UNIT–1 : Introduction to Cloud Computing (PYQ Solution)

AKTU · Cloud Computing (BCAM051) · Solutions by **Chirag Jain**, Ycotes

SECTION A Answers (≈70 Words Each)

Q1. Define cloud computing and list its key characteristics. (Jan 2026, Jan 2025, 2024, Feb 2024)

Cloud computing is a model for delivering computing resources (servers, storage, databases, networking, software) over the Internet on a pay-as-you-use basis (NIST: ubiquitous, convenient, on-demand network access to a shared pool of configurable resources). Key characteristics: on-demand self-service, broad network access, resource pooling (multi-tenancy), rapid elasticity, measured service, plus scalability, high availability and cost efficiency.

Q2. Why is elasticity a key feature of cloud platforms? How does it contribute to cost optimisation? (Jan 2026, Feb 2024)

Elasticity is the ability to automatically add or remove resources as demand changes. Users scale out during peaks and scale in when demand falls, so they never run idle capacity. Because billing is per use, they pay only for what they consume, avoid over-provisioning for peak load and avoid performance loss from under-provisioning, which lowers cost and keeps service quality.

Q3. What is on-demand provisioning and how does it benefit cloud users? (Jan 2026)

On-demand provisioning means resources such as VMs, storage or databases are allocated automatically and instantly when requested, without human intervention by the provider. Benefits: fast deployment (minutes instead of

weeks), no up-front investment, pay-per-use billing, easy scaling and freedom to experiment.

Q4. Trace the evolution of cloud computing. Evaluate the challenges faced during the evolution. (Jan 2025)

Evolution: mainframes with time-sharing (1960s) → client-server and clusters → distributed and grid computing (1990s) → utility computing and virtualization → web services/SaaS (Salesforce, 1999) → modern cloud (AWS 2006, Azure, Google Cloud).

Challenges: limited network bandwidth and latency, immature virtualization, security and privacy worries, lack of standards and interoperability (vendor lock-in), reliability and SLA guarantees, and data-protection regulations.

Q5. Compare parallel computing and distributed computing. (2024)

Parallel computing	Distributed computing
Many processors in **one machine** with shared memory	Many **independent computers** linked by a network
Tightly coupled, one task split into subtasks	Loosely coupled, nodes cooperate and may run different tasks
Goal: speed (performance)	Goal: resource sharing, scalability, fault tolerance
Communication through shared memory	Communication by message passing
Example: multi-core GPU	Example: Hadoop cluster, cloud

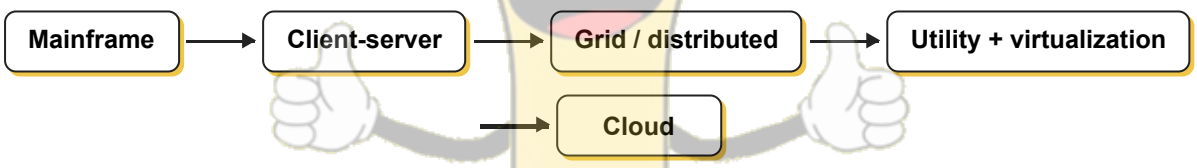
SECTION B Answers (≈250 Words Each)

Q1. Define the need of cloud computing. Explain the evolution of cloud computing and its relationship with distributed computing. (Jan 2026, 2024)

Need: traditional IT requires large capital spending, long procurement, over-provisioning for peak load, and in-house maintenance. Cloud computing gives **low cost, on-demand scalability, fast deployment, global access, automatic updates, backup and disaster recovery**, so businesses can focus on their core work.

Evolution:

1. **Mainframes (1960s):** centralised, time-shared computing via dumb terminals.
2. **Client-server and clusters (1980s-90s):** PCs and networked servers.
3. **Distributed and grid computing:** many computers cooperate on large tasks.
4. **Utility computing and virtualization (2000s):** computing sold like electricity; VMs allow sharing hardware.
5. **Cloud computing (2006 onward):** AWS EC2/S3, Google App Engine, Microsoft Azure; SaaS, PaaS, IaaS.



