countries: a systematic review and meta-analysis*. *Journal of Empirical Research in Vocational Education and Training*, 5(3), 1-14. doi: 10.1186/1877-6345-5-3.
 46. Velte, A. T., Velte, J. V., & Elsenpeter, R. (2010). *Cloud Computing: A Practical Approach*. New York: The McGraw-Hill.
 47. Westlake. (2022). *Benefits of cloud computing for businesses*. Retrieved from <https://www.westlake-it.co.uk/news/2022/05/30/benefits-of-cloud-computing-for-businesses/>
 48. Wikipedia. (2022). *Border Gateway Protocol*. Retrieved from https://en.wikipedia.org/wiki/Border_Gateway_Protocol
 49. Wired. (2020). *Data Centers Aren't Devouring the Planet's Electricity—Yet*. Retrieved from <https://www.wired.com/story/data-centers-not-devouring-planet-electricity-yet/>



APPENDIX

Appendix 1: Case of the 2012 US Presidential Campaign and how AWS supported Obama

In this unit, we will look at how Amazon's cloud computing technology allowed President Obama's 2012 Presidential campaign to avoid an IT investment that would have run into the tens of millions of dollars.

A look at our case study:

The campaign's IT team used AWS to build, launch, run, and grow their apps. Following the election, they backed everything up to Amazon S3 and scaled way, far down. They created and operated over 200 AWS apps that could handle millions of people. On the final four days of the campaign, one of these applications, the campaign call tool, handled 7,000 concurrent users and placed over two million calls.

Why use AWS?

Here are 3 key aspects that influenced why AWS would be used as the cloud computing provider in the Obama campaign:

1. Security and Compliance

Elections draw some of the world's most aggressive information security threats. When it comes to election technology, information security is a major priority. AWS understands election administrators' responsibilities and meets or exceeds security and compliance standards at every level of our customers' cloud journey. AWS prioritizes data security, and our worldwide infrastructure is developed and managed in accordance with security best practices.

2. Voter Engagement

In 2018, all millennials (those aged 18 to 29) were eligible to vote in the United States for the first time. Millennials prefer online transactions and have high expectations for tailored customer experiences. AWS offered building blocks that can be quickly assembled to support virtually any secure workload for targeted outreach.

3. Elections Management:

Elections Management refers to back-office duties like voter registration that serve as operational efficiency drivers across multiple linked systems, applications, and local organizations spanning counties and districts. AWS offers a number of database services to help with voter registration. These fully managed systems may be launched in minutes with a few clicks. Furthermore, the AWS Database Migration Service facilitates a simple and cost-effective transition to the AWS Cloud.

How it was done:

- The primary registry of voter file information was a database hosted on Amazon RDS. This database combined data from various sources (including www.barackobama.com and donor information from the finance team) to provide campaign managers with a dynamic, fully integrated picture of what was going on.
- This collection of databases enabled campaign workers to target and segment prospective voters, shift marketing resources based on near real-time feedback on the effectiveness of specific ads, and power a donation system that raised more than \$1 billion (the 30th largest ecommerce site in the world).

The Obama campaign's apps are equivalent in extent and complexity to those seen in the largest companies and data-rich startups.

To give a point-by-point example of how the election campaign made use of apps that were available on the AWS cloud platform and performed tasks both complex and massive in scale:

- Vertica and Elastic MapReduce are used to model massive amounts of data.
- Multi-channel media management via TV, print, online, mobile, radio, and email with dynamic production, targeting, retargeting, and multi-variant testing, similar to what you'd find in a competent digital media agency.
- Coordination and collaboration of volunteers, contributors, and supporters on a social level.
- Large-scale transaction processing.
- Voter abuse prevention and protection, including incident collection and volunteer deployment.
- A comprehensive information distribution system for campaign news, polling, topic information, voter registration, and more.

Since the 2016 U.S. presidential election, Amazon Web Services has quietly increased its presence in state and local elections; more than 40 states now use one or more of Amazon's election offerings as do America's two major political parties, Democratic presidential candidate Joe Biden, and the federal agency in charge of enforcing federal campaign finance laws.

While it does not handle voting on election day, according to company documents and interviews, AWS now runs state and county election websites, stores voter registration rolls and ballot data, facilitates overseas voting by military personnel, and helps provide live election-night results.

Nonetheless, Amazon's growing presence in the elections industry may jeopardize what many officials regard as a strength of the US voting system: decentralization.

While most security experts agree that while Amazon's cloud is likely to be much more difficult to hack than the systems it is replacing, putting data from multiple jurisdictions on a single system raises the possibility that a single major breach could be disastrous. "It makes Amazon a more attractive target

for hackers" and "increases the difficulty of dealing with an insider attack," said Chris Vickery, director of cyber risk research at cybersecurity startup Upguard.

The privatization of voting infrastructure is part of a larger trend that has swept across nearly every aspect of government in America, from parking tickets to prisons, and has continued under the Trump administration.

According to companies that partner with both firms for government contracts, Azure, AWS's main competitor, has a sizable government business and offers some election services, but it has not focused on them and lags behind Amazon.

Questions to consider:

1. What are the advantages of putting elections on a cloud platform?
2. How is decentralization considered a threat?
3. Read and comment on how AWS used sentiment analysis to reflect on the inauguration speeches of Obama vs Trump and the conclusions made:

<https://medium.com/@szekelygergoo/use-aws-to-compare-inauguration-speeches-of-obama-and-trump-670068ea39d5>

Appendix 2: Code Snippets

Usecase: Chatbot for students in EDU institution

Importance of Natural Language Understanding (NLU) cannot be highlighted enough but that is the main reason why this thesis is even considering to be enrolled. From technology perspective Microsoft has really great service to offer. Language Understanding Service (LUIS) is one of the best NLU solution on the market. However, every Microsoft service that is somehow in relation with NLU is coupled to LUIS in the background. With LUIS it is easy to add language understanding to any app. It is designed to identify valuable information in conversations, LUIS interprets user goals (intents) and distils valuable information from sentences (entities), for a high quality, nuanced language model. LUIS integrates seamlessly with the Azure Bot Service, making it easy to create a sophisticated bot.

```

{
  "query": "Book me a flight to Cairo",
  "topScoringIntent": {
    "intent": "BookFlight",
    "score": 0.9887482
  },
  "intents": [
    {
      "intent": "BookFlight",
      "score": 0.9887482
    },
    {
      "intent": "None",
      "score": 0.04272597
    }
  ],
  "entities": [
    {
      "entity": "cairo",
      "type": "Location",
      "startIndex": 20,
      "endIndex": 24,
      "score": 0.956781447
    }
  ]
}

```

Figure 0.1. LUIS in Action

For example, for query like “Book me a flight to Cario”, LUIS is capable to turn results in a JSON form. where could be found valuable information like BookFlight as an Intent with 98% of accuracy and entities like Cairo as a Location entity with 95 % of accuracy. Even if Bots and NLU are pretty mature technologies there are still possibilities to some Students questions stays unanswered or misunderstood. These situations should be well treated, and students should have another possible option to fulfil their request. One of common approaches for that situation are quick replies. Quick replies are small buttons or menus which have already prepared and predicted possible questions which can be written but also chosen by pressing the right predicted question.

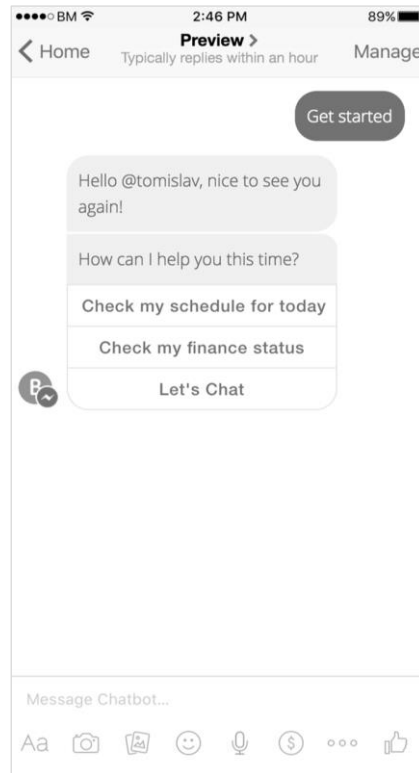


Figure 0.2. Quick Replies

Other possible solution is to offer to chat or call Student Office Desk Staff directly but that should be only in rear cases. The main idea of Student Service Support Chatbot is to decrease amount of student calls to minimum.

Use case: Digital asset certification using distributed ledger/blockchain

Application Modules

This type of application is intended for private blockchain. This means that each educational institution should have its own stream that only the people in the institution have the authority to store the diploma. All streams are stored in the main book that is distributed to all nodes, ie educational institutions in this example. The more nodes in the chain, the better, because the chain becomes ever stronger and safer.

The application consists of three modules:

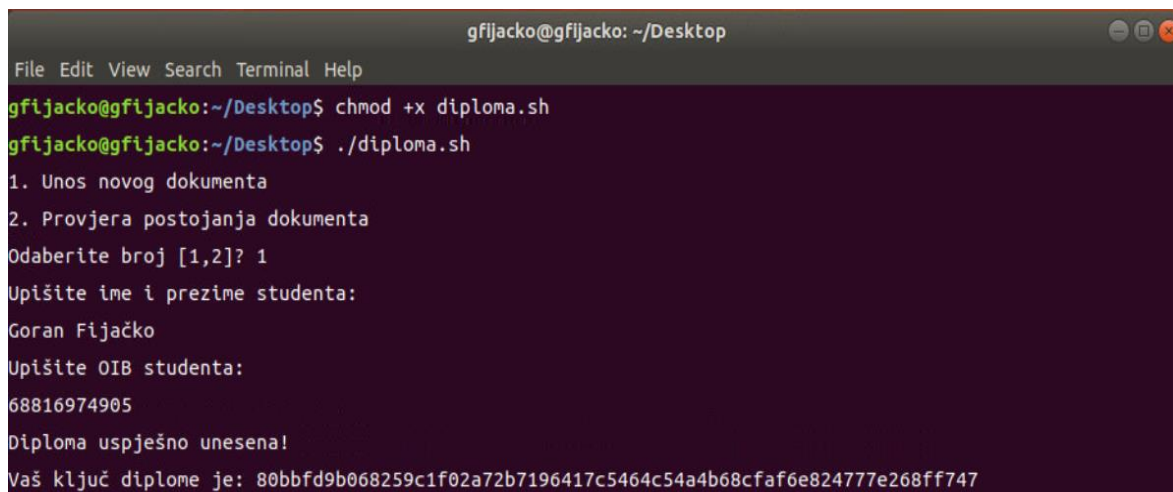
1. Module for entering a diploma
2. Diploma check module
3. Diploma printing module

The first module is for entering a diploma. It switches the entered data into a hexadecimal form and stores them in the chain and returns the transaction ID (txid) back. Transaction ID is a private key that

is awarded to a graduate student because it can be used to check the diploma data in the chain. The Diploma Check Module, combined with the OIB and Transaction ID, sends a query to the chain and verifies whether there is a record in the chain. Thereafter, it gives a positive or negative answer, depending on whether there is a really required degree in the chain and whether it complies with the OIB entered. The diploma printing module prints a diploma on the screen in PDF format. All of the modules listed in this example are displayed in the command-line text interface, ie in the Ubuntu Operating System Terminal. They can also be programmed into a web application and used in WEB browsers.

User Roles

After the student successfully completes the faculty and defends his graduate thesis, the faculty system reports that the student has graduated. With this application and the module for entering the diploma, an authorized person at the university will enter the name, last name and OIB graduate student and this information will be stored in the chain. As feedback, he receives a Transaction ID which gives the student and enrolls on the original print diploma. It can also be printed in the form of a bar code whose scan is the value of the Transaction ID.



```
gfijacko@gfijacko: ~/Desktop
File Edit View Search Terminal Help
gfijacko@gfijacko:~/Desktop$ chmod +x diploma.sh
gfijacko@gfijacko:~/Desktop$ ./diploma.sh
1. Unos novog dokumenta
2. Provjera postojanja dokumenta
Odaberite broj [1,2]? 1
Upišite ime i prezime studenta:
Goran Fijačko
Upišite OIB studenta:
68816974905
Diploma uspješno unesena!
Vaš ključ diplome je: 80bbfd9b068259c1f02a72b7196417c5464c54a4b68cfaf6e824777e268ff747
```

Figure 0.3. Showing the module for entering the diploma

The student gets his deserved diploma and his private key diploma, which in this case is 80bbfd9b068259c1f02a72b7196417c5464c54a4b68cfaf6e824777e268ff747. He then reports for a job and after a call from the employer goes to the job interview. The employer asks for a degree to check his qualifications. The procedure is currently being conducted so that the employer contacts the educational institution to verify the validity of the diploma, most often in writing. This process is long-lasting and consumes a lot of resources. But in this case, the employer gets a diploma with a private key. The employer then appoints the OIB person applying for a job and the public key in the application. This way in a fraction of a second returns the information on the validity of the diploma.

```
Molim Vas odaberite opciju:
1. Unos novog dokumenta
2. Provjera postojanja dokumenta
Odaberite broj [1,2]? 2
OIB:
68816974905
Ključ:
80bbfd9b068259c1f02a72b7196417c5464c54a4b68cfaf6e824777e268ff747
Goran Fijačko diplomirao na Visokom učilištu Algebra, smjer Multimedija, 15.10.2018. u Zagrebu.
Prikazati diplomu?
1. Da
2. Ne
Odaberite broj [1,2]? 1
Diploma će se prikazati u PDF-u!
```

Figure 0.4. Show the diploma verification module

After the application's confirmation is answered, the screen prints. The name and surname of the student, educational institution, orientation, date and place of graduation are written in the print. The employer eventually has the option of printing a copy of the diploma for his own archive. If you choose a print option, the diploma will be generated and opened in PDF format.

