Answer Sheet

Q1) You would like to share some documents with public users accessing an S3 bucket over the Internet. What are two valid methods of granting public read permissions so you can share the documents? (choose 2)

✅ Grant public read access to the objects when uploading

Explanation:-Access policies define access to resources and can be associated with resources (buckets and objects) and users You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. Bucket policies can be used to grant permissions to objects You can define permissions on objects when uploading and at any time afterwards using the AWS Management Console. You cannot use a bucket ACL to grant permissions to objects within the bucket. You must explicitly assign the permissions to each object through an ACL attached as a subresource to that object Using an EC2 instance as a bastion host to share the documents is not a feasible or scalable solution You can configure an S3 bucket as a static website and use CloudFront as a front-end however this is not necessary just to share the documents and imposes some constraints on the solution. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

⚪ Share the documents using CloudFront and a static website

✅ Use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket granting read access to public anonymous users

Explanation:-Access policies define access to resources and can be associated with resources (buckets and objects) and users You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. Bucket policies can be used to grant permissions to objects You can define permissions on objects when uploading and at any time afterwards using the AWS Management Console. You cannot use a bucket ACL to grant permissions to objects within the bucket. You must explicitly assign the permissions to each object through an ACL attached as a subresource to that object Using an EC2 instance as a bastion host to share the documents is not a feasible or scalable solution You can configure an S3 bucket as a static website and use CloudFront as a front-end however this is not necessary just to share the documents and imposes some constraints on the solution. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

⚪ Grant public read on all objects using the S3 bucket ACL

Q2)

A Solutions Architect is designing an authentication solution using the AWS STS that will provide temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users).

What supported sources are available to the Architect for users? (choose 2)

✅ OpenID Connect

Explanation:-The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users) Federation can come from three sources: - Federation (typically AD) - Federation with Mobile Apps (e.g. Facebook, Amazon, Google or other Open ID providers) - Cross account access (another AWS account) The question has asked for supported sources for users. Cognito user pools contain users, but identity pools do not You cannot use STS with local users on a PC or an EC2 instance References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

⚪ EC2 instance

✅ Another AWS account

Explanation:-The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users) Federation can come from three sources: - Federation (typically AD) - Federation with Mobile Apps (e.g. Facebook, Amazon, Google or other Open ID providers) - Cross account access (another AWS account) The question has asked for supported sources for users. Cognito user pools contain users, but identity pools do not You cannot use STS with local users on a PC or an EC2 instance References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

⚪ A local user on a user's PC

Q3)

You are building an application that will collect information about user behavior. The application will rapidly ingest large amounts of dynamic data and requires very low latency. The database must be scalable without incurring downtime.

Which database would you recommend for this scenario?

⚪ RDS with MySQL

✅ DynamoDB

Explanation:-Amazon Dynamo DB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability Push button scaling means that you can scale the DB at any time without incurring downtime DynamoDB provides low read and write latency RDS uses EC2 instances so you have to change your instance type/size in order to scale compute vertically RedShift uses EC2 instances as well so you need to choose your instance type/size for scaling compute vertically, but you can also scale horizontally by adding more nodes to the cluster Rapid ingestion of dynamic data is not an ideal use case for RDS or RedShift References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/

⚪ RedShift

⚪ RDS with Microsoft SQL

Q4) A Solutions Architect is building a complex application with several back-end APIs. The architect is considering using Amazon API Gateway. With Amazon API Gateway what are features that assist with creating and managing APIs? (Choose 2)

✅ You can define plans that meter and restrict third-party developer access to APIs

Explanation:-Metering – define plans that meter and restrict third-party developer access to APIs Lifecycle Management – Operate multiple API versions and multiple stages for each version simultaneously so that existing applications can continue to call previous versions after new API versions are published References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-api-gateway/

⚪ Flexible message delivery over multiple transport protocols

⚪ You can define the maintenance window or AWS will schedule a 30 minute window

✅ You can operate multiple API versions and multiple stages for each version simultaneously

Explanation:-Metering – define plans that meter and restrict third-party developer access to APIs Lifecycle Management – Operate multiple API versions and multiple stages for each version simultaneously so that existing applications can continue to call previous versions after new API versions are published References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-api-gateway/

Q5)

Your company would like to restrict the ability of most users to change their own passwords whilst continuing to allow a select group of users within specific user groups.

What is the best way to achieve this? (choose 2)

✅ Under the IAM Password Policy deselect the option to allow users to change their own passwords

Explanation:-A password policy can be defined for enforcing password length, complexity etc. (applies to all users) You can allow or disallow the ability to change passwords using an IAM policy and you should attach this to the group that contains the users, not to the individual users themselves You cannot use an IAM role to perform this function The AWS STS is not used for controlling password policies. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

✅ Create an IAM Policy that grants users the ability to change their own password and attach it to the groups that contain the users

Explanation:-A password policy can be defined for enforcing password length, complexity etc. (applies to all users) You can allow or disallow the ability to change passwords using an IAM policy and you should attach this to the group that contains the users, not to the individual users themselves You cannot use an IAM role to perform this function The AWS STS is not used for controlling password policies References: https://digitalcloud.training/certification-

training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

- Create an IAM Role that grants users the ability to change their own password and attach it to the groups that contain the users
- Disable the ability for all users to change their own passwords using the AWS Security Token Service

---

Q6)

A colleague from your company's IT Security team has notified you of an Internet-based threat that affects a certain port and protocol combination. You have conducted an audit of your VPC and found that this port and protocol combination is allowed on an Inbound Rule with a source of 0.0.0.0/0. You have verified that this rule only exists for maintenance purposes and need to make an urgent change to block the access.

What is the fastest way to block access from the Internet to the specific ports and protocols?

- You don't need to do anything; this rule will only allow access to VPC based resources
- ✅ Update the security group by removing the rule

Explanation:-Security group membership can be changed whilst instances are running Any changes to security groups will take effect immediately You can only assign permit rules in a security group, you cannot assign deny rules If you delete the security you will remove all rules and potentially cause other problems You do need to make the update, as it's the VPC based resources you're concerned about. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- Delete the security group
- Add a deny rule to the security group with a higher priority

---

Q7)

You are a Solutions Architect at Digital Cloud Training. One of your clients has requested that you design a solution for distributing load across a number of EC2 instances across multiple AZs within a region. Customers will connect to several different applications running on the client's servers through their browser using multiple domain names and SSL certificates. The certificates are stored in AWS Certificate Manager (ACM).

What is the optimal architecture to ensure high availability, cost effectiveness, and performance?

- Launch a single ALB and bind multiple SSL certificates to multiple secure listeners
- ✅ Launch a single ALB and bind multiple SSL certificates to the same secure listener. Clients will use the Server Name Indication (SNI) extension

Explanation:-You can use a single ALB and bind multiple SSL certificates to the same listener With Server Name Indication (SNI) a client indicates the hostname to connect to. SNI supports multiple secure websites using a single secure listener You cannot have the same port in multiple listeners so adding multiple listeners would not work. Also, when using standard HTTP/HTTPS the port will always be 80/443 so you must be able to receive traffic on the same ports for multiple applications and still be able to forward to the correct instances. This is where host-based routing comes in With host-based routing you can route client requests based on the Host field (domain name) of the HTTP header allowing you to route to multiple domains from the same load balancer (and share the same listener) You do not need multiple ALBs and it would not be cost-effective. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Launch multiple ALBs and bind separate SSL certificates to each ELB
- Launch a single ALB, configure host-based routing for the domain names and bind an SSL certificate to each routing rule

---

Q8)

A Linux instance running in your VPC requires some configuration changes to be implemented locally and you need to run some commands.

Which of the following can be used to securely connect to the instance?

- EC2 password
- ✅ Key pairs

Explanation:-A key pair consists of a public key that AWS stores, and a private key file that you store For Windows AMIs, the private key file is required to obtain the password used to log into your instance For Linux AMIs, the private key file allows you to securely SSH into your instance The "EC2 password" might refer to the operating system password. By default you cannot login this way to Linux and must use a key pair. However, this can be enabled by setting a password and updating the /etc/ssh/sshd_config file You cannot login to an EC2 instance using certificates/public keys, References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

- Public key
- SSL/TLS certificate

Q9)

One of your EC2 instances runs an application process that saves user data to an attached EBS volume. The EBS volume was attached to the EC2 instance after it was launched and is unencrypted. You would like to encrypt the data that is stored on the volume as it is considered sensitive however you cannot shutdown the instance due to other application processes that are running.

What is the best method of applying encryption to the sensitive data without any downtime?

- Create an encrypted snapshot of the current EBS volume. Restore the snapshot to the EBS volume
- ✅ Create and mount a new encrypted EBS volume. Move the data to the new volume and then delete the old volume

Explanation:-There is no direct way to change the encryption state of a volume. Either create an encrypted volume and copy data to it or take a snapshot, encrypt it, and create a new encrypted volume from the snapshot.

- Unmount the volume and enable server-side encryption. Re-mount the EBS volume
- Leverage the AWS Encryption CLI to encrypt the data on the volume

---

Q10)

You are a Solutions Architect at Digital Cloud Training. A client has requested a design for a highly-available, fault tolerant architecture for the web and app tiers of a three-tier application. The requirements are as follows:

- Web instances will be in a public subnet and app instances will be in a private subnet

- Connections to EC2 instances should be automatically distributed across AZs

- A minimum of 12 web server EC2 instances must be running at all times

- A minimum of 6 app server EC2 instances must be running at all times

- The failure of a single availability zone (AZ) must not affect The availability of The application or result in a reduction of capacity beneath The stated requirements

Which of the following design options would be the most suitable and cost-effective solution?

- ✅ One Auto Scaling Group using 3 AZs and a minimum of 18 EC2 instances behind an Internet facing ALB for the web layer. One Auto Scaling Group using 3 AZs and a minimum of 9 EC2 instances behind an internal-only ALB for the app layer

Explanation:-Simple scaling maintains a current number of instances, you can manually change the ASGs min/desired/max and attach/detach instances. Auto Scaling will try to distribute EC2 instances evenly across AZs In this scenario you must have a minimum of 12 instances running in the event of an AZ failure, therefore with 18 instances across 3 AZs if one AZ fails you still have enough instances ELBs can be Internet-facing or internal-only. Remember that internet-facing ELBs have public IPs, whereas internal-only ELBs have private IPs have public IPs. Therefore, you must have 2 ELBs, one for the web layer and one for the app layer. Otherwise the web layer would have to hairpin the traffic back to the public IP of the ELB rather than forwarding it to the internal ELB and this is not a supported configuration. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- One Auto Scaling Group using 3 AZs and a minimum of 12 EC2 instances behind an Internet facing ALB for the web layer. One Auto Scaling Group using 3 AZs and a minimum of 6 EC2 instances behind an internal-only ALB for the app layer

- One Auto Scaling Group with a minimum of 18 EC2 instances for the web layer. One Auto Scaling Group using 3 AZs and a minimum of 9 EC2 instances for the app layer. A single Internet-facing ALB using 3 AZs and two target groups for the web and app layers
- One Auto Scaling Group with a minimum of 12 EC2 instances for the web layer. One Auto Scaling Group using 3 AZs and a minimum of 6 EC2 instances for the app layer. A single Internet-facing ALB using 3 AZs and two target groups for the web and app layers

---

Q11)

A customer has asked you to recommend the best solution for a highly available database. The database is a relational OLTP type of database and the customer does not want to manage the operating system the database runs on. Failover between AZs must be automatic.

