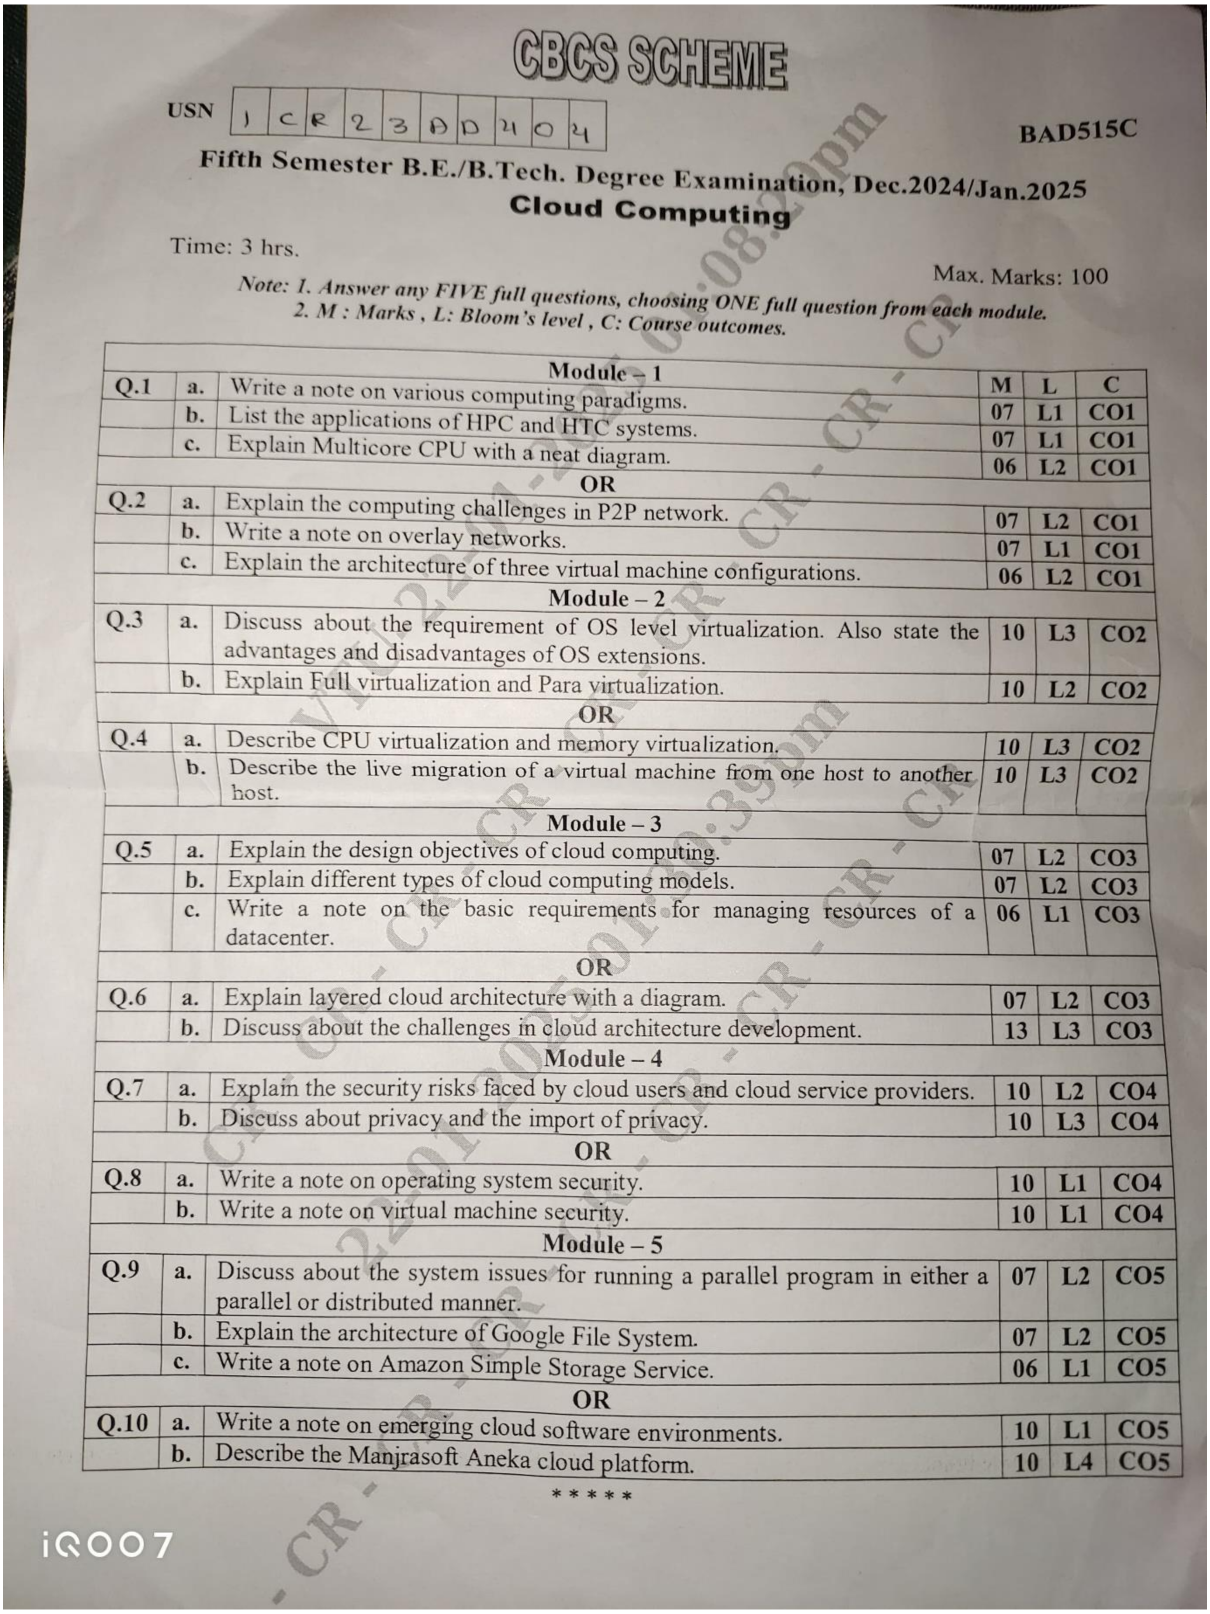

CBCS SCHEME

USN 1 C R 2 3 A D 4 0 4

BAD515C

Fifth Semester B.E./B.Tech. Degree Examination, Dec.2024/Jan.2025

Cloud Computing

Time: 3 hrs. Max. Marks: 100

Note: 1. Answer any FIVE full questions, choosing ONE full question from each module.
2. M : Marks , L: Bloom's level , C: Course outcomes.

		Module – 1	M	L	C
Q.1	a.	Write a note on various computing paradigms.	07	L1	CO1
	b.	List the applications of HPC and HTC systems.	07	L1	CO1
	c.	Explain Multicore CPU with a neat diagram.	06	L2	CO1
		OR			
Q.2	a.	Explain the computing challenges in P2P network.	07	L2	CO1
	b.	Write a note on overlay networks.	07	L1	CO1
	c.	Explain the architecture of three virtual machine configurations.	06	L2	CO1
		Module – 2			
Q.3	a.	Discuss about the requirement of OS level virtualization. Also state the advantages and disadvantages of OS extensions.	10	L3	CO2
	b.	Explain Full virtualization and Para virtualization.	10	L2	CO2
		OR			
Q.4	a.	Describe CPU virtualization and memory virtualization.	10	L3	CO2
	b.	Describe the live migration of a virtual machine from one host to another host.	10	L3	CO2
		Module – 3			
Q.5	a.	Explain the design objectives of cloud computing.	07	L2	CO3
	b.	Explain different types of cloud computing models.	07	L2	CO3
	c.	Write a note on the basic requirements for managing resources of a datacenter.	06	L1	CO3
		OR			
Q.6	a.	Explain layered cloud architecture with a diagram.	07	L2	CO3
	b.	Discuss about the challenges in cloud architecture development.	13	L3	CO3
		Module – 4			
Q.7	a.	Explain the security risks faced by cloud users and cloud service providers.	10	L2	CO4
	b.	Discuss about privacy and the import of privacy.	10	L3	CO4
		OR			
Q.8	a.	Write a note on operating system security.	10	L1	CO4
	b.	Write a note on virtual machine security.	10	L1	CO4
		Module – 5			
Q.9	a.	Discuss about the system issues for running a parallel program in either a parallel or distributed manner.	07	L2	CO5
	b.	Explain the architecture of Google File System.	07	L2	CO5
	c.	Write a note on Amazon Simple Storage Service.	06	L1	CO5
		OR			
Q.10	a.	Write a note on emerging cloud software environments.	10	L1	CO5
	b.	Describe the Manjrasoft Aneka cloud platform.	10	L4	CO5