For ease of use of the application after release to production, it is a better choice to use it as a WEB application. This means that everything shown will be moved to a web server and the application will access the https protocol (e.g., via URL <https://www.diplome.hr>) in web browsers. This means that users only need an internet connection and an account in the application to quickly and securely check the validity of the diploma.

Usecase: Remote-controlled smart devices in smart home

In order to interpret the effect of the ambient conditions of the stores to customer behavior we can use IoT sensors to measure brightness, temperature and humidity and determine/control their influence on the customer basket. This involves determining thresholds for unfavorable brightness, unpleasant temperature and inadequate humidity levels. The technological solution should be deployed in the form of a decision support system that can analyze the mutual relationships between IoT collected data, specific product groups and overall transactions in the store. Part of the decision support system should be able to control technical conditions in automated manner through interoperable interface embedded into existing air-conditioning systems. Since ambient conditions are usually not equal in the entire store because some products can require different conditions (e.g. frozen food has different acceptable ambient temperature range than other food), we should include into analytical data sets the store zone that identify particular area requiring specific environmental conditions.

The proposed data points are divided into two granularity levels: shop visit and product bought. The data points are collected from existing transactional databases and the data store containing IoT sensors real-time data.

Transactional data source tables are as follows in below figures.

Environment

Field Name	Data Type	Description (Optional)
Time	AutoNumber	Time stamp (granulation level is arbitrary)
StoreAreaID	Number	The area of the store
Temperature	Number	Mean temperature of the store area
Brightness	Number	Mean brightness of the store area
Humidity	Number	Mean humidity of the store area

Transactions

Field Name	Data Type	Description (Optional)
VisitID	AutoNumber	
ProductID	Short Text	Bought product during the visit
Quantity	Number	Quantity of the bought product

StoreAreas

Field Name	Data Type	Description (Optional)
StoreAreaID	AutoNumber	
StoreAreaName	Short Text	The name of the area of the store

Products

Field Name	Data Type	Description (Optional)
ProductID	AutoNumber	
ProductName	Short Text	The name of the product
ProductCategory	Short Text	The category of the product
ProductSubcategory	Short Text	The subcategory of the product
ProductWeight	Number	The weight of the one unit of the product
StoreAreaID	Short Text	The area in which the product is located in the store
Price	Number	Price of one unit of the product

Visits

Field Name	Data Type	Description (Optional)
VisitID	AutoNumber	
EnteringTime	Date/Time	Time when customer entered the store
ExitingTime	Date/Time	Time when customer arrived to the cash register

Figure 0.5. Transnational data source

On a figure below, the ETL relations are shown:

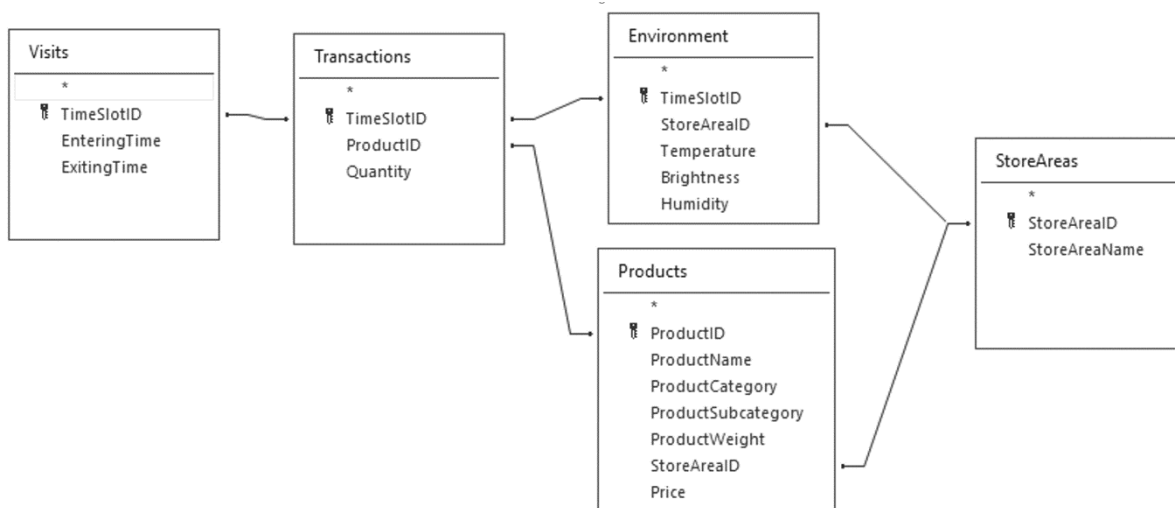


Figure 0.6. ETL relations

The variables available to analyse shop visits after applying ETL procedure are:

Field Name	Data Type	Description (Optional)
TimeSlotID	AutoNumber	
EnteringTime	Date/Time	
ExitingTime	Date/Time	
ProductID	Number	
Quantity	Number	
Environment_StoreAreaID	Number	
Temperature	Number	
Brightness	Number	
Humidity	Number	
ProductName	Short Text	
ProductCategory	Short Text	
ProductSubcategory	Short Text	
ProductWeight	Number	
Products_StoreAreaID	Number	
Price	Number	
StoreAreaName	Short Text	

Figure 0.7. The variables after applying ETL procedure

As target variables for machine learning we can now derive:

- Number of items (N)– number of different products purchased by a customer in one store visit (i.e., number of items in shopping basket);
- Weight of purchases (W) – weight of all products purchased by a customer in one store visit
- Quantity of items (Q) – quantity of items of all products (summed across all types of products) purchased by a customer in one store visit.

Other group of possible target variables – retail business indicators - are described separately in the next section.

Use case: Automation of tasks using cloud based services

To demonstrate how to carry out an MBA programming language R was used, and, in particular, the arules package, along with some code included as a proof-of-concept. The example used is available

at arulesViz Vignette and use a data set of grocery sales that contains 9,835 individual transactions with 169 items. First step was to look at the items in the transactions and, in particular, plot the relative frequency of the 25 most frequent items. This is equivalent to the support of these items where each itemset contains only the single item. The bar plot illustrates the groceries that are frequently bought at this store, and it is notable that the support of even the most frequent items is relatively low (for example, the most frequent item occurs in only around 2.5% of transactions). These insights were used to inform the minimum threshold when running the Apriori algorithm; for example, we know that in order for the algorithm to return a reasonable number of rules we'll need to set the support threshold at well below 0.025.




Figure 0.8. Bar plot of the support of the 25 most frequent items bought

By setting a support threshold of 0.001 and confidence of 0.5, we can run the Apriori algorithm and obtain a set of 5,668 results. These threshold values are chosen so that the number of rules returned is high, but this number would reduce if we increased either threshold or support. Experimenting is recommended with these thresholds to obtain the most appropriate values. Whilst there are too many rules to be able to look at them all individually, we can look at the five rules with the largest lift in below table.

Table 0.1. The five rules with the largest lift

Rule	Support	Confidence	Lift
{instant food products,soda}=>{hamburger meat}	0.001	0.632	19.00
{soda, popcorn}=>{salty snacks}	0.001	0.632	16.70
{flour, baking powder}=>{sugar}	0.001	0.556	16.41
{ham, processed cheese}=>{white bread}	0.002	0.633	15.05
{whole milk, instant food products}=>{hamburger meat}	0.002	0.500	15.04

These rules seem to make intuitive sense. For example, the first rule might represent the sort of items purchased for a BBQ, the second for a movie night and the third for baking. Rather than using the thresholds to reduce the rules down to a smaller set, it is usual for a larger set of rules to be returned so that there is a greater chance of generating relevant rules. Alternatively, we can use visualisation techniques to inspect the set of rules returned and identify those that are likely to be useful. Using the *arulesViz* package, rules by confidence, support and lift are plotted. This plot illustrates the relationship between the different metrics. The optimal rules are those that lie on what's known as "support-confidence boundary". Essentially, they lie on the right hand border of the plot where either support, confidence or both are maximised. The plot function in the *arulesViz* package has a useful interactive function that allows you to select individual rules (by clicking on the associated data point), which means the rules on the border can be easily identified.




Figure 0.9. A scatter plot of the confidence, support and lift metrics

There are lots of other plots available to visualize the rules, but one other figure that we would recommend exploring is the graph-based visualization of the top ten rules in terms of lift (more than ten rules can be included, but these types of graphs can easily get cluttered). In this graph the items grouped around a circle represent an itemset and the arrows indicate the relationship in rules. For example, the purchase of sugar is associated with purchases of flour and baking powder. The size of the circle represents the level of confidence associated with the rule and the color the level of lift (the larger the circle and the darker the grey the better).

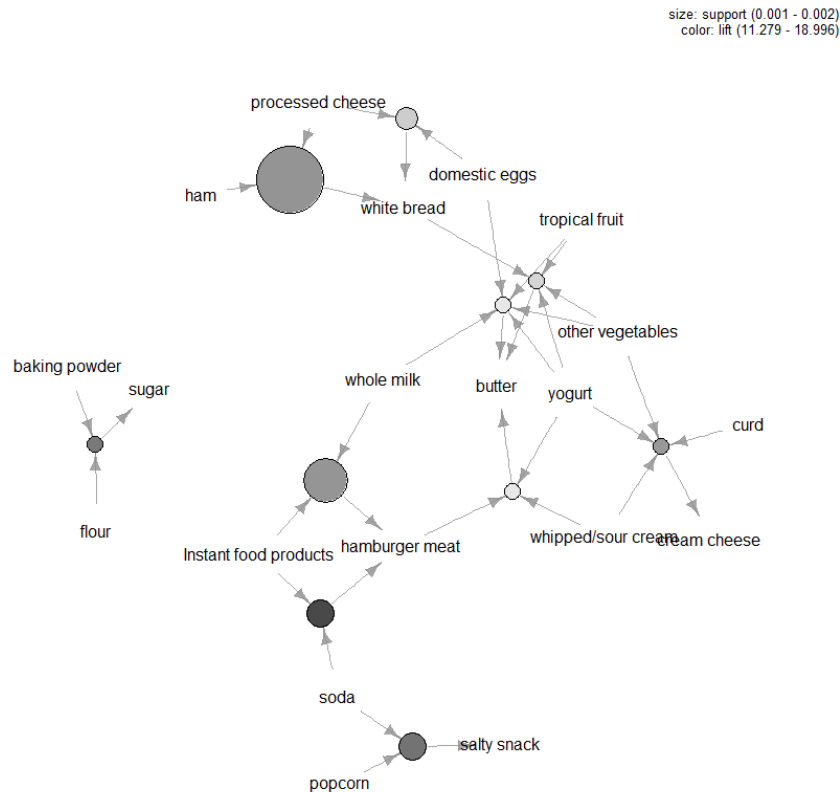


Figure 0.10. Graph-based visualisation of the top ten rules in terms of lift

Market Basket Analysis is a useful tool for retailers who want to better understand the relationships between the products that people buy. There are many tools that can be applied when carrying out MBA and the trickiest aspects to the analysis are setting the confidence and support thresholds in the Apriori algorithm and identifying which rules are worth pursuing. Typically the latter is done by measuring the rules in terms of metrics that summarize how interesting they are, using visualization techniques and also more formal multivariate statistics. Ultimately the key to MBA is to extract value from your transaction data by building up an understanding of the needs of your consumers. This type of information is invaluable if you are interested in marketing activities such as cross-selling or targeted campaigns.

R Code

```
library("arules")
library("arulesViz")

#Load data set:
data("Groceries")
summary(Groceries)

#Look at data:
inspect(Groceries[1])
LIST(Groceries)[1]

#Calculate rules using apriori algorithm and specifying support and confidence thresholds:
rules = apriori(Groceries, parameter=list(support=0.001, confidence=0.5))

#Inspect the top 5 rules in terms of lift:
inspect(head(sort(rules, by="lift"),5))

#Plot a frequency plot:
itemFrequencyPlot(Groceries, topN = 25)

#Scatter plot of rules:
library("RColorBrewer")
plot(rules,control=list(col=brewer.pal(11,"Spectral")),main="")

#Rules with high lift typically have low support.

#The most interesting rules reside on the support/confidence border which can be clearly seen in this plot.

#Plot graph-based visualisation:
subrules2 <- head(sort(rules, by="lift"), 10)
plot(subrules2, method="graph",control=list(type="items",main=""))
```

Usecase: Water supply management using distance readers in water supply networks

The LoRa protocol is a modulation of wireless data transmission based on existing Chirp Spread Spectrum (CSS) technology. With its characteristics, it belongs to the group of low power consumption and large coverage area (LPWAN) protocols. Looking at the OSI model, it belongs to the first, physical

layer. The history of the LoRa protocol begins with the French company Cycleo, whose founders created a new physical layer of radio transmission based on the existing CSS modulation. Their goal was to provide wireless data exchange for water meters, electricity and gas meters. In 2012, Semtech acquired Cycleo and developed chips for client and access devices. Although CSS modulation had hitherto been applied to military radars and satellite communications, LoRa had simplified its application, eliminating the need for precise synchronization, with the introduction of a very simple way of encoding and decoding signals. In this way, the price of chips became acceptable for widespread use. LoRa uses unlicensed frequency spectrum for its work, which means that its use does not require the approval or lease of a concession from the regulator. These two factors, low cost and free use, have made this protocol extremely popular in a short period of time.

