100 CROSS-QUESTIONS

& ANSWERS

AWS Cloud Security | Topics 1-18

Trap questions, follow-ups & scenario drills - the questions an interviewer asks AFTER your first answer

Boond boond se ghada bharta hai - one question a drop, the pot fills.

HOW TO USE THIS

Cover the answer. Say your answer OUT LOUD in one or two sentences. Then compare. A cross-question tests the WHY behind a fact, the exception to a rule, or a broken scenario - so if your answer matches the fact but misses the reasoning, count it as wrong and re-read that topic. Target: 20 questions/day = done in 5 days, matching your study plan.

CLOUD BASICS (Q1-Q15)

Q1. You said cloud is pay-as-you-go. Then why do many companies get HIGHER bills than on-prem?

A. Elasticity without governance = waste: idle instances, over-provisioned sizes, forgotten volumes, data-transfer (egress) charges. Cloud is cheaper only when actively managed - cost optimisation is an ongoing job, not a default.

Q2. If AWS manages the infrastructure, why do cloud breaches still happen?

A. Because ~99% of cloud security failures are customer misconfiguration - public S3 buckets, leaked access keys, 0.0.0.0/0 security groups - not AWS infrastructure being hacked. Shared responsibility: your side fails, not theirs.

Q3. Is cloud ALWAYS more secure than on-premises?

A. No. AWS gives better physical security and tooling, but a misconfigured cloud account is weaker than a well-run data center. Security depends on the operator, not the location.

Q4. Scalability vs elasticity - are they the same?

A. No. Scalability = ability to grow. Elasticity = automatic grow AND shrink with demand. Elastic implies scalable; scalable does not imply elastic.

Q5. On-demand self-service is a benefit. How is it also a security risk?

A. Anyone with valid credentials can create resources in seconds - shadow IT, and cryptomining within minutes of a key leak. Speed of provisioning = speed of attack.

Q6. In SaaS the provider manages everything. So what can still go wrong on YOUR side?

A. Weak passwords, no MFA, over-shared documents/links, misconfigured permissions, insider misuse. Identity and data are always the customer's responsibility - in every model.

Q7. A client says: 'Move everything to cloud to save money.' What do you cross-question?

A. Which workloads? Steady 24/7 loads can be cheaper on-prem or on reserved capacity; spiky loads suit cloud. Also ask about migration cost, egress fees, and compliance constraints.

Q8. Multi-tenancy: my data sits on hardware shared with strangers. Risk?

A. Isolation is enforced by the hypervisor (AWS Nitro). Cross-tenant attacks are largely theoretical; if compliance demands it, use Dedicated Instances/Hosts. The practical risk remains your own misconfiguration.

Q9. RDS is PaaS and AWS patches the engine. Can you skip security work?

A. No - you still own network access (SG), IAM, encryption choice, credential management, and SQL injection in your own application code.

Q10. Why is IaaS more flexible but riskier than PaaS?

A. You control the OS - install anything, but also inherit patching, hardening, host firewall. More control = more responsibility = larger misconfiguration surface.

Q11. Someone says 'Lambda is SaaS because there are no servers.' Correct them.

A. Lambda is FaaS (serverless compute, under the PaaS umbrella) - you still write, secure, and IAM-scope the code. SaaS is a finished application like Gmail.

Q12. In which service model is a DATA breach the provider's responsibility?

A. Trick question - data responsibility is always yours in all three models. The provider answers only for breaches of the layers they manage.

Q13. What is 'measured service' and its security use?

A. Usage metering. Billing anomalies are a compromise signal - a sudden EC2 spike in a region you never use usually means stolen keys running miners.

Q14. Vertical vs horizontal scaling - which does cloud prefer and why?

A. Horizontal (add instances). It matches elasticity and load balancing and removes single points of failure. Vertical hits hardware ceilings and usually needs downtime.

Q15. What NEW security problem does hybrid cloud create?

A. The link itself (VPN/Direct Connect) becomes an attack path, identity must federate across two worlds, and inconsistent controls between on-prem and cloud create gaps attackers pivot through.

AWS BASICS - GLOBAL INFRASTRUCTURE (Q16-Q23)

Q16. Console vs CLI vs SDK - which is 'more secure'?

A. None inherently - all call the same SigV4-signed APIs gated by the same IAM. Security depends on credential handling: long-lived CLI keys on a laptop are riskier than a console session with MFA.

Q17. What is SigV4 and why should a security engineer care?

