USN										

Internal Assessment Test 1 – Sept. 2025

Subject/Code: CLOUD COMPUTING/ BAD515C				
Date:26-09-2025	Duration: 90 mins	Max. Marks: 50	Semester: 05	Branch: AI&DS

Sl.	Answer any FIVE FULL Questions	Marks	CO	RBT
1	Define cloud computing. How is the Cloud Computing Reference Model structured? Draw a diagram to support your explanation. Cloud Computing: Cloud computing is the delivery of computing services like servers, storage, databases, networking, software, and applications over the Internet ("the cloud"). It allows users to access resources on-demand, pay only for what they use, and scale services as needed without managing physical hardware.	10	CO1	L1

	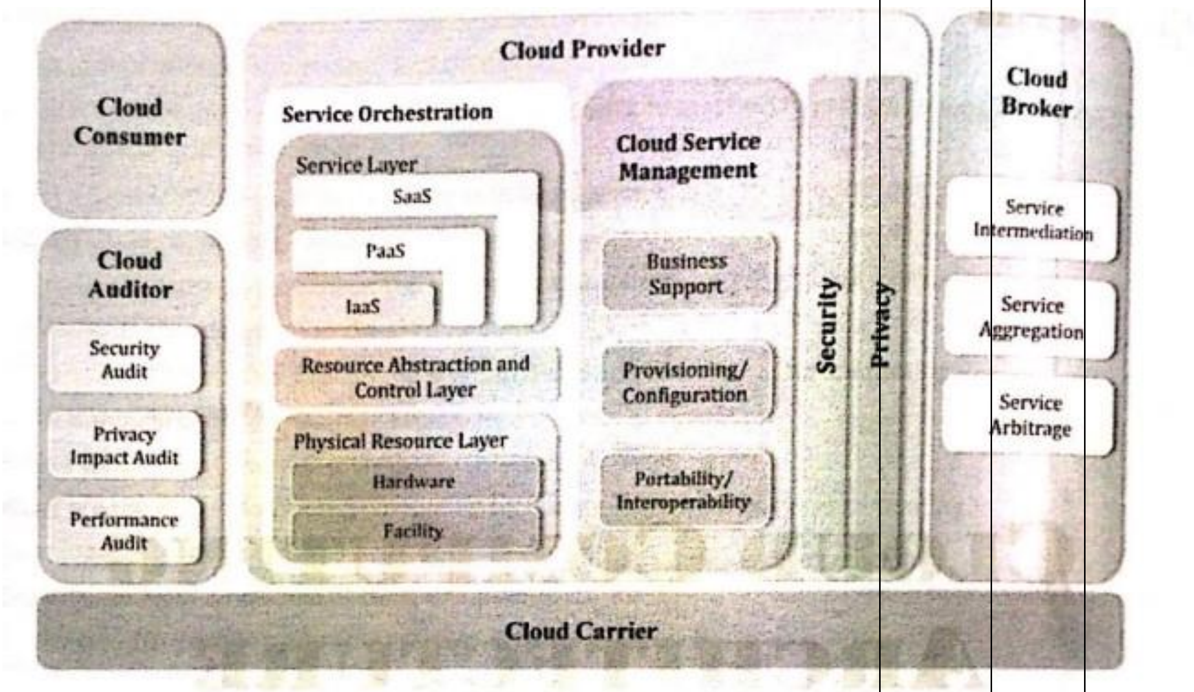 **Figure 2.1 Cloud Computing Reference Model** Cloud Computing Reference Model: • The Cloud Computing Reference Model is structured into five layers that show how cloud services are delivered • Cloud Application Layer – Provides end-user services such as Gmail, Google Docs, Office 365, etc. (SaaS). • Cloud Software Environment Layer – Offers platforms for developers (PaaS), e.g., Google App Engine, AWS Elastic Beanstalk. • Cloud Software Infrastructure Layer – Provides virtualized resources like servers, storage, and networks (IaaS), e.g., AWS EC2, Microsoft Azure. • Software Kernel Layer – Manages the operating systems, hypervisors, and basic resource allocation. • Hardware & Firmware Layer – Physical servers, storage devices, and networking hardware form the base.			
2	Analyze the trade-offs between performance, security, and energy efficiency in distributed cloud systems. Trade-offs between Performance, Security, and Energy Efficiency in Distributed Cloud Systems Performance vs. Security • High performance requires faster access, low latency, and quick data transfer. • Strong security (encryption, firewalls, access control) can add overhead, slowing down performance. • Example: Encrypting all data increases security but reduces speed. Performance vs. Energy Efficiency • High performance needs more computing power, more servers, and higher processing speed. • This increases energy usage and reduces efficiency.	10	CO1	L2