The EBYTE E32 (868T20D) module was used to create the project. The module is based on the Semtech SX1276 chip. The maximum output power of the module is 100 mW, and the manufacturer has declared a range of up to 3 km using a 5dBi antenna without obstacles, at a transfer rate of 2.4 kbps. This module does not have an integrated LoRaWAN protocol, but is designed for direct communication (P2P). If it is to be used for LoRaWAN, then the protocol needs to be implemented on a microcontroller. Communication between the module and the microcontroller is realized through the UART interface (serial port) and two control terminals which are used to determine the state of operation of the module. The module will return feedback via the AUX statement.

LoRaWAN is a software protocol based on the LoRa protocol. Unlike the patent-bound LoRa transmission protocol, LoRaWAN is an open industry standard operated by the nonprofit LoRa Alliance. The protocol uses an unlicensed ISM area (Industry, Science and Medicine) for its work. In Europe, LoRaWAN uses the ISM part of the spectrum that covers the range between 863 - 870 MHz [4]. This range is divided into 15 channels of different widths. For a device to be LoRaWAN compatible, it must be able to use at least the first five channels of 125 kHz and support transmission speeds of 0.3 to 5 kbps. Due to the protection against frequency congestion, the operating cycle of the LoRaWAN device is very low and the transmission time must not exceed 1% of the total operation of the device.

In addition to defining the type of devices and the way they communicate via messages, the LoRaWAN protocol also defines the appearance of the network itself [5]. It consists of end devices, usually various types of sensors in combination with LoRaWAN devices. The sensors appear to central transceivers or concentrators. One sensor can respond to multiple hubs which improves the resilience and range of the network. Hubs are networked to servers that process incoming messages. One of the tasks of the server is to recognize multiple received messages and remove them. Central transceivers must be able to receive a large number of messages using multi-channel radio transceivers and adaptive mode, adapting to the capabilities of the end device. The security of the LoRaWAN network is ensured by authorizing the sensor to the central transceiver, and messages can be encrypted between the sensor and the application server via AES encryption.



MQTT is a simple messaging protocol. It is located in the application layer of the TCP / IP model (5-7 OSI models). It was originally designed for messaging in M2M systems (direct messaging between machines). Its main advantage is the small need for network and computer resources. For these reasons, it has become one of the primary protocols in the IoT world. This protocol is based on the principle of subscriptions to messages and their publication through intermediaries. An intermediary, commonly called a broker, is a server that receives and distributes messages to clients who may be publishers of messages or may be subscribed to them in order to receive them. The two clients will never communicate with each other.

The most important segment of the sensor platform is its reliability. To make sure that an accident occurs in time, we must first ensure the reliability of the platform. Precisely for this reason, in the solution proposed in this paper, periodic reporting from the sensor platform to the system is set. The device will report periodically every 12 hours, and this is taken care of by the alarm system on the microcontroller. Namely, STM32F411 is equipped with a clock that monitors real time (RTC), and offers the ability to set two independent alarms. In this case, one of them is in charge of waking up the process that sends periodic messages with the current state of the measured water flow through the meter.

Before the software implementation of the measurement, it should be noted that the pulse given by the sensor at its output voltage is 5 V. Although the used microcontroller will tolerate this voltage at its input, it is better to lower it to the declared input value of 3.3 V. Such voltage is obtained by two resistors, one with a value of 10 k Ω and the other of 22 k Ω , connected in a simple voltage divider [9]. The connection method is clearly shown in the diagram. The flow volume measurement itself is done by monitoring the number of pulses sent by the water sensor via a standard time counter. Each pulse will be registered by the microcontroller as an interrupt. When pulses appear, it is possible to measure the flow and report it via LoRa radio transmission.

The frequency of the timer is set to 1 MHz via a divider. By comparing the number of clock cycles between the two interrupts, one can very easily obtain the pulse frequency given by the water flow sensor. Knowing the pulse frequency and pulse characteristic, the water flow can be calculated using pre-defined procedure.

The first measured flow value greater than zero sets the sensor platform to an alarm state. As long as there is a flow, periodic advertising will take place every 15 minutes instead of every 12 hours. Five minutes after the flow stops, the device will sound the end of the alarm, and the next call will be made regularly after 12 hours or earlier in the event of a new alarm. The alarm system works internally in such a way that the last measured value of the water flow is read every 5 seconds. This value, together with the current counter time, is continuously stored by the measurement process in the form of a time and flow structure. The read value is stored in a field the size of three elements. If after three readings all three elements in the field are equal, it can be determined that there was no flow in the last 15 seconds and the device exits the alarm state. The system waits another five minutes before announcing



the end of the alarm over the LoRa connection. If the flow occurs again within these five minutes, the system will act as if the alarm has not stopped, that is, it will send a flow message after 15 minutes.




Figure 0.11. Water flow sensor connection diagram

LoRa notifications are intentionally delayed so that in the event of a constant occurrence and interruption of the flow, they would not often send radio messages.

Real life experience

During the measurement, the circuit is supplied with 5 V DC. This is the recommended operating voltage for the LoRa module and water flow sensor used, while the microcontroller can be powered by 5 V or 3.3 V. In this measurement, the first goal is to show that the peak current value will not reach a value greater than 300 mA, which is the maximum that the microcontroller circuit can withstand. This data allows us to power the entire circuit through the microcontroller using the built-in USB port and thus simplify the appearance of the entire sensor. The second goal is to reduce power consumption in order to prolong the autonomy of the sensor operation as much as possible. As an external power supply, a laboratory power supply R-SPS3010 from Nice-power was used, which can provide a stable operating voltage from 0 to 30 V with a current of up to 10 A. The universal measuring instrument UT139B from UNI-T is connected in series. It is set to measure milliamperes during the measurement, keeping the maximum measured value on the screen.

Range measurement

The range was measured from the Zagreb settlement of Vrbani 3, which is located next to Lake Jarun. This location, gives us an insight into what range can be expected in urban and what in rural conditions. Namely, from the central transceiver to the north there is a very urban part with many residential buildings and dense traffic infrastructure, while on the south side is Lake Jarun and the Sava River, which are mostly green areas, smaller forests, and only a few lower buildings. The limiting factor is the position of the antenna of the central transceiver, which was located on the first floor of a residential building, approximately 4 m above ground level and surrounded by buildings. When measuring on the side of

the central transceiver, an omnidirectional antenna with a gain of 3.5 dBi was used, which is stationary placed on the outside of the window of a residential building. On the sensor side, for mobility, a smaller antenna with 2 dBi gain was used. The signal was sent in the open "out of hand". The position of each measurement was recorded via a GPS device on a mobile device and later transferred to Google Earth. In Google Earth, it is possible to import recorded measuring points and measure the distance between them and the antenna of the central transceiver. According to the manufacturer's specification, the maximum range that can be expected from these modules is 3 km in almost ideal conditions with a 5 dBi antenna. In order to somehow approach this distance despite the unfavorable measurement position, the data transfer rate was reduced from the standard module settings from 2.4 kbps to 300 bps. Due to the small amount of data that needs to be transmitted, this is not a limiting factor in practice, and due to the low transmission speed, a smaller amount of errors was obtained when recognizing the received signal and increased success in receiving messages over long distances. In figure below the measured range of the fabricated LoRa system is shown. The position of the central transceiver is shown with an asterisk, while the points from which the signal from the sensor managed to reach it are shown in green. Red dots indicate places where it was not possible to communicate between the sensor and the central transceiver. As expected, the largest range of 3393 m was achieved to the southeast, where apart from a couple of residential buildings near the antenna, there were no additional obstacles. Towards the southwest, the obtained result was 2773 m. However, according to the urban part of the city, the maximum achieved range was 982 m to the east, and to the north it was only 860 m.



Figure 0.12. Central transceiver antenna position and measuring range

According to the specification, the maximum consumption of the used module is 130 mA. The measured consumption of the water flow sensor is 4 mA. The maximum current that can be conducted through the sensor board development board is 300 mA, and the circuit on the development platform used is designed so that the Vbus USB terminal and the 5 V terminals of the circuit are on the same bus. From this we can conclude that the entire interface with the sensor and the LoRa module can be powered by the USB interface.

However, it is necessary to optimize the consumption so that the circuit can run on a commercially available battery for as long as possible. Table shows the current measurements during the operation

of the microcontroller. Here, the microcontroller operated with a maximum operating clock of 96 MHz and without any power optimization. Data are given separately for each element to make it easier to track optimization.



Co-funded by
the European Union

2021-1-SI01-KA220-VET-000034641

Table 0.2. Circuit current without optimization

Connected system components	Current [mA]	State
Microcontroller	26.65	Wait
Microcontroller	26.88	Event stop
Microcontroller + LoRa Module	39.16	Wait
Microcontroller + LoRa Module	121.5	Signal send
Microcontroller + LoRa Module + Sensor	42.51	Wait
Microcontroller + LoRa Module + Sensor	125.7	Signal send

As the flow sensor does not have the possibility of optimization, in Table the values of the current flowing through it are singled out and at the end of each step they will only be added to the obtained results. Table shows that by reducing the operating clock, the current decreased by 11 mA, which is a decrease of slightly more than 40% in the consumption of microprocessors.

Table 0.3. Current through the water sensor

Current [mA]	State
3.35	Idle
4.03	Flow

The first step of optimization is to lower the processor clock to 48 MHz.


Table 0.4 Current with reduced microprocessor clock speed

Connected system components	Current [mA]	State
Microcontroller	15.50	Wait
Microcontroller	15.91	Event stop
Microcontroller + LoRa Module	28.15	Wait

As the LoRa module on the sensor platform is not used for receiving messages, there is no need to keep it constantly active. Fortunately, this module has a mode in which it shuts down its radio transceiver. By changing the code on the microcontroller, an operating mode was introduced where the radio transceiver is turned on only when necessary. With this procedure, the total current through the

microcontroller and the LoRa module dropped to 17.7 mA in standby mode. The STM32F411 microcontroller has various energy saving functions. One of them is a sleep state in which we stop the processor clock completely and listen only to interruptions coming from external devices or clocks. As FreeRTOS was used in the paper, instead of directly sending the microprocessor to sleep, FreeRTOS tickless mode was used. In it, FreeRTOS stops working and puts the microprocessor to sleep. This lowers the current through the circuit consisting of the microcontroller and the LoRa module to 5.87 mA in standby mode, with the total current through the entire circuit now being only 9.22 mA in standby mode. Measuring the current strength has successfully shown how it is possible to use a USB port to power the entire circuit. Also, in several interventions on the program code of the microprocessor, it was possible to lower the current from 42.51 mA to 9.22 mA, which is a difference of 78%. This is very important because waiting time is the state in which the circuit is located almost all the time. Using a portable USB charger (power bank) with a capacity of 10000 mAh (the most common value at the time of writing), with such consumption can be counted on approximately 40 days of autonomous operation of the sensor. Radio signal acquisition showed very good results considering the power and position of the antenna. This measurement is an indication of how even without a great search for the ideal antenna position, a quite decent range can be achieved with a device that has the output power of an average home Wi-Fi system. The maximum measured distance was 3393 m in terms of measurements from ground level and without optical visibility. There is also a large difference in the behaviour of LoRa radio protocols between urban and rural areas. While in an uninhabited area the range exceeded the manufacturer's specifications, in places with several residential buildings, the range dropped sharply. It can be concluded that for the purpose of reporting adverse events in rural and remote areas, LoRa LPWAN is an excellent solution. Smaller range in the urban area is very easy to compensate with more densely placed central transceivers.





Karta


Mapa



Očitanja

2012 siječanj	542	2.236
veljača	978	2.228
ožujak	527	1.935
travanj	563	2.037
svibanj	331	2.578
lipanj	322	2.022
srpanj	573	2.062
kolovoz	386	5.374
rujan	378	2.252
listopad	1.219	1.947
studeni	464	2.325
prosinac	327	1.976
2013 siječanj	288	2.125
veljača	275	2.157
ožujak	202	1.799
travanj	329	2.256
svibanj	250	1.812
lipanj	424	1.952
srpanj	507	2.036
kolman9	402	2.542

Mjesečna potrošnja



Baždarenje




Figure 0.13. LoRa LPWAN

Use case: Rule-based phishing website classification

Before training the classification model it was necessary to choose the features that are relevant and useful for the classification process. To evaluate the features, we used the ranking of features based on the following methods:

1. Information gain that ranks features based on the calculated information gain relative to the classification class, numerical features are first discretized.
2. Gain ratio ranks features based on the calculated gain ratio. Gain ratio is calculated as Information gain divided by the entropy of the feature for which the ratio is computed.



3. Symmetrical uncertainty is a measure that eliminates redundant and meaningless features, which have no interconnectivity with other features.
4. Relief method was proposed by Kira and Rendell and is used for the selection of statistically relevant features, it is resistant to noise and the interdependence of features.

Features are evaluated in a way that is randomly sampled from a given set of instances and take nearest neighbours that belong to the class. If the neighbours are aligned with instances the weighting factor increases, in contrast if the closest neighbours are different the weighting factor decreases.




Figure 0.14. Comparison of different methods for feature selection

If we look at the ranked features, we see that the features that dominate the dataset are:

- post data, which shows us whether the client has filled/not filled the fake
- form in the Lino system
- session change, which shows us if user, during the session, has changed the
- session identifier or not
- session duration, duration of the session in seconds
- robots, which shows us whether the user accessed/not accessed robots.txt file, which defines the rules of robot conduct.

