



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Module:

ACCA7027 Cloud Engineering Principles and Practice

MSc Cybersecurity and Digital Forensics

University of Wales Trinity Saint David

AKASH NATUBHAI LUHARIYA

STUDENT ID: 2535266

Submission: 18th August 2026

Table of Contents

- 1. Introduction**
- 2. Virtualisation and Its Role in Cloud Computing**
 - 2.1 Principles of Virtualisation
 - 2.2 Advantages of Virtualisation in Cloud Computing
 - 2.3 Critical evaluation
- 3. Containerisation and Docker**
 - 3.1 Principles of Containerisation
 - 3.2 Docker as a Containerisation Platform
 - 3.3 Advantages of Containerisation in Cloud Computing
 - 3.4 Dockerfile Creation
 - 3.5 Creating and Executing a Docker Image
 - 3.6 Publishing Docker Images to Docker Hub
 - 3.7 Container Volume
 - 3.8 Networking for Containers
 - 3.9 End to End Docker Walk Through
 - 3.10 Criticism of Docker and Containerization
- 4. Load Balancing, Scalability and Application Gateways**
 - 4.1 Load Balancing
 - 4.2 Scalability
 - 4.3 Application Gateways
 - 4.4 Cloud Provider Strategy: Azure
 - 4.5 Load Balancing and Scaling Relationship
 - 4.6 Combined Benefits for the Client
- 5. Driving Factors in Cloud Computing and Serverless Computing**
 - 5.1 Driving Factors in the Development of Cloud Computing
 - 5.1.1 Cost Efficiency
 - 5.1.2 Flexibility and Agility
 - 5.1.3 Scalability and Elasticity
 - 5.2 Serverless Computing and Function as a Service
 - 5.3 Comparison between Serverless and Containers
 - 5.4 Criticism Serverless Computing
- 6. Practical Demonstration and Video Evidence**
- 7. Critical Discussion**
- 8. Conclusion**

References

1. Introduction

Cloud computing has brought about changes in the ways organisations obtain and use computing resources. Cloud computing allows users to have access to computing resources such as servers, storage, networks, and applications through on-demand provisioning of computing resources based on varying demands [10], [20].

There are various technologies that support today's cloud computing environments. The technology of virtualisation enables partitioning of physical hardware into several virtual machines that allow sharing of resources. Containerisation is another form of abstraction that packages applications and their dependences to provide portable environments. Docker is an example of containerisation software that facilitates application development and scalability [5], [6], [8].

The cloud applications should also react to varying demands appropriately. Load balancing allows distribution of requests among resources while scalability allows adding resources when there is more workload. Application gateways can allow for intelligent routing and securing of traffic at application layer. Serverless computing using platforms like AWS Lambda and Azure Functions have allowed execution of event-based code without having to deal with servers [11], [14], [15].

The following report provides a critical analysis of virtualization, Docker containerization, images, volumes, and networking; load balancing, scalability, and application gateways; and drivers behind the use of cloud services. It will also provide a comparison between serverless computing and containerization. The practical activities using Docker and videos will show real-life implementations of these concepts. [4]

2. Virtualisation and Its Role in Cloud Computing

2.1 Principles of Virtualisation

Virtualisation is a basic technology used in cloud infrastructure due to its ability to abstract physical computing resources and create multiple isolated virtual environments on them. The difference between virtualisation and deployment of applications is that while without virtualisation each physical server is allocated to one application, with the help of virtualisation, computing resources of one physical machine can be distributed among several virtual machines (VMs). Virtual machines are managed by a hypervisor which allocates physical resources for each of them such as computing power, amount of memory, amount of storage and interfaces with networks. In cloud environment, abstraction allows infrastructure providers to distribute computing resources dynamically according to user needs [1], [2], [4].

The principle of resource abstraction can be exemplified by a physical server which has a multi-core processor, 64 GB of RAM and considerable amount of storage. Without virtualisation, an organisation installs one OS and application on this server. If application uses only part of the resources, then there will be unused capacity. Using virtualisation, one physical server can run several VMs which have their own shares of CPU, memory, and storage resources. For example, VM1 can contain a web server, VM2 an application server and VM3 a developer's environment. All workloads will be isolated but will use the same hardware resources.

Figure 1: Virtualisation of shared physical hardware




Figure 1. Several virtual machines running on the resources of one physical server with the help of hypervisor.

The resource sharing ability is especially useful in cloud computing because IaaS providers build large data centres which consist of considerable amounts of physical resources. According to research by Shi and Lin, virtualisation allows IaaS providers to pool physical resources and present them to users in form of virtual machines [4].

2.2 Advantages of Virtualisation in Cloud Computing

The first advantage is high efficiency of resource usage. One physical server may support several loads and there is no need in separate hardware. Good placement of VMs is important since poor allocation will lead to wasted resources, increased power consumption, and higher costs [4], [11]. Thus, cloud vendors apply resource management techniques to know where to place VMs and to migrate workloads between hosts.

Another advantage is isolation. Every VM works in a logically isolated environment and applications running on one VM have no access to processes and filesystem of other VM. It makes VMs suitable for hosting several customers or applications in shared infrastructure. Isolation is not full, however. Research shows that performance interference occurs when VMs share common physical hardware such as CPU, memory, or storage [1], [2]. Therefore, cloud vendors must carefully manage resources to provide good service.