* * * * *

1a. Various computing paradigms:

Computing paradigm Distinctions

- Centralized computing
  - All computer resources are centralized in on
    physical system.
  - All resources (processors, memory, and storage) are fully shared and tightly coupled within one integrated OS
- Parallel computing -

- all processors are either tightly coupled wi centralized shared memory or loosely coupled with distributed memory.
- Interprocessor communication is accomplishe through shared memory or via message passing.

* Distributed computing -

- A distributed system consists of multiple autonomous computers, each having its own private memory, communicating

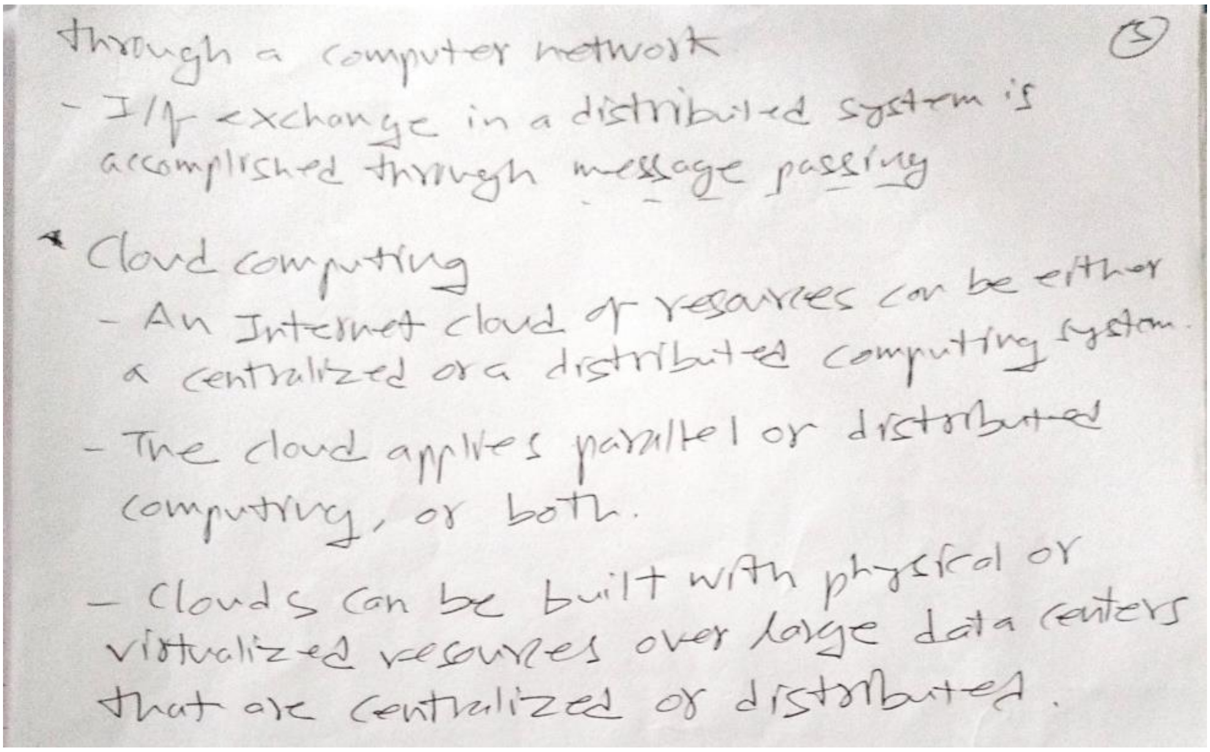

through a computer network 3

- I/f exchange in a distributed system is accomplished through message passing

* Cloud computing
  - An Internet cloud of resources can be either a centralized or a distributed computing system.
  - The cloud applies parallel or distributed computing, or both.
  - Clouds can be built with physical or virtualized resources over large data centers that are centralized or distributed.