Which of the below options would you suggest to the customer?

- Use DynamoDB
- ✅ Use RDS in a Multi-AZ configuration

Explanation:-Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. With RDS you can configure Multi-AZ which creates a replica in another AZ and synchronously replicates to it (DR only) RedShift is used for analytics OLAP not OLTP If you install a DB on an EC2 instance you will need to manage to OS yourself and the customer wants it to be managed for them DynamoDB is a managed database of the NoSQL type. NoSQL DBs are not relational DBs. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

- Install a relational database on EC2 instances in multiple AZs and create a cluster
- Use RedShift in a Multi-AZ configuration

Q12)

You are troubleshooting a connectivity issue where you cannot connect to an EC2 instance in a public subnet in your VPC from the Internet.

Which of the configuration items in the list below would you check first? (choose 2)

- ✅ The subnet has “Auto-assign public IPv4 address" set to “Yes"

Explanation:-Public subnets are subnets that have: “Auto-assign public IPv4 address?? set to “Yes?? which will assign a public IP The subnet route table has an attached Internet Gateway The instance will also need to a security group with an inbound rule allowing the traffic EC2 instances always have a private IP address assigned. When using a public subnet with an Internet Gateway the instance needs a public IP to be addressable from the Internet NAT gateways are used to enable outbound Internet access for instances in private subnets. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

- There is a NAT Gateway installed in the subnet
- The subnet route table has an attached NAT Gateway
- ✅ The security group attached to the EC2 instance has an inbound rule allowing the traffic

Explanation:-Public subnets are subnets that have: “Auto-assign public IPv4 address?? set to “Yes?? which will assign a public IP The subnet route table has an attached Internet Gateway The instance will also need to a security group with an inbound rule allowing the traffic EC2 instances always have a private IP address assigned. When using a public subnet with an Internet Gateway the instance needs a public IP to be addressable from the Internet NAT gateways are used to enable outbound Internet access for instances in private subnets. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

Q13)

You would like to provide some on-demand and live streaming video to your customers. The plan is to provide the users with both the media player and the media files from the AWS cloud. One of the features you need is for the content of the media files to begin playing while the file is still being downloaded.

What AWS services can deliver these requirements? (choose 2)

- ✅ Use CloudFront with a Web and RTMP distribution

Explanation:-For serving both the media player and media files you need two types of distributions: - A web distribution for the media player - An RTMP distribution for the media files RTMP: - Distribute streaming media files using Adobe Flash Media Server’s RTMP protocol - Allows an end user to begin playing a media file before the file has finished downloading from a CloudFront edge location - Files must be stored in an S3 bucket (not an EBS volume or EC2 instance). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-cloudfront/

- Use CloudFront with an RTMP distribution
- Store the media files on an EC2 instance
- ✅ Store the media files in an S3 bucket

Explanation:-For serving both the media player and media files you need two types of distributions: - A web distribution for the media player - An RTMP distribution for the media files RTMP: - Distribute streaming media files using Adobe Flash Media Server’s RTMP protocol - Allows an end user to begin playing a media file before the file has finished downloading from a CloudFront edge location - Files must be stored in an S3 bucket (not an EBS volume or EC2 instance). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-cloudfront/

Q14)

There is a new requirement to implement in-memory caching for a Financial Services application due to increasing read-heavy load. The data must be stored persistently. Automatic failover across AZs is also required.

Which two items from the list below are required to deliver these requirements? (choose 2)

- ✅ ElastiCache with the Redis engine

Explanation:-Redis engine stores data persistently Memached engine does not store data persistently Redis engine supports Multi-AZ using read replicas in another AZ in the same region You can have a fully automated, fault tolerant ElastiCache-Redis implementation by enabling both cluster mode and multi-AZ failover Memcached engine does not support Multi-AZ failover or replication. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-elasticache/

- ElastiCache with the Memcached engine
- ✅ Multi-AZ with Cluster mode and Automatic Failover enabled

Explanation:-Redis engine stores data persistently Memached engine does not store data persistently Redis engine supports Multi-AZ using read replicas in another AZ in the same region You can have a fully automated, fault tolerant ElastiCache-Redis implementation by enabling both cluster mode and multi-AZ failover Memcached engine does not support Multi-AZ failover or replication. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-elasticache/

- Multiple nodes placed in different AZs

Q15) A Solutions Architect is designing a data archive strategy using Amazon Glacier. The Architect needs to explain the features of the service to his manager, which statements about Glacier are correct? (choose 2)

- ✅ Glacier objects are visible through S3 only

Explanation:-Glacier objects are visible through S3 only (not Glacier directly) The contents of an archive that has been uploaded cannot be modified Uploading archives is synchronous Downloading archives is asynchronous Retrieval can take a few hours. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

- The contents of an archive can be modified after uploading
- ✅ Uploading archives is synchronous; downloading archives is asynchronous

Explanation:-Glacier objects are visible through S3 only (not Glacier directly) The contents of an archive that has been uploaded cannot be modified Uploading archives is synchronous Downloading archives is asynchronous Retrieval can take a few hours. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

- Retrieval is immediate

Q16)

The association between a poll-based source and a Lambda function is called the event source mapping. Event sources maintain the mapping configuration except for stream-based services such as _____ and _____ for which the configuration is made on the Lambda side and Lambda performs the polling.

Fill in the blanks from the options below (choose 2)

- ✅ DynamoDB

Explanation:-Event sources are mapped to Lambda functions Event sources maintain the mapping configuration except for stream-based services (e.g. DynamoDB, Kinesis) for which the configuration is made on the Lambda side and Lambda performs the polling This question is really just asking you to identify which of the listed services are stream-based services. DynamoDB and Kinesis are both used for streaming data. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-route-53/

- S3
- IoT Button
- ✅ Kinesis

Explanation:-Event sources are mapped to Lambda functions Event sources maintain the mapping configuration except for stream-based services (e.g. DynamoDB, Kinesis) for which the configuration is made on the Lambda side and Lambda performs the polling This question is really just asking you to identify which of the listed services are stream-based services. DynamoDB and Kinesis are both used for streaming data. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-route-53/

Q17)

The data scientists in your company are looking for a service that can process and analyze real-time, streaming data. They would like to use standard SQL queries to query the streaming data.

Which combination of AWS services would deliver these requirements?

- DynamoDB and EMR
- ✅ Kinesis Data Streams and Kinesis Data Analytics

Explanation:-Kinesis Data Streams enables you to build custom applications that process or analyze streaming data for specialized needs Amazon Kinesis Data Analytics is the easiest way to process and analyze real-time, streaming data. Kinesis Data Analytics can use standard SQL queries to process Kinesis data streams and can ingest data from Kinesis Streams and Kinesis Firehose but Firehose cannot be used for running SQL queries DynamoDB is a NoSQL database that can be used for storing data from a stream but cannot be used to process or analyze the data or to query it with SQL queries. Elastic Map Reduce (EMR) is a hosted Hadoop framework and is not used for analytics on streaming data. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-kinesis/

- ElastiCache and EMR
- Kinesis Data Streams and Kinesis Firehose

Q18)

You are a Solutions Architect at a media company and you need to build an application stack that can receive customer comments from sporting events. The application is expected to receive significant load that could scale to millions of messages within a short space of time following high-profile matches.

As you are unsure of the load required for the database layer what is the most cost-effective way to ensure that the messages are not dropped?

- Use RDS Auto Scaling for the database layer which will automatically scale as required
- ✅ Create an SQS queue and modify the application to write to the SQS queue. Launch another application instance the polls the queue and writes messages to the database

Explanation:-Amazon Simple Queue Service (Amazon SQS) is a web service that gives you access to message queues that store messages waiting to be processed. SQS offers a reliable, highly-scalable, hosted queue for storing messages in transit between computers and is used for distributed/decoupled applications This is a great use case for SQS as the messages you don't have to over-provision the database layer or worry about messages being dropped RDS Auto Scaling does not exist. With RDS you have to select the underlying EC2 instance type to use and pay for that regardless of the actual load on the DB With DynamoDB there are now 2 pricing options: - Provisioned capacity has been around forever and is one of the incorrect answers to this question. With provisioned capacity you have to specify the number of read/write capacity units to provision and pay for these regardless of the load on the database. - With the the new On-demand capacity mode DynamoDB is charged based on the data reads and writes your application performs on your tables. You do not need to specify how much read and write throughput you expect your application to perform because DynamoDB instantly accommodates your workloads as they ramp up or down. it might be a good solution to this question but is not an available option. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/application-integration/amazon-sqs/

- Write the data to an S3 bucket, configure RDS to poll the bucket for new messages
- Use DynamoDB and provision enough write capacity to handle the highest expected load

---

Q19)

You are a Solutions Architect at Digital Cloud Training. A large multi-national client has requested a design for a multi-region, multi-master database. The client has requested that the database be designed for fast, massively scaled applications for a global user base. The database should be a fully managed service including the replication.

Which AWS service can deliver these requirements?

- RDS with Multi-AZ
- S3 with Cross Region Replication
- ✅ DynamoDB with Global Tables and Cross Region Replication

Explanation:-Cross-region replication allows you to replicate across regions: - Amazon DynamoDB global tables provides a fully managed solution for deploying a multi-region, multi-master database - When you create a global table, you specify the AWS regions where you want the table to be available - DynamoDB performs all of the necessary tasks to create identical tables in these regions, and propagate ongoing data changes to all of them RDS with Multi-AZ is not multi-master (only one DB can be written to at a time), and does not span regions S3 is an object store not a multi-master database There is no such thing as EBS replication. You could build your own database stack on EC2 with DB-level replication but that is not what is presented in the answer. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/

- EC2 instances with EBS replication

Q20)

The application development team in your company has a new requirement for the deployment of a container solution. You plan to use the AWS Elastic Container Service (ECS). The solution should include load balancing of incoming requests across the ECS containers and allow the containers to use dynamic host port mapping so that multiple tasks from the same service can run on the same container host.

Which AWS load balancing configuration will support this?