Flexible and scalable environment is the third advantage. New VM can be created using existing image or template without need of buying new server and its installation. VMs can also be migrated between physical servers in case of need of rebalancing. The research on VM migration shows its importance for dynamic allocation of resources, compliance with service-level requirements and lower operating costs [3]. These features correspond to the cloud computing requirements in terms of fast allocation of resources.

2.3 Critical evaluation

Although virtualisation brings many advantages, it still causes some challenges related to management and performance of system. The software itself consumes some resources and several VMs competing for common hardware may demonstrate performance variations [1], [2], [3]. In addition, inefficient allocation of VMs leads to high power consumption and low economy of consolidation [1], [2], [3]. The problem of security also should be managed since vulnerabilities or errors of configuration at the level of hypervisor or host may affect several virtual loads.

Still, the advantages of virtualisation make it very suitable for cloud infrastructure. Consolidation of loads, isolation of customers, dynamic allocation of resources and VM migration provide technical base for scalable IaaS platforms. Most importantly, virtualisation should not be confused with cloud computing. The former is the technology of creating abstract computing environment while cloud computing is combination of virtualisation and other technologies like automation, networking, storage, and service management.

Virtualization can be regarded as one of the most essential technologies that enable cloud computing, but not necessarily the only definition of cloud computing. The most valuable role it plays is converting a static physical environment to a dynamic resource pool of logical resources.

3. Containerisation and Docker

3.1 Principles of Containerisation

Containerisation is a form of operating-system-level virtualisation whereby the application and its dependencies are bundled in an isolated execution environment known as a container. Unlike a traditional virtual machine, a container does not need a complete guest operating system in most cases. Containers share the operating-system kernel but are isolated from processes, file system and network of other applications running on the host OS. Containerisation is therefore relatively light-weight and facilitates efficient creation, start and replication of application workloads.

Differences between virtual machines and containers are significant when it comes to cloud computing. While a virtual machine has a full guest operating system, a container includes the application, runtime, libraries, and configuration necessary for the application. Therefore, it is possible to have many containers on the same host because there will be reduced operating system overhead. Recent literature suggests that containerisation is a key technology for cloud-native environments due to portability, efficiency, scalability, and automated deployment [5], [6], [8], [18].

Containerisation is mostly used with microservices architectures. The idea is to develop separate independent services instead of one large application. For instance, an e-commerce application may include services for authentication, product management, payments, and order processing. Each service is packaged in its own container and can be deployed and scaled separately. Literature analysing Docker-based microservice architectures reveals that containers provide isolation and flexibility while allowing to deploy services independently [5], [7], [19].

The basic structure can be illustrated as follows:

- Physical/Cloud Infrastructure -> Host Operating System -> Container Runtime -> Containers -> Applications

Unlike virtualised infrastructure, which is structured in a following way:

- Physical Infrastructure -> Hypervisor -> Virtual Machines -> Guest Operating Systems -> Applications

This is one of the reasons why containers are so popular in modern cloud environments. Nevertheless, containers cannot be regarded as smaller virtual machines as they have unique isolation and security properties [1], [2], [18].

3.2 Docker as a Containerisation Platform

Docker offers an environment for creating, running, and deploying containerised applications. The key elements include Docker files, images, containers, registries, volumes, and networks.

The Docker image is a packaged and immutable template for the creation of containers. It includes file system layers and other configurations necessary to run an application. Container is an active instance of an image. A Dockerfile includes all the steps that will be used to build an image. The instructions in Docker documentation include FROM, COPY, RUN, CMD, ENTRYPOINT, ENV, EXPOSE, USER, and WORKDIR [5], [6].

In such case, the following relationship can be established:

Dockerfile → Build → Image → Run → Container

This distinction between image and container is crucial. The image can be built only once but later used to produce multiple similar containers. In this way, reproducibility is ensured.

3.3 Advantages of Containerisation in Cloud Computing

There are several benefits offered by containerisation that are specifically related to cloud computing.


Portability

The application with all its dependencies is encapsulated in the image, and thus, differences between development, testing, and production environments decrease. It is possible to transfer a container image between compatible environments without manually recreating application dependencies.

3.4 Dockerfile Creation

A practical example can be created using an HTML webpage hosted on NGINX. The project comprises an index.html file and a Dockerfile.

for example, the HTML application could be:

A screenshot of a code editor window with multiple tabs. The active tab is 'Dockerfile' and it contains the following text:


```
FROM nginx:alpine
LABEL maintainer="akash"
COPY ./site/ /usr/share/nginx/html/
EXPOSE 80
CMD ["nginx", "-g", "daemon off;"]
```

The editor interface includes a menu bar (File, Edit, View), a toolbar with various icons, and a status bar at the bottom showing 'Ln 9, Col 35 | 127 characters', 'Plain text', '100%', 'Windows (CRLF)', and 'UTF-8'. The Windows taskbar is visible at the very bottom of the image.

The FROM keyword indicates the base image. Here, the NGINX Alpine image is used to provide a web-server environment. The COPY keyword will copy the student's HTML file to the directory that is used by NGINX. The EXPOSE keyword specifies the port number that is used by the application. Notably, according to Docker documentation, EXPOSE does not expose the port but rather publishes it during runtime using keywords such as -p [5], [6].

This example shows the usefulness of the Infrastructure as Code methodology since the web server is installed automatically and the file is copied without manual intervention.