1b. Application of HPC and HTC systems

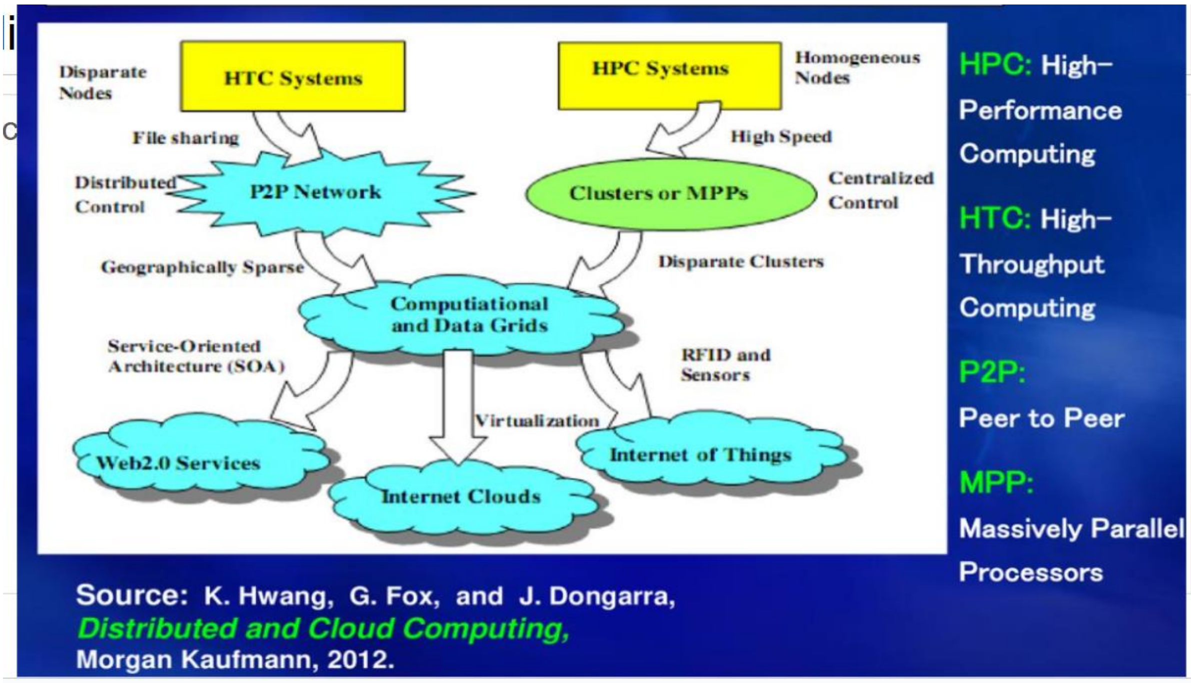

HPC: High-Performance Computing

HTC: High-Throughput Computing

P2P: Peer to Peer

MPP: Massively Parallel Processors

Source: K. Hwang, G. Fox, and J. Dongarra, *Distributed and Cloud Computing*, Morgan Kaufmann, 2012.

