UNIT–4 : Resource Management and Security in Cloud (PYQ Solution)

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

SECTION A Answers (≈70 Words Each)

Q1. What is cloud storage and how is it different from traditional storage systems? (Jan 2026)

Cloud storage stores data on remote, provider-managed servers accessed over the Internet, billed by usage (S3, Google Drive). Differences: traditional storage uses owned local hardware (DAS, SAN, NAS) with up-front cost and fixed capacity, while cloud storage is elastic, pay-per-use, replicated, accessible anywhere and maintained by the provider; traditional storage gives more local control and works offline.

Q2. Discuss Identity and Access Management (IAM) in cloud environments. (Jan 2025)

IAM is the framework of policies and tools that ensures the right users and services have the right access to the right resources. It covers authentication (passwords, MFA, SSO), authorisation (roles, policies, least privilege), user and group management, federation and auditing. In the cloud (AWS IAM, Azure AD) it prevents unauthorised access and supports compliance.

Q3. Discuss encryption techniques for securing cloud resources. (Jan 2025)

Encryption converts data into unreadable ciphertext. Techniques: symmetric (AES, one shared key, fast, for bulk data), asymmetric (RSA/ECC, key pairs, for key exchange and signatures), hashing (SHA-256 for integrity), TLS/SSL for data in transit, encryption at rest (disk and database, key management

service/KMS), and advanced options like homomorphic encryption and client-side encryption where the customer holds the keys.

Q4. Define VM security. (2024)

VM security is the set of measures that protect virtual machines, the hypervisor and the data they hold from attack. It covers VM isolation, hypervisor hardening, patching of guest OS, secure images and snapshots, protection against VM escape, VM sprawl and side-channel attacks, network micro-segmentation, and access control to management consoles.

Q5. List the security threats in cloud computing. (2024)

- **Data breaches and data loss**
- **Account hijacking and insider threats**
- **Insecure APIs and interfaces**
- **Denial of Service (DoS/DDoS)**
- **Malware injection and VM escape attacks**
- **Misconfiguration and weak IAM**
- **Shared-technology vulnerabilities, vendor lock-in and compliance violations**

Q6. Explain autoscaling in cloud services. (Feb 2024)

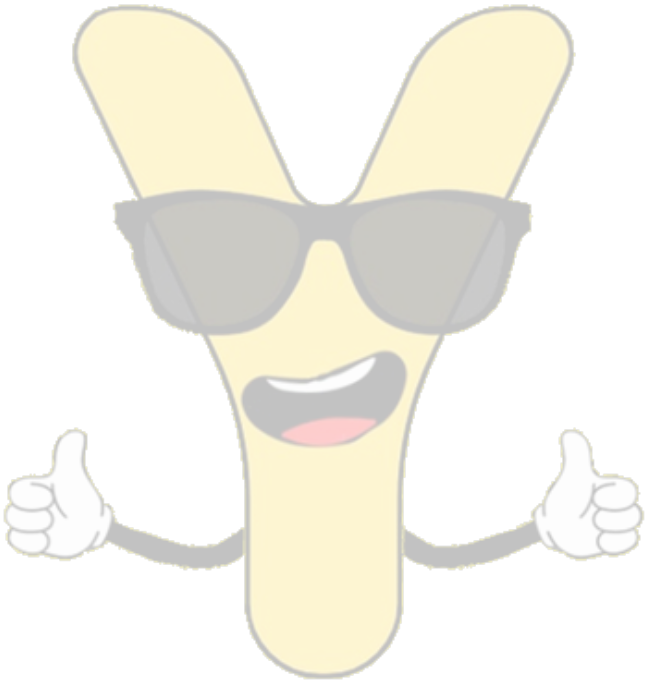
Autoscaling automatically adjusts the number of compute resources (VM instances) according to load, using rules or metrics such as CPU > 70% (scale out) or < 30% (scale in). A load balancer distributes traffic to new instances. Benefits: consistent performance during peaks, lower cost during low demand and high availability. Examples: AWS Auto Scaling, Azure VM Scale Sets.