3.5 Creating and Executing a Docker Image



```

Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

PS C:\Users\Aastha> cd "C:\Users\akash\Cloud Engineering Assignment"
PS C:\Users\akash\Cloud Engineering Assignment> docker build -t cloud-assignment:1.0 .
[+] Building 1.2s (7/7) FINISHED
=> [internal] load build definition from Dockerfile                                docker:desktop-linux
=> => transferring dockerfile: 173B                                              0.0s
=> [internal] load metadata for docker.io/library/nginx:alpine                  0.0s
=> [internal] load .dockerignore                                                 0.6s
=> => transferring context: 2B                                                  0.0s
=> [internal] load build context                                                 0.0s
=> => transferring context: 61B                                                 0.0s
=> [1/2] FROM docker.io/library/nginx:alpine@sha256:4a73073bd557c65b759505da037898b61f1be6cbcc3c2c3aeac22d2a470c 0.1s
=> => resolve docker.io/library/nginx:alpine@sha256:4a73073bd557c65b759505da037898b61f1be6cbcc3c2c3aeac22d2a470c 0.1s
=> CACHED [2/2] COPY ./site/ /usr/share/nginx/html/                            0.0s
=> exporting to image                                                            0.2s
=> => exporting layers                                                            0.0s
=> => exporting manifest sha256:aaa95d2d9080e7c5a02dbd9d21dee7d2afaf4cbb813c92a34bce5b34bffb559 0.0s
=> => exporting config sha256:d5168d1420674c1ee104a814c401f8b87882f96f19f9b8d809e2bd339a487e97 0.0s
=> => exporting attestation manifest sha256:5c2b59f7581be08c718a5eb1dd3bd5d5fcd9ad2cd20684399e211ee596525b2b 0.1s
=> => exporting manifest list sha256:bd7abbe34cbc988c09a196c4d800b2c039e3bbe42728f783c30dd41aa91dcb45 0.0s
=> => naming to docker.io/library/cloud-assignment:1.0                         0.0s
=> => unpacking to docker.io/library/cloud-assignment:1.0                     0.0s

View build details: docker-desktop://dashboard/build/desktop-linux/desktop-linux/7diq2f1f90ylue0ppkc5mbyig
PS C:\Users\akash\Cloud Engineering Assignment>

```

This gives practical proof that the Dockerfile has built a good image and that the image can be executed as a working web application.

The container status can be checked with:

- docker PS

The container log can be viewed with:

- docker logs cloud-web

Next, the lifecycle management can be conducted as:

- docker stop cloud-web.
- docker rm cloud-web

Here we see how different the image and container are: while the container is shut down, the image remains as a reusable artefact.

3.6 Publishing Docker Images to Docker Hub

It is particularly important when using containers if it is possible to transfer the images between different environments. Docker Hub is a service which provides the ability to store, manage and share images [5], [6].

Initially, the developer logs in to the service:

Now, Docker documentation suggests that labelling and pushing are the ways to add the local image to a Docker Hub repository [5], [6]. The image will be downloaded later into another

environment:

```

>> => exporting config sha256:d5168d1420674clee104a814c401f8b87882f96f19f9b8d809e2bd339a487e97      0.0s
>> => exporting attestation manifest sha256:5c2b59f7581be08c718a5eb1dd3bd5d5fcd9ad2cd20684399e211ee596525b2b 0.1s
>> => exporting manifest list sha256:bd7abbe34cbc988c09a196c4d800b2c039e3bbe42728f783c30dd41aa91dcb45 0.0s
>> => naming to docker.io/library/cloud-assignment:1.0      0.0s
>> => unpacking to docker.io/library/cloud-assignment:1.0    0.0s

View build details: docker-desktop://dashboard/build/desktop-linux/desktop-linux/7diq2f1f90ylue0ppkc5mbyig
PS C:\Users\akash\Cloud Engineering Assignment> docker login
Authenticating with existing credentials... [Username: akash1662]

Info → To login with a different account, run 'docker logout' followed by 'docker login'

Login Succeeded
PS C:\Users\akash\Cloud Engineering Assignment> docker tag cloud-assignment:1.0 akash1662/cloud-assignment:1.0
PS C:\Users\akash\Cloud Engineering Assignment> docker push akash1662/cloud-assignment:1.0
invalid reference format
PS C:\Users\akash\Cloud Engineering Assignment> docker push akash1662/cloud-assignment:1.0
invalid reference format
PS C:\Users\akash\Cloud Engineering Assignment> docker push akash1662/cloud-assignment:1.0
The push refers to repository [docker.io/akash1662/cloud-assignment]
c6a14fcffc1e: Pushed
e8720549761f: Pushed
d0008c891db4: Pushed
390dc935348d: Pushed
3cd534fe98c6: Pushed
46519e7231d2: Pushed
62bec68d7c31: Pushed
1223f016b4e4: Pushed
55afalecc21d: Pushed
46f977ee452f: Pushed
1.0: digest: sha256:bd7abbe34cbc988c09a196c4d800b2c039e3bbe42728f783c30dd41aa91dcb45 size: 856
PS C:\Users\akash\Cloud Engineering Assignment>
  
```


In practice, it means that the image is detached from the developer's computer and is ready for usage by another developer or cloud service provider. In case of CI/CD, for example, the image could be automatically built, evaluated, and moved to production environment.

Nevertheless, when it comes to publishing images, one needs to pay attention to security issues. Companies should not put passwords, API keys, or any other secrets inside their images. Besides, the images should be scanned for any known vulnerabilities. The container image is not necessarily secure just because it is isolated [17].