Evolution of cloud computing

Relationship with distributed computing: cloud computing is built on distributed computing, in which independent computers coordinate through a network, appearing to the user as a single system. The cloud adds **virtualization, resource pooling, automation, elasticity and pay-per-use billing** on top of distributed systems, and uses their fault tolerance and scalability to serve many tenants.

Conclusion: the cloud is the commercial culmination of distributed, grid and utility computing.

Q2. Analyze the underlying principles of parallel computing in cloud environments and how it improves cloud performance. Derive equations

for workload balancing for 10,000 tasks. (Jan 2025, Feb 2024)

Principles of parallel computing:

- **Decomposition:** split a problem into independent subtasks.
- **Concurrency:** run subtasks on several processors/VMs at once.
- **Communication and synchronisation:** exchange results among tasks.
- **Load balancing:** spread work so no node is idle or overloaded.
- **Scalability:** performance grows as nodes are added.

Contribution to cloud applications: shorter response time, higher throughput (MapReduce, batch jobs), better resource utilisation and fault tolerance (a failed task is re-run elsewhere).

Performance equations: Speed-up S = T1/Tp. **Amdahl's law:** S(p) = 1 / (f + (1 − f)/p), where f is the serial fraction. For f = 0.05 and p = 10: S = 1/(0.05 + 0.095) = 6.9. Efficiency E = S/p = 0.69.

Workload balancing for N = 10,000 tasks on n nodes:

- Homogeneous nodes: tasks per node = N/n. With n = 10, each node gets **1,000** tasks.
- Heterogeneous nodes with speeds s_i: tasks_i = N × s_i / Σ s_j. For speeds 1, 2, 2, 5 (sum 10): **1,000, 2,000, 2,000, 5,000** tasks, so all nodes finish together.
- Completion time T = max(tasks_i / s_i) + communication overhead. Balanced: T = N/Σs = 1,000 time units.
- Imbalance factor = (max load − average load)/average load; a good scheduler keeps it near 0.

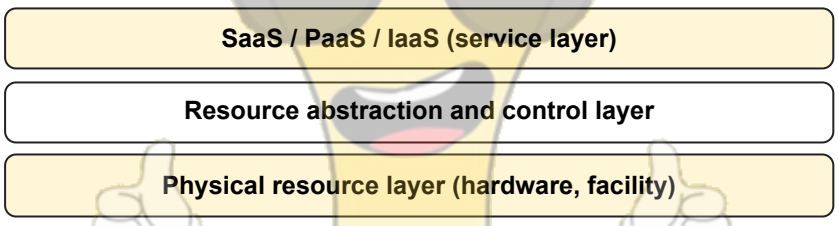
Conclusion: parallelism yields performance only if the serial fraction is small and the load is balanced.

SECTION C Answers (≈500 Words Each)

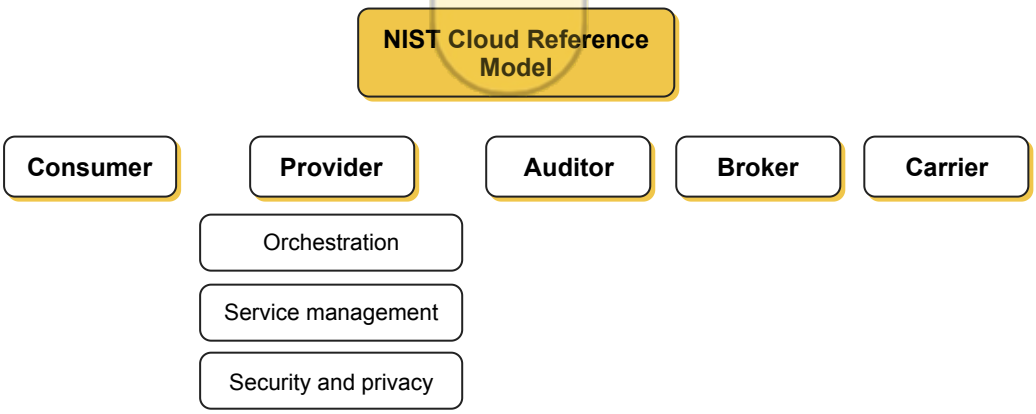
Q1. Explain the NIST cloud reference architecture and its significance in standardisation. Describe the cloud computing reference model with a diagram. (Jan 2026, 2024)