1c. Multicore CPU

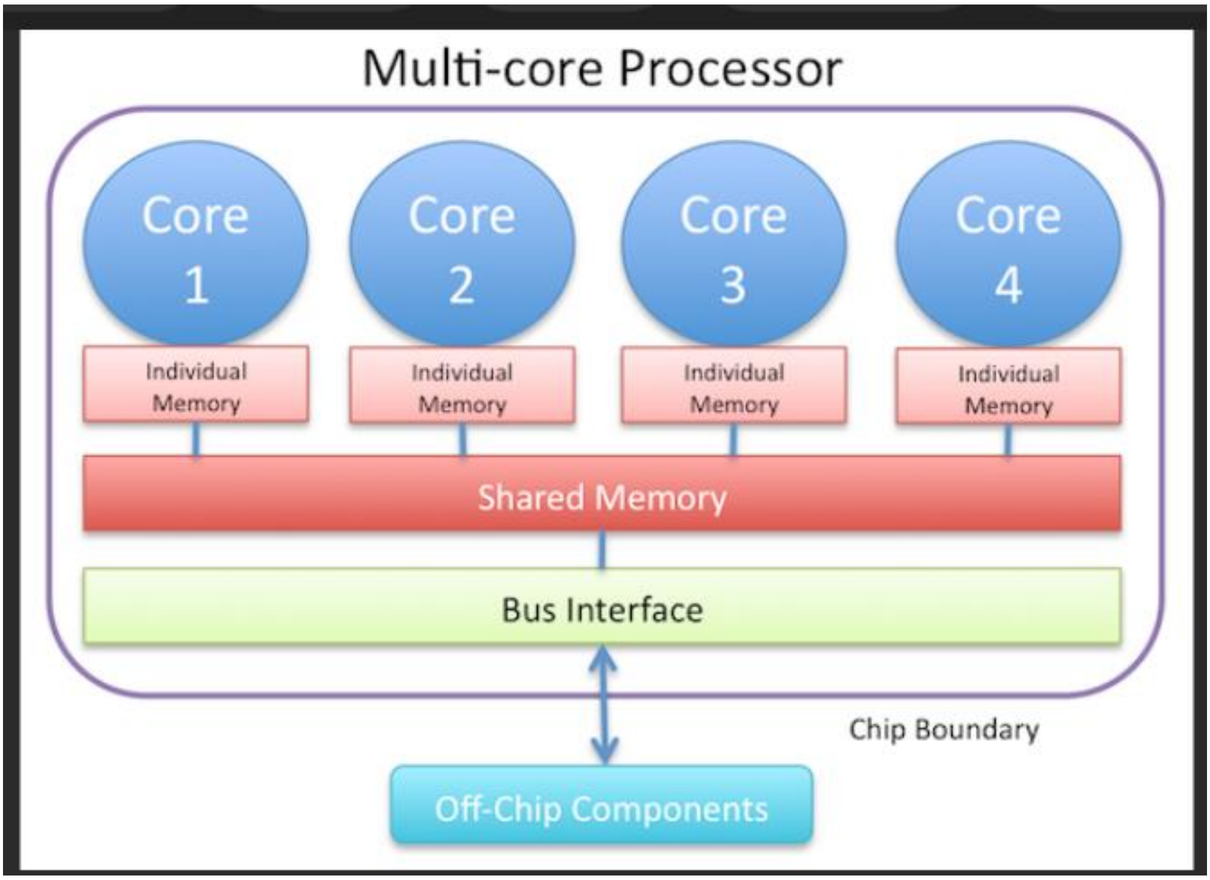

3a. OS level virtualization

OS-level virtualization, also known as containerization, is a technique that allows multiple isolated instances of an operating system to run on the same machine.

How it works

- The kernel of the operating system allows for multiple isolated user-space instances.
- These instances are called containers, virtual environments (VEs), virtual private servers (VPSes), or jails.
- Programs running in a container can only see the container's contents and devices.
- The virtualized operating system handles the users and their requests individually.

Benefits

- **Fast deployment**: Containers start guests in seconds.
- **Less resource requirement**: Guests in containers share the OS with the host or even no OS.
- **Improved security, manageability and availability**: OS-level virtualization has been widely used to improve these aspects.

Use cases

- OS-level virtualization is used in cloud computing to run numerous programs for many users on the same machine simultaneously.
- OS-level virtualization can be used to consolidate physical and virtual resources in data centers.