3.7 Container Volume

The container is usually considered replaceable or ephemeral, which becomes a challenge when dealing with the situation where the application produces some information that must outlive the life cycle of the container. Solve this issue, Docker introduced volumes that provide persistent storage that is independent of the specific container.

Volume can be created using:



```

e8720549761f: Pushed
d0008c891db4: Pushed
390dc935348d: Pushed
3cd534fe98c6: Pushed
46519e7231d2: Pushed
62bec68d7c31: Pushed
1223f016b4e4: Pushed
55afalecc21d: Pushed
46f977ee452f: Pushed
1.0: digest: sha256:bd7abbe34cbc988c09a196c4d800b2c039e3bbe42728f783c30dd41aa91dcb45 size: 856
PS C:\Users\akash\Cloud Engineering Assignment> docker volume create cloud-data
cloud-data
PS C:\Users\akash\Cloud Engineering Assignment> docker ls
docker: unknown command: docker ls

Run 'docker --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker volume ls
DRIVER      VOLUME NAME
local       cloud-data
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name-volume -p 8081:80 -v cloud-data:/usr/share/html cloud-assignment
:1.0

What's next:
  Debug this container error with Gordon → docker ai "help me fix this container error"
unknown flag: --name-volume

Usage: docker run [OPTIONS] IMAGE [COMMAND] [ARG...]

Run 'docker run --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name cloud-assignment-container -p 8081:80 -v cloud-data:/usr/share/html cloud-assignment:1.0
482b4df75e9e9bf89e7d10fde5a64a40acde6073a866248e1a3907ca2601b0d6
PS C:\Users\akash\Cloud Engineering Assignment>

```

The Docker documentation differentiates between the volumes and the writable layer of the container and provides tools to mount the volumes in containers [6], [7].

For example, let us take an instance of the database container:

Database Container | v /var/lib/data | v Docker Volume | v persistent data

If the container that houses the database is deleted and re-created, the volume can then be attached to this new container. Application and data lifecycles are thus managed independently.

But a cloud system cannot consist only of creating volumes. An organisation needs to think about backups, replication, availability, access control, and recovery. A Docker volume running on a single host cannot be considered available [4], [10].

3.8 Networking for Containers

Networking facilitates communication between containers and external systems. The networking options in Docker include bridge networks and other drivers based on the deployment environment [10].

User-defined networks can be created with the help of:

- docker network creates cloud-network.

The network can be inspected with:

- docker network ls
- and:
- docker network inspect cloud-network.

A container can be attached to the network during creation:

```

local      cloud-data
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name-volume -p 8081:80 -v cloud-data:/usr/share/html cloud-assignment:1.0

What's next:
  Debug this container error with Gordon → docker ai "help me fix this container error"
unknown flag: --name-volume

Usage: docker run [OPTIONS] IMAGE [COMMAND] [ARG...]

Run 'docker run --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name cloud-assignment-container -p 8081:80 -v cloud-data:/usr/share/html cloud-assignment:1.0
482b4df75e9e9bf89e7d10fde5a64a40acde6073a866248e1a3907ca2601b0d6
PS C:\Users\akash\Cloud Engineering Assignment> docker network create cloud-network
c4d6a63e2f0f769dd9a29bb100652b4f6a8daa465fb059fd21b5bf79882e440c
PS C:\Users\akash\Cloud Engineering Assignment> docker network run -d --name web-network --network cloud-network -p 8082:80 cloud-assignment:1.0
unknown shorthand flag: 'd' in -d

Usage: docker network

Run 'docker network --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name web-network --network cloud-network -p 8082:80 cloud-assignment:1.0
b002090183bb3c4ff32e2b849910d4ea9d73e7f02a91244ebe49b9fd33b4e5cf
PS C:\Users\akash\Cloud Engineering Assignment> docker network ls
NETWORK ID        NAME                DRIVER              SCOPE
2b2d25f727f4      bridge              bridge              local
c4d6a63e2f0f      cloud-network        bridge              local
6aa29206010b      host                host                local
0636923c0264      none                null                local
PS C:\Users\akash\Cloud Engineering Assignment>

```


The documentation from Docker says that “services running on containers attached to a user-defined network may communicate with each other without needing to publish all of the services’ ports” [10].

The web service can communicate with the API, while the API communicates with the database. Only services that need to be externally accessible need to publish appropriate ports.

This separation improves architectural organisation and can reduce unnecessary network exposure [10], [17].

3.9 End to End Docker Walk Through

The significance of each exercise individually increases in importance when integrated into a single process.



```

:1.0

What's next:
  Debug this container error with Gordon → docker ai "help me fix this container error"
unknown flag: --name-volume

Usage: docker run [OPTIONS] IMAGE [COMMAND] [ARG...]

Run 'docker run --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name cloud-assignment-container -p 8081:80 -v cloud-data:/usr/share/html cloud-assignment:1.0
482b4df75e9e9bf89e7d10fde5a64a40acde6073a866248e1a3907ca2601b0d6
PS C:\Users\akash\Cloud Engineering Assignment> docker network create cloud-network
c4d6a63e2f0f769dd9a29bb100652b4f6a8daa465fb059fd21b5bf79882e440c
PS C:\Users\akash\Cloud Engineering Assignment> docker network run -d --name web-network --network cloud-network -p 8082:80 cloud-assignment:1.0
unknown shorthand flag: 'd' in -d

Usage: docker network

Run 'docker network --help' for more information
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d --name web-network --network cloud-network -p 8082:80 cloud-assignment:1.0
b002090183bb3c4ff32e2b849910d4ea9d73e7f02a91244eba49b9fd33b4e5cf
PS C:\Users\akash\Cloud Engineering Assignment> docker network ls
NETWORK ID        NAME                DRIVER              SCOPE
2b2d25f727f4      bridge              bridge              local
c4d6a63e2f0f      cloud-network        bridge              local
6aa29206010b      host                host                local
0636923c0264      none                null                local
PS C:\Users\akash\Cloud Engineering Assignment> docker run -d -p 8080:80 akash1662/cloud-assignment:1.0
814c383c994a11dc5d84bceb19bf134408a0a69e18545e51871ccb0c6de4ed3c
PS C:\Users\akash\Cloud Engineering Assignment>
  