A. The request-signing scheme on every AWS API call: proves identity, integrity, and timestamp (blocks replay attacks). Invalid signature = rejected. It is why API security reduces to credential security.

Q18. Client stores Indian payment data in us-east-1. Problem?

A. Data residency - RBI requires payment data stored in India (ap-south-1). Region choice is a legal/compliance decision, not just latency.

Q19. Are IAM users regional?

A. No - IAM is global (also Route 53, CloudFront). EC2, VPC, EBS are regional. S3 buckets are regional but the bucket NAMESPACE is global.

Q20. One AZ fails. Does your app survive? Cross-question your own design.

A. Only if: instances in >=2 AZs, load balancer spanning AZs, data replicated (Multi-AZ DB, snapshots), and no single AZ-locked dependency (like one EBS volume).

Q21. Multi-AZ vs Multi-Region - one line each.

A. Multi-AZ = high availability against a data-center failure. Multi-Region = disaster recovery against a whole-region failure, plus latency and compliance drivers.

Q22. Why is ap-south-1a in your account NOT the same physical AZ as ap-south-1a in mine?

A. AWS randomizes AZ letters per account to spread load. For cross-account coordination use AZ IDs (e.g., aps1-az1), which are physical.

Q23. How do edge locations improve security, not just speed?

A. CloudFront/Route 53 absorb DDoS at 400+ edge POPs far from your origin, hide the origin IP, terminate TLS, and host AWS WAF rules.

VPC, SUBNET, ROUTE TABLE, IGW (Q24-Q45)

Q24. What exactly makes a subnet 'public'?

A. One thing only: its route table contains 0.0.0.0/0 -> IGW. Not a name, not a checkbox, not a setting on the subnet itself.

Q25. Instance in a public subnet but no public IP - reachable from internet?

A. No. It needs BOTH a public IP and the IGW route. The IGW performs 1:1 NAT; without a public IP there is no mapping to translate.

Q26. Can two VPCs have the same CIDR block?

A. Yes - they are isolated networks. It only becomes a problem when you try to peer/connect them: overlapping CIDRs cannot route to each other.

Q27. Why does AWS reserve 5 IPs per subnet?

A. Network address (.0), VPC router (.1), DNS (.2), reserved future use (.3), and broadcast (last). So a /24 gives 251 usable, not 254.

Q28. How many usable IPs in a /28 subnet?

A. 16 total minus 5 reserved = 11. /28 is also the smallest subnet size AWS allows.

Q29. Route table: 10.0.0.0/16 -> local, 0.0.0.0/0 -> IGW. Packet to 10.0.5.9 goes where?

A. Local. Longest-prefix match: /16 is more specific than /0. Destination route specificity decides, never rule order.

Q30. Can you delete the 'local' route?

A. No. It is implicit and immutable - it guarantees intra-VPC routing always works.

Q31. Can one subnet have two route tables?

A. No - exactly one route table per subnet. The reverse is fine: one route table can be associated with many subnets.

Q32. Private-subnet instance needs OS updates from the internet. Add an IGW route?

A. No - that makes it public. Put a NAT Gateway in a public subnet; the private route table gets 0.0.0.0/0 -> NAT. Outbound-only connectivity.

Q33. NAT Gateway vs IGW - the core difference?

A. IGW is two-way (inbound possible for instances with public IPs). NAT is outbound-only and hides instances behind the NAT's Elastic IP.

Q34. Attacker compromises the web EC2 in the public subnet. What limits lateral movement to the DB?

A. Layers: DB in private subnet (no IGW path), DB security group allowing only the web tier's SG on the DB port, NACLs, and least-privilege IAM. Blast-radius control by design.

Q35. Does a VPC span regions?

A. No - a VPC is regional and spans AZs within that region. Cross-region connectivity needs peering or Transit Gateway.

Q36. What is the bandwidth limit of an Internet Gateway?

A. Practically none - it is horizontally scaled and redundant; AWS imposes no bandwidth constraint and there is no availability risk for you to manage.

Q37. Why put the load balancer in the public subnet and the app in private?

A. The LB becomes the only internet-facing surface: a single choke point for TLS, WAF, and logging, while app instances stay unreachable directly. Minimal attack surface.

Q38. Two subnets, same VPC, different AZs - can they talk by default?

A. Yes - the implicit local route covers the entire VPC CIDR. To block it you need NACL/SG rules; routing cannot separate them.

Q39. What is an ENI, and why do security people care?

A. Elastic Network Interface - the virtual NIC. Security groups actually attach to ENIs. One instance can hold multiple ENIs; a dual-homed instance can accidentally bridge two networks.