- ✅ Use an Application Load Balancer (ALB) and map the ECS service to the ALB

Explanation:-It is possible to associate a service on Amazon ECS to an Application Load Balancer (ALB) for the Elastic Load Balancing (ELB) service An Application Load Balancer allows dynamic port mapping. You can have multiple tasks from a single service on the same container instance. The Classic Load Balancer requires that you statically map port numbers on a container instance. You cannot run multiple copies of a task on the same instance, because the ports would conflict An NLB does not support host-based routing (ALB only), and this would not help anyway. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ecs/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Use a Classic Load Balancer (CLB) and create a static mapping of the ports
- Use a Network Load Balancer (NLB) and host-based routing
- You cannot run multiple copies of a task on the same instance, because the ports would conflict

---

Q21) To improve security in your AWS account you have decided to enable multi-factor authentication (MFA). You can authenticate using an MFA device in which two ways? (choose 2)

- Locally to EC2 instances
- ✅ Through the AWS Management Console

Explanation:-You can authenticate using an MFA device in the following ways: Through the AWS Management Console – the user is prompted for a user name, password and authentication code Using the AWS API – restrictions are added to IAM policies and developers can request temporary security credentials and pass MFA parameters in their AWS STS API requests Using the AWS CLI by obtaining temporary security credentials from STS (aws sts get-session-token). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

- Using a key pair
- ✅ Using the AWS API

Explanation:-You can authenticate using an MFA device in the following ways: Through the AWS Management Console – the user is prompted for a user name, password and authentication code Using the AWS API – restrictions are added to IAM policies and developers can request temporary security credentials and pass MFA parameters in their AWS STS API requests Using the AWS CLI by obtaining temporary security credentials from STS (aws sts get-session-token). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-iam/

Q22)

An application that was recently moved into the AWS cloud has been experiencing some authentication issues. The application is currently configured to authenticate to an on-premise Microsoft Active Directory Domain Controller via a VPN connection. Upon troubleshooting the issues, it seems that latency across the VPN connection is causing authentication to fail. Your company is very cost sensitive at the moment and the administrators of the Microsoft AD do not want to manage any additional directories. You need to resolve the issues quickly.

What is the best solution to solve the authentication issues taking cost considerations into account?

- Create an AWS Direct Connect connection to reduce the latency between your company and AWS
- Use the AWS Active Directory Service for Microsoft Active Directory and join your existing on-premise domain
- ✅ Install an additional Microsoft Active Directory Domain Controller for your existing domain on EC2 and configure the application to authenticate to the local DC

Explanation:-Direct Connect is an incorrect option as it can take months to provision and a quick resolution has been requested The best answer is to Install an additional Microsoft Active Directory Domain Controller for your existing domain on EC2: - When you build your own you can join an existing on-premise Active Directory domain/directory (replication mode) - You must establish a VPN (on top of Direct Connect if you have it) - Replication mode is less secure than establishing trust relationships AWS Microsoft AD does not support replication mode where replication to an on-premise AD takes place The option to use the AWS Active Directory Service for Microsoft Active Directory and create a new domain is incorrect as it involves creating a new directory which the administrators don't want. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-directory-service/

- Use the AWS Active Directory Service for Microsoft Active Directory and create a new domain. Establish a trust relationship with your existing on-premise domain

---

Q23)

You are designing an identity, authorization and access management solution for the AWS cloud. The features you need include the ability to manage user accounts and group memberships, create and apply group policies, securely connect to Amazon EC2 instances, and provide Kerberos-based single sign-on (SSO). You do not need to establish trust relationships with other domains, use DNS dynamic update, implement schema extensions or use other advanced directory features.

What would be the most cost-effective solution?

- ✅ Use AWS Simple AD

Explanation:-AWS Simple AD is an inexpensive Active Directory-compatible service with common directory features. It is a standalone, fully managed, directory on the AWS cloud. Simple AD is generally the least expensive option and the best choice for less than 50000 users and don't need advanced AD features. It is powered by SAMBA 4 Active Directory compatible server AD Connector is a directory gateway for redirecting directory requests to an Active Directory service. As you only require simple features and are looking for cost-effectiveness this would not be the best option as you must maintain an Active Directory service The AWS Directory Service for Microsoft AD is a fully managed AWS service on AWS infrastructure.It is the best choice if you have more than 5000 users and/or need a trust relationship set up. In this case you don't need those features and it would be more expensive so isn't the best options Amazon Cloud Directory enables you to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions, it is not used for authentication use cases. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-directory-service/

- Use AWS Directory Service for Microsoft AD
- Use Amazon Cloud Directory
- Use AD Connector

Q24)

You work for a company that produces TV commercials. You are planning to run an advertising campaign during a major political event that will be watched by millions of people over several days. It is expected that your website will receive large bursts of traffic following commercial breaks. You have performed an analysis and determined that you will need up to 150 EC2 web instances to process the traffic which is within the client's budgetYou need to ensure you deliver a high quality and consistent user experience whilst not exceeding the client's budget.

How would you design a highly available and elastic solution?

- Create an Auto Scaling Group across multiple AZs with a desired capacity of 150 EC2 instances. Launch an Application Load Balancer and specify the same AZs as the ASG and pre-warm the ALB by contacting AWS prior to the event
- Create an Auto Scaling Group across multiple AZs with a desired capacity of 150 EC2 instances. Launch an Application Load Balancer and specify the same AZs as the ASG
- Create an Auto Scaling Group across multiple AZs with a maximum capacity of 150 EC2 instances. Launch an Application Load Balancer and specify the same AZs as the ASG
- ✅ Create an Auto Scaling Group across multiple AZs with a maximum capacity of 150 EC2 instances. Launch an Application Load Balancer and specify the same AZs as the ASG and pre-warm the ALB by contacting AWS prior to the event

Explanation:-For this solution you must provide an elastic solution that can scale quickly with demand up to the client's budget limit. Therefore, as the analysis shows you will need up to 150 EC2 instances, which is within the client's budget you should set the ASG with a maximum capacity of 150 EC2 instances so it cannot exceed the budget. If you're anticipating a fast increase in load you can contact AWS and instruct them to pre-warm (provision) additional ELB nodes, this will ensure that the nodes will be ready when needed. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

Q25) For operational access to your AWS environment you are planning to setup a bastion host implementation. Which of the below are AWS best practices for setting up bastion hosts? (choose 2)

- ✅ Deploy in 2 AZs and use an Auto Scaling group to ensure that the number of bastion host instances always matches the desired capacity you specify during launch

Explanation:-You can configure EC2 instances as bastion hosts (aka jump boxes) in order to access your VPC instances for management. Bastion hosts are deployed in public (not private) subnets within your VPC. You can use the SSH or RDP protocols to connect to bastion hosts You need to configure a security group with the relevant permissions to allow the SSH or RDP protocols. You can also use security group rules to restrict the IP addresses/CIDRs

that can access the bastion host. Bastion hosts can use auto-assigned public IPs or Elastic IPs It is a best practice is to deploy Linux bastion hosts in two AZs, use Auto Scaling (set to 1 to just replace)and Elastic IP addresses Setting the security rule to allow from the 0.0.0.0/0 source would allow any host on the Internet to access your bastion. It's a security best practice to restrict the sources to known (safe) IP addresses or CIDR blocks. You would not want to allow unrestricted access to ports on the bastion host. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

- Bastion hosts are deployed in the private subnets of the VPC
- ✅ Elastic IP addresses are associated with the bastion instances to make it easier to remember and allow these IP addresses from on-premises firewalls

Explanation:-You can configure EC2 instances as bastion hosts (aka jump boxes) in order to access your VPC instances for management. Bastion hosts are deployed in public (not private) subnets within your VPC. You can use the SSH or RDP protocols to connect to bastion hosts You need to configure a security group with the relevant permissions to allow the SSH or RDP protocols. You can also use security group rules to restrict the IP addresses/CIDRs that can access the bastion host. Bastion hosts can use auto-assigned public IPs or Elastic IPs It is a best practice is to deploy Linux bastion hosts in two AZs, use Auto Scaling (set to 1 to just replace)and Elastic IP addresses Setting the security rule to allow from the 0.0.0.0/0 source would allow any host on the Internet to access your bastion. It's a security best practice to restrict the sources to known (safe) IP addresses or CIDR blocks. You would not want to allow unrestricted access to ports on the bastion host. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

- Access to the bastion hosts is configured to 0.0.0.0/0 for ingress in security groups

---

Q26)

An application running on an external website is attempting to initiate a request to your company's website on AWS using API calls. A problem has been reported in which the requests are failing with an error that includes the following text:

"Cross-Origin Request Blocked: The Same Origin Policy disallows reading the remote resource"You have been asked to resolve the problem, what is the most likely solution?

- The IAM policy does not allow access to the API
- The ACL on the API needs to be updated
- ✅ Enable CORS on the APIs resources using the selected methods under the API Gateway

Explanation:-Can enable Cross Origin Resource Sharing (CORS) for multiple domain use with Javascript/AJAX: - Can be used to enable requests from domains other the APIs domain - Allows the sharing of resources between different domains - The method (GET, PUT, POST etc) for which you will enable CORS must be available in the API Gateway API before you enable CORS - If CORS is not enabled and an API resource received requests from another domain the request will be blocked - Enable CORS on the APIs resources using the selected methods under the API Gateway IAM policies are not used to control CORS and there is no ACL on the API to update This error would display whether using SSL/TLS or not. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-api-gateway/

- The request is not secured with SSL/TLS

Q27)

You are an entrepreneur building a small company with some resources running on AWS. As you have limited funding you're extremely cost conscious.

Which AWS service can send you alerts via email or SNS topic when you are forecast to exceed your funding capacity so you can take action?

- Cost Explorer
- ✅ AWS Budgets

Explanation:-AWS Budgets gives you the ability to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. Budget alerts can be sent via email and/or Amazon Simple Notification Service (SNS) topic The AWS Cost Explorer is a free tool that allows you to view charts of your costs The AWS Billing Dashboard can send alerts when you're bill reaches certain thresholds but you must use AWS Budgets to created custom budgets that notify you when you are forecast to exceed a budget The AWS Cost and Usage report tracks your AWS usage and provides estimated charges associated with your AWS account but does not send alerts References: https://aws.amazon.com/aws-cost-management/aws-budgets/

- AWS Billing Dashboard
- Cost ; Usage reports

Q28) A company is in the process of deploying an Amazon Elastic Map Reduce (EMR) cluster. Which of the statements below accurately describe the EMR service? (choose 2)

- ✅ EMR utilizes a hosted Hadoop framework running on Amazon EC2 and Amazon S3

Explanation:-Amazon EMR is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. EMR utilizes a hosted Hadoop framework running on Amazon EC2 and Amazon S3. EMR uses Apache Hadoop as its distributed data processing engine which is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing EMR launches all nodes for a given cluster in the same Amazon EC2 Availability Zone EFS is a fully-managed service that makes it easy to set up and scale file storage in the Amazon Cloud. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-emr/

- EMR makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), and other processing
- ✅ EMR launches all nodes for a given cluster in the same Amazon EC2 Availability Zone

Explanation:-Amazon EMR is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. EMR utilizes a hosted Hadoop framework running on Amazon EC2 and Amazon S3. EMR uses Apache Hadoop as its distributed data processing engine which is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing EMR launches all nodes for a given cluster in the same Amazon EC2 Availability Zone EFS is a fully-managed service that makes it easy to set up and scale file storage in the Amazon Cloud References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-emr/

- EMR is a fully-managed service that makes it easy to set up and scale file storage in the Amazon Cloud

---

Q29)

As a SysOps engineer working at Digital Cloud Training, you are constantly trying to improve your processes for collecting log data. Currently you are collecting logs from across your AWS resources using CloudWatch and a combination of standard and custom metrics. You are currently investigating how you can optimize the storage of log files collected by CloudWatch.