```

Cloud Engineering Assignment

localhost:8080

Set Google Chrome as your default browser and pin it to your taskbar

Set as default

Hello from Docker!



This workflow explains the relevance of containerization to cloud computing. The application is developed only once, packaged along with all dependencies in an image, and thereafter can be made available using the registry. The cloud computing infrastructure can leverage this image to instantiate multiple application instances. Further application instances can be created based on increased demand; hence containerization becomes compatible with cloud load balancing and scalability practices [11], [12], [18].

The above workflow shows separation of concerns where the Dockerfile is used to build the application environment, images are the deployable items, containers are live instances,

Docker Hub helps in the distribution process, volumes ensure persistence, while networking allows service interaction.

In cloud-native systems, such separation helps automate the entire process in a repeatable manner. From research on Docker and cloud native systems, it can be inferred that the usage of containerization is increasing in scalable distributed applications and microservices architectures [5], [7], [11], [19]. But the value of containerization gets magnified when containers are part of other activities like orchestration, monitoring, security, and automation.

3.10 Criticism of Docker and Containerization

While there are many benefits to using Docker, there are some considerations about the technology before adopting containerization.

Containerization is a step forward when compared to the legacy server-based approach in general. While virtualization is all about making hardware abstractions, containerization makes application abstractions. This idea has been put to practice using Docker images, Dockerfiles, Docker registries, volumes, and networking. Combining them with native cloud services for scalability, load balancing, and orchestration enables the development of more portable, repeatable, and adaptive applications [1], [5], [8].

4. Load Balancing, Scalability and Application Gateways

4.1 Load Balancing

Load balancing is the task of routing incoming application requests across different computing resources including virtual machines, containers, or application instances. The main idea is to avoid overloading a particular resource when others are idle. In cloud computing, good load balancing will increase resource utilization, response time, throughput, and availability [12], [13] of an application. This especially true for web applications since the number of users can be varied over time.

For example, an online application of some university can start with using two web servers. If there are thousands of students accessing the website simultaneously before the deadline for some assignment, then all their requests to a single server will lead to higher CPU usage, higher response time and may cause service interruption. The load balancer will spread these requests between the servers.




Figure 8. Incoming client requests distributed across multiple application instances by a cloud load balancer.

4.2 Scalability

Scalability is the ability of the application to manage changes in workload by expanding or reducing the computing resources available. There are two basic techniques of scaling – vertical scaling and horizontal scaling.

Vertical scaling involves expanding the resources of one node: for instance, expanding virtual machine memory from 4 GB to 16 GB. The approach is simple but limited by the upper capacity of a particular node and dependent on the type of platform.

Horizontal scaling means increasing the amount of application instances:

➤ 2 instances → 4 instances → 8 instances

The technique is especially suitable for cloud-native applications because they could start and stop instances based on workload. For example, AWS EC2 Auto Scaling allows one to expand or reduce the number of instances based on the needs and integrate them into Elastic Load Balancing [11], [12].

Nevertheless, scalability is not just an addition of the number of resources unlimitedly. More instances mean more infrastructure costs. Besides, the application should have the ability to work properly in case of adding more instances, especially when it comes to statefulness, session and persistent data [11], [18].

4.3 Application Gateway

The application gateway offers a way to manage and route web traffic via application-level. An application gateway can look at HTTP data and making routing decisions on URL paths and host headers whereas a simple network load balancer relies on IP address and port numbers for routing [10], [12].

Microsoft Azure Application Gateway works on OSI Level 7 and features application-aware routing, SSL/TLS termination, autoscaling, zone redundancy and integrates with Web Application Firewall functionalities [4]. E.g., requests for:

- www.example.com/images

can be routed to image processing service, while for:

- www.example.com/orders

to an order management service.

4.4 Cloud Provider Strategy: Azure

Several load balancing strategies are offered by Azure, which is suitable for various needs. Application Gateway strategy is suitable for HTTP/HTTPS applications due to a combination of Layer 7 routing with load balancing and other features [4].

The capacity of the Azure Application Gateway v2 is scalable; it increases or decreases depending on traffic changes [5]. It is also capable of covering availability zones, providing increased fault tolerance [5]. Azure represents a good example of how load balancing and scaling can work together.

Architecture of an online store application:

Figure 9: Azure application architecture combining gateway, load balancing, and scaling




Figure 9. Application Gateway distributing customer traffic across scalable application instances.

4.5 Load Balancing and Scaling Relationship

Load balancing and scaling deal with different but related issues. Whereas scaling provides or takes away application capacity, load balancing helps distribute traffic among this capacity.