Q7. Differentiate between monolithic and microservices architecture.

(Feb 2024)

Monolithic	Microservices
Single, tightly coupled codebase and deployable unit	Many small, independent services
Scaled as a whole	Each service scaled independently
One technology stack	Different stacks per service
One failure can crash everything	Failures isolated
Simple to start, hard to change at scale	Complex to operate but easy to evolve

---

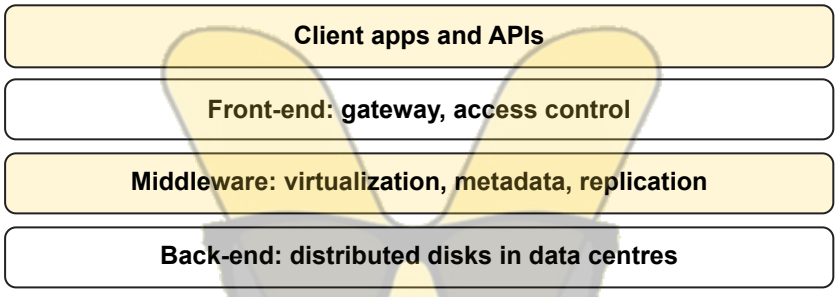


SECTION B Answers (≈250 Words Each)

Q1. Discuss cloud storage architecture and the advantages of Storage-as-a-Service. (Jan 2026)

Cloud storage architecture has these layers:

1. **Front-end:** client interface, REST/HTTP APIs, gateways and authentication.
2. **Middleware / management layer:** storage virtualization, metadata servers, data placement, replication, deduplication, encryption and billing.
3. **Back-end:** distributed physical storage across many commodity servers and disks in several data centres.



Cloud storage architecture

Data is stored as **object, file or block** storage, split into chunks, **replicated (typically 3 copies)** or erasure-coded, and spread across nodes and zones so a failure does not cause loss.

Advantages of STaaS:

- **Cost saving:** no hardware purchase; pay per GB.
- **Elastic scalability** on demand.
- **High durability and availability** through replication.
- **Remote access** and easy collaboration.
- **Backup and disaster recovery** built in.
- **No maintenance**, with automatic upgrades.
- **Security features**: encryption, access control.

Conclusion: distributed layered design gives cloud storage its reliability and scale.

Q2. Discuss inter-cloud resource management strategies. Simulate a scenario where multiple clouds collaborate during a high-demand

period. (Jan 2025)

Inter-cloud resource management coordinates resources among several cloud providers so that a cloud short of capacity can use others'.

Strategies:

- **Federation:** clouds agree to share resources under common policies.
- **Cloud bursting:** a private cloud overflows to a public one.
- **Multi-cloud / broker based:** a broker selects the cheapest or fastest provider.
- **Load balancing and migration** of VMs across clouds, with SLA-aware scheduling, pricing/market-based allocation, and standard APIs for interoperability.

Simulation: three clouds each with capacity 1,000 units. During a festival sale, demand is A = 1,200, B = 400, C = 500 (total 2,100).

- **Without collaboration:** A serves 1,000 and rejects 200. Served = 1,000 + 400 + 500 = 1,900, so success = 1,900/2,100 = **90.5%**.
- **With collaboration:** A forwards the 200 excess units, 100 to B and 100 to C (spare capacity 600 and 500). Served = **2,100 (100%)**; utilisation A 100%, B 50%, C 60%; average = 2,100/3,000 = **70%**.
- Gain: +9.5 percentage points in served demand and zero SLA violations. Cost: transfer and inter-cloud pricing.

Conclusion: collaboration absorbs peaks without each cloud over-provisioning.