Which of the following are valid options for storing CloudWatch log files? (choose 2)

- ✅ CloudWatch Logs

Explanation:-Valid options for storing logs include: - CloudWatch Logs - Centralized logging system (e.g. Splunk) - Custom script and store on S3 RedShift, EFS and EBS are not valid options for storing CloudWatch log files. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/management-tools/amazon-cloudwatch/

- EFS
- ✅ Splunk

Explanation:-Valid options for storing logs include: - CloudWatch Logs - Centralized logging system (e.g. Splunk) - Custom script and store on S3 RedShift, EFS and EBS are not valid options for storing CloudWatch log files. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/management-tools/amazon-cloudwatch/

- EBS

Q30)

Your company uses Amazon Glacier to store files that must be retained for compliance reasons and are rarely accessed. An auditor has requested access to some information that is stored in a Glacier archive. You have initiated an archive retrieval job.

Which factors are important to know about the process from this point? (choose 2)

- There is a charge if you delete data within 90 days
- ✅ Following retrieval, you have 24 hours to download your data

Explanation:-There is a charge if you delete data within 90 days – however we are not talking about deleting data here, just retrieving it Retrieved data is available for 24 hours by default (can be changed) Amazon Glacier must complete a job before you can get its output Glacier automatically encrypts data at rest using AES 256 symmetric keys and supports secure transfer of data over SSL Retrieved data will not be encrypted if it was uploaded unencrypted You do not need an MFA device to access the retrieved files. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

- ✅ Amazon Glacier must complete a job before you can get its output

Explanation:-There is a charge if you delete data within 90 days – however we are not talking about deleting data here, just retrieving it Retrieved data is available for 24 hours by default (can be changed) Amazon Glacier must complete a job before you can get its output Glacier automatically encrypts data at rest using AES 256 symmetric keys and supports secure transfer of data over SSL Retrieved data will not be encrypted if it was uploaded unencrypted You do not need an MFA device to access the retrieved files. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

- The retrieved data will always be encrypted

Q31)

A company is considering using EC2 Reserved Instances to reduce cost. The Architect involved is concerned about the potential limitations in flexibility of using RIs instead of On-Demand instances.

Which of the following statements about RIs are useful to the Architect? (choose 2)

- ✅ RIs can be sold on the Reserved Instance Marketplace

Explanation:-Capacity is reserved for a term of 1 or 3 years Standard = commitment of 1 or 3 years, charged whether it's on or off Scheduled = reserved for specific periods of time, accrue charges hourly, billed in monthly increments over the term (1 year) Scheduled RIs match your capacity reservation to a predictable recurring schedule RIs are used for steady state workloads and predictable usage Ideal for applications that need reserved capacity Upfront payments can reduce the hourly rate Can switch AZ within the same region Can change the instance size within the same instance type Instance type modifications are supported for Linux only Cannot change the instance size of Windows RIs Billed whether running or not Can sell reservations on the AWS marketplace Can be used in Auto Scaling Groups Can be used in Placement Groups. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

- There is a fee charged for any RI modifications
- You cannot launch RIs using Auto Scaling Groups
- ✅ You can use RIs in Placement Groups

Explanation:-Capacity is reserved for a term of 1 or 3 years Standard = commitment of 1 or 3 years, charged whether it's on or off Scheduled = reserved for specific periods of time, accrue charges hourly, billed in monthly increments over the term (1 year) Scheduled RIs match your capacity reservation to a predictable recurring schedule RIs are used for steady state workloads and predictable usage Ideal for applications that need reserved capacity Upfront payments can reduce the hourly rate Can switch AZ within the same region Can change the instance size within the same instance type Instance type modifications are supported for Linux only Cannot change the instance size of Windows RIs Billed whether running or not Can sell reservations on the AWS marketplace Can be used in Auto Scaling Groups Can be used in Placement Groups. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

Q32)

Your company has recently formed a partnership with another company. Both companies have resources running in the AWS cloud and you would like to be able to access each other's resources using private IP addresses. The resources for each company are in different AWS regions and you need to ensure that fully redundant connectivity is established.

You have established a VPC peering connection between the VPCs, what steps need to be taken next to establish connectivity and resource sharing between the VPCs across regions? (choose 2)

- Establish an IPSec VPN between the VPCs
- Establish redundant Direct Connect connections between the VPCs
- ✅ Manually add routes to each VPCs routing tables as required to enable IP connectivity

Explanation:-Peering connections can be created with VPCs in different regions (available in most regions now). Data sent between VPCs in different regions is encrypted (traffic charges apply). You must update route tables to configure routing. You must also update the inbound and outbound rules for VPC security group to reference security groups in the peered VPC When creating a VPC peering connection with another account you need to enter the account ID and VPC ID from the other account You do not use an IPSec VPN or Direct Connect to establish VPC peering, the connections are internal to AWS using the AWS network infrastructure BGP routing configuration is required for Direct Connect but not for VPC peering. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- ✅ Update Security Group rules to allow resource sharing

Explanation:-Peering connections can be created with VPCs in different regions (available in most regions now). Data sent between VPCs in different regions is encrypted (traffic charges apply). You must update route tables to configure routing. You must also update the inbound and outbound rules for VPC security group to reference security groups in the peered VPC When creating a VPC peering connection with another account you need to enter the account ID and VPC ID from the other account You do not use an IPSec VPN or Direct Connect to establish VPC peering, the connections are internal to AWS using the AWS network infrastructure BGP routing configuration is required for Direct Connect but not for VPC peering. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

Q33)

Several websites you run on AWS use multiple Internet-facing Elastic Load Balancers (ELB) to distribute incoming connections to EC2 instances running web applications. The ELBs are configured to forward using either TCP (layer 4) or HTTP (layer 7) protocols. You would like to start recording the IP addresses of the clients that connect to your web applications.

Which ELB features will you implement with which protocols? (choose 2)

- X-Forwarded-For request header and TCP
- ✅ X-Forwarded-For request header and HTTP

Explanation:-Proxy protocol for TCP/SSL carries the source (client) IP/port information X-forwarded-for for HTTP/HTTPS carries the source IP/port information In both cases the protocol carries the source IP/port information right through to the web server. If you were happy to just record the source connections on the load balancer you could use access logs. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- ✅ Proxy Protocol and TCP

Explanation:-Proxy protocol for TCP/SSL carries the source (client) IP/port information X-forwarded-for for HTTP/HTTPS carries the source IP/port information In both cases the protocol carries the source IP/port information right through to the web server. If you were happy to just record the source connections on the load balancer you could use access logs References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Proxy Protocol and HTTP

Q34)

Your company has offices in several locations around the world. Each office utilizes resources deployed in the geographically closest AWS region. You would like to implement connectivity between all of the VPCs so that you can provide full access to each other's resources. As you are security conscious you would like to ensure the traffic is encrypted and does not traverse the public Internet. The topology should be many-to-many to enable all VPCs to access the resources in all other VPCs.

How can you successfully implement this connectivity using only AWS services? (choose 2)

- Use software VPN appliances running on EC2 instances

✅ Use inter-region VPC peering

Explanation:-Peering connections can be created with VPCs in different regions (available in most regions now) Data sent between VPCs in different regions is encrypted (traffic charges apply) You cannot do transitive peering so a hub and spoke architecture would not allow all VPCs to communicate directly with each other. For this you need to establish a mesh topology A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services, it does not provide full VPC to VPC connectivity Using software VPN appliances to connect VPCs together is not the best solution as it is cumbersome, expensive and would introduce bandwidth and latency constraints (amongst other problems). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

✅ Implement a fully meshed architecture

Explanation:-Peering connections can be created with VPCs in different regions (available in most regions now) Data sent between VPCs in different regions is encrypted (traffic charges apply) You cannot do transitive peering so a hub and spoke architecture would not allow all VPCs to communicate directly with each other. For this you need to establish a mesh topology A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services, it does not provide full VPC to VPC connectivity Using software VPN appliances to connect VPCs together is not the best solution as it is cumbersome, expensive and would introduce bandwidth and latency constraints (amongst other problems). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

⚪ Implement a hub and spoke architecture

Q35)

The company you work for is currently transitioning their infrastructure and applications into the AWS cloud. You are planning to deploy an Elastic Load Balancer (ELB) that distributes traffic for a web application running on EC2 instances. You still have some application servers running on-premise and you would like to distribute application traffic across both your AWS and on-premises resources.

How can this be achieved?

✅ Provision a Direct Connect connection between your on-premises location and AWS and create a target group on an ALB to use IP based targets for both your EC2 instances and on-premises servers

Explanation:-The ALB (and NLB) supports IP addresses as targets Using IP addresses as targets allows load balancing any application hosted in AWS or on-premises using IP addresses of the application back-ends as targets You must have a VPN or Direct Connect connection to enable this configuration to work You cannot use instance ID based targets for on-premises servers and you cannot mix instance ID and IP address target types in a single target group The CLB does not support IP addresses as targets. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/ https://aws.amazon.com/blogs/aws/new-application-load-balancing-via-ip-address-to-aws-on-premises-resources/

⚪ Provision a Direct Connect connection between your on-premises location and AWS and create a target group on an ALB to use Instance ID based targets for both your EC2 instances and on-premises server

⚪ Provision an IPSec VPN connection between your on-premises location and AWS and create a CLB that uses cross-zone load balancing to distributed traffic across EC2 instances and on-premises servers

⚪ This cannot be done, ELBs are an AWS service and can only distributed traffic within the AWS cloud

---

Q36)

You are undertaking a project to make some audio and video files that your company uses for onboarding new staff members available via a mobile application. You are looking for a cost-effective way to convert the files from their current formats into formats that are compatible with smartphones and tablets. The files are currently stored in an S3 bucket.