In case an application consists of two servers and the workload grows, autoscaling adds two more servers. As a result, a load balancer will be able to spread the traffic between four servers. In the opposite situation, the reduction in workload allows eliminating excess instances along with automatic adjustment of traffic distribution.

This connection is illustrated by AWS guides, which state that instances launched by EC2 Auto Scaling may be automatically added to Elastic Load Balancing and removed when they are terminated [11], [12]. At the same time, the scientific literature mentions load balancing as a critical element for maintaining performance and resource utilization in cloud computing environment [12], [13].

Thus, scaling without load balancing may result in underutilization of added resources, and load balancing without scaling cannot generate extra capacity if all existing resources are busy.

4.6 Combined Benefits for the Client

In combination with application gateways, load balancing and scaling become more reliable solutions [11], [12]. An application gateway offers intelligent application routing as well as

certain security features. The load balancer distributes requests among healthy application instances, and scaling regulates the number of application instances depending on demand.

For a client, this could offer:

- Increased availability: Failed application instances will not get any requests.
- Increased performance: Requests will be distributed rather than concentrated on one application instance.
- Scalability: Additional resources may be engaged during the peak period of load.
- Cost savings: Unnecessary application instances may be eliminated due to low demand.
- Increased security: Application gateway can be equipped with web application firewall features.
- Greater flexibility: Various application instances could be routed based on specific HTTP request characteristics.

Thus, the combination becomes a full-scale solution for traffic management. The gateway decides on how traffic should be managed; the load balancer defines where the traffic should be distributed; the scaling controls how many application instances should be available. The proposed combination is especially relevant to cloud applications with volatile demand.

5. Driving Factors in Cloud Computing and Serverless Computing

5.1 Driving Factors in the Development of Cloud Computing

Cloud computing development was triggered by such economic, technological, and organizational factors as the need for higher flexibility and efficiency. The existing IT infrastructure implies that organizations should buy servers, networks, and storage to install applications. It leads to high expenditures and the necessity to forecast the demands properly. With cloud computing, it is possible to get computing capacity as a service, hence avoiding the necessity of buying lots of permanent infrastructure [20].

5.1.1 Cost Efficiency

Costs are among the driving factors for implementing cloud computing technology. With traditional data centre infrastructure, it is necessary to buy enough infrastructure to cope with peak loads even though maximum load is needed infrequently. Hence, hardware can be idle for a considerable period [20]. With cloud computing, it becomes possible to manage resources according to needs.

5.1.2 Flexibility and Agility

Thanks to cloud platforms, it is possible to quickly provision computing resources without the need to acquire new hardware. Businesses can build development environments, deploy software and experiment with new services without buying extra physical servers. It

contributes to organizational flexibility and helps enterprises to adjust to changes in customer demands.

The other benefit of cloud computing is that it enables experimentation. Computing resources can be provisioned temporarily for purposes of development or testing and then terminated. Thus, new projects do not require a lot of investment from organizations, which helps them to respond to changing business needs [20].

5.1.3 Scalability and Elasticity

Scalability is another driver of cloud adoption. The capacity of a cloud application can vary significantly between regular and peak times. For example, the workload of an online store can increase during a major sales event. In response to such a situation, cloud infrastructure increases the capacity of an application during peak hours and decreases it afterwards [11], [12].

The use of cloud computing provides better resource utilization as businesses do not necessarily have to maintain their maximum capacity. Moreover, customer experience is also improved through this approach [11], [20].


5.2 Serverless Computing and Function as a Service

Serverless computing is an abstraction in cloud computing. Serverless computing is not about the absence of the servers; rather, the underlying infrastructure is managed by the cloud provider, while the developer focuses on the application code. Function as a Service (FaaS) is one of the most popular serverless architectures where applications are broken down into functions that run based on events [14], [15].

AWS Lambda is an example of FaaS. AWS defines Lambda as compute service that allows running code without provisioning or managing servers and features automatic scaling and pay-per-execution billing [14], [15]. Triggers for execution include HTTP requests, file uploads, and messages.

Azure Functions is another serverless solution. According to Microsoft, Azure Functions is an event-driven platform for building serverless solutions that allow file processing, scheduled executions, creating web APIs, handling database events, and messaging [14], [15]. Its Flex Consumption hosting option supports event-driven scalability and pay-as-you-go pricing [5].

A typical serverless workflow can therefore be represented as:



For instance, when a customer uploads an image, the storage service may trigger a serverless function to process the image and save the output. There is no need to maintain the server ready to always receive uploads.

5.3 Comparison between Serverless and Containers

Factor	Containerisation	Serverless/FaaS
Deployment unit	Container image	Function
Infrastructure management	Customer/platform responsibility	Mainly cloud provider
Runtime control	High	Lower
Portability	High	Lower
Scaling	Platform/orchestrator dependent	Automatic
Long-running workloads	Well suited	More workload-dependent
Event-driven workloads	Suitable	Highly suitable
Operational overhead	Higher	Lower

Containerization is an option when developers need full control over the application runtime. Docker allows developers to package the applications and their dependencies in an image that

can then be deployed across all similar environments. Thus, containers are good for long running services and applications requiring specific runtime environment settings.

Serverless is a more abstract solution. Developer deploys functions, cloud provider deploys the execution infrastructure and manages capacity issues. That may lead to reduced operational burden and increased speed of software development. Systematic review of serverless computing indicates that reduced infrastructural management is among key features of the approach, along with performance, testing, debugging, portability, and deployment issues [14], [15], [16].