Q3. Explain resource provisioning and resource provisioning methods. (2024)

Resource provisioning is the process of **allocating VMs, storage and network resources to applications** according to user demand and SLA, and releasing them when no longer needed.

Methods:

1. **Static (advance) provisioning:** a fixed amount is allocated in advance for the contract period. Simple, but wastes resources if demand varies.
2. **Dynamic provisioning:** resources are added or removed at run time as load changes (autoscaling). Efficient and cost saving.
3. **User self-provisioning:** customers request resources through a web portal/API with a credit card, immediately.

4. **Reserved vs on-demand vs spot:** reserved instances for long, steady use (discount), on-demand for short-term, **spot** for flexible cheap workloads.
5. **Predictive vs reactive:** predictive uses forecasts; reactive responds to thresholds.



Resource provisioning cycle

Challenges: over-provisioning (cost), under-provisioning (SLA violation) and provisioning delay.

Q4. Discuss the importance of Identity and Access Management (IAM) in cloud security. How does IAM manage user permissions and access control? (Feb 2024)

Importance: in the cloud, resources are reachable over the Internet, so **identity is the new perimeter**. IAM prevents unauthorised access, limits damage from stolen credentials, enforces compliance and provides audit trails.

How IAM works:

- **Authentication:** verifies identity using passwords, **multi-factor authentication**, certificates, SSO and federation (SAML, OAuth).
- **Authorisation and access control: role-based access control (RBAC)** groups permissions into roles (admin, developer, viewer); attribute-based and policy-based rules define who can do what on which resource; **least privilege** gives only needed rights.
- **User lifecycle:** creating, changing and deleting accounts and groups; temporary credentials.
- **Monitoring and audit:** logging (CloudTrail) of every access to detect misuse.

Example: an AWS IAM policy allows a "Developer" role to read and write S3 bucket "app-data" but not delete it.

Conclusion: IAM gives centralised, fine-grained control of who can access cloud resources.

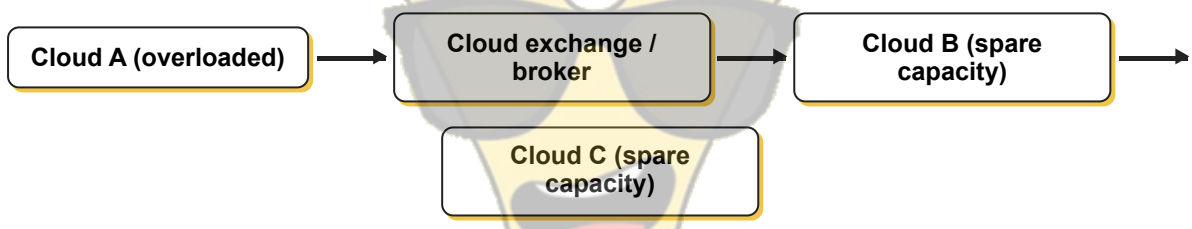
SECTION C Answers (≈500 Words Each)

Q1. Explain inter-cloud resource management and the challenges in global resource sharing. Analyze global exchange of cloud resources and propose a model to optimise allocation. (Jan 2026, Jan 2025, 2024)

Inter-cloud resource management is the coordination of resources across multiple independent clouds (federation, multi-cloud, cloud brokers) so providers can **lend, rent or exchange capacity** and users can obtain services across clouds.

Architecture:

- **Cloud exchange/market:** a directory and broker where providers publish capacity and price and users find the best match.
- **Cloud coordinator:** exports local services, monitors load and negotiates with peers.
- **Cloud broker:** selects providers and handles SLAs.



Global exchange of cloud resources

Strategies: load balancing, VM migration, cloud bursting, SLA-aware scheduling, price-based (auction) allocation and disaster failover.

Challenges: interoperability and standards (different APIs, VM formats), **security and trust** across domains, **SLA and pricing** differences, **data movement cost and latency**, **legal/regulatory** issues (data sovereignty), monitoring and billing consistency, and vendor lock-in.