What AWS service can help with converting the files?

⚪ MediaConvert

⚪ Data Pipeline

✅ Elastic Transcoder

Explanation:-Amazon Elastic Transcoder is a highly scalable, easy to use and cost-effective way for developers and businesses to convert (or “transcode??) video and audio files from their source format into versions that will playback on devices like smartphones, tablets and PCs MediaConvert converts file-based content for broadcast and multi-screen delivery Data Pipeline helps you move, integrate, and process data across AWS compute and storage resources, as well as your on-premises resources Rekognition is a deep learning-based visual analysis service. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/media-services/amazon-elastic-transcoder/

⚪ Rekognition

Q37)

A company uses CloudFront to provide low-latency access to cached files. An Architect is considering the implications of using CloudFront Regional Edge Caches.

Which statements are correct in relation to this service? (choose 2)

✅ Regional Edge Caches are enabled by default for CloudFront Distributions

Explanation:-Regional Edge Caches are located between origin web servers and global edge locations and have a larger cache than any individual edge location, so your objects remain in cache longer at these locations. Regional Edge caches aim to get content closer to users and are enabled by default for CloudFront Distributions (so you don't need to update your distributions) There are no additional charges for using Regional Edge Caches You can write to regional edge caches too. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-cloudfront/ https://aws.amazon.com/about-aws/whats-new/2016/11/announcing-regional-edge-caches-for-amazon-cloudfront/

⚪ There are additional charges for using Regional Edge Caches

✅ Regional Edge Caches have larger cache-width than any individual edge location, so your objects remain in cache longer at these locations

Explanation:-Regional Edge Caches are located between origin web servers and global edge locations and have a larger cache than any individual edge location, so your objects remain in cache longer at these locations. Regional Edge caches aim to get content closer to users and are enabled by default for CloudFront Distributions (so you don't need to update your distributions) There are no additional charges for using Regional Edge Caches You can write to regional edge caches too. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-cloudfront/ https://aws.amazon.com/about-aws/whats-new/2016/11/announcing-regional-edge-caches-for-amazon-cloudfront/

⚪ Regional Edge Caches are read-only

Q38)

The company you work for has a presence across multiple AWS regions. As part of disaster recovery planning you are formulating a solution to provide a regional DR capability for an application running on a fleet of Amazon EC2 instances that are provisioned by an Auto Scaling Group (ASG). The applications are stateless and read and write data to an S3 bucket. You would like to utilize the current AMI used by the ASG as it has some customizations made to it.

What are the steps you might take to enable a regional DR capability for this application? (choose 2)

✅ Enable cross region replication on the S3 bucket and specify a destination bucket in the DR region

Explanation:-There are two parts to this solution. First you need to copy the S3 data to each region (as the instances are stateless), then you need to be able to deploy instances from an ASG using the same AMI in each regions. - CRR is an Amazon S3 feature that automatically replicates data across AWS Regions. With CRR, every object uploaded to an S3 bucket is automatically replicated to a destination bucket in a different AWS Region that you choose, this enables you to copy the existing data across to each region - AMIs of both Amazon EBS-backed AMIs and instance store-backed AMIs can be copied between regions. You can then use the copied AMI to create a new launch configuration (remember that you cannot modify an ASG launch configuration, you must create a new launch configuration) There's no such thing as Multi-AZ for an S3 bucket (it's an RDS concept) Changing permissions on an AMI doesn't make it usable from another region, the AMI needs to be present within each region to be used. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

⚪ Enable multi-AZ for the S3 bucket to enable synchronous replication to the DR region

⚪ Modify the permissions of the AMI so it can be used across multiple regions

✅ Copy the AMI to the DR region and create a new launch configuration for the ASG that uses the AMI

Explanation:-There are two parts to this solution. First you need to copy the S3 data to each region (as the instances are stateless), then you need to be able to deploy instances from an ASG using the same AMI in each regions. - CRR is an Amazon S3 feature that automatically replicates data across AWS Regions. With CRR, every object uploaded to an S3 bucket is automatically replicated to a destination bucket in a different AWS Region that you choose, this enables you to copy the existing data across to each region - AMIs of both Amazon EBS-backed AMIs and instance store-backed AMIs can be copied between regions. You can then use the copied AMI to create a new launch configuration (remember that you cannot modify an ASG launch configuration, you must create a new launch configuration) There's no such thing as Multi-AZ for an S3 bucket (it's an RDS concept) Changing permissions on an AMI doesn't make it usable from another region, the AMI needs to be present within each region to be used. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

Q39)

An application hosted in your VPC uses an EC2 instance with a MySQL DB running on it. The database uses a single 1TB General Purpose SSD (GP2) EBS volume. Recently it has been noticed that the database is not performing well and you need to improve the read performance.

What are two possible ways this can be achieved? (choose 2)

- Add multiple EBS volumes in a RAID 1 array
- ✓ Add multiple EBS volumes in a RAID 0 array

Explanation:-RAID 0 = 0 striping – data is written across multiple disks and increases performance but no redundancy RAID 1 = 1 mirroring – creates 2 copies of the data but does not increase performance, only redundancy SSD, Provisioned IOPS – I01 provides higher performance than General Purpose SSD (GP2) and you can specify the IOPS required up to 50 IOPS per GB and a maximum of 32000 IOPS RDS read replicas cannot be created from EC2 instances Creating an active/passive cluster doesn't improve read performance as the passive node is not servicing requests. This is use for fault tolerance. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

- Add an RDS read replica in another AZ
- ✓ Use a provisioned IOPS volume and specify the number of I/O operations required

Explanation:-RAID 0 = 0 striping – data is written across multiple disks and increases performance but no redundancy RAID 1 = 1 mirroring – creates 2 copies of the data but does not increase performance, only redundancy SSD, Provisioned IOPS – I01 provides higher performance than General Purpose SSD (GP2) and you can specify the IOPS required up to 50 IOPS per GB and a maximum of 32000 IOPS RDS read replicas cannot be created from EC2 instances Creating an active/passive cluster doesn't improve read performance as the passive node is not servicing requests. This is use for fault tolerance References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

Q40)

Your company is reviewing their information security processes. One of the items that came out of a recent audit is that there is insufficient data recorded about requests made to a few S3 buckets. The security team requires an audit trail for operations on the S3 buckets that includes the requester, bucket name, request time, request action, and response status.

Which action would you take to enable this logging?

- Create a CloudTrail trail that audits S3 bucket operations
- Enable S3 event notifications for the specific actions and setup an SNS notification
- ✓ Enable server access logging for the S3 buckets to save access logs to a specified destination bucket

Explanation:-Server access logging provides detailed records for the requests that are made to a bucket. To track requests for access to your bucket, you can enable server access logging. Each access log record provides details about a single access request, such as the requester, bucket name, request time, request action, response status, and an error code, if relevant For capturing IAM/user identity information in logs you would need to configure AWS CloudTrail Data Events (however this does not audit the bucket operations required in the question) Amazon S3 event notifications can be sent in response to actions in Amazon S3 like PUTs, POSTs, COPYs, or DELETEs.S3 event notifications records the request action but not the other requirements of the security team CloudWatch metrics do not include the bucket operations specified in the question References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/ https://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html

- Create a CloudWatch metric that monitors the S3 bucket operations and triggers an alarm

---

Q41) A colleague has asked you some questions about how AWS charge for DynamoDB. He is interested in knowing what type of workload DynamoDB is best suited for in relation to cost and how AWS charges for DynamoDB? (choose 2)

- ✓ DynamoDB is more cost effective for read heavy workloads

Explanation:-DynamoDB charges: - DynamoDB is more cost effective for read heavy workloads - It is priced based on provisioned throughput (read/write) regardless of whether you use it or not NOTE: With the DynamoDB Auto Scaling feature you can now have DynamoDB dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. However, this is relatively new and may not yet feature on the exam. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

- DynamoDB is more cost effective for write heavy workloads
- ✓ Priced based on provisioned throughput (read/write) regardless of whether you use it or not

Explanation:-DynamoDB charges: - DynamoDB is more cost effective for read heavy workloads - It is priced based on provisioned throughput (read/write) regardless of whether you use it or not NOTE: With the DynamoDB Auto Scaling feature you can now have DynamoDB dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. However, this is relatively new and may not yet feature on the exam. See the link below for more details References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

- You provision for expected throughput but are only charged for what you use

---

Q42)

You are a Solutions Architect at Digital Cloud Training. One of your clients runs an application that writes data to a DynamoDB table. The client has asked how they can implement a function that runs code in response to item level changes that take place in the DynamoDB table.

What would you suggest to the client?

- Enable server access logging and create an event source mapping between AWS Lambda and the S3 bucket to which the logs are written
- ✓ Enable DynamoDB Streams and create an event source mapping between AWS Lambda and the relevant stream

Explanation:-DynamoDB Streams help you to keep a list of item level changes or provide a list of item level changes that have taken place in the last 24hrs. Amazon DynamoDB is integrated with AWS Lambda so that you can create triggers—pieces of code that automatically respond to events in DynamoDB Streams If you enable DynamoDB Streams on a table, you can associate the stream ARN with a Lambda function that you write. Immediately after an item in the table is modified, a new record appears in the table's stream. AWS Lambda polls the stream and invokes your Lambda function synchronously when it detects new stream records An event source mapping identifies a poll-based event source for a Lambda function. It can be either an Amazon Kinesis or DynamoDB stream. Event sources maintain the mapping configuration except for stream-based services (e.g. DynamoDB, Kinesis) for which the configuration is made on the Lambda side and Lambda performs the polling You cannot configure DynamoDB as a Kinesis Data Streams producer You can write Lambda functions to process S3 bucket events, such as the object-created or object-deleted events. For example, when a user uploads a photo to a bucket, you might want Amazon S3 to invoke your Lambda function so that it reads the image and creates a thumbnail for the photo . However, the questions asks for a solution that runs code in response to changes in a DynamoDB table, not an S3 bucket A local secondary index maintains an alternate sort key for a given partition key value, it does not record item level changes References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.Lambda.html https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-lambda/

- Create a local secondary index that records item level changes and write some custom code that responds to updates to the index
- Use Kinesis Data Streams and configure DynamoDB as a producer

---

Q43)

Your company is starting to use AWS to host new web-based applications. A new two-tier application will be deployed that provides customers with access to data records. It is important that the application is highly responsive and retrieval times are optimized. You're looking for a persistent data store that can provide the required performance.

From the list below what AWS service would you recommend for this requirement?