Thus, the compromise is between control and convenience. Containers offer greater control and portability, serverless – reduced responsibility for infrastructural management.

5.4 Criticism of Serverless Computing

Serverless computing offers numerous benefits but creates certain limitations. One limitation is cold start latency, when an idle function may require an environment for execution to be started before a request is processed [16]. It is a problem that affects applications needing consistently fast performance.

Vendor lock-in is another challenge. Serverless applications rely on provider-specific storage, message queues, authentication, and monitoring services. Multi-cloud serverless computing research states vendor lock-in as one of the critical problems since the important level of abstraction makes switching providers harder [14], [15].

Serverless still requires security measures to be applied. While the infrastructure management is delegated to the provider, developers must care about application permissions, dependencies, secrets, and configuration. According to the research, serverless applications present their own security issues because applications include multiple functions communicating with cloud services [15], [17].

Therefore, serverless computing cannot be considered as a universal alternative for containers. Containers may prove to be better suited when runtime control, portability, and long execution time are needed. Serverless computing is well suited for event-based, sporadic, and fluctuating workloads [14], [15], [16].

In general, development of cloud computing technologies is based on increasing flexibility, elasticity, automation, and abstraction. Virtualization abstracts physical hardware, containers.

6. Practical Demonstration and Video Evidence

The practical demonstration offers the evidence of the entire process of development of a Docker application. In the first place, the Dockerfile is used for creation of an image of a simple web application. The image is constructed with the use of “docker build” command and run as a container with port binding that makes it possible to access the HTML webpage from the browser [5], [6]. The next step includes the creation of Docker volume to show how data persistence works and user-defined network to show container-to-container communication [7], [10]. Finally, the image is tagged and uploaded to the Docker Hub, which means that one can distribute his/her container image on another machine [5], [6].

Video: <https://youtu.be/MuvthF70Ss0?si=KVj1zAs-tBoMPaJW>

7. Critical Discussion

The technologies reviewed above show that cloud computing brings numerous advantages but at the same time involves compromises. While containerization enhances deployment consistency, portability, and resource utilization, containers still require efficient image management, network management, storage, and security management. According to research, containerized architecture allows enhancing data-centre efficiency by avoiding over-allocation in comparison with the traditional applications deployment [1], [2], [8]. Nevertheless, containerization does not eliminate infrastructure complexity as orchestration, monitoring and security become crucial aspects of the operation.

In an equivalent way, load balancing and scalability are not separate but complementary properties. Load balancing distributes tasks among resources while scalability offers extra resources in case of demand increase. The research on Docker-based cloud environments shows that proper scheduling and load balancing increase workload distribution in containerized resources [7], [11], [12]. Thus, only an increase in application instance numbers does not ensure better performance without effective workload distribution [11], [12], [13].

The higher level of abstraction offered by serverless computing simplifies infrastructure management. However, this advantage comes together with certain compromises related to security, portability, and control. Marin *et al.* emphasize particular security issues arising in serverless architecture which include complex function interaction and access control [15], [17]. Thus, serverless technology is not always better than containers.

In conclusion, there is no universal best cloud architecture. Container is recommended when the runtime control and portability are significant factors whereas serverless computing architecture provides value where workloads are event driven and vary [14], [15]. Hence, the combination of both may offer a good compromise between control, scalability, flexibility, and efficiency. The right technology should be chosen based on workload and other criteria like cost and performance.

8. Conclusion

As shown throughout this report, cloud computing represents a holistic system whereby virtualisation, containerisation, networking, scalability and serverless technologies meet various infrastructure and application needs. While virtualisation makes resource sharing efficient, Docker containerisation allows for portable and reproducible application environments. The practical demonstration, in turn, showed how Dockerfiles, images, containers, volumes, networking and image publishing may be used in contemporary application development [1], [5], [6], [10].

In addition, it has been found that load balancing and scalability complement each other, as the former ensures proper distribution of load, while the latter increases capacity. Azure Application Gateway further expands on these abilities through the provision of application-layer routing, autoscaling and zone redundancy [11], [12]. Lastly, the research on the cloud adoption reasons proved cost efficiency, flexibility, elasticity, and minimal infrastructure

management among important motivating factors. Serverless platforms like AWS Lambda and Azure Functions abstract even more from runtime, but still containers remain a better option where needed [1], [14], [15].

In summary, there is no single optimal cloud technology. Instead, the efficient design of cloud architecture requires choosing and integrating appropriate technologies depending on the workload, performance, security, cost, and availability considerations [1], [11], [15], [17], [20]. This justifies a hybrid approach that allows for the coexistence of containers, and serverless computing technologies.