- Example: Running servers at peak load improves response but consumes more electricity.

Performance Metrics and Scalability

- **Performance Metrics**:
  - Measured in **MIPS, Tflops, TPS**, job response time, and network latency.
  - QoS, dependability, and security resilience are also performance factors.

- **Amdahl's Law** → Speedup limited by sequential portion (α).
  - Even with infinite processors, speedup = 1/α.
  - Example: If α=0.25 → max speedup = 4.
- **Fixed Workload Problem** → Efficiency drops as cluster size grows (many nodes idle).
- **Gustafson's Law** → Scale **problem size** with number of processors → maintains efficiency.
  - Example: With 256 nodes, efficiency ≈ 75%.
  - Rule: Use **Amdahl** for fixed workload, **Gustafson** for scaled workload.

- **High Availability (HA)**:
  - Defined as: Availability = MTTF / (MTTF + MTTR).
  - Goal → **No single point of failure**.
  - Strategies → Redundancy, fault isolation, quick recovery.
- **System Comparisons**:
  - **SMP/MPP** → Lower availability (centralized).
  - **NUMA** → Better with multiple OSes.
  - **Clusters & Clouds** → HA with failover, but availability drops as system grows.
  - **Grids** → Higher availability due to fault isolation.
  - **P2P** → Low availability, peers unreliable.

Security vs. Energy Efficiency

- Stronger security requires extra processes (encryption, intrusion detection, monitoring).
- These consume additional computing resources and power.
- Example: Continuous security scans protect data but use more energy.

Balancing All Three

- Cloud providers must carefully balance performance, security, and energy efficiency.
- Too much focus on one can weaken the other two.
- Optimizations like virtualization, efficient algorithms, green data centers, and adaptive load balancing help achieve balance.

	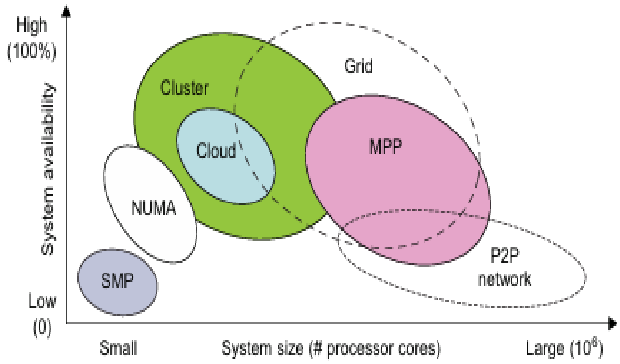 **FIGURE 1.24** Estimated system availability by system size of common configurations in 2010.			
3	Describe the classification of virtualization and discuss the techniques used for hardware virtualization. 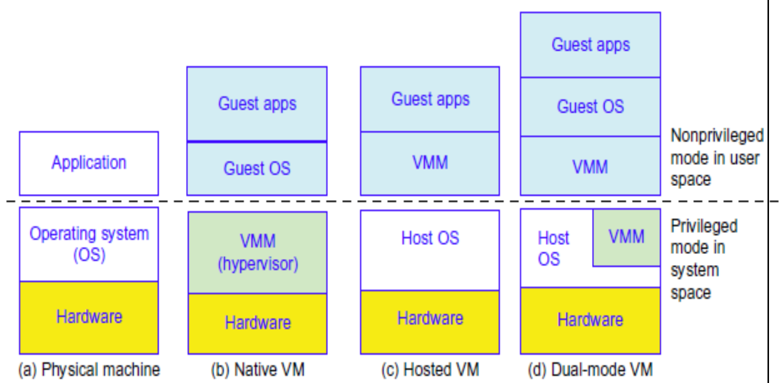 Classification of Virtualization: Virtualization is the process of creating a virtual version of hardware, OS, storage, or network. It can be classified as: Hardware Virtualization – Virtual machines created on physical hardware using hypervisors. Operating System Virtualization – Multiple isolated user spaces on a single OS (e.g., containers like Docker). Server Virtualization – Dividing a physical server into multiple logical servers. Storage Virtualization – Pooling multiple storage devices to appear as one. Network Virtualization – Combining network resources into a single virtual network. Desktop Virtualization – Running a desktop environment on a remote server, accessed via thin clients. Techniques Used for Hardware Virtualization: Full Virtualization The hypervisor completely simulates hardware. Guest OS runs unmodified as if on real hardware. Example: VMware, VirtualBox. Para-Virtualization Guest OS is modified to interact with the hypervisor. Faster than full virtualization because it reduces overhead. Example: Xen.	10	CO2	L2