- ElastiCache with the Memcached engine
- ✅ ElastiCache with the Redis engine

Explanation:-ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. The in-memory caching provided by ElastiCache can be used to significantly improve latency and throughput for many read-heavy application workloads or compute-intensive workloads There are two different database engines with different characteristics as per below: Memcached - Not persistent - Cannot be used as a data store - Supports large nodes with multiple cores or threads - Scales out and in, by adding and removing nodes Redis - Data is persistent - Can be used as a datastore - Not multi-threaded - Scales by adding shards, not nodes Kinesis Data Streams is used for processing streams of data, it is not a persistent data store RDS is not the optimum solution due to the requirement to optimize retrieval times which is a better fit for an in-memory data store such as ElastiCache References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-elasticache/

- Kinesis Data Streams
- RDS in a multi-AZ configuration

Q44)

You are a Solutions Architect at Digital Cloud Training. A client from a large multinational corporation is working on a deployment of a significant amount of resources into AWS. The client would like to be able to deploy resources across multiple AWS accounts and regions using a single toolset and template.

You have been asked to suggest a toolset that can provide this functionality?

- Use a CloudFormation template that creates a stack and specify the logical IDs of each account and region
- ✅ Use a CloudFormation StackSet and specify the target accounts and regions in which the stacks will be created

Explanation:-AWS CloudFormation StackSets extends the functionality of stacks by enabling you to create, update, or delete stacks across multiple accounts and regions with a single operation Using an administrator account, you define and manage an AWS CloudFormation template, and use the template as the basis for provisioning stacks into selected target accounts across specified regions. An administrator account is the AWS account in which you create stack sets A stack set is managed by signing in to the AWS administrator account in which it was created. A target account is the account into which you create, update, or delete one or more stacks in your stack set Before you can use a stack set to create stacks in a target account, you must set up a trust relationship between the administrator and target accounts A regular CloudFormation template cannot be used across regions and accounts. You would need to create copies of the template and then manage updates You do not need to use a third-party product such as Terraform as this functionality can be delivered through native AWS technology References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/management-tools/aws-cloudformation/ https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html

- Use a third-party product such as Terraform that has support for multiple AWS accounts and regions
- This cannot be done, use separate CloudFormation templates per AWS account and region

---

Q45)

Your client is looking for a fully managed directory service in the AWS cloud. The service should provide an inexpensive Active Directory-compatible service with common directory features. The client is a medium-sized organization with 4000 users.

As the client has a very limited budget it is important to select a cost-effective solution.What would you suggest?

- AWS Active Directory Service for Microsoft Active Directory
- ✅ AWS Simple AD

Explanation:-Simple AD is an inexpensive Active Directory-compatible service with common directory features. It is a standalone, fully managed, directory on the AWS cloud and is generally the least expensive option. It is the best choice for less than 5000 users and when you don't need advanced AD features Active Directory Service for Microsoft Active Directory is the best choice if you have more than 5000 users and/or need a trust relationship set up. It provides advanced AD features that you don't get with SimpleAD Amazon Cognito is an authentication service for web and mobile apps AWS Single Sign-On (SSO) is a cloud SSO service that makes it easy to centrally manage SSO access to multiple AWS accounts and business applications References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/security-identity-compliance/aws-directory-service/

- Amazon Cognito
- AWS Single Sign-On

Q46)

You have been asked to implement a solution for capturing, transforming and loading streaming data into an Amazon RedShift cluster. The solution will capture data from Amazon Kinesis Data Streams.

Which AWS services would you utilize in this scenario? (choose 2)

- ✅ Kinesis Data Firehose for capturing the data and loading it into RedShift

Explanation:-For this solution Kinesis Data Firehose can be used as it can use Kinesis Data Streams as a source and can capture, transform, and load streaming data into a RedShift cluster. Kinesis Data Firehose can invoke a Lambda function to transform data before delivering it to destinations Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), and other processing, this solution does not involve video streams AWS Data Pipeline is used for processing and moving data between compute and storage services. It does not work with streaming data as Kinesis does Elastic Map Reduce (EMR) is used for processing and analyzing data using the Hadoop framework. It is not used for transforming streaming data. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-kinesis/

- Kinesis Video Streams for capturing the data and loading it into RedShift
- EMR for transforming the data
- ✅ Lambda for transforming the data

Explanation:-For this solution Kinesis Data Firehose can be used as it can use Kinesis Data Streams as a source and can capture, transform, and load streaming data into a RedShift cluster. Kinesis Data Firehose can invoke a Lambda function to transform data before delivering it to destinations Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), and other processing, this solution does not involve video streams AWS Data Pipeline is used for processing and moving data between compute and storage services. It does not work with streaming data as Kinesis does Elastic Map Reduce (EMR) is used for processing and analyzing data using the Hadoop framework. It is not used for transforming streaming data. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-kinesis/

Q47)

You are creating a design for a web-based application that will be based on a web front-end using EC2 instances and a database back-end. This application is a low priority and you do not want to incur costs in general day to day management.

Which AWS database service can you use that will require the least operational overhead?

- RDS
- RedShift
- EMR
- ✅ DynamoDB

Explanation:-Out of the options in the list, DynamoDB requires the least operational overhead as there are no backups, maintenance periods, software updates etc. to deal with RDS, RedShift and EMR all require some operational overhead to deal with backups, software updates and maintenance periods References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-dynamodb/

Q48)

A new Big Data application you are developing will use hundreds of EC2 instances to write data to a shared file system. The file system must be stored redundantly across multiple AZs within a region and allow the EC2 instances to concurrently access the file

system. The required throughput is multiple GB per second.

From the options presented which storage solution can deliver these requirements?

- Amazon EBS using multiple volumes in a RAID 0 configuration
- ✅ Amazon EFS

Explanation:-Amazon EFS is the best solution as it is the only solution that is a file-level storage solution (not block/object-based), stores data redundantly across multiple AZs within a region and you can concurrently connect up to thousands of EC2 instances to a single filesystem Amazon EBS volumes cannot be accessed by concurrently by multiple instances Amazon S3 is an object store, not a file system Amazon Storage Gateway is a range of products used for on-premises storage management and can be configured to cache data locally, backup data to the cloud and also provides a virtual tape backup solution References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-efs/

- Amazon S3
- Amazon Storage Gateway

Q49) A company has deployed Amazon RedShift for performing analytics on user data. When using Amazon RedShift, which of the following statements are correct in relation to availability and durability? (choose 2)

- ✅ RedShift always keeps three copies of your data

Explanation:-RedShift always keeps three copies of your data and provides continuous/incremental backups Corrections: Single-node clusters do not support data replication Manual backups are not automatically deleted when you delete a cluster References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/

- Single-node clusters support data replication
- ✅ RedShift provides continuous/incremental backups

Explanation:-RedShift always keeps three copies of your data and provides continuous/incremental backups Corrections: Single-node clusters do not support data replication Manual backups are not automatically deleted when you delete a cluster References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/

- RedShift always keeps five copies of your data

Q50)

You are planning to launch a RedShift cluster for processing and analyzing a large amount of data. The RedShift cluster will be deployed into a VPC with multiple subnets.

Which construct is used when provisioning the cluster to allow you to specify a set of subnets in the VPC that the cluster will be deployed into?

- DB Subnet Group
- Subnet Group
- Availability Zone (AZ)
- ✅ Cluster Subnet Group

Explanation:-You create a cluster subnet group if you are provisioning your cluster in your virtual private cloud (VPC) A cluster subnet group allows you to specify a set of subnets in your VPC When provisioning a cluster you provide the subnet group and Amazon Redshift creates the cluster on one of the subnets in the group A DB Subnet Group is used by RDS A Subnet Group is used by ElastiCache Availability Zones are part of the AWS global infrastructure, subnets reside within AZs but in RedShift you provision the cluster into Cluster Subnet Groups. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/ https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-cluster-subnet-groups.html

---

Q51)

A Solutions Architect is designing the system monitoring and deployment layers of a serverless application. The system monitoring layer will manage system visibility through recording logs and metrics and the deployment layer will deploy the application stack and manage workload changes through a release management process.The Architect needs to select the most appropriate AWS services for these functions.

Which services and frameworks should be used for the system monitoring and deployment layers? (choose 2)

- Use AWS Lambda to package, test, and deploy the serverless application stack
- ✅ Use Amazon CloudWatch for consolidating system and application logs and monitoring custom metrics

Explanation:-AWS Serverless Application Model (AWS SAM) is an extension of AWS CloudFormation that is used to package, test, and deploy serverless applications With Amazon CloudWatch, you can access system metrics on all the AWS services you use, consolidate system and application level logs, and create business key performance indicators (KPIs) as custom metrics for your specific needs AWS Lambda is used for executing your code as functions, it is not used for packaging, testing and deployment. AWS Lambda is used with AWS SAM AWS X-Ray lets you analyze and debug serverless applications by providing distributed tracing and service maps to easily identify performance bottlenecks by visualizing a request end-to-end References: https://docs.aws.amazon.com/lambda/latest/dg/serverless_app.html https://digitalcloud.training/certification-training/aws-solutions-architect-associate/management-tools/amazon-cloudwatch/

- Use AWS X-Ray to package, test, and deploy the serverless application stack
- ✅ Use AWS SAM to package, test, and deploy the serverless application stack

Explanation:-AWS Serverless Application Model (AWS SAM) is an extension of AWS CloudFormation that is used to package, test, and deploy serverless applications With Amazon CloudWatch, you can access system metrics on all the AWS services you use, consolidate system and application level logs, and create business key performance indicators (KPIs) as custom metrics for your specific needs AWS Lambda is used for executing your code as functions, it is not used for packaging, testing and deployment. AWS Lambda is used with AWS SAM AWS X-Ray lets you analyze and debug serverless applications by providing distributed tracing and service maps to easily identify performance bottlenecks by visualizing a request end-to-end References: https://docs.aws.amazon.com/lambda/latest/dg/serverless_app.html https://digitalcloud.training/certification-training/aws-solutions-architect-associate/management-tools/amazon-cloudwatch/

Q52)

You have just created a new security group in your VPC. You have not yet created any rules.

Which of the statements below are correct regarding the default state of the security group? (choose 2)

- ✅ There are no inbound rules and traffic will be implicitly denied

Explanation:-Custom security groups do not have inbound allow rules (all inbound traffic is denied by default) Default security groups do have inbound allow rules (allowing traffic from within the group) All outbound traffic is allowed by default in both custom and default security groups Security groups act like a stateful firewall at the instance level. Specifically security groups operate at the network interface level of an EC2 instance. You can only assign permit rules in a security group, you cannot assign deny rules and there is an implicit deny rule at the end of the security group. All rules are evaluated until a permit is encountered or continues until the implicit deny. You can create ingress and egress rules References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- ✅ There is an outbound rule that allows all traffic to all IP addresses

Explanation:-Custom security groups do not have inbound allow rules (all inbound traffic is denied by default) Default security groups do have inbound allow rules (allowing traffic from within the group) All outbound traffic is allowed by default in both custom and default security groups Security groups act like a stateful firewall at the instance level. Specifically security groups operate at the network interface level of an EC2 instance. You can only assign permit rules in a security group, you cannot assign deny rules and there is an implicit deny rule at the end of the security group. All rules are evaluated until a permit is encountered or continues until the implicit deny. You can create ingress and egress rules. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- There is an outbound rule allowing traffic to the Internet Gateway
- There are is an inbound rule that allows traffic from the Internet Gateway

---

Q53)

You need to setup a distribution method for some static files. The requests will be mainly GET requests and you are expecting a high volume of GETs often exceeding 2000 per second. The files are currently stored in an S3 bucket.