Optimisation model: minimise Σ cost_i × x_i subject to Σ x_i = demand, x_i ≤ spare capacity_i and latency/SLA constraints. Greedy rule: allocate overflow to the cheapest cloud with spare capacity.

Efficiency calculation: clouds A, B, C each have capacity 1,000; demand 1,200, 400, 500. Without sharing, 1,900 of 2,100 demand is served (**90.5%**) with 200 units rejected. With the model, A sends 200 units to B and C; all 2,100 are served (**100%**), average utilisation rises from 63.3% (1,900/3,000) to **70%** and the served demand improves by 10.5% relative.

Conclusion: inter-cloud sharing raises utilisation and reliability but needs standards and trust frameworks.

Q2. Describe security challenges in cloud computing with examples. Evaluate them with a case study of a data breach in a SaaS platform. Discuss security governance. (Jan 2026, Jan 2025, 2024)

Security challenges:

- **Data breaches and data loss:** unauthorised access to stored data (misconfigured S3 buckets).
- **Multi-tenancy risks:** one tenant exploiting shared infrastructure (VM escape, side channels).
- **Account hijacking and insider threats:** stolen credentials, malicious admins.
- **Insecure APIs.**
- **DoS/DDoS attacks** affecting availability.
- **Compliance and legal issues:** data location, GDPR, audits.
- **Loss of control and vendor lock-in:** dependence on the provider's security.
- **Shadow IT and misconfiguration.**

Case study (generic SaaS breach): a SaaS CRM provider stored customer data with a weak admin API key committed in a public code repository. Attackers used it to query the tenant database and exported records of thousands of customers.
Causes: poor secret management, no MFA, excessive privilege, lack of monitoring.
Impact: privacy violation, regulatory fines, loss of trust and costly notification.
Lessons: rotate and vault secrets, enforce MFA and least privilege, encrypt data, monitor logs and have an incident response plan.

Security governance: the policies, roles and processes by which an organisation directs and controls cloud security.

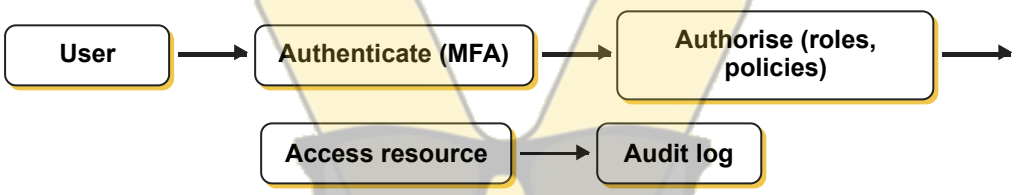
- Define **security policies, standards and roles** (CISO, data owners).
- **Risk assessment and compliance management** (ISO 27001, SOC 2).
- **SLA and contract review**, shared-responsibility definition.
- **Audit, monitoring, incident response and business continuity.**
- **Vendor management** and data classification.

Shared responsibility	Provider	Customer
Physical and hypervisor security	Yes	No
Data, identities, configuration	No	Yes

Conclusion: cloud security needs technology, governance and shared responsibility.

Q3. Explain IAM and its importance in securing cloud-based systems. (Jan 2026)

Identity and Access Management (IAM) controls who (users, applications, devices) can access which cloud resources and what they may do. Core functions are **authentication** (passwords, MFA, SSO, federation), **authorisation** (RBAC, policies, least privilege), **account lifecycle management** and **auditing/logging**.



IAM flow

Importance:

- Protects against **unauthorised access and credential theft**.
- Enforces **least privilege** and separation of duties.
- Centralised control across many services and accounts.
- Meets **compliance** requirements through audit trails.
- Enables **secure federation** with external identity providers and temporary credentials.
- Reduces insider risk.

Example: in AWS IAM, a role "ReadOnlyAuditor" can only list and read resources, while "Admin" can modify them, and MFA is required for sensitive actions.