<table>
<tr>
<td></td>
<td>Hardware-Assisted Virtualization
Modern CPUs provide special instructions (Intel VT-x, AMD-V) to support virtualization.
Improves performance and efficiency.</td>
<td></td>
<td></td>
<td></td>
</tr>
<tr>
<td>4</td>
<td>

Explain virtualization and discuss its pros and cons. Describe the different hypervisors.

Virtualization:

Virtualization is the process of creating a virtual version of resources like hardware, operating systems, storage, or networks. It allows multiple virtual machines (VMs) to run on a single physical system by sharing resources.

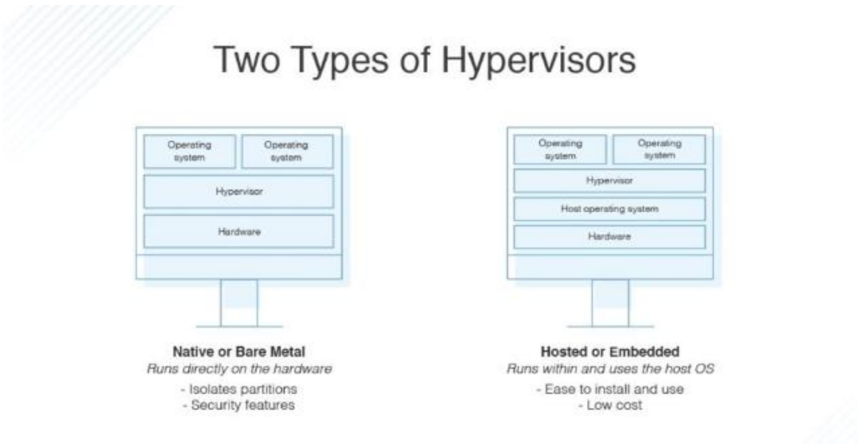

Pros of Virtualization:

Efficient Resource Utilization – Better use of CPU, memory, and storage.

Cost Saving – Reduces hardware and maintenance costs.

Flexibility – Easy to create, modify, and migrate virtual machines.

Isolation – Each VM runs independently, reducing system crashes.

Scalability – Supports rapid scaling in cloud computing.

Cons of Virtualization:

Performance Overhead – VMs may run slower than physical machines.

Security Risks – Attacks on hypervisors can affect all VMs.

Management Complexity – Requires skilled administration.

High Initial Setup Cost – Needs powerful servers and software licenses.

Types of Hypervisors:

Type 1 – Bare-metal Hypervisor

Runs directly on hardware without a host OS.

High performance and efficiency.

Examples: VMware ESXi, Microsoft Hyper-V, Xen.