According to AWS best practices, what can you do to optimize performance?

✅ Integrate CloudFront with S3 to cache the content

Explanation:-Amazon S3 automatically scales to high request rates. For example, your application can achieve at least 3,500 PUT/POST/DELETE and 5,500 GET requests per second per prefix in a bucket. There are no limits to the number of prefixes in a bucket If your workload is mainly sending GET requests, in addition to the preceding guidelines, you should consider using Amazon CloudFront for performance optimization. By integrating CloudFront with Amazon S3, you can distribute content to your users with low latency and a high data transfer rate Transfer Acceleration is used to accelerate object uploads to S3 over long distances (latency) Cross-region replication creates a replica copy in another region but should not be used for spreading read requests across regions. There will be 2 S3 endpoints and CRR is not designed for 2 way sync so this would not work well ElastiCache is used for caching database content not S3 content References: https://docs.aws.amazon.com/AmazonS3/latest/dev/request-rate-perf-considerations.html https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-s3/

- Use ElastiCache to cache the content
- Use S3 Transfer Acceleration
- Use cross-region replication to spread the load across regions

---

Q54)

Your company has started using the AWS CloudHSM for secure key storage. A recent administrative error resulted in the loss of credentials to access the CloudHSM. You need access to data that was encrypted using keys stored on the hardware security module.

How can you recover the keys that are no longer accessible?

- Restore a snapshot of the CloudHSM

✅ There is no way to recover your keys if you lose your credentials

Explanation:-Amazon does not have access to your keys or credentials and therefore has no way to recover your keys if you lose your credentials. References: https://aws.amazon.com/cloudhsm/faqs/

- Log a case with AWS support and they will use MFA to recover the credentials
- Reset the CloudHSM device and create a new set of credentials

---

Q55)

You have implemented the AWS Elastic File System (EFS) to store data that will be accessed by a large number of EC2 instances. The data is sensitive and you are working on a design for implementing security measures to protect the data. You need to ensure that network traffic is restricted correctly based on firewall rules and access from hosts is restricted by user or group.

How can this be achieved with EFS? (choose 2)

- Use Network ACLs to control the traffic

✅ Use EFS Security Groups to control network traffic

Explanation:-You can control who can administer your file system using IAM. You can control access to files and directories with POSIX-compliant user and group-level permissions. POSIX permissions allows you to restrict access from hosts by user and group. EFS Security Groups act as a firewall, and the rules you add define the traffic flow You cannot use AWS WAF to protect EFS data using users and groups You do not use IAM to control access to files and directories by user and group, but you can use IAM to control who can administer the file system configuration You use EFS Security Groups to control network traffic to EFS, not Network ACLs References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-efs/ https://aws.amazon.com/efs/features/

- Use AWS Web Application Firewall (WAF) to protect EFS

✅ Use POSIX permissions to control access from hosts by user or group

Explanation:-You can control who can administer your file system using IAM. You can control access to files and directories with POSIX-compliant user and group-level permissions. POSIX permissions allows you to restrict access from hosts by user and group. EFS Security Groups act as a firewall, and the rules you add define the traffic flow You cannot use AWS WAF to protect EFS data using users and groups You do not use IAM to control access to files and directories by user and group, but you can use IAM to control who can administer the file system configuration You use EFS Security Groups to control network traffic to EFS, not Network ACLs References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/storage/amazon-efs/ https://aws.amazon.com/efs/features/

Q56)

You have recently enabled Access Logs on your Application Load Balancer (ALB). One of your colleagues would like to process the log files using a hosted Hadoop service.

What configuration changes and services can be leveraged to deliver this requirement?

✅ Configure Access Logs to be delivered to S3 and use EMR for processing the log files

Explanation:-Access Logs can be enabled on ALB and configured to store data in an S3 bucket. Amazon EMR is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. EMR utilizes a hosted Hadoop framework running on Amazon EC2 and Amazon S3 Neither Kinesis or EC2 provide a hosted Hadoop service You cannot configure access logs to be delivered to DynamoDB References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/analytics/amazon-emr/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Configure Access Logs to be delivered to DynamoDB and use EMR for processing the log files
- Configure Access Logs to be delivered to EC2 and install Hadoop for processing the log files
- Configure Access Logs to be delivered to S3 and use Kinesis for processing the log files

---

Q57)

You are designing the disk configuration for an EC2 instance. The instance needs to support a MapReduce process that requires high throughput for a large dataset with large I/O sizes. You need to provision the most cost-effective storage solution option.

What EBS volume type will you select?

- EBS General Purpose SSD
- EBS Provisioned IOPS SSD
- EBS General Purpose SSD in a RAID 1 configuration

✅ EBS Throughput Optimized HDD

Explanation:-EBS Throughput Optimized HDD is good for the following use cases (and is the most cost-effective option: Frequently accessed, throughput intensive workloads with large datasets and large I/O sizes, such as MapReduce, Kafka, log processing, data warehouse, and ETL workloads Throughput is measured in MB/s, and includes the ability to burst up to 250 MB/s per TB, with a baseline throughput of 40 MB/s per TB and a maximum throughput of 500 MB/s per volume The SSD options are more expensive References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

Q58)

An EC2 instance you manage is generating very high packets-per-second and performance of the application stack is being impacted. You have been asked for a resolution to the issue that results in improved performance from the EC2 instance.

What would you suggest?

✅ Use enhanced networking

Explanation:-Enhanced networking provides higher bandwidth, higher packet-per-second (PPS) performance, and consistently lower inter-instance latencies. If your packets-per-second rate appears to have reached its ceiling, you should consider moving to enhanced networking because you have likely reached the upper thresholds of the VIF driver. It is only available for certain instance types and only supported in VPC. You must also launch an HVM AMI with the appropriate drivers AWS currently supports enhanced networking capabilities using SR-IOV. SR-IOV provides direct access to network adapters, provides higher performance (packets-per-second) and lower latency You do not need to create a RAID 1 array (which is more for redundancy than performance anyway) A placement group is used to increase network performance between instances. In this case there is only a single instance so it won't help Adding multiple IP addresses is not a way to increase performance of the instance as the same amount of bandwidth is available to the Elastic Network Interface (ENI). References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/ https://aws.amazon.com/premiumsupport/knowledge-center/enable-configure-enhanced-networking/

- Add multiple Elastic IP addresses to the instance
- Create a placement group and put the EC2 instance in it
- Configure a RAID 1 array from multiple EBS volumes

Q59)

A web application you manage receives order processing information from customers and places the messages on an SQS queue. A fleet of EC2 instances are configured to pick up the messages, process them, and store the results in a DynamoDB table. The current configuration has been resulting in a large number of empty responses to ReceiveMessage requests.

You would like to update the configuration to eliminate empty responses to reduce operational overhead. How can this be done?

- Configure Short Polling to eliminate empty responses by reducing the length of time a connection request remains open
- Use a FIFO (first-in-first-out) queue to preserve the exact order in which messages are sent and received
- Use a Standard queue to provide at-least-once delivery, which means that each message is delivered at least once
- ✅ Configure Long Polling to eliminate empty responses by allowing Amazon SQS to wait until a message is available in a queue before sending a response

Explanation:-The correct answer is to use Long Polling which will eliminate empty responses by allowing Amazon SQS to wait until a message is available in a queue before sending a response The problem does not relate to the order in which the messages are processed in and there are no concerns over messages being delivered more than once so it doesn't matter whether you use a FIFO or standard queue Long Polling: - Uses fewer requests and reduces cost - Eliminates false empty responses by querying all servers - SQS waits until a message is available in the queue before sending a response - Requests contain at least one of the available messages up to the maximum number of messages specified in the ReceiveMessage action - Shouldn't be used if your application expects an immediate response to receive message calls - ReceiveMessageWaitTime is set to a non-zero value (up to 20 seconds) - Same charge per million requests as short polling Changing the queue type would not assist in this situation Short Polling: - Does not wait for messages to appear in the queue - It queries only a subset of the available servers for messages (based on weighted random execution) - Short polling is the default - ReceiveMessageWaitTime is set to 0 - More requests are used, which implies higher cost References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/application-integration/amazon-sqs/ https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-long-polling.html

---

Q60)

You are running an Auto Scaling Group (ASG) with an Elastic Load Balancer (ELB) and a fleet of EC2 instances. Health checks are configured on the ASG to use EC2 status checks The ELB has determined that an EC2 instance is unhealthy and has removed it from service. However, you noticed that the instance is still running and has not been terminated by the ASG.

What would be an explanation for this?

- ✅ The ELB health check type has not been selected for the ASG and so it is unaware that the instance has been determined to be unhealthy by the ELB and has been removed from service

Explanation:-If using an ELB it is best to enable ELB health checks as otherwise EC2 status checks may show an instance as being healthy that the ELB has determined is unhealthy. In this case the instance will be removed from service by the ELB but will not be terminated by Auto Scaling Connection draining is not the correct answer as the ELB has taken the instance out of service so there are no active connections The health check grace period allows a period of time for a new instance to warm up before performing a health check More information on ASG health checks: By default uses EC2 status checks Can also use ELB health checks and custom health checks ELB health checks are in addition to the EC2 status checks If any health check returns an unhealthy status the instance will be terminated With ELB an instance is marked as unhealthy if ELB reports it as OutOfService A healthy instance enters the InService state If an instance is marked as unhealthy it will be scheduled for replacement If connection draining is enabled, Auto Scaling waits for in-flight requests to complete or timeout before terminating instances The health check grace period allows a period of time for a new instance to warm up before performing a health check (300 seconds by default) References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

- The ASG is waiting for the cooldown timer to expire before terminating the instance
- The health check grace period has not yet expired
- Connection draining is enabled and the ASG is waiting for in-flight requests to complete

---

Q61)

Your company runs a web-based application that uses EC2 instances for the web front-end and RDS for the database back-end. The web application writes transaction log files to an S3 bucket and the quantity of files is becoming quite large. You have determined that it is acceptable to retain the most recent 60 days of log files and permanently delete the rest.

What can you do to enable this to happen automatically?