Q40. 0.0.0.0/0 vs ::/0 - difference and the audit trap?

A. All IPv4 vs all IPv6. A rule 'open to the world' must be checked for BOTH; many audits find the IPv4 rule and miss the IPv6 twin.

Q41. Can a route send 0.0.0.0/0 to another EC2 instance instead of a gateway?

A. Yes - the NAT-instance / firewall-appliance pattern (requires disabling source/dest check on that instance). Legacy; NAT Gateway replaced it for most uses.

Q42. Can a subnet's CIDR be larger than the VPC's CIDR?

A. No. Every subnet must be a subset of the VPC CIDR, and subnets cannot overlap each other.

Q43. VPC A and B are peered. A reaches B, but B cannot reach A. Why?

A. Peering routes must be added in BOTH VPCs' route tables. B is missing its return route (or B's SG/NACL blocks the traffic). Peering itself is not routing.

Q44. Is the default VPC production-safe?

A. No - public subnets, auto-assign public IPs, permissive defaults. Build a custom VPC; many organisations delete the default VPC as a hardening step.

Q45. How would you DETECT someone attaching an unexpected IGW to a private VPC?

A. CloudTrail events CreateInternetGateway/AttachInternetGateway feeding an alert, or an AWS Config rule flagging IGW attachments. (Preview of your next topics - detection.)

SECURITY GROUP vs NACL (Q46-Q60)

Q46. SG allows inbound 443. Do you need an outbound rule for the response?

A. No - SGs are stateful. Connection tracking auto-allows return traffic of an allowed connection.

Q47. NACL allows ONLY inbound 443. Does HTTPS work?

A. No. NACLs are stateless: the response leaves on an ephemeral port (1024-65535) and needs an explicit OUTBOUND allow. The most classic NACL trap.

Q48. What are ephemeral ports and who chooses them?

A. Temporary client-side ports for return traffic, picked by the client OS - Linux ~32768-60999, Windows 49152-65535. Allow 1024-65535 in stateless rules to cover all.

Q49. NACL: rule 100 DENY 1.2.3.4, rule 200 ALLOW all. Is 1.2.3.4 blocked?

A. Yes - rules evaluate lowest number first, first match wins. Swap the numbers and the same IP gets in.

Q50. Can you write that DENY in a Security Group instead?

A. No - SGs have no DENY, only allow (plus implicit deny). NACL is the only VPC-native way to blacklist an IP at the network layer.

Q51. SG has: allow 22 from 10.0.0.0/16 AND allow 22 from 0.0.0.0/0. Effective result?

A. Open to the world. SG rules are a union - most permissive wins. There is no 'more specific rule overrides' logic in SGs.

Q52. You delete ALL rules from a security group. What happens?

A. All inbound denied (implicit deny); if outbound rules were removed too, outbound is denied as well. The instance is network-isolated.

Q53. Default NACL vs a NEW custom NACL - default behavior?

A. Default NACL: allow all in/out. Freshly created custom NACL: DENY everything until you add rules. Opposite defaults - classic exam trap.

Q54. DB SG allows 3306 'from sg-web'. Web instance IPs change. Does the rule break?

A. No - SG referencing tracks group MEMBERSHIP, not IPs. That is exactly why it beats CIDR-based rules for app-to-app traffic.

Q55. Can a SG rule reference a SG in another VPC?

A. Only across a VPC peering connection in the same region. You cannot reference arbitrary SGs in unrelated VPCs.

Q56. One instance, five security groups - how do rules combine?

A. Union of all ALLOW rules across every attached SG (default limit ~5 SGs per ENI). One permissive SG poisons the whole set.

Q57. Inbound was allowed by NACL and SG, but the response is blocked by the SG outbound rules. Possible?

A. No - trick question. A stateful SG never blocks the return traffic of a connection it allowed in. If responses fail, suspect the NACL.

Q58. Two instances in the SAME subnet talk to each other. Does the NACL filter it?

A. No - NACLs filter only traffic crossing the subnet boundary. Intra-subnet traffic is checked by SGs alone.

Q59. 'Port 22 open to 0.0.0.0/0, just temporarily for debugging.' Cross-examine that.

A. Scanners find it in minutes and brute-force starts immediately. Use SSM Session Manager (no inbound ports), or at minimum your own IP /32 with a time-boxed removal.

Q60. A packet is blocked inbound. Which control rejected it FIRST - NACL or SG?