</td>
<td>10</td>
<td>CO2</td>
<td>L2</td>
</tr>
</table>

	Type 2 – Hosted Hypervisor Runs on top of a host operating system. Easier to set up but less efficient due to OS overhead. Examples: Oracle VirtualBox, VMware Workstation.			
5	Describe the various computing platforms and technologies, highlighting their key features. **Various Computing Platforms and Technologies with Key Features:** 1. Mainframe Computing • Powerful centralized computers used for large-scale transaction processing. • Features: High reliability, supports thousands of users, large storage, strong security. 2. Client-Server Computing • Work is divided between client (request) and server (response). • Features: Centralized management, easy resource sharing, scalable, suitable for enterprise applications. 3. Cluster Computing • Group of interconnected computers working together as a single system. • Features: High performance, fault tolerance, cost-effective compared to supercomputers. 4. Grid Computing • Uses distributed resources from multiple locations to solve large problems. • Features: Resource sharing, large-scale computation, heterogeneous systems support. 5. Cloud Computing • Provides on-demand services (IaaS, PaaS, SaaS) via the internet. • Features: Scalability, pay-as-you-go, accessibility from anywhere, elasticity. 6. Edge Computing • Processing done near data source (IoT devices, sensors) instead of central servers. • Features: Low latency, real-time processing, reduces bandwidth usage. 7. Fog Computing • Extension of cloud to local devices between cloud and edge. • Features: Decentralized, improves performance for IoT, supports mobility and location awareness.	10	CO3	L2
6	Your company wants to deploy infrastructure on a cloud provider. Draw and explain the IaaS Reference Model to illustrate the services and layers involved.	10	CO3	L3

	1.			

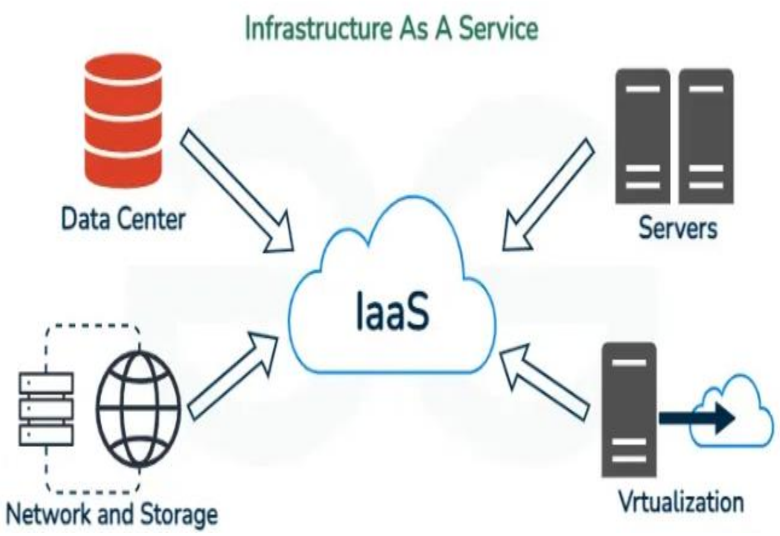

IaaS (Infrastructure as a Service) Reference Model:

IaaS provides virtualized computing resources (servers, storage, and networking) over the internet. Users can deploy and manage their own applications without worrying about the underlying hardware.

Layers of IaaS Reference Model:

1. Physical Layer
   - Consists of actual hardware: servers, storage devices, and networking equipment.
2. Virtualization Layer
   - Uses hypervisors to virtualize hardware into multiple Virtual Machines (VMs).
   - Ensures isolation, scalability, and efficient resource utilization.
3. Infrastructure Layer
   - Provides virtual servers, storage, and networking resources to users.
   - Example: AWS EC2, Azure Virtual Machines.
4. Management Layer
   - Handles monitoring, provisioning, billing, load balancing, and security.
   - Provides APIs, dashboards, and automation tools.
5. Application Layer (User Layer)
   - End-users or organizations access services through web portals, CLI, or APIs.
   - They install and run their applications on provisioned VMs.

CI CCI HOD