Q4. Explain VM security issues and preventive measures in cloud environments. (Jan 2026)

VM security issues:

- **VM escape:** a guest breaks out and attacks the hypervisor or other VMs.

- **Hypervisor vulnerabilities** (a single compromise exposes all VMs).
- **VM sprawl:** unmanaged, unpatched VMs.
- **Insecure VM images and snapshots** containing secrets or malware.
- **Side-channel and cross-VM attacks** in multi-tenancy.
- **Inter-VM traffic** invisible to physical firewalls.
- **VM theft/copying** of image files and unauthorised migration.

Preventive measures:

1. Harden and **patch** the hypervisor and guest OSs.
2. Strong **isolation**, dedicated hosts for sensitive workloads.
3. **Secure, scanned VM images**; restrict snapshot access.
4. **Virtual firewalls and micro-segmentation**, IDS/IPS.
5. Strict **access control and MFA** for management consoles.
6. **Encrypt** disks and VM images; use secure boot.
7. **Monitoring, logging and lifecycle management** to avoid sprawl.

Q5. Explain the role of security standards in cloud computing and some widely adopted standards. (Feb 2024)

Role of standards: they give **common requirements and best practices**, create trust between providers and customers, enable compliance with law, support **auditing and certification**, and improve interoperability and risk management.

Widely adopted standards:

- **ISO/IEC 27001:** information security management system; **27017** (cloud security controls) and **27018** (PII protection in cloud).
- **SOC 1/2/3 (AICPA):** audit reports on controls and security.
- **NIST SP 800-53, NIST cloud guidelines:** security controls.
- **CSA CCM / STAR:** Cloud Security Alliance controls and registry.
- **PCI DSS:** payment card data.
- **HIPAA:** healthcare data (USA).
- **GDPR:** data protection (EU), and India's DPDP Act.
- **FedRAMP:** US government cloud authorisation.
- **TLS, SAML, OAuth, OpenID:** protocol standards for secure communication and identity.

Conclusion: adopting recognised standards demonstrates assurance and simplifies compliance.

Q6. Discuss the methods of resource provisioning in cloud environments. How do on-demand and reserved instances differ? (2024, Feb 2024)

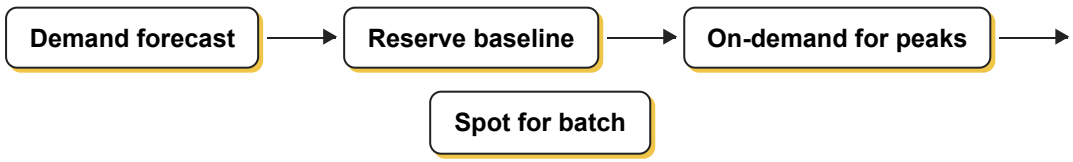
Resource provisioning methods:

- **Static provisioning:** capacity fixed in advance for predictable load; may waste or fall short.
- **Dynamic provisioning (autoscaling):** capacity changes in real time with load, using reactive thresholds or predictive models.
- **Self-service provisioning:** the user requests resources through a portal or API.
- **Advance reservation and spot provisioning:** capacity booked beforehand, or idle capacity bought cheaply.

On-demand vs reserved instances:

Feature	On-demand	Reserved
Commitment	None	1 or 3 years
Price	Higher per hour	Up to 40-70% cheaper
Payment	Pay per second/hour	Upfront or partial/monthly
Flexibility	Start and stop anytime	Fixed capacity commitment
Best for	Short, unpredictable workloads, testing	Steady, long-running workloads

Example: a server at ₹8/hour on demand costs 8 × 720 = ₹5,760/month; a reserved rate of ₹4.5/hour costs ₹3,240/month (44% less) if used continuously. A mix of reserved (baseline) and on-demand (peaks) is usually cheapest.



Provisioning mix