The **NIST Cloud Computing Reference Architecture** (SP 500-292) is a neutral, vendor-independent model describing **five actors**:

1. **Cloud Consumer:** person or organisation that uses cloud services.
2. **Cloud Provider:** supplies the service. Activities: **service deployment, service orchestration** (service layer SaaS/PaaS/IaaS over resource abstraction and control layer over the physical layer), **cloud service management** (business support, provisioning, portability and interoperability), **security and privacy**.
3. **Cloud Auditor:** independently assesses security, privacy and performance.
4. **Cloud Broker:** manages use and delivery; offers **intermediation, aggregation and arbitrage**.
5. **Cloud Carrier:** provides connectivity and transport (network, telecom).



Provider's service orchestration layers



NIST actors

Significance in standardisation:

- Gives a **common vocabulary** and taxonomy for providers, buyers and regulators.

- Helps **compare and evaluate** cloud services and security requirements.
- Promotes **interoperability and portability**, reducing vendor lock-in.
- Supports **security, auditing and compliance** (roles of auditor, security layer).
- Guides cloud **design and procurement** in government and industry.

Conclusion: the reference model provides a standard map of cloud actors and activities that helps the industry build compatible and secure services.

Q2. Compare traditional IT resource provisioning with cloud-based provisioning. Provide a cost analysis model for a small business migrating to the cloud. (Jan 2025)

Aspect	Traditional	Cloud
Investment	Large CapEx (servers, space, power)	OpEx, pay-per-use
Setup time	Weeks to months	Minutes
Capacity	Fixed, sized for peak, often idle	Elastic, scales with demand
Maintenance	In-house staff	Provider manages hardware
Reliability/DR	Costly to build	Built-in redundancy, backups
Risk	Under- or over-provisioning	Vendor lock-in, recurring bills

Cost model: TCO (5 years) = Capital cost + Σ yearly running cost. Assume the business needs 4 servers.

Traditional: hardware ₹8,00,000; space, power, cooling and licences ₹1,50,000/yr; one administrator ₹3,00,000/yr.
TCO = 8,00,000 + 5 × (1,50,000 + 3,00,000) = 8,00,000 + 22,50,000 = **₹30,50,000**.

Cloud: 4 VMs × ₹6,000/month = ₹24,000/month = ₹2,88,000/yr; storage and bandwidth ₹48,000/yr; part-time cloud admin ₹1,20,000/yr. Yearly = ₹4,56,000.
TCO = 5 × 4,56,000 = **₹22,80,000**.

Saving = 30,50,000 − 22,80,000 = **₹7,70,000 (about 25%)**, with zero up-front spending. Further saving comes from turning off idle VMs. Break-even is immediate because there is no capital outlay.

Conclusion: cloud provisioning is cheaper and faster for small businesses with variable demand.

Q3. Evaluate the importance of elasticity in cloud computing. Give a numerical model for traffic scaling from 10,000 to 1 million users in a video streaming service. (Jan 2025)

Importance of elasticity: it matches capacity to demand automatically (autoscaling), **avoids over-provisioning cost and under-provisioning failures**, keeps response time and SLA steady, and enables handling of unpredictable events such as a match or movie premiere.

Numerical model. Assumptions: one streaming server handles **5,000 concurrent users**; cost ₹10 per server-hour.

- At 10,000 users: servers = 10,000/5,000 = **2**.
- At 1,000,000 users: servers = 1,000,000/5,000 = **200**.

Daily pattern: peak of 1 million users for 4 hours; for the other 20 hours, about 20,000 users (4 servers).