A. NACL - it sits at the subnet edge and is evaluated before the packet ever reaches the ENI where the SG applies.

EC2 (Q61-Q72)

Q61. When you create a key pair, what is stored where?

A. AWS keeps only the PUBLIC key (injected into the instance's authorized_keys). You download the private key once; AWS cannot recover or re-issue it.

Q62. You lost the .pem file. How do you regain access?

A. Options: SSM Session Manager (if agent + role exist), EC2 Instance Connect, or the surgical route - stop instance, detach root EBS, attach to a rescue instance, add a new key to authorized_keys, reattach.

Q63. Why is an IAM role better than access keys stored on the instance?

A. Keys are long-lived and copyable - they leak into AMIs, snapshots, and repos. A role delivers short-lived STS credentials via metadata: auto-rotated, nothing on disk.

Q64. What is 169.254.169.254?

A. The link-local Instance Metadata Service (IMDS) endpoint - serves instance info and the role's temporary credentials. Reachable only from inside the instance.

Q65. Explain the Capital One breach as one attack chain.

A. SSRF flaw in a web app -> app tricked into calling the IMDSv1 credentials URL -> role credentials returned to attacker -> credentials used to list and copy S3 data. IMDSv2 plus a least-privilege role would have broken the chain twice.

Q66. HOW exactly does IMDSv2 stop that SSRF?

A. It requires a PUT request to obtain a session token, then that token as a header on every GET. Typical SSRF can only trigger plain GETs and cannot inject headers - the credentials URL now fails.

Q67. Are user-data scripts secret?

A. No - readable via the API (DescribeInstanceAttribute) and instance metadata. Never place passwords or keys in user data; pull them at runtime from Secrets Manager or SSM Parameter Store.

Q68. An instance app needs to read S3. Where does that permission live?

A. In an IAM role (instance profile) attached to the instance, with a policy allowing s3:GetObject on the specific bucket ARN. Not in the security group - SGs do network, IAM does API.

Q69. Your t3.micro slows down under sustained load. Attack?

A. No - burstable family economics. CPU credits: baseline performance plus credits for bursts; exhausted credits mean throttling (or extra cost in unlimited mode). A capacity symptom, not a compromise.

Q70. Stop vs terminate - what happens to data and IPs?

A. Stop: EBS persists, instance-store data lost, private IP kept, public IP released (unless Elastic IP). Terminate: instance destroyed; root EBS deleted by default (DeleteOnTermination=true).

Q71. Can two instances share one EBS volume?

A. Only io1/io2 Multi-Attach, same AZ, with a cluster-aware filesystem - a niche case. The standard answer: no; use EFS for shared file storage.

Q72. You build an AMI from a hardened instance. Any leak risk?

A. Yes - an AMI captures the whole disk: leftover keys, shell history, logs, cached secrets. Sanitize before imaging, and audit that no AMI is accidentally shared public.

EBS (Q73-Q80)

Q73. Encrypt an EXISTING unencrypted volume in place?

A. Impossible. The path is: snapshot -> copy the snapshot WITH encryption enabled -> create a new volume from the encrypted copy -> swap it in. Guaranteed interview question.

Q74. Snapshot of an encrypted volume - is it encrypted? Can it be made public?

A. Encrypted with the same key lineage, and KMS-encrypted snapshots CANNOT be made public - only shared with specific accounts that are also granted key access.

Q75. Snapshots are incremental. Why does deleting a middle snapshot not break restores?

A. Only changed blocks are stored, but AWS silently preserves any block still referenced by later snapshots. Every snapshot can independently restore the full volume.

Q76. Move an EBS volume from AZ-a to AZ-b?

A. Volumes are AZ-locked. Snapshot it (regional), then create a new volume from the snapshot in AZ-b.

Q77. gp3 vs gp2 - the one interview point?

A. gp3 decouples IOPS/throughput from size (3000 IOPS baseline, ~20% cheaper). gp2 ties IOPS to size (3 IOPS/GB). Migrating gp2 -> gp3 = same performance, lower cost.

Q78. A database needs a consistent 50,000 IOPS. Which volume type?

A. io2 (Block Express for the extreme end). gp3 caps at 16,000 IOPS, so it cannot serve this.

Q79. 'EBS encryption by default' - what does it actually cover?

A. All NEW volumes and snapshots in that region auto-encrypt: data at rest, data moving between instance and volume, and snapshots. Existing unencrypted volumes are untouched.

Q80. An attacker gains ec2:CreateSnapshot and ModifySnapshotAttribute. Risk?