Aforementioned features were selected manually, we ranked all features according to the score of the feature selection method. We selected most significant features for our classification models, in our case the top five features.

Classification Model Selection for Bot from Human Differentiation

A prerequisite for using supervised learning methods and selecting the optimal subset of features is a labelled dataset. Selected features should contribute to the generalization of some classes, ie. for each class they should be able to make a unique behavioural profile. To evaluate the performance of the classification method we used the K-fold cross validation method. For our purposes we used $k = 10$ parts - the relevant literature states that $k = 10$ parts are a optimum number for estimating errors.

Decision Tree C 4.5

Firstly, for classification purposes, we evaluated a decision tree algorithm C 4.5, which is an upgrade of the classic algorithm ID3. Both algorithms are the result of research made by Ross Quinlan. C 4.5 uses a dataset for learning to create a redundant tree. In the case of using similar data, in the learning and validation.




Figure 0.15. Pruned tree, using the full set of features

Classifier has good results but when we use an independent validation set the classifier usually produces bad results. After building a redundant tree, the tree is converted to the IF/THEN rules and the algorithm calculates the best conditions for classification accuracy we remove the IF conditions if they do not reduce the classification accuracy. Pruning is done from the leaves to the root of the tree and is based on the pessimistic estimation of errors; errors are related to the percentage of incorrectly classified cases in the training dataset. Based on the difference in accuracy of rules and standard deviation taken from the binomial distribution we define a certain upper limit of confidence that is usually 0.25, based on which the trees are pruned. For building our models with C 4.5 we set the confidence threshold for pruning to 0.25 and the minimum number of instances per leaf is 2.

	Class	TP Rate	FP Rate	F-measure	AUC
C 4.5 Experiment #1	Human	0.177	0	0.301	0.773
	Robot	1	0.823	0.972	0.773
C 4.5 Experiment #2	Human	0.793	0.002	0.872	0.985
	Robot	0.998	0.207	0.992	0.985
SVM Experiment #1	Human	0.265	0	0.419	0.801
	Robot	1	0.735	0.979	0.801
SVM Experiment #2	Human	0.962	0.006	0.942	0.976
	Robot	0.998	0.042	0.997	0.978

Figure 0.16. Classification results for C 4.5 and SVM, experiment 1 uses only selected features. Experiment 2 uses selected features plus Country and ASN of client

Prior to the classification we removed the class instance of unknown visitors because they represented human attempts to attack with manually entered vales or using non-existing browsers. Method C 4.5 has resulted in the pruned tree showed, which is the same with the optimal selection of features and using the full set of features. It is important to note that the algorithm C 4.5 is very good at choosing features by using heuristics in creating and deleting subtrees.

If we look at the mentioned results, we can see that for the given features (Experiment 1) we have a classification accuracy of 94.5% and a perfect rate of correct positives for robots (TP rate). Classifier badly classifies human visitors (TPR = 0.177), and degrades the classification ability of the robot where the rate is false positive rate is high - 0.823. Looking at F-measure we can say that a good classifier detects correctly robots while wrongly classifies human visitors and commonly (> 80%) declares them robots. We tested the C 4.5 classier with two additional features - client's country and ASN of service provider. This features were resolved from the IP address using the aforementioned GeolP database. This subset (C 4.5 Experiment 2) is shown in mentioned results. We reduced the number of false positives for the class Robot to 0.207, thus the result of classification for class Human was better 0.793.

Support Vector Machine

SVM is an algorithm that finds the maximum margin of separation between classes, while defining the margin as distance between the critical points which are closest to the surface of separation. Points nearest to the surface are called support vectors, the margin M can be seen as the width of the separation between the surfaces. Calculating the support vector is an optimization problem which can be solved using different optimization algorithms. The trick used in calculating SVM is using different kernel functions, which move unsolvable or inadequate problems to a higher dimension, where it can be solved. In our experiments we trained our SVM models with the sequential minimal optimization algorithm using a linear kernel $K(x, y) = \langle x, y \rangle$ where $_ = 1.0-12$ and tolerance is set to 0.001, previously training data was normalised. For Experiment 1 features SVM performs better than C 4.5, the precision

was 95.8%. Human visitors are still a problem, although SVM has much higher rate of true positives (26.5%). Increased detection rate of human visitor gives a lower rate of wrong detection of robots (73.5%). F measure is very good for robots and much better for human visitors (even better than method C 4.5) - but still too low to be used - (0.419). With additional features Country and ASN (Experiment 2) we obtained a false rate for both classes under 5%. The true positive rate was also high for class Human 0.962 and for class Robot 0.998. We can conclude that with this subset of features and with regular retraining to avoid the concept drift, this model is feasible for every day use.

Usecase: Data loss prevention cloud-based system

Comparison of DLP solutions available on the market

Comparison of DLP solutions available on the market, based on the Gartner Gartner® Releases 2022 Market Guide for Data Loss Prevention: Key Takeaways.

Symantec

Based in Mountain View, California, Symantec has been on the DLP market since the acquisition of Vontu in 2007. Symantec has recently released Symantec Data Loss Prevention 15.0 and has component products for DLP Enforce, DLP IT Analytics, cloud storage (supports more than 65 cloud applications), Cloud Prevent for Microsoft Office 365, DLP for endpoint, DLP for network and DLP storage, as well as third-party security technology DLP API support, such as content retrieval, reporting, and FlexResponse for encrypting content or DRM application. Symantec continues to invest in DLP technology and improves its data protection business unit. In 2016, Symantec has made the acquisition of Blue Coat, which gives it the option of purchasing Elastica and Perspecsy for the Blue Coat, for which there are integration of DLP policies through the two-way REST API between Elastica and Symantec DLP. Symantec is a convenient choice for organizations that require advanced detection techniques and integration with CASB for a unique data protection policy.

Advantages

Symantec offers the most advanced detection techniques in the market with advanced functionality such as form recognition, image analysis, and handwriting recognition that can cover a wide range of data loss scenarios. Symantec supports a hybrid deployment model for several of its DLP products where Detection Servers installed on AWS, Azure, or Rackspace connect to a local DLP Enforce platform. Symantec's SmartResponse system offers a wide range of administrative flexibility based on content actions that conform to the DLP rule. Its Vector Machine Learning (VML) DLP enables users to learn the DLP system by providing both positive and negative sample content. This could be useful if traditional matchmaking methods are not sufficient to match content correctly.

Weaknesses

Symantec clients expressed frustration when purchasing or updating Data Insight plugins for Symantec DLP, which is now owned by Veritas. Make sure your Symantec DLP vendor can also sell Veritas Data Insight if you are interested in this add-on. Monitoring and detecting sensitive data in cloud applications requires DLP endpoint detection and the required Symantec CASB connectors to achieve full functionality. Clients express concern over the overall cost of implementing Symantec DLP, compared to competing products.

Digital Guardian

Established in 2002, Digital Guardian (formerly Verdasys) is headquartered in Waltham, Massachusetts. Access to Digital Guardian DLP is primarily via the DLP endpoint, with strong partnerships for DLP product network integration and DLP detection by October 2015, when Code Green Networks (CGN) was acquired through acquisition. Since then, it has launched it as a line of Digital Guardian Network DLP products. The Digital Guardian endpoint covers DLP, advanced threat protection, and endpoint detection and response (EDR) in a single agent installed on desktop computers, laptops and servers running on Windows, Linux, and Mac OS X, as well as support for VDI environments. The Digital Guardian Network DLP and the Digital Guardian Discovery product cover DLP networks, cloud data protection, and data discovery, and are offered as hardware, software applications, and / or virtual apps. During 2016, Digital Guardian worked on simplifying and integrating management capabilities between its DLP endpoints and assets from CGN acquisitions. Digital Guardian also has an existing partnership with Fidelis Cybersecurity Network DLP. Several Gartner clients recently talked about this partnership, and Gartner believes that, in addition to existing joint customers, the partnership will continue to reduce and eventually stop. Digital Guardian is a suitable choice for organizations with strong concerns about the legislation, particularly in the health sector and financial services, as well as the organization with the requirements of AD protection of intellectual property. Digital Guardian is also a good choice for organizations that require uniformity of DLP rules to work equally well in all Windows, Mac OS X and Linux operating systems.

Advantages

Clients report faster implementation times and successful projects when using the Digital Guardian product in combination with managed digital guardian services. Digital Guardian has integration with wider security products, including threat intelligence, network sandbox, User and Entity Analysis (UEBA), Cloud Data Protection, and Security Event Management (SIEM, including IBM QRadar and Splunk applications). Customers like the option of modular licensing for the DLP endpoint, with support for Windows, Mac OS X and Linux, and feature endpoints that can be licensed in any combination of device visibility and control, DLP and advanced threat protection. Digital Guardian's vision shows a strong understanding of technology, security, threats, and trends in the industry that will shape their bids.

Weaknesses

Digital Guardian does not have a common policy for end points and network products. The Digital Guardian Agent can not differentiate between personal and business accounts for Microsoft OneDrive. However, it can prevent the use of personal Microsoft OneDrive applications. Customers expressed concern about the speed of integration of the acquired CGN. Structured data indexing is not supported by the Digital Guardian endpoint agent, but this feature is available through CGN agent.

Forcepoint

In 2015, Raytheon and Vista Equity Partners completed a joint venture that combines Websense, a portfolio company Vista Equity and Raytheon Cyber Products. In 2016, the company gained two lines of Intel Security - Stonesoft and Sidewinder firewalls through acquisition - and restarted the combined company as Forcepoint. Raytheon has already held municipal share Forcepoint's, and Vista Equity Partners holds a minority stake. Headquartered in Austin, Texas, Forcepoint has been a leader in the DLP product market, previously known as Raytheon-Websense, for several years. The Forcepoint DLP product line includes Forcepoint DLP Discover, Forcepoint DLP Gateway, Forcepoint Cloud Applications, and Forcepoint DLP Endpoint. During the years of delivery of DLP and integrated DLP modules for its secure web and e-mail gateway products, Forcepoint has created an outstanding DLP packet for network coverage, endpoints and data discovery (both client and cloud), with special attention to the protection of intellectual property and the implementation of the compliance policy with the regulations. Forcepoint is a suitable choice for organizations with requirements for legal compliance and intellectual property protection or organizations that want to implement DLP virtual devices in the Azure public cloud infrastructure.

Advantages

Forcepoint DLP Endpoint can automatically encrypt / decrypt files via Microsoft RMS without removing RMS protection based on end-to-end data, motion data, and discovery rules. Forcepoint provides over 350 predefined rules and embedded component UEBA for additional security analytical features that perform incident risk rating, identify threats from internal users, point out endangered endpoints, and calculate data theft risk indicators to identify the most vulnerable users and activities. Indexing structured data, especially data-indexing support in Salesforce, cites clients as the key differentiator factor.

Weaknesses

Clients reported problems with technical support for indexing structured data. If you need to index structured data in the database, make sure you thoroughly test it on live data in your specific database environment. Raytheon's involvement in the defense market will help strengthen Forcepoint with additional intelligence and products. However, there is no success of security vendors owned by defense structures that have succeeded in commercial markets. Forcepoint's relevance in some geographic areas can be problematic because of Raytheon's strong American loyalty. Some Gartner clients have noted this complaint and see if this is causing your concern in your organization.

Intel Security (today: McAfee)

Over the past few years, Intel has changed its investment in and from various product lines several times and has not sufficiently considered these changes inside and outside the company. This has caused exhaustion of employees at alarming rates, many of which have been launched by new security companies or are employed by competitive security vendors. Historically, in many of Intel's security products, there has been a chronic lack of investment.

Intel's security approach was to integrate acquisitions with McAfee's ePolicy Orchestrator (McAfee ePO) policy management system, alert monitoring, and link security events between DLP Event Ends, Network Transfers, and Restricted Data on Storage Data in the Organization. The DLP 10.0 release has brought further improvements to DLP and updates to DLP online products in 2016 highlighted McAfee's renewed focus on data protection. Intel Security is a good choice for organizations that have significant resources invested in McAfee ePO and want a unique supplier that can provide DLP, device control and encryption.

Advantages

DLP integration within McAfee Web Gateway proxy supports decrypting and re-encrypting site traffic site traffic, including email service providers and cloud storage products. The capture database can index and store all visible network and endpoint components. Clients reported this useful to test new rules, forensic analysis of events that occurred prior to policy making and after-event investigation. It also supports e-discovery and legacy retention, as well as integration directly with the software Guidance Software and AccessData. McAfee DLP includes the basic level of data classification on the DLP 10 endpoint for Windows and Mac OS X and can still be firmly integrated with Titus and Bold James for various data classification options. DLP endpoint rules are aware of locations and may have different responses and content remedies when they are online when they are offline. The Federation of Security Innovations (SIA) is still robust and is a good way for Intel Security customers to maximize their DLP investment due to proven and tested integration of data product classifications, DRIF and UEBA suppliers.

Weaknesses

McAfee DLP supports native API integration with Cloud Data Box but support for other cloud applications and cloud storage support are missing. Intel Security has made some improvements to DLP Agent 10 on Mac OS X, but still lacks support for email, web, and cloud. Linux is not supported. Customers report that the configuration of DLP rules can be complex and disadvantageous compared to other DLP products. The future success of Intel Security in the DLP market will depend on their performance while acting as a company and whether the focus can be on data security tasks over a longer period of time.

