

**1) A company is storing an access key with an access key ID and secret access key in a text file on a custom Amazon Machine Image (AMI). The company uses the access key to access Amazon DynamoDB tables from instances created from the AMI. The security team has mandated a more secure solution.**

**Which solution will meet the security team's mandate?**

- A) Put the access key in an Amazon S3 bucket and retrieve the access key on boot from the instance.
- B) Pass the access key to the instances through instance user data.
- C) Obtain the access key from a key server launched in a private subnet.
- D) Create an IAM role with permissions to access the table and launch all instances with the new role.

**2) A company is developing a highly available web application using stateless web servers.**

**Which AWS services or features are suitable for storing session state data? (Select TWO.)**

- A) Amazon CloudWatch
- B) Amazon DynamoDB
- C) Elastic Load Balancing
- D) Amazon ElastiCache
- E) AWS Storage Gateway

**3) Company salespeople upload their sales figures daily. A solutions architect needs a durable storage solution for these documents that also protects against users accidentally deleting important documents.**

**Which action will protect against unintended user actions?**

- A) Store data in an Amazon Elastic Block Store (Amazon EBS) volume and create snapshots once a week.
- B) Store data in an Amazon S3 bucket and enable versioning.
- C) Store data in two Amazon S3 buckets in different AWS Regions.
- D) Store data in Amazon EC2 instance storage.

**4) An application requires a highly available relational database with an initial storage capacity of 8 TB. The database will grow by 8 GB every day. To support expected traffic, at least eight read replicas will be required to handle database reads.**

**Which AWS service will meet these requirements?**

- A) Amazon DynamoDB
- B) Amazon S3
- C) Amazon Aurora
- D) Amazon Redshift

**5) A solutions architect is designing a critical business application with a relational database that runs on an Amazon EC2 instance. It requires a single Amazon Elastic Block Store (Amazon EBS) volume that can support up to 32,000 IOPS.**

**Which EBS volume type can meet the performance requirements of this application?**

- A) EBS Provisioned IOPS SSD (io1)
- B) EBS Throughput Optimized HDD (st1)
- C) EBS General Purpose SSD (gp2)
- D) EBS Cold HDD (sc1)

**6) A web application allows users to upload orders to an Amazon S3 bucket. The resulting Amazon S3 events trigger an AWS Lambda function that inserts a message to an Amazon SQS queue. A single Amazon EC2 instance reads messages from the queue, processes them, and stores them in an Amazon DynamoDB table partitioned by unique order ID. Next month, traffic is expected to increase by a factor of 10 and a solutions architect is reviewing the architecture for possible scaling problems.**

**Which component is MOST likely to need re-architecting to be able to scale to accommodate the new traffic?**

- A) Lambda function
- B) SQS queue
- C) EC2 instance
- D) DynamoDB table

**7) An application saves its logs to an Amazon S3 bucket. A user wants to keep the logs for 1 month for troubleshooting purposes, and then purge the logs.**

**What feature will enable this?**

- A) Adding a bucket policy to the S3 bucket
- B) Configuring lifecycle configuration rules for the S3 bucket
- C) Creating an IAM policy for the S3 bucket
- D) Enabling cross-origin resource sharing (CORS) on the S3 bucket

**8) An application running on Amazon EC2 instances processes sensitive information stored in Amazon S3. The information is accessed over the internet. The security team is concerned that the internet connectivity to Amazon S3 is a security risk.**

**Which solution will resolve the security concern?**

- A) Access the data through an internet gateway.
- B) Access the data through a VPN connection.
- C) Access the data through a NAT gateway.
- D) Access the data through a VPC endpoint for Amazon S3.

**9) A company is building an Amazon Redshift cluster in its shared services VPC. The cluster will host sensitive data.**

**How can the company control which networks can access the cluster?**

- A) Run the cluster in a different VPC and connect through VPC peering.
- B) Create a database user inside the Amazon Redshift cluster for users on the network only.
- C) Define a cluster security group for the cluster that allows access from the allowed networks.
- D) Only allow access to networks that connect with the shared services network using VPN.

**10) A solutions architect is designing an online shopping application running in a VPC on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The application tier must read and write data to a customer-managed database cluster. There should be no access to the database from the internet, but the cluster must be able to obtain software patches from the internet.**

**Which VPC design meets these requirements?**

- A) Public subnets for both the application tier and the database cluster
- B) Public subnets for the application tier and private subnets for the database cluster
- C) Public subnets for the application tier and NAT gateway, and private subnets for the database cluster
- D) Public subnets for the application tier, and private subnets for the database cluster and NAT gateway

**Answers**

- 1) D – IAM roles for EC2 instances allow applications running on the instance to access AWS resources without having to create and store any access keys. Any solution involving the creation of an access key then introduces the complexity of managing that secret.
- 2) B, D – Both DynamoDB and ElastiCache provide high-performance storage of key-value pairs. CloudWatch and Elastic Load Balancers are not storage services. Storage Gateway is a storage service, but it is a hybrid storage service that enables on-premises applications to use cloud storage.
- 3) B – If a versioned object is deleted, then it can still be recovered by retrieving the final version. Response A would lose any changes committed since the previous snapshot. Storing the data in two S3 buckets would provide slightly more protection, but a user could still delete the object from both buckets. EC2 instance storage is ephemeral and should never be used for data requiring durability.
- 4) C – Aurora is a relational database that will automatically scale to accommodate data growth. Amazon Redshift does not support read replicas and will not automatically scale. DynamoDB is a NoSQL service, not a relational database. Amazon S3 is object storage, not a relational database.
- 5) A – EBS Provisioned IOPS SSD provides sustained performance for mission-critical low-latency workloads. EBS General Purpose SSD can provide bursts of performance up to 3,000 IOPS and have a maximum baseline performance of 16,000 IOPS. The two HDD options are lower cost, high throughput volumes.
- 6) C – A single EC2 instance will not scale and is a single point of failure in the architecture. A much better solution would be to have EC2 instances in an Auto Scaling group across two Availability Zones read messages from the queue. The other responses are all managed services that can be configured to scale or will scale automatically.
- 7) B – Lifecycle configuration allows lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. Bucket policies and IAM define access to objects in an S3 bucket. CORS enables clients in one domain to interact with resources in a different domain.
- 8) D – VPC endpoints for Amazon S3 provide secure connections to S3 buckets that do not require a gateway or NAT instances. NAT gateways and internet gateways still route traffic over the internet to the public endpoint for Amazon S3. There is no way to connect to Amazon S3 using a VPN.
- 9) C – A security group can grant access to traffic from the allowed networks using the CIDR range for each network. VPC peering and VPN are connectivity services and cannot control traffic for security. Amazon Redshift user accounts address authentication and authorization at the user level and have no control over network traffic.
- 10) C – The online application must be in public subnets to allow access from clients' browsers. The database cluster must be in private subnets to meet the requirement that there be no access from the internet. A NAT gateway is required to give the database cluster the ability to download patches from the internet. NAT gateways must be deployed in public subnets.