3b. full virtualization and para virtualization

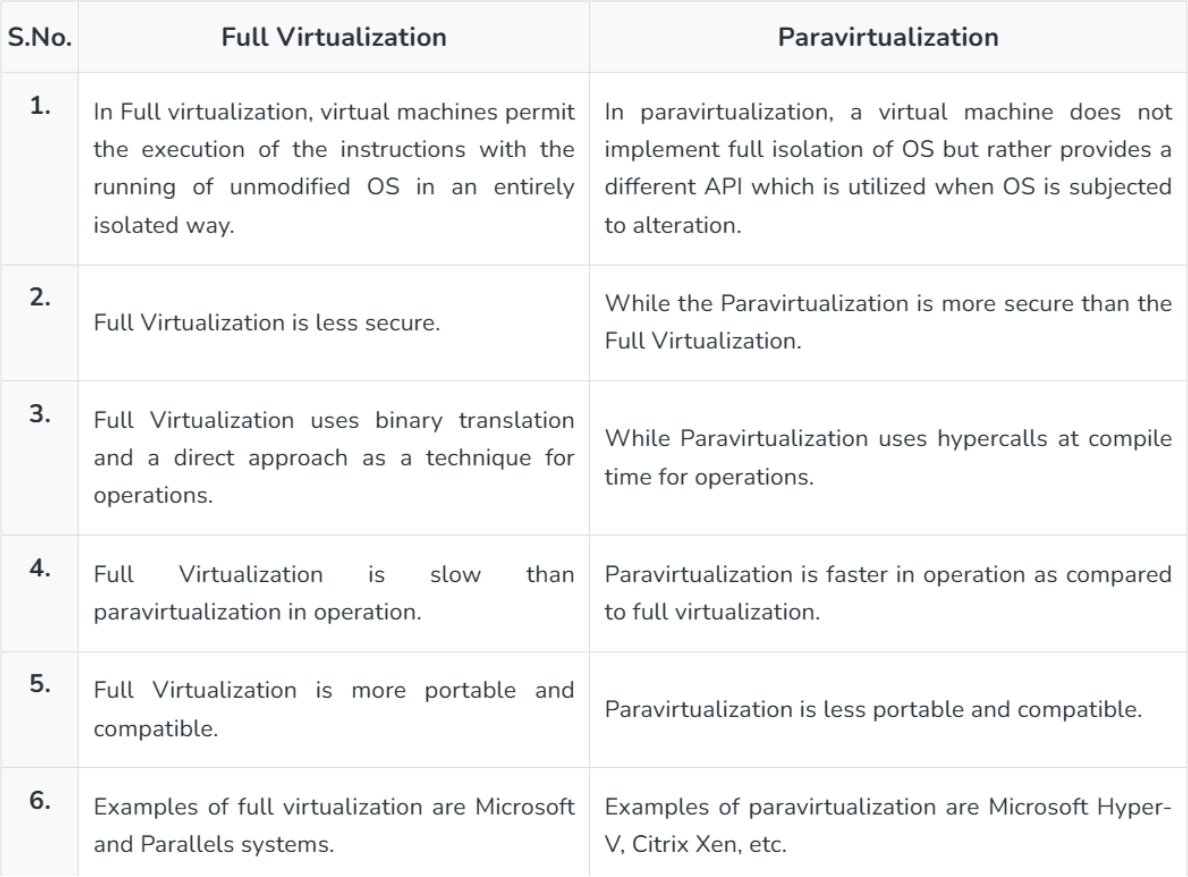

S.No.	Full Virtualization	Paravirtualization
1.	In Full virtualization, virtual machines permit the execution of the instructions with the running of unmodified OS in an entirely isolated way.	In paravirtualization, a virtual machine does not implement full isolation of OS but rather provides a different API which is utilized when OS is subjected to alteration.
2.	Full Virtualization is less secure.	While the Paravirtualization is more secure than the Full Virtualization.
3.	Full Virtualization uses binary translation and a direct approach as a technique for operations.	While Paravirtualization uses hypercalls at compile time for operations.
4.	Full Virtualization is slow than paravirtualization in operation.	Paravirtualization is faster in operation as compared to full virtualization.
5.	Full Virtualization is more portable and compatible.	Paravirtualization is less portable and compatible.
6.	Examples of full virtualization are Microsoft and Parallels systems.	Examples of paravirtualization are Microsoft Hyper-V, Citrix Xen, etc.

5a. The design objectives of cloud computing include cost savings, scalability, flexibility, and innovation.

Cost savings

- Eliminates the need for physical hardware and infrastructure investments
- Uses commodity hardware and open source software to minimize costs

Scalability

- Allows organizations to scale resources up or down based on demand
- Supports smooth expansion of functional modules

Flexibility

- Provides access to resources and data from any location with an internet connection
- Allows for changes to the configuration and topology as needed

Innovation

- Enables rapid development and deployment of applications
- Supports real-time collaboration among team members
- Fosters a culture of continuous improvement and agility

5b. The main types of cloud computing models are public, private, and hybrid clouds. Each model can use a variety of services, including Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

Cloud computing models

- **Public cloud**: A cloud environment that uses IT infrastructure that's not owned by the end user.
- **Private cloud**: A cloud environment that's dedicated to a single user or group.
- **Hybrid cloud**: A combination of public and private cloud elements that are connected over the internet.