Usecase: Dynamic web site hosting



Co-funded by
the European Union

2021-1-SI01-KA220-VET-000034641

Inspired by: <https://www.linkedin.com/pulse/host-dynamic-website-aws-sara-mostafa/>

How to deploy a dynamic website with AWS by uploading your website content into S3 bucket, create an EC2 instance to host web app on it as in this scenario EC2 acts like a public server all people from the world can visit this server.

Amazon S3 (Simple Storage Service) is a service offered by AWS for object storage through a web service interface. It can be used to store or retrieve any amount of data such as documents, images, videos, etc. S3 bucket is a resource in Amazon S3. It is a container where files and folders can be uploaded.

Amazon EC2 (Elastic Compute Cloud) is a service offered by AWS. It is considered as a virtual server.

IAM (Identity and access management) Role is used to give permission to service to do something on another service.

LAMP web server can be used to host a static website or deploy a dynamic PHP application that reads and writes information to a database.

Steps

Step 1: Create S3 Bucket

You will need to create an S3 bucket to put your website's files and folders. To do this, login into your AWS management console and click on Services on the top navbar. From the Services drop-down, select S3 from the Storage section. This should display the S3 dashboard.

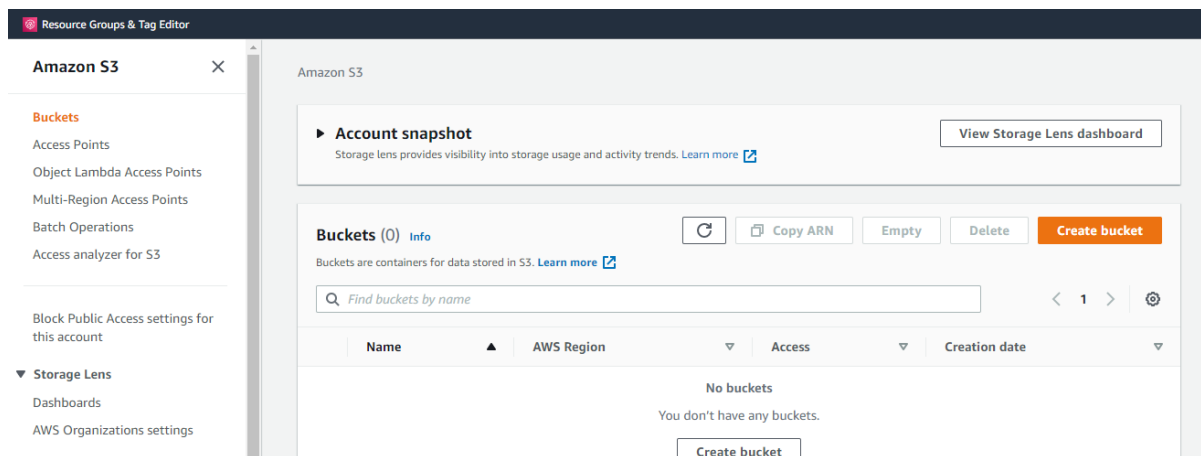


Figure 0.17. Creating an S3 bucket – first step

From the S3 dashboard, click on Create bucket. Give the bucket a unique name, the name you choose must be globally unique. Next, choose your preferred AWS Region from the drop-down.

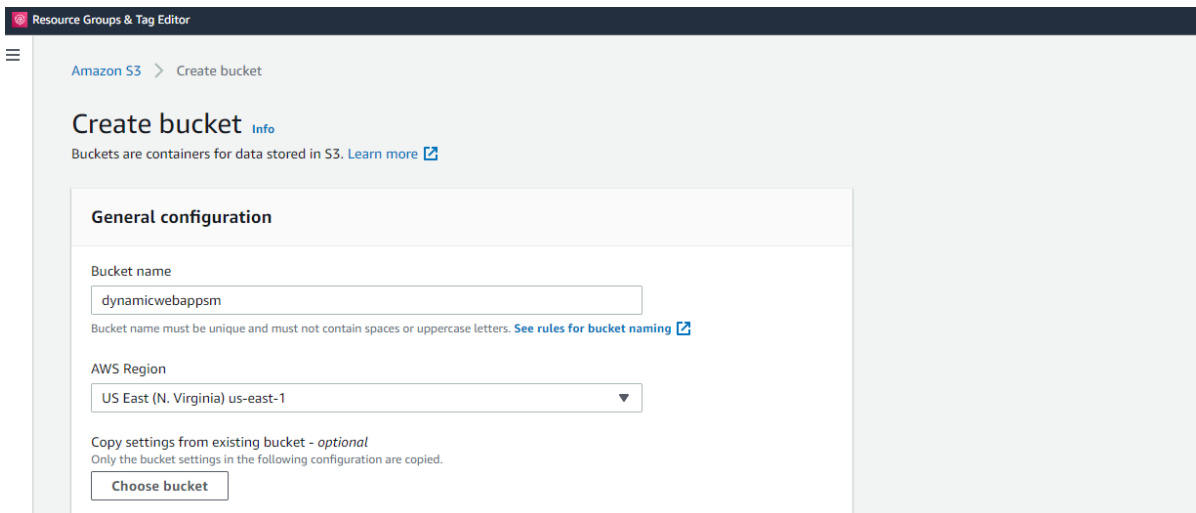


Figure 0.18. Creating an S3 bucket – second step

Under Block Public Access settings for this bucket section, check the Block all public access checkbox. This is done to make the bucket not accessible to the public.

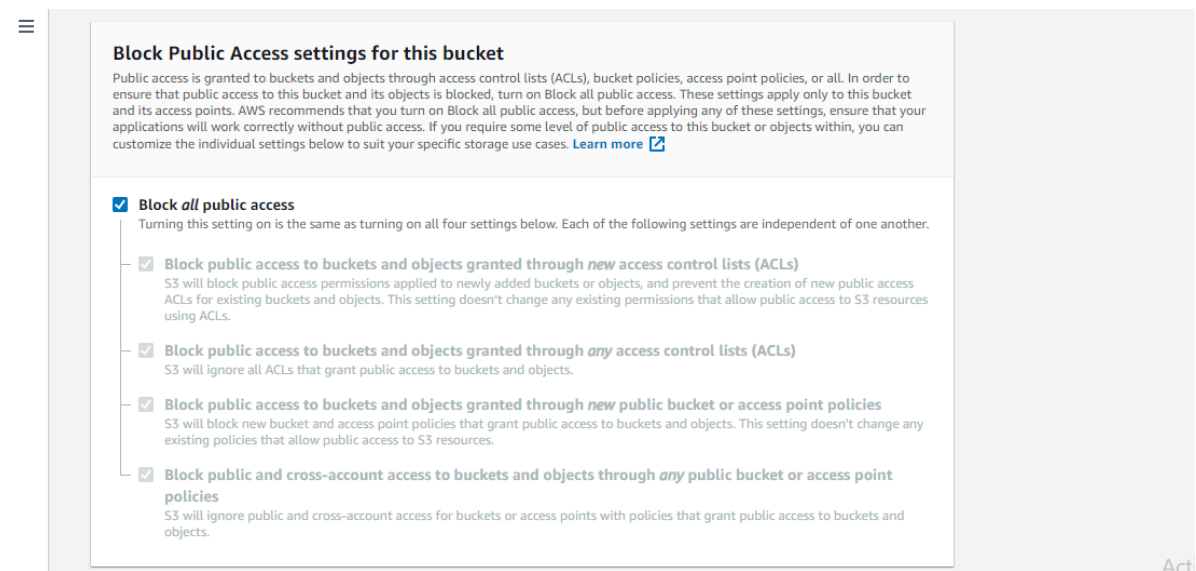


Figure 0.19. Creating an S3 bucket – third step

Click on Disable for Bucket Versioning. You can also Add tag to the bucket for easy identification.

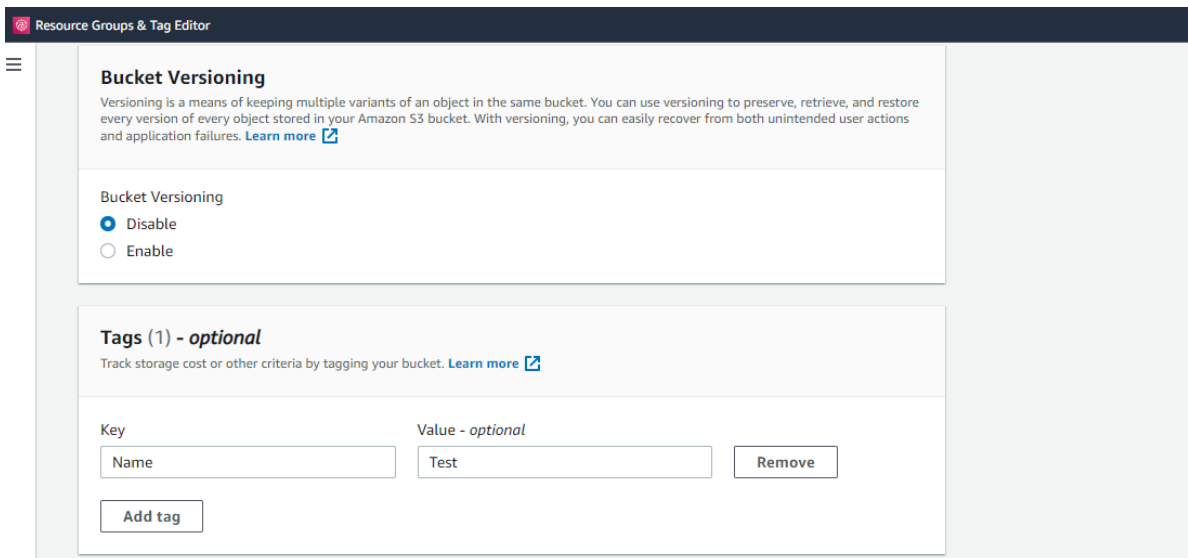


Figure 0.20. Creating an S3 bucket – fourth step

Under Default encryption section, click on Enable for Server-side encryption. Then check Amazon S3 Key (SSE-S3).

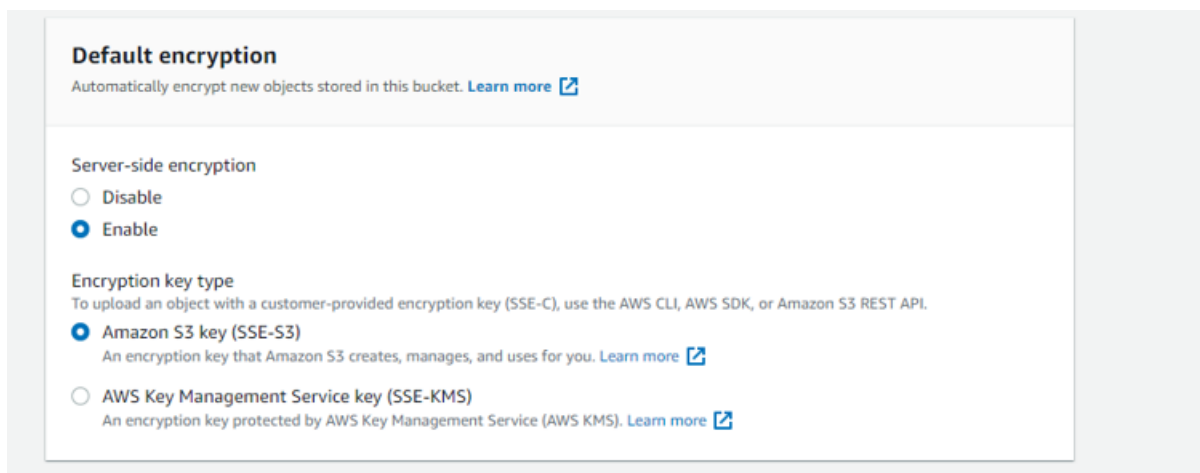


Figure 0.21. Creating an S3 bucket – fifth step

Then click on Create bucket.

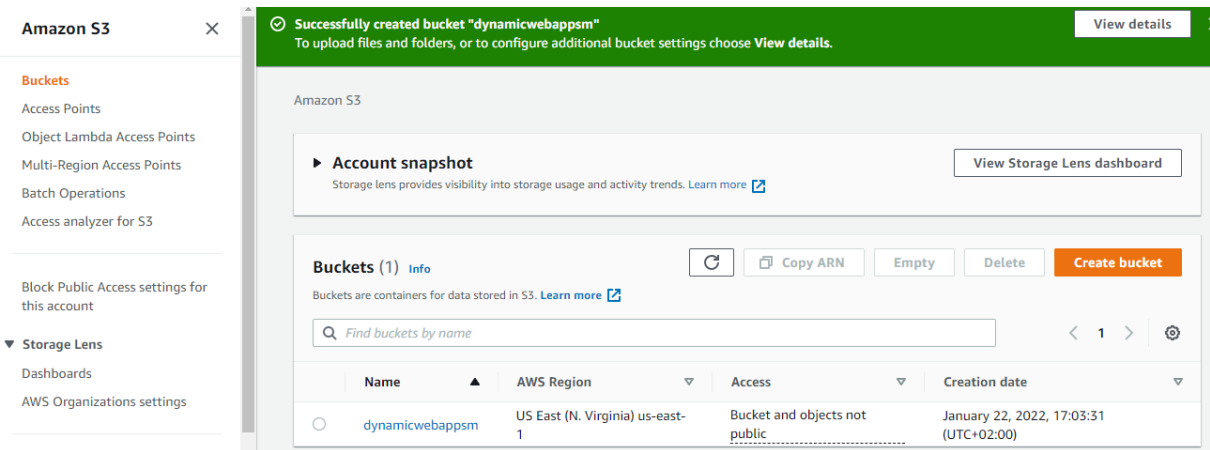


Figure 0.22. Creating an S3 bucket – sixth step

Step 2: Upload web files to S3 bucket

After creating the bucket, you need to upload your website’s files and folders into it.