References

- [1] W. Felter, A. Ferreira, R. Rajamony, and J. Rubio, "An updated performance comparison of virtual machines and Linux containers," in *Proc. 2015 IEEE Int. Symp. Performance Analysis of Systems and Software (ISPASS)*, 2015, pp. 171–172, doi: 10.1109/ISPASS.2015.7095802.
- [2] R. K. Barik, R. K. Lenka, K. R. Rao, and D. Ghose, "Performance analysis of virtual machines and containers in cloud computing," in *Proc. 2016 Int. Conf. Computing, Communication and Automation (ICCCA)*, 2016, pp. 1204–1210, doi: 10.1109/CCAA.2016.7813925.
- [3] F. Zhang, G. Liu, X. Fu, and R. Yahyapour, "A survey on virtual machine migration: Challenges, techniques, and open issues," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 2, pp. 1206–1243, 2018, doi: 10.1109/COMST.2018.2794881.
- [4] L. Mashayekhy, M. M. Nejad, and D. Grosu, "Physical machine resource management in clouds: A mechanism design approach," *IEEE Trans. Cloud Comput.*, vol. 3, no. 3, 2015, doi: 10.1109/TCC.2014.2369419.
- [5] D. Jaramillo, D. V. Nguyen, and R. Smart, "Leveraging microservices architecture by using Docker technology," in *Proc. SoutheastCon 2016*, 2016, doi: 10.1109/SECON.2016.7506647.
- [6] S. Singh and N. Singh, "Containers & Docker: Emerging roles & future of Cloud technology," in *Proc. 2016 2nd Int. Conf. Applied and Theoretical Computing and Communication Technology (iCATccT)*, 2016, doi: 10.1109/ICATCCT.2016.7912109.
- [7] D. N. Jha, S. Garg, P. P. Jayaraman, and R. Buyya, "A holistic evaluation of Docker containers for interfering microservices," in *Proc. 2018 IEEE Int. Conf. Services Computing (SCC)*, 2018, doi: 10.1109/SCC.2018.00012.
- [8] K. Kumar and M. Kurhekar, "Economically efficient virtualization over cloud using Docker containers," in *Proc. 2016 IEEE Int. Conf. Cloud Computing in Emerging Markets (CCEM)*, 2016, pp. 95–100, doi: 10.1109/CCEM.2016.24.
- [9] B. G. Xavier, T. C. Ferreto, and L. C. Jersak, "Time provisioning evaluation of KVM, Docker and Unikernels in a Cloud Platform," in *Proc. 2016 16th IEEE/ACM Int. Symp.*

Cluster, Cloud and Grid Computing (CCGrid), 2016, pp. 277–280, doi: 10.1109/CCGrid.2016.86.

[10] R. Jain and S. Paul, “Network virtualization and software defined networking for cloud computing: A survey,” *IEEE Commun. Mag.*, vol. 51, no. 11, pp. 24–31, 2013, doi: 10.1109/MCOM.2013.6658648.

[11] A. Al-Dhuraibi, F. Paraiso, N. Djarallah, and P. Merle, “Elasticity in cloud computing: State of the art and research challenges,” *IEEE Trans. Services Comput.*, 2017, doi: 10.1109/TSC.2017.2711009.

[12] M. S. Al Reshan, D. Syed, N. Islam, A. Shaikh, M. Hamdi, M. A. Elmagzoub, G. Muhammad, and K. H. Talpur, “A fast converging and globally optimized approach for load balancing in cloud computing,” *IEEE Access*, vol. 11, pp. 11390–11404, 2023, doi: 10.1109/ACCESS.2023.3241279.

[13] “A survey of load balancing in cloud computing: Challenges and algorithms,” in *Proc. IEEE*, 2013. IEEE Xplore document no. 6472470.

Important: For [13], IEEE's indexed page confirms the publication, but the accessible search metadata did not expose the author list. I would **not invent the authors**. If your university requires full bibliographic metadata, use IEEE Xplore's “Cite This” output for this one.

[14] N. Ekwe-Ekwe and L. Amos, “The state of FaaS: An analysis of public Functions-as-a-Service providers,” in *Proc. 2024 IEEE 17th Int. Conf. Cloud Computing (CLOUD)*, 2024, pp. 430–438, doi: 10.1109/CLOUD62652.2024.00055.

[15] C. Lin, N. Mahmoudi, S. Fan, and H. Khazaei, “Fine-grained performance and cost modeling and optimization for FaaS applications,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 34, no. 1, pp. 180–194, 2023, doi: 10.1109/TPDS.2022.3214783.

[16] P. Vahidinia, B. Farahani, and F. Shams Aliee, “Mitigating cold start problem in serverless computing: A reinforcement learning approach,” *IEEE Internet Things J.*, 2022, doi: 10.1109/JIOT.2022.3165127.

[17] K. German and O. Ponomareva, “An overview of container security in a Kubernetes cluster,” in *Proc. 2023 IEEE Ural-Siberian Conf. Biomedical Engineering, Radioelectronics and Information Technology (USBREIT)*, 2023, pp. 283–285, doi: 10.1109/USBREIT58508.2023.10158865.

[18] A. K. Verma, R. Patel, P. Rai, and S. Patel, “Performance comparison between Docker container and virtual machine microservices,” in *Proc. 2024 IEEE 13th Int. Conf. Communication Systems and Network Technologies (CSNT)*, 2024, pp. 718–723, doi: 10.1109/CSNT60213.2024.10546015.

[19] G. Liu, B. Huang, Z. Liang, M. Qin, H. Zhou, and Z. Li, “Microservices: Architecture, container, and challenges,” in *Proc. 2020 IEEE 20th Int. Conf. Software Quality, Reliability and Security Companion (QRS-C)*, 2020, pp. 629–635, doi: 10.1109/QRS-C51114.2020.00107.

[20] B. Martens, M. Walterbusch, and F. Teuteberg, “Costing of cloud computing services: A total cost of ownership approach,” in *Proc. 2012 45th Hawaii Int. Conf. System Sciences (HICSS)*, 2012, pp. 1563–1572, doi: 10.1109/HICSS.2012.186.