Cloud computing services

- **IaaS**: A service that provides virtualized computing resources like servers, storage, and networking.
- **PaaS**: A service that provides a platform for developers to build, deploy, and manage applications.
- **SaaS**: A service that delivers software applications over the internet.

Benefits of cloud computing

- Reduced costs for managing and maintaining on-premises infrastructure
- Scalable resources like RAM and flexible bandwidth
- The ability to adjust storage and processing capabilities
- The ability to operate without large upfront investments

5c. Basic requirements for managing data center resources include: reliable storage infrastructure, robust network connectivity, effective cooling systems, robust security measures, capacity planning to anticipate demand, disaster recovery and continuity planning, regular hardware maintenance, and skilled staff with expertise in IT operations and cybersecurity to ensure optimal performance, data integrity, and uptime.

Key aspects to consider:

- **Physical infrastructure:**
  Servers, storage devices, network switches, cooling systems, power distribution units, and physical security measures.
- **Network management:**
  Monitoring network performance, managing bandwidth, ensuring redundancy, and secure routing.
- **Data storage:**
  Implementing reliable storage solutions with redundancy to prevent data loss.
- **Security:**
  Access controls, intrusion detection, data encryption, and regular security audits.
- **Capacity planning:**
  Forecasting future resource needs and scaling infrastructure accordingly.
- **Monitoring and alerting:**
  Real-time monitoring of system health, performance metrics, and alerts for potential issues.
- **Disaster recovery:**
  Developing plans to restore operations in case of a major outage or disaster.
- **Energy efficiency:**
  Optimizing power usage and cooling systems to minimize energy consumption.
- **Staff expertise:**
  Skilled technicians with knowledge of data center hardware, software, network management, and security protocols.

7a. Cloud users and cloud service providers face a number of security risks, including data breaches, account hijacking, and malware.

Security risks for cloud users

- **Data breaches**
  Unauthorized access to sensitive data, which can be caused by malicious attacks, human error, or system vulnerabilities
- **Account hijacking**
  Cybercriminals gain unauthorized access to a user's account and use it to access other accounts
- **Data loss**

Accidental deletion, hardware failure, or malicious attacks can lead to data loss
Security risks for cloud service providers

- **Denial of service attacks**
  Can shut down cloud services, making them temporarily unavailable to users
- **Malware attacks**
  Criminals use malicious software to gain access or cause damage to a computer or network
- **Cloud misconfiguration**
  Errors in cloud service configuration that can lead to security vulnerabilities
- **Regulation and compliance issues**
  Cloud service providers need to meet regulatory requirements and industry standards
  Other cloud security challenges Poor identity and access management (IAM), Insider threats, Evolving attack surfaces, Limited visibility, and Shortage of security experts.

7b.

"Privacy" refers to an individual's right to control how their personal information is collected, used, and shared, essentially the freedom to be left alone or to determine when and how information about them is disclosed to others; the "importance of privacy" lies in its role in protecting personal autonomy, safeguarding sensitive information, and maintaining individual dignity across various aspects of life like personal relationships, health, and online activities.

Key points about privacy and its importance:

- **Individual control:**
  Privacy allows individuals to decide what information about themselves is shared and with whom, promoting personal agency and self-determination.
- **Psychological well-being:**
  Maintaining privacy can contribute to a sense of security and comfort, enabling individuals to express themselves freely without fear of judgment or unwanted intrusion.
- **Protection of sensitive data:**
  Privacy safeguards personal information like medical records, financial details, and contact information from unauthorized access or misuse.
- **Social trust:**
  Respecting privacy fosters trust within communities and relationships, allowing for open communication without fear of information being shared inappropriately.
- **Legal implications:**
  Many jurisdictions have enacted data privacy laws, like the European Union's GDPR, to protect individuals' privacy rights and hold organizations accountable for data breaches.
  Different aspects of privacy:
- **Physical privacy:** The right to be alone in physical spaces without unwanted observation.
- **Information privacy:** The right to control how personal data is collected, stored, and used.
- **Communication privacy:** The right to private conversations and correspondence.
- **Digital privacy:** The right to control personal information shared online.