From the S3 dashboard, click on the name of the bucket you just created. On the Objects tab, you can see that the bucket is currently empty, click on the Upload button.

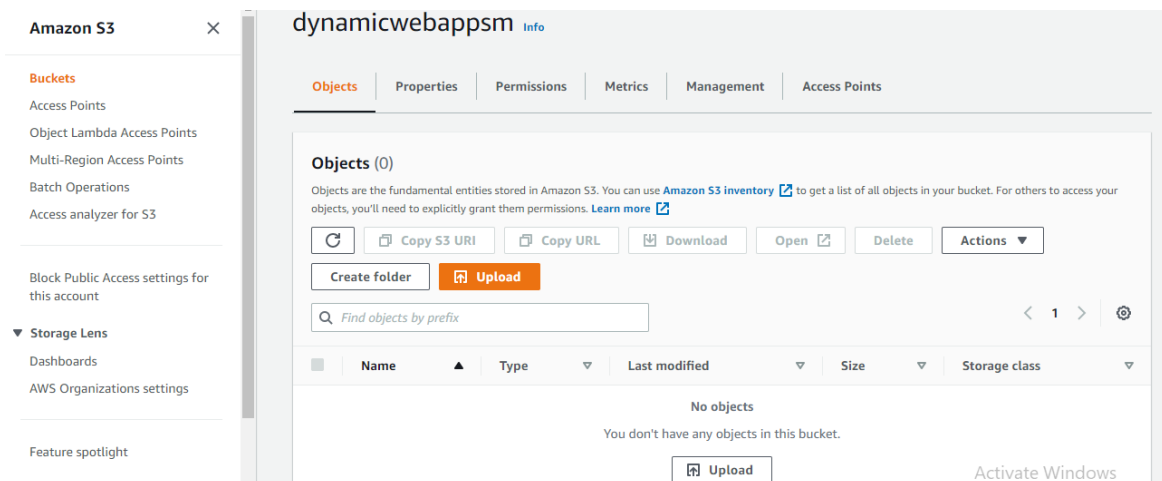


Figure 0.23. Upload web files to S3 bucket – first step

This should take you to the Upload page.

Name	Date modified	Type	Size
assets	1/8/2022 7:10 PM	File folder	
css	1/8/2022 7:10 PM	File folder	
js	1/8/2022 7:10 PM	File folder	
index.html	11/29/2020 7:04 AM	Microsoft Edge H...	30 KB

Figure 0.24. Upload web files to S3 bucket – second step

Drag and drop files and folders you want to upload here, or choose **Add files**, or **Add folders**.

Files and folders (14 Total, 438.2 KB) Remove Add files Add folder

All files and folders in this table will be uploaded.

Find by name < 1 2 >

<input type="checkbox"/>	Name ▲	Folder ▼	Type ▼	Size ▼
<input type="checkbox"/>	avataaars.svg	assets/img/	image/svg+xml	12.5 KB
<input type="checkbox"/>	cabin.png	assets/img/portfolio/	image/png	35.7 KB
<input type="checkbox"/>	cake.png	assets/img/portfolio/	image/png	16.7 KB
<input type="checkbox"/>	circus.png	assets/img/portfolio/	image/png	27.3 KB
<input type="checkbox"/>	contact_me.js	assets/mail/	text/javascript	3.6 KB
<input type="checkbox"/>	contact_me.php	assets/mail/	-	1.1 KB
<input type="checkbox"/>	favicon.ico	assets/img/	image/x-icon	22.9 KB
<input type="checkbox"/>	game.png	assets/img/portfolio/	image/png	25.3 KB
<input type="checkbox"/>	index.html	-	text/html	29.9 KB
<input type="checkbox"/>	jqBootstrapValidation.js	assets/mail/	text/javascript	35.3 KB

Figure 0.25. Upload web files to S3 bucket – third step

After the necessary files and folders have been added, scroll down and click on Upload. The uploading should be done in a few minutes depending on your network and content size. Also, please do not close the tab while the upload process is going on.

Step 3: Create IAM Role

Now, EC2 want to pull code from S3. So you want to create IAM Role to give EC2 permission to access S3. To do this, from the Services drop-down, select IAM from the Security, Identity& Compliance section. From the IAM dashboard, click on Roles. Then Click on Create role.

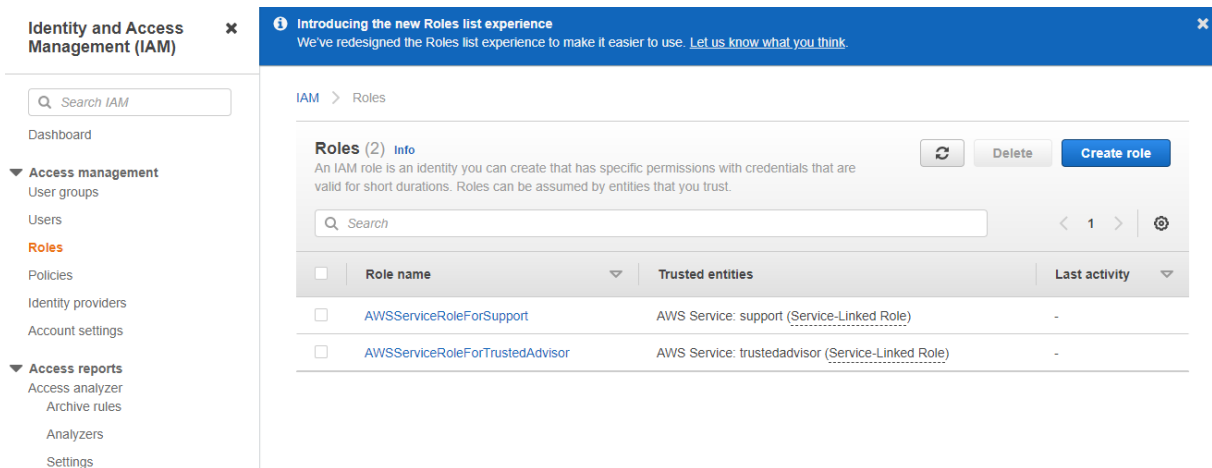


Figure 0.26. Create IAM Role – first step

Choose EC2 and click Next: Permissions.

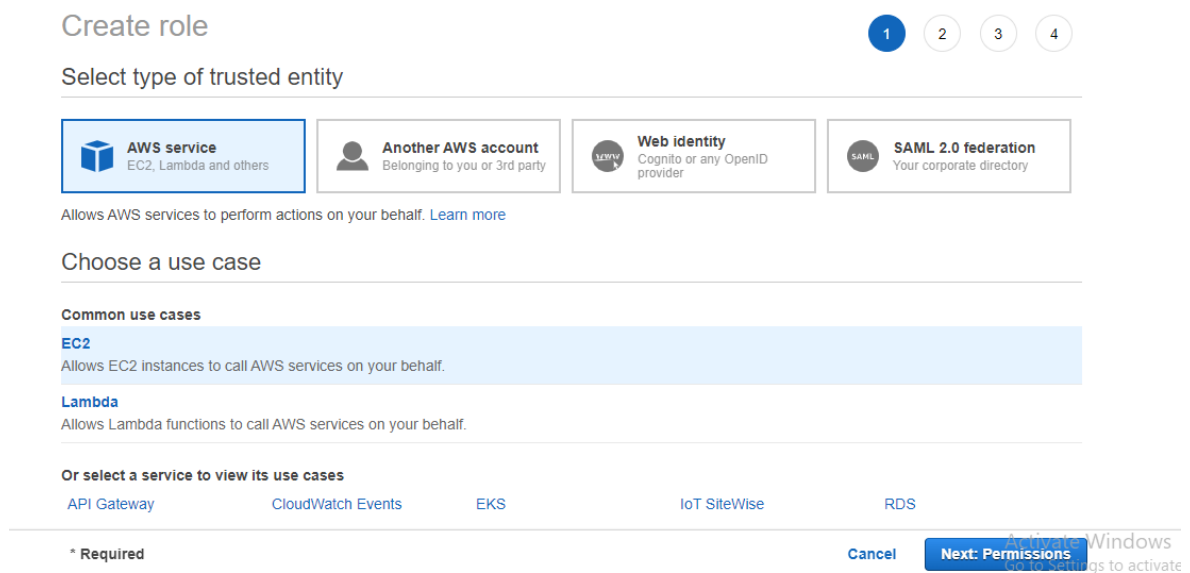


Figure 0.27. Create IAM Role – second step

Search for S3 and check AmazonS3FullAccess. Then click Next: Tags.

▼ Attach permissions policies

Choose one or more policies to attach to your new role.

Create policy ↻

Filter policies Showing 9 results

	Policy name	Used as
<input type="checkbox"/>	▶ AmazonDMSRedshiftS3Role	None
<input checked="" type="checkbox"/>	▶ AmazonS3FullAccess	None
<input type="checkbox"/>	▶ AmazonS3ObjectLambdaExecutionRolePolicy	None
<input type="checkbox"/>	▶ AmazonS3OutpostsFullAccess	None
<input type="checkbox"/>	▶ AmazonS3OutpostsReadOnlyAccess	None
<input type="checkbox"/>	▶ AmazonS3ReadOnlyAccess	None
<input type="checkbox"/>	▶ IVSRecordToS3	None
<input type="checkbox"/>	▶ QuickSightAccessForS3StorageManagementAnalyticsReadOnly	None

* Required Cancel Previous **Next: Tags**

Figure 0.28. Create IAM Role – third step

Click on Next: Review.

Create role 1 2 3 4

Add tags (optional)

IAM tags are key-value pairs you can add to your role. Tags can include user information, such as an email address, or can be descriptive, such as a job title. You can use the tags to organize, track, or control access for this role. [Learn more](#)

Key	Value (optional)	Remove
<input type="text" value="Add new key"/>	<input type="text"/>	

You can add 50 more tags.

Cancel Previous **Next: Review**

Figure 0.29. Create IAM Role – fourth step

Give the role name and description. Then click on Create role.

Create role



Review

Provide the required information below and review this role before you create it.

Role name*
Use alphanumeric and '+,=, @, -' characters. Maximum 64 characters.

Role description
Maximum 1000 characters. Use alphanumeric and '+,=, @, -' characters.

Trusted entities AWS service: ec2.amazonaws.com

Policies  [AmazonS3FullAccess](#)

Permissions boundary Permissions boundary is not set

* Required

Cancel

Previous

Create role

Figure 0.30. Create IAM Role – fifth step

Now, the role has been created successfully.

Identity and Access Management (IAM)

Introducing the new Roles list experience
We've redesigned the Roles list experience to make it easier to use. [Let us know what you think.](#)

✓ The role **ec2s3role** has been created.

IAM > Roles

Roles (3) Info Refresh Delete Create role

An IAM role is an identity you can create that has specific permissions with credentials that are valid for short durations. Roles can be assumed by entities that you trust.

<input type="checkbox"/>	Role name	Trusted entities	Last activity
<input type="checkbox"/>	AWSServiceRoleForSupport	AWS Service: support (Service-Linked Role)	-
<input type="checkbox"/>	AWSServiceRoleForTrustedAdvisor	AWS Service: trustedadvisor (Service-Linked Role)	-
<input type="checkbox"/>	ec2s3role	AWS Service: ec2	-

Figure 0.31. Create IAM Role – sixth step

Step 4: Create an EC2 instance

You will need to create an EC2 instance to install apache (/var/www/html) and copy the content of S3 to html directory. To do this, from the Services drop-down, select EC2 from the Compute section. This should display the EC2 dashboard. From the EC2 dashboard, click on Launch instance.

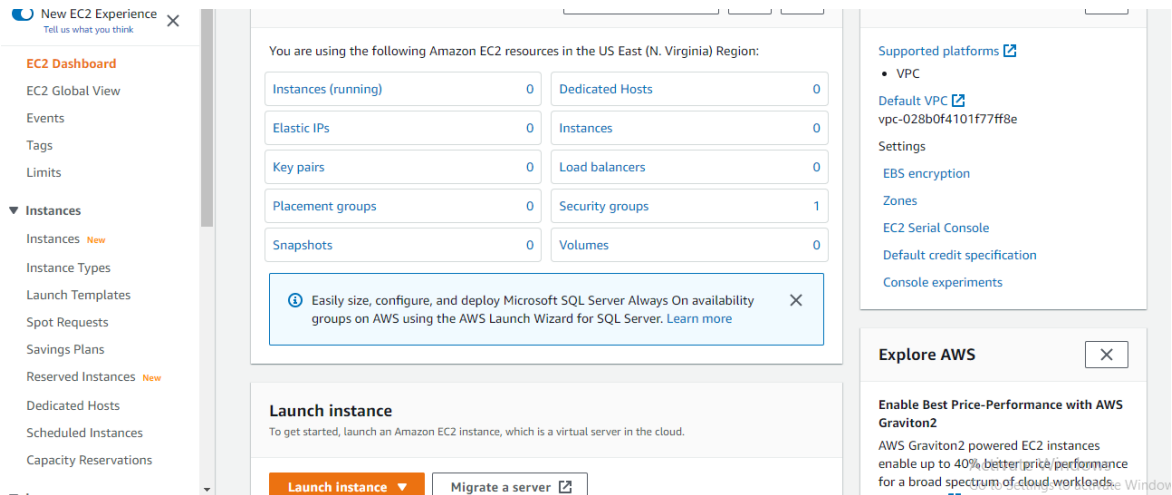


Figure 0.32. Create an EC2 instance – first step

For AMI, choose Quick Start and click on Select for Amazon Linux (Free tier eligible).