- **Static provisioning for peak:** 200 servers × 24 h = 4,800 server-hours × ₹10 = **₹48,000/day**.
- **Elastic provisioning:** (200 × 4) + (4 × 20) = 800 + 80 = 880 server-hours × ₹10 = **₹8,800/day**.
- **Saving** = (48,000 − 8,800)/48,000 = **81.7%**.

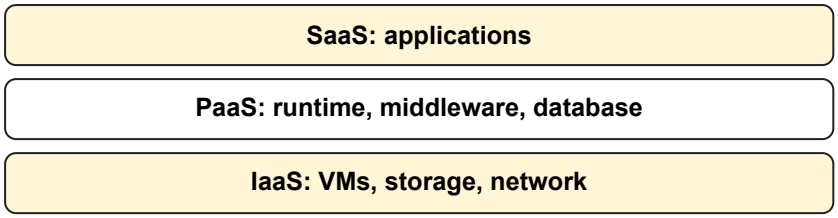
Scaling policy: add servers when CPU utilisation > 70% and remove when < 30%, with a scale-out lag of a few minutes (a load balancer distributes requests).

Conclusion: elasticity gives performance at peak and economy off-peak.

Q4. Discuss cloud computing delivery models with advantages and disadvantages. (2024)

Cloud delivery (service) models:

- **IaaS (Infrastructure as a Service):** virtual machines, storage, networks (AWS EC2, Azure VMs). The user manages OS, middleware and applications.
- **PaaS (Platform as a Service):** a platform to develop and deploy applications without managing infrastructure (Google App Engine, Heroku).
- **SaaS (Software as a Service):** ready-made applications through a browser (Gmail, Salesforce, Office 365).



Cloud service models

Model	Advantages	Disadvantages
IaaS	Full control, flexible, pay-per-use, no hardware	User handles OS and security patching, needs skill
PaaS	Faster development, automatic scaling, no infrastructure work	Vendor lock-in, limited runtime choices
SaaS	No installation, automatic updates, accessible anywhere	Less customisation, data held by the provider, needs Internet

Common advantages: low cost, scalability, rapid deployment. **Common disadvantages:** dependence on Internet, security and privacy risks, and downtime outside the user's control.

Q5. Explain on-demand provisioning and how it enhances scalability of cloud resources. (Feb 2024)

On-demand provisioning is the automatic allocation of computing resources at the moment a user requests them, through a self-service portal or API, without waiting for the provider's staff; resources are released and billing stops when no longer needed.

Types: static (pre-sized), **dynamic** (adjusts to load) and **user self-provisioning** (customer buys resources directly).

How it enhances scalability:

- **Vertical scaling:** increase CPU/RAM of a VM on request.
- **Horizontal scaling:** launch more VM instances behind a load balancer within minutes.
- **Elastic scale-in** when load drops, saving cost.
- A **resource pool** and virtualization make capacity appear unlimited, so applications grow smoothly without redesign.
- Autoscaling policies (CPU threshold) trigger provisioning automatically.

Benefits: no capital cost, fast time-to-market and better SLA compliance.

Q6. Explain the importance of Distributed Computing in cloud systems. List the essential characteristics of cloud computing as defined by NIST. (Feb 2024)

Distributed computing uses multiple networked computers that coordinate by message passing to work as one system. Its importance in the cloud:

- **Scalability:** add nodes to increase capacity.
- **Fault tolerance and availability:** work shifts when a node fails.
- **Resource sharing** and utilisation across data centres.
- **Performance:** parallel execution (MapReduce) and geographic proximity to users.
- **Cost effectiveness** using commodity hardware.
  Cloud storage, big data processing and microservices all rest on distributed principles.

NIST essential characteristics:

1. **On-demand self-service:** provision resources without human interaction.
2. **Broad network access:** available through standard devices over the network.
3. **Resource pooling:** multi-tenant pool dynamically assigned.
4. **Rapid elasticity:** scale out and in quickly.
5. **Measured service:** metering and pay-per-use billing.