9a. When running a parallel program, whether on a single machine with multiple cores (parallel) or across a network of computers (distributed), common system issues include: load balancing, communication overhead, synchronization issues, race conditions, memory management, data consistency, fault tolerance, and network latency; all of which can significantly impact the performance and reliability of your parallel application.

Key System Issues:

- **Load Balancing:**

Uneven distribution of workload across available processors, leading to some processors being idle while others are overloaded, impacting overall efficiency.

- **Communication Overhead:**
  The time spent transferring data between processors, which can become significant in distributed systems with high network latency.
- **Synchronization Issues:**
  Coordinating access to shared data between multiple threads or processes to avoid race conditions where the result depends on the order of execution.
- **Race Conditions:**
  When multiple threads try to access and modify shared data simultaneously, potentially leading to inconsistent results.
- **Memory Management:**
  Efficiently allocating and managing memory across multiple processors, especially in shared memory systems, can be challenging.
- **Data Consistency:**
  Ensuring that all copies of data across different processors are consistent, especially when updates are made in parallel.
- **Fault Tolerance:**
  Designing a system to handle processor or network failures and continue running despite disruptions.
- **Network Latency:**
  Delays in data transmission over the network, which can significantly impact performance in distributed systems.

Factors Affecting System Issues:

- **Program Design:** How well the problem is parallelized and tasks are divided among processors.
- **Hardware Architecture:** The capabilities of the processors and network infrastructure.
- **Parallel Programming Model:** The chosen paradigm for parallel programming (e.g., threads, message passing).

Mitigating System Issues:

- **Efficient Task Decomposition:**
  Carefully dividing the problem into independent, manageable tasks for parallel execution.
- **Load Balancing Techniques:**
  Dynamically adjusting workload distribution among processors to ensure even utilization.
- **Synchronization Primitives:**
  Using mechanisms like mutexes, semaphores, and critical sections to control access to shared data.
- **Optimized Communication Patterns:**
  Minimizing unnecessary data transfers and using efficient communication protocols.
- **Data Replication Strategies:**
  Replicating data strategically to improve access times and fault tolerance.

9b.

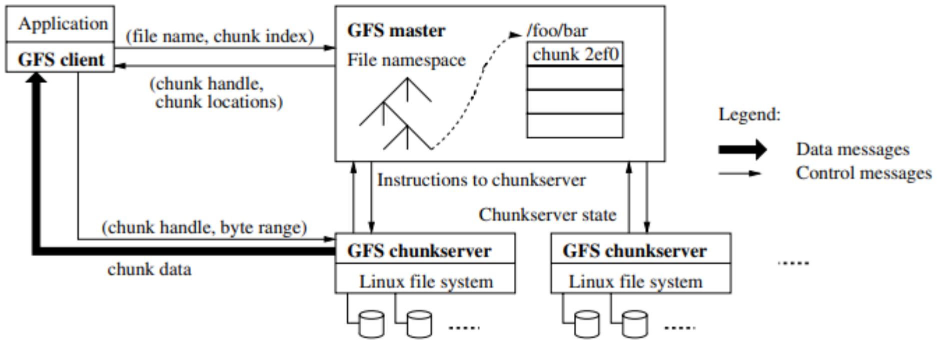

Figure 1: GFS Architecture

https://static.googleusercontent.com/media/research.google.com/en//archive/gfs-sosp2003.pdf

9c.

What is Amazon S3?

Amazon Simple Storage Service (Amazon S3) is an object storage service offering industry-leading scalability, data availability, security, and performance. Millions of customers of all sizes and industries store, manage, analyze, and protect any amount of data for virtually any use case, such as data lakes, cloud-native applications, and mobile apps. With cost-effective storage classes and easy-to-use management features, you can optimize costs, organize and analyze data, and configure fine-tuned access controls to meet specific business and compliance requirements.

Benefits

Scalability

Durability and availability

Security and data protection

Lowest price and highest performance