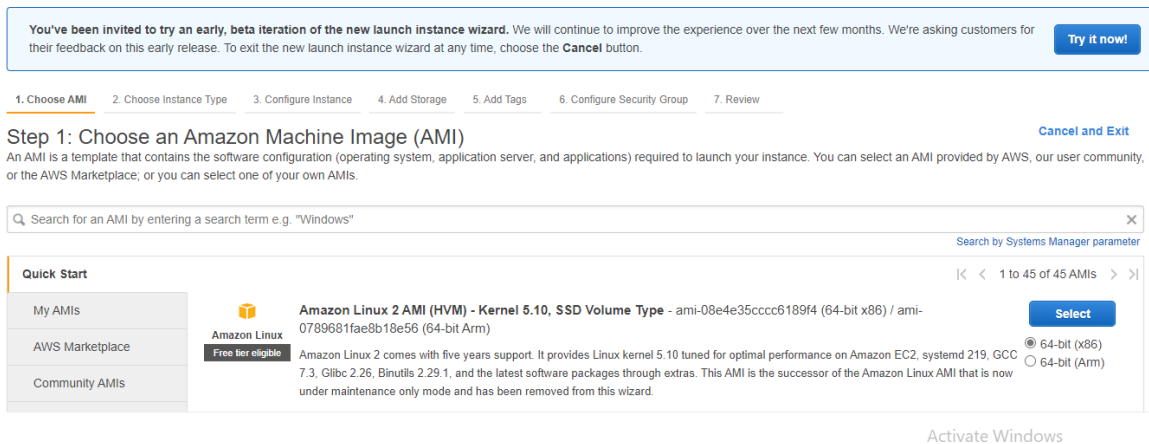


Figure 0.33. Create an EC2 instance – second step

For an instance type, choose t2.micro (Free tier eligible). And click on Next: Configure Instance Details.

Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. [Learn more](#) about instance types and how they can meet your computing needs.

Filter by: All instance families Current generation Show/Hide Columns

Currently selected: t2.micro (- ECUs, 1 vCPUs, 2.5 GHz, -, 1 GiB memory, EBS only)

	Family	Type	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance	IPv6 Support
<input type="checkbox"/>	t2	t2.nano	1	0.5	EBS only	-	Low to Moderate	Yes
<input checked="" type="checkbox"/>	t2	t2.micro <small>Free tier eligible</small>	1	1	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.small	1	2	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.medium	2	4	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.large	2	8	EBS only	-	Low to Moderate	Yes

Cancel Previous **Review and Launch** Next: Configure Instance Details

Figure 0.34. Create an EC2 instance – third step

Determine 1 for Number of instances, default vpc for Network and Default in us-east-1a for Subnet.

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances [Launch into Auto Scaling Group](#)

Purchasing option Request Spot instances

Network [Create new VPC](#)

Subnet [Create new subnet](#)
4091 IP Addresses available

Auto-assign Public IP

Hostname type

DNS Hostname Enable IP name IPv4 (A record) DNS requests
 Enable resource-based IPv4 (A record) DNS requests
 Enable resource-based IPv6 (AAAA record) DNS requests

Figure 0.35. Create an EC2 instance – fourth step

Choose ec2s3role or whatever you named for IAM role. And Terminate for Shutdown behaviour. Then click on Next: Add Storage.

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

Placement group Add instance to placement group

Capacity Reservation

Domain join directory [Create new directory](#)

IAM role [Create new IAM role](#)

Shutdown behavior

Stop - Hibernate behavior Enable hibernation as an additional stop behavior

Enable termination protection Protect against accidental termination

Monitoring Enable CloudWatch detailed monitoring
[Additional charges apply.](#)

Tenancy
[Additional charges will apply for dedicated tenancy.](#)

Elastic Inference Add an Elastic Inference accelerator

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)

Figure 0.36. Create an EC2 instance – fifth step

Click on Next: Add Tags.

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encryption
Root	/dev/xvda	snap-0c03ce90cef384dca	8	General Purpose SSD (gp2)	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypt

[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Shared file systems

You currently don't have any file systems on this instance. Select "Add file system" button below to add a file system.

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Tags](#)

Figure 0.37. Create an EC2 instance – sixth step

You can add tag Name: DynamicSite. Then click on Next: Configure Security Group.

Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. A copy of a tag can be applied to volumes, instances or both. Tags will be applied to all instances and volumes. [Learn more](#) about tagging your Amazon EC2 resources.

Key (128 characters maximum)	Value (256 characters maximum)	Instances (i)	Volumes (i)	Network Interfaces (i)
Name	DynamicSite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Add another tag (Up to 50 tags maximum)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Configure Security Group](#)

Figure 0.38. Create an EC2 instance – seventh step

Select Create a new security group. Give it Name: DynamicWebsiteSG and description: SG for DynamicWebApp. For SSH rule select My IP for Source. Click on Add Rule and select HTTP for Type and Anywhere for Source. Last rule select HTTPS for Type and Anywhere for Source. Click on Review and Launch.

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group: Create a new security group
 Select an existing security group

Security group name:
 Description:

Type (i)	Protocol (i)	Port Range (i)	Source (i)	Description (i)
SSH	TCP	22	My IP 197.42.126.153/32	e.g. SSH for Admin Desktop
HTTP	TCP	80	Anywhere 0.0.0.0, ::0	e.g. SSH for Admin Desktop
HTTPS	TCP	443	Anywhere 0.0.0.0, ::0	e.g. SSH for Admin Desktop

Add Rule

[Cancel](#) [Previous](#) [Review and Launch](#)

Figure 0.39. Create an EC2 instance – eight step

Click on Launch.



Step 7: Review Instance Launch

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	197.42.126.153/32	
HTTP	TCP	80	0.0.0.0/0	
HTTP	TCP	80	:::0	
HTTPS	TCP	443	0.0.0.0/0	
HTTPS	TCP	443	:::0	

Figure 0.40. Create an EC2 instance – ninth step

Select Create a new key pair and RSA for type. Give it name WebServerKey and click on Download Key Pair. Note: You should download the key to can ssh on EC2. Click on Launch Instances.

Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance. Amazon EC2 supports ED25519 and RSA key pair types.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Create a new key pair

Key pair type
 RSA ED25519

Key pair name
WebServerKey

Download Key Pair

You have to download the private key file (*.pem file) before you can continue. Store it in a secure and accessible location. You will not be able to download the file again after it's created.

Cancel Launch Instances

Figure 0.41. Create an EC2 instance – tenth step

Now, instance is launching successfully.

Launch Status

Your instances are now launching
The following instance launches have been initiated: [i-095e1941ebb94afb2](#) [View launch log](#)

Get notified of estimated charges
Create billing alerts to get an email notification when estimated charges on your AWS bill exceed an amount you define (for example, if you exceed the free usage tier).

How to connect to your instances

Your instances are launching, and it may take a few minutes until they are in the **running** state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.

Click **View Instances** to monitor your instances' status. Once your instances are in the **running** state, you can **connect** to them from the Instances screen. [Find out](#) how to connect to your instances.

▼ Here are some helpful resources to get you started

- [How to connect to your Linux instance](#)
- [Amazon EC2: User Guide](#)
- [Learn about AWS Free Usage Tier](#)
- [Amazon EC2: Discussion Forum](#)

[Activate Windows](#)

Figure 0.42. Create an EC2 instance – eleventh step

Click on Review Instance and wait Status check will be 2/2 checks passed.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone
DynamicSite	i-095e1941ebb94afb2	Running	t2.micro	2/2 checks passed	No alarms	us-east-1a

Figure 0.43. Create an EC2 instance - twelve step

Step 5: SSH with MobaXterm

Now, you want to connect to EC2 by using MobaXterm. First you should copy public IPv4 address of EC2 instance.

Instance ID	Public IPv4 address	Private IPv4 addresses
i-095e1941ebb94afb2 (DynamicSite)	35.173.198.68 open address	172.31.92.121

Instance state	Private IP DNS name (IPv4 only)	Public IPv4 DNS
Running	ip-172-31-92-121.ec2.internal	ec2-35-173-198-68.compute-1.amazonaws.com open address

Instance type	Elastic IP addresses	Answer private resource DNS name IPv4 (A)
t2.micro	-	ip-172-31-92-121.ec2.internal

AWS Compute Optimizer finding	IAM Role	VPC ID
Opt-in to AWS Compute Optimizer for recommendations. Learn more	ec2s3role	vpc-028b0f4101f77ff8e

Subnet ID
subnet-057363357afd00871

Figure 0.44 Connecting to EC2 by using MobaXterm - first step

Open MobaXterm and start a new remote session by clicking on Session.

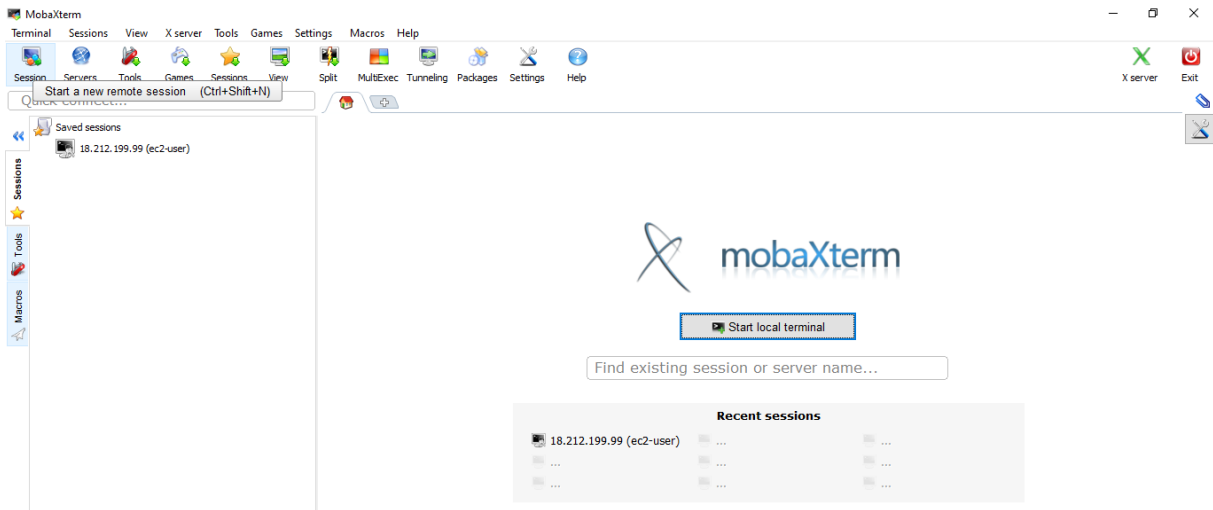


Figure 0.45. Connecting to EC2 by using MobaXterm - second step

Click on SSH. Paste IP of your EC2 For example:(3.86.76.216). And ec2-user for Specify username. Click on Advanced SSH settings, check Use private key and browse location of key. Click OK.

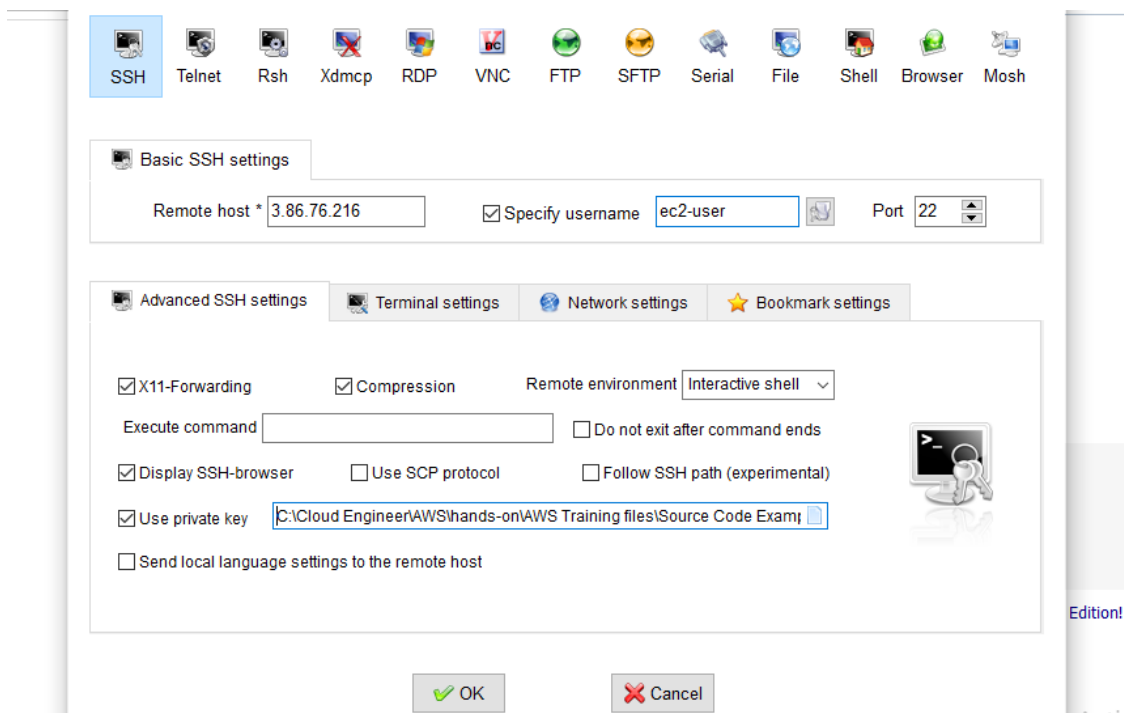


Figure 0.46. Connecting to EC2 by using MobaXterm – third step

Now, you connected to EC2 successfully.