A. They snapshot your disks and share the snapshot to their own account - full data theft without ever logging into the instance. Why snapshot permissions belong in least-privilege reviews.

S3 (Q81-Q90)

Q81. Bucket policy grants public read, but Block Public Access is ON. Result?

A. Blocked. BPA is the override kill-switch - public grants in policies/ACLs are simply ignored while it is on.

Q82. Your IAM policy allows s3:GetObject; the bucket policy explicitly denies you. Result?

A. Denied. Explicit deny beats any allow, across all policy types. No exceptions.

Q83. Policy grants s3:ListBucket on arn:aws:s3:::mybucket/* - does listing work?

A. No. ListBucket is a BUCKET-level action and needs the bucket ARN without /*. Object actions (GetObject) need the /* form. The #1 cause of mysterious AccessDenied.

Q84. 'My bucket name is random - nobody will find it.' Destroy that argument.

A. Bucket names leak via DNS, TLS certificates, logs, and code repos, and enumeration tooling is public. Obscurity is not access control; only policy is.

Q85. SSE-S3 vs SSE-KMS - when is KMS worth paying for?

A. KMS gives you your own key, a key policy (separate access gate), per-use audit in CloudTrail, and an instant kill-switch (disable the key). Use it for regulated or sensitive data.

Q86. Does encryption at rest protect against a leaked presigned URL?

A. No. Encryption is transparent to any AUTHORIZED request - and a valid presigned URL IS an authorized request until it expires.

Q87. A 7-day presigned URL leaked on day 1. Options?

A. Invalidate the signer: rotate/revoke the credentials or strip the permissions that signed it; or move/delete the object. You cannot expire one URL selectively - so design with short expiries.

Q88. Versioning is ON. An attacker deletes an object. Is data gone?

A. No - a delete marker is added; prior versions survive. Remove the marker to restore. MFA Delete or Object Lock protects even the versions themselves.

Q89. Bucket ACL vs bucket policy - modern answer?

A. ACLs are legacy, coarse-grained, and now disabled by default (Object Ownership: bucket-owner enforced). Use JSON bucket policies with conditions; keep ACLs off.

Q90. Force HTTPS-only access to a bucket - how?

A. Bucket policy with an explicit Deny where Condition aws:SecureTransport = false. Deny beats everything, so plain HTTP dies.

IAM & IAM POLICY (Q91-Q100)

Q91. One attached policy allows an action, another explicitly denies it. Result?

A. Deny. Evaluation order of strength: explicit deny > explicit allow > implicit deny. There is no policy precedence by attachment or order.

Q92. Why roles instead of IAM users for applications?

A. Roles issue temporary STS credentials - auto-expiring, auto-rotating, nothing stored, every assumption logged. A leaked user access key lives until a human rotates it.

Q93. What does a permission boundary do that a normal policy cannot?

A. It GRANTS nothing - it caps. Effective permissions = identity policy INTERSECT boundary. Used to let teams create IAM entities without escalating past the ceiling.

Q94. Why is iam:PassRole with Resource:* privilege escalation?

A. The user can pass ANY role - including an admin role - to a service they control (EC2, Lambda), then operate as that role. Always scope PassRole to specific role ARNs.

Q95. Can a group contain another group? Can a role join a group?

A. No and no. Groups hold users only - no nesting, no roles. A frequent one-word-trap in interviews.

Q96. Root account hardening - name three beyond 'enable MFA'.

A. Hardware MFA key (not just virtual), DELETE root access keys entirely, and a CloudTrail alarm on any root login. Bonus: verified security contact e-mail.

Q97. Policy: Action s3:* on arn:aws:s3:::mybucket (no /*). What is wrong - twice?

A. Under-scoped: object actions fail without the /* resource. Over-scoped: it still grants dangerous bucket actions like DeleteBucket and PutBucketPolicy. Both broken in one line.

Q98. Condition aws:MultiFactorAuthPresent - the subtle trap?

A. For access-key API calls the key is ABSENT, not 'false'. A plain Bool test misbehaves; the robust pattern is Deny with BoolIfExists: aws:MultiFactorAuthPresent = false.

Q99. Trust policy vs permission policy on a role?

A. Trust policy = WHO may assume the role (Principal). Permission policy = WHAT the role can do. Both required - and a trust policy with Principal:* means anyone can assume your role.

Q100. How do you FIND unused permissions to trim toward least privilege?

A. IAM Access Advisor ('service last accessed' data) plus IAM Access Analyzer policy generation from real CloudTrail activity. Least privilege is iterative pruning, never a one-shot guess.