Figure 0.47. Connecting to EC2 by using MobaXterm - fourth step

Step 6: Install a LAMP web server on Amazon Linux 2

The following procedures help you install an Apache web server with PHP and MariaDB. To ensure that all of your software packages are up to date, perform a quick software update on your instance.

```
sudo yum update -y
```

Install the lamp-mariadb10.2-php7.2 and php7.2 Amazon Linux Extras repositories to get the latest versions of the LAMP MariaDB and PHP packages for Amazon Linux 2.

```
sudo amazon-linux-extras install -y lamp-mariadb10.2-php7.2 php7.2
```

Now, you can install the Apache web server, MariaDB, and PHP software packages.

```
sudo yum install -y httpd mariadb-server
```

Start the Apache web server.

```
sudo systemctl start httpd
```

Use the systemctl command to configure the Apache web server to start at each system boot.

```
sudo systemctl enable httpd
```

You can verify that httpd is on by running

```
sudo systemctl is-enabled httpd
```

Now, you want to copy content of website from S3 to directory `/var/www/html` in EC2 . Make sure you copy your S3 bucket name.

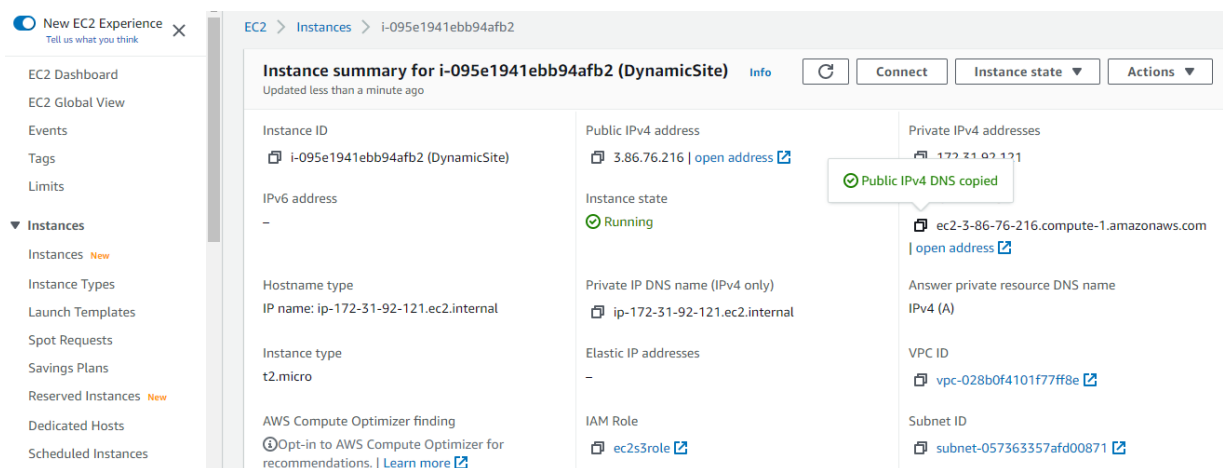
```
sudo aws s3 cp s3://dynamicwebappsm --region us-east-1 /var/www/html/ --recursive
```

To verify that content is copied to `/var/www/html` .

```
cd /var/www/html
```

```
ls
```

Copy public IPv4 DNS and paste it in a new tab.



The screenshot displays the AWS Management Console interface for an EC2 instance. The instance is named 'i-095e1941ebb94afb2 (DynamicSite)' and is currently in a 'Running' state. A green notification bubble with a checkmark and the text 'Public IPv4 DNS copied' is overlaid on the console. The public IPv4 address is listed as 3.86.76.216. The console also shows other details such as the instance type (t2.micro), IAM role (ec2s3role), and VPC ID (vpc-028b0f4101f77ff8e).

Figure 0.48. Installing a LAMP web server on Amazon Linux 2

Congratulations, you have deployed a dynamic website on EC2 successfully.

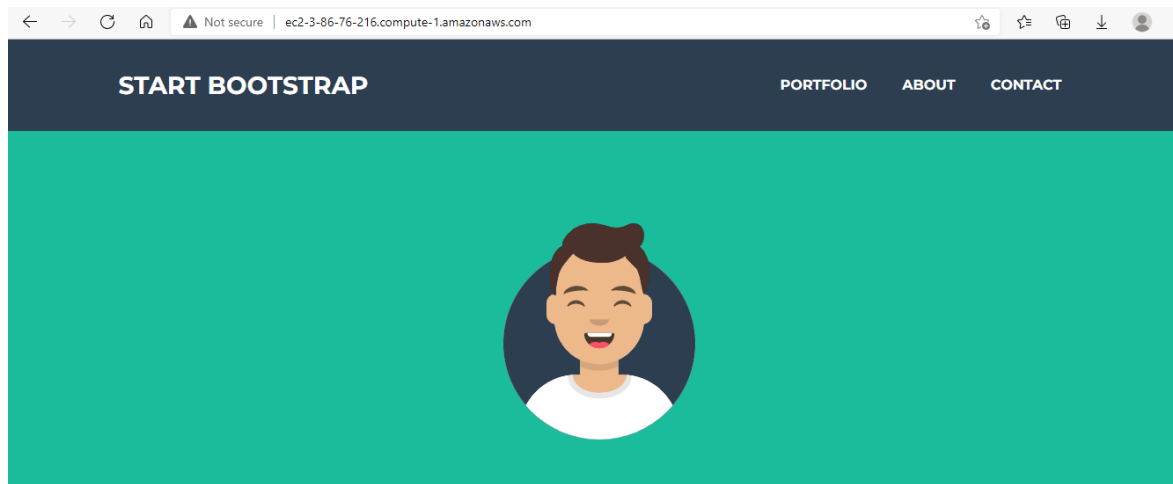


Figure 0.49. Successful deployment a dynamic website on EC2

Usecase: Host a static website using AWS (or other clouds)

Step-by-Step Guide

Basic Configurations

1. Go to the S3 console and create a new bucket with default settings.
2. Go the properties of your bucket and choose the option "Static website hosting."
3. Enable the option "Use this bucket to host a website."
4. Provide the names of the HTML to be displayed as the homepage and the HTML file that will be displayed in case an error occurs on your site.

Optionally, provide redirection rules if you want to route requests conditionally according to specific object key names, prefixes in the request, or response codes to some other object in the same bucket or external URL.

Static website hosting

Endpoint : <http://staticwebsitehostingexample.s3-website-eu-west-1.amazonaws.com>

Use this bucket to host a website [Learn more](#)

Index document [i](#)

Error document [i](#)

Redirection rules (optional) [i](#)

Redirect requests [Learn more](#)

Disable website hosting

Figure 0.50. Host a static website using AWS – first step

Now, go the Permissions section of your bucket and add the following into your Bucket Policy section:

Overview Properties Permissions Management

Access Control List Bucket Policy CORS configuration

Bucket policy editor ARN: [arn:aws:s3:::staticwebsitehostingexample](#)

Type to add a new policy or edit an existing policy in the text area below.

```

1  {
2  "Version": "2018-07-15",
3  "Statement": [
4  {
5  "Sid": "PublicReadForGetBucketObjects",
6  "Effect": "Allow",
7  "Principal": "*",
8  "Action": ["s3:GetObject"],
9  "Resource": ["arn:aws:s3:::your-bucket-name/*"]
10 }
11 ]
12 }

```

Figure 0.51. Host a static website using AWS – second step

Replace *your-bucket-name* with the name of your bucket.

To enable your S3 static website to respond to requests like GET and POST coming from an external application hosted on a certain domain, you would need to configure CORS in your bucket settings. To do this, add the following into the CORS configuration section of Permissions:

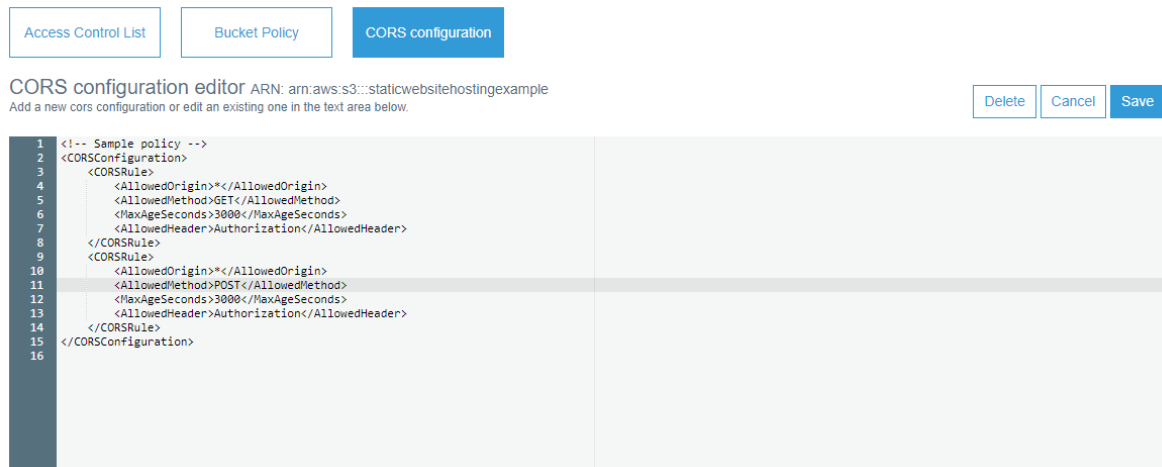


Figure 0.52. Host a static website using AWS – second step

Upload your code. For this tutorial, create two simple HTML files by the name of index.html and error.html and upload them to the bucket.

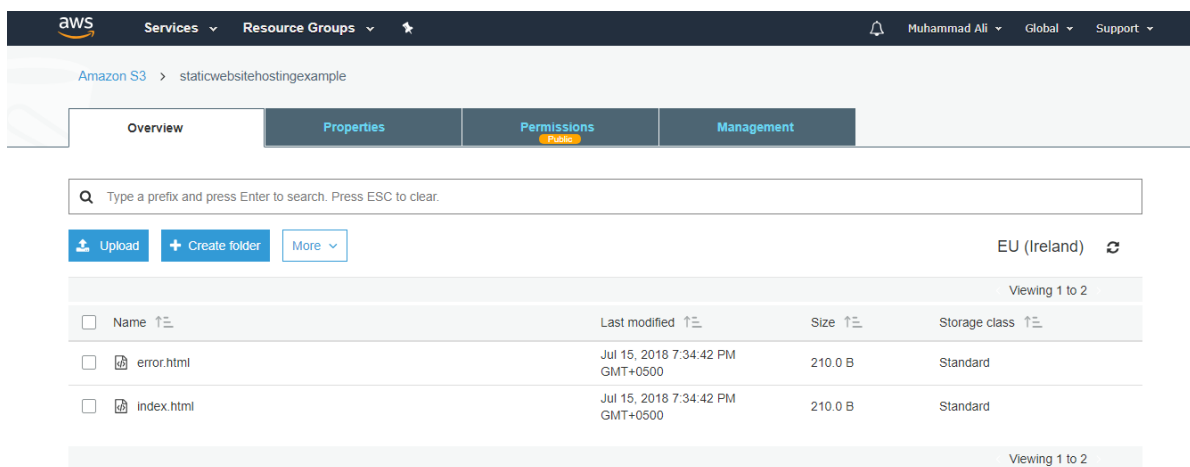


Figure 0.53. Host a static website using AWS – third step

To launch and test the site, the endpoint can be retrieved from Properties > Static website hosting.

Enrich Your Website by Adding Dynamic Behaviour

You can use a combination of HTML5 and CSS3 to graphically enrich your website. You can also use jQuery Ajax to call an API (microservice) and dynamically fetch data from a data source and display it on your website. Similarly, by invoking API endpoints using Ajax, you can store any kind of user's data back to your data source, like any other web application. If your requirement is to use AWS only for all your development needs, you can use a combination of API Gateway and Lambda to build APIs, a tutorial for which can be found here.

CORS Settings in API Gateway Endpoints

It is important to note that when developing APIs (microservices) using an API Gateway and Lambda, make sure to do the following:

Enable CORS in the API gateway at the time of creating a new resource.

New Child Resource

Use this page to create a new child resource for your resource.

Configure as proxy resource ⓘ

Resource Name*

Resource Path*

You can add path parameters using brackets. For example, the resource path **{username}** represents a path parameter called 'username'. Configuring **/{proxy+}** as a proxy resource catches all requests to its sub-resources. For example, it works for a GET request to **/foo**. To handle requests to **/**, add a new ANY method on the **/** resource.

Enable API Gateway CORS ⓘ

* Required

[Cancel](#) [Create Resource](#)

Figure 0.54. Host a static website using AWS – fourth step

When writing the lambda function (which you will integrate with the API Gateway endpoint to provide functionality to your microservice), make sure to add an additional parameter into the response header by the name of **Access-Control-Allow-Origin** with the value ******