- Use an S3 lifecycle policy to move the log files that are more than 60 days old to the GLACIER storage class
- Write a Ruby script that checks the age of objects and deletes any that are more than 60 days old
- ✅ Use an S3 lifecycle policy with object expiration configured to automatically remove objects that are more than 60 days old

Explanation:-Moving logs to Glacier may save cost but the questions requests that the files are permanently deleted Object Expiration allows you to schedule removal of your objects after a defined time period Using Object Expiration rules to schedule periodic removal of objects eliminates the need to build processes to identify objects for deletion and submit delete requests to Amazon S3 References: https://aws.amazon.com/about-aws/whats-new/2011/12/27/amazon-s3-announces-object-expiration/ https://aws.amazon.com/about-aws/whats-new/2011/12/27/amazon-s3-announces-object-expiration/

- Use an S3 bucket policy that deletes objects that are more than 60 days old

---

Q62)

A DynamoDB table you manage has a variable load, ranging from sustained heavy usage some days, to only having small spikes on others. The load is 80% read and 20% write. The provisioned throughput capacity has been configured to account for the heavy load to ensure throttling does not occur.You have been asked to find a solution for saving cost.

What would be the most efficient and cost-effective solution?

- Use DynamoDB DAX to increase the performance of the database
- ✅ Create a DynamoDB Auto Scaling scaling policy

Explanation:-DynamoDB auto scalinguses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. This is the most efficient and cost-effective solution Manually adjusting the provisioned throughput is not efficient Using AWS Lambda to modify the provisioned throughput is possible but it would be more cost-effective to use DynamoDB Auto Scaling as there is no cost to using it DynamoDB DAX is an in-memory cache that increases the performance of DynamoDB. However, it costs money and there is no requirement to increase performance References: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

- Create a CloudWatch alarm that notifies you of increased/decreased load, and manually adjust the provisioned throughput
- Create a CloudWatch alarm that triggers an AWS Lambda function that adjusts the provisioned throughput

---

Q63)

As a Solutions Architect at Digital Cloud Training you are helping a client to design a multi-tier web application architecture. The

client has requested that the architecture provide low-latency connectivity between all servers and be resilient across multiple locations. The client uses Microsoft SQL Server for existing databases.

The client has a limited budget for staff costs and does not need to access the underlying operating system

What would you recommend as the most efficient solution?

- Amazon EC2 instances with Microsoft SQL Server and data replication between two different AZs
- Amazon EC2 instances with Microsoft SQL Server and data replication within an AZ
- Amazon RDS with Microsoft SQL Server
- ✅ Amazon RDS with Microsoft SQL Server in a Multi-AZ configuration

Explanation:-As the client does not need to manage the underlying operating system and they have a limited budget for staff, they should use a managed service such as RDS. Multi-AZ RDS creates a replica in another AZ and synchronously replicates to it which enables the required resilience across multiple locations With EC2 you have full control at the operating system layer (not required) and can install your own database. However, you would then need to manage the entire stack and therefore staff costs would increase so this is not the best solution. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-rds/

Q64)

An organization in the agriculture sector is deploying sensors and smart devices around factory plants and fields. The devices will collect information and send it to cloud applications running on AWS.

Which AWS service will securely connect the devices to the cloud applications?

- AWS DMS
- AWS Glue
- AWS Lambda
- ✅ AWS IoT Core

Explanation:-An organization in the agriculture sector is deploying sensors and smart devices around factory plants and fields. The devices will collect information and send it to cloud applications running on AWS

Q65)

A development team needs to run up a few lab servers on a weekend for a new project. The servers will need to run uninterrupted for a few hours.

Which EC2 pricing option would be most suitable?

- Reserved
- ✅ On-Demand

Explanation:-Spot pricing may be the most economical option for a short duration over a weekend but you may have the instances terminated by AWS and there is a requirement that the servers run uninterrupted On-Demand pricing ensures that instances will not be terminated and is the most economical option Reserved pricing provides a reduced cost for a contracted period (1 or 3 years), and is not suitable for ad hoc requirements Dedicated instances run on hardware that's dedicated to a single customer and are more expensive than regular On-Demand instances References: https://aws.amazon.com/ec2/pricing/

- Dedicated instances
- Spot

Q66)

A security officer has requested that all data associated with a specific customer is encrypted. The data resides on Elastic Block Store (EBS) volumes.

Which of the following statements about using EBS encryption are correct? (choose 2)

- All instance types support encryption
- All attached EBS volumes must share the same encryption state
- ✅ Data in transit between an instance and an encrypted volume is also encrypted

Explanation:-All EBS types and all instance families support encryption Not all instance types support encryption There is no direct way to change the encryption state of a volume Data in transit between an instance and an encrypted volume is also encrypted You can have encrypted and non-encrypted EBS volumes on a single instance. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

- ✅ There is no direct way to change the encryption state of a volume

Explanation:-All EBS types and all instance families support encryption Not all instance types support encryption There is no direct way to change the encryption state of a volume Data in transit between an instance and an encrypted volume is also encrypted You can have encrypted and non-encrypted EBS volumes on a single instance References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ebs/

Q67)

You are putting together an architecture for a new VPC on AWS. Your on-premise data center will be connected to the VPC by a hardware VPN and has public and VPN-only subnets.

The security team has requested that all traffic that hits the public subnets on AWS must be directed over the VPN to the corporate firewall.

How can this be achieved?

- ✅ In the public subnet route table, add a route for your remote network and specify the virtual private gateway as the target

Explanation:-Route tables determine where network traffic is directed. In your route table, you must add a route for your remote network and specify the virtual private gateway as the target. This enables traffic from your VPC that's destined for your remote network to route via the virtual private gateway and over one of the VPN tunnels. You can enable route propagation for your route table to automatically propagate your network routes to the table for you You must select the virtual private gateway (AWS side of the VPN) not the customer gateway (customer side of the VPN) in the target in the route table NAT Gateways are used to enable Internet access for EC2 instances in private subnets, they cannot be used to direct traffic to VPG You must create the route table rule in the route table attached to the public subnet, not the VPN-only subnet References: https://docs.aws.amazon.com/vpc/latest/userguide/VPC_VPN.html

- In the public subnet route table, add a route for your remote network and specify the customer gateway as the target
- Configure a NAT Gateway and configure all traffic to be directed via the virtual private gateway
- In the VPN-only subnet route table, add a route that directs all Internet traffic to the virtual private gateway

---

Q68)

One of your clients has requested advice on the correct choice of Elastic Load Balancer (ELB) for an application they are planning to deploy on AWS. The application requires extremely high throughput and extremely low latencies. The connections will be made using the TCP protocol and the ELB must support load balancing to multiple ports on an instance.

Which ELB would you suggest the client uses?

- Classic Load Balancer
- ✅ Network Load Balancer

Explanation:-The Network Load Balancer operates at the connection level (Layer 4), routing connections to targets – Amazon EC2 instances, containers and IP addresses based on IP protocol data. It is architected to handle millions of requests/sec, sudden volatile traffic patterns and provides extremely low

latencies. It provides high throughput and extremely low latencies and is designed to handle traffic as it grows and can load balance millions of requests/second. NLB also supports load balancing to multiple ports on an instance The CLB operates using the TCP, SSL, HTTP and HTTPS protocols. It is not the best choice for requirements of extremely high throughput and low latency and does not support load balancing to multiple ports on an instance The ALB operates at the HTTP and HTTPS level only (does not support TCP load balancing) Route 53 is a DNS service, it is not a type of ELB (though you can do some types of load balancing with it) References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Application Load Balancer
- Route 53

Q69)

An RDS database is experiencing heavy read traffic. You are planning on creating read replicas.

When using Amazon RDS with Read Replicas, which of the deployment options below are valid? (choose 2)

- Cross-subnet
- ✅ Within an Availability Zone

Explanation:-Read Replicas can be within an AZ, Cross-AZ and Cross-Region Read replicas are used for read heavy DBs and replication is asynchronous. Read replicas are for workload sharing and offloading Read replicas cannot be cross-continent, cross-subnet or cross-facility. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-rds/

- Cross-Continent
- ✅ Cross-Availability Zone

Explanation:-Read Replicas can be within an AZ, Cross-AZ and Cross-Region Read replicas are used for read heavy DBs and replication is asynchronous. Read replicas are for workload sharing and offloading Read replicas cannot be cross-continent, cross-subnet or cross-facility. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-rds/

Q70)

A developer is writing code for AWS Lambda and is looking to automate the release process.

Which AWS services can be used to automate the release process of Lambda applications? (choose 2)

- AWS Glue
- AWS OpsWorks
- ✅ AWS CodeDeploy

Explanation:-You can automate your serverless application's release process using AWS CodePipeline and AWS CodeDeploy The following AWS services can be used to fully automate the deployment process: You use CodePipeline to model, visualize, and automate the steps required to release your serverless application You use AWS CodeDeploy to gradually deploy updates to your serverless applications You use CodeBuild to build, locally test, and package your serverless application You use AWS CloudFormation to deploy your application References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-lambda/ https://docs.aws.amazon.com/lambda/latest/dg/build-pipeline.html

- ✅ AWS CodePipeline

Explanation:-You can automate your serverless application's release process using AWS CodePipeline and AWS CodeDeploy The following AWS services can be used to fully automate the deployment process: You use CodePipeline to model, visualize, and automate the steps required to release your serverless application You use AWS CodeDeploy to gradually deploy updates to your serverless applications You use CodeBuild to build, locally test, and package your serverless application You use AWS CloudFormation to deploy your application References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-lambda/ https://docs.aws.amazon.com/lambda/latest/dg/build-pipeline.html

---

Q71)

In your VPC you have several EC2 instances that have been running for some time. You have logged into an instance and need to determine a few pieces of information including what IAM role is assigned, the instance ID and the names of the security groups that are assigned to the instance.

From the options below, what would be a source of this information?

- Parameters
- ✅ Metadata

Explanation:-Instance metadata is data about your instance that you can use to configure or manage the running instance and is available at http://169.254.169.254/latest/meta-data Tags are used to categorize and label resources Parameters are used in databases User data is used to configure the system at launch time and specify scripts. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/ https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html#instancedata-data-categories

- Tags
- User data

Q72)

The application development team in your company have developed a Java application and saved the source code in a .war file. They would like to run the application on AWS resources and are looking for a service that can handle the provisioning and management of the underlying resources it will run on.

What AWS service would allow the developers to upload the Java source code file and provide capacity provisioning and infrastructure management?

- AWS CloudFormation
- AWS CodeDeploy
- AWS OpsWorks
- ✅ AWS Elastic Beanstalk

Explanation:-AWS Elastic Beanstalk can be used to quickly deploy and manage applications in the AWS Cloud. Developers upload applications and Elastic Beanstalk handles the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring Elastic Beanstalk supports applications developed in Go, Java, .NET, Node.js, PHP, Python, and Ruby, as well as different platform configurations for each language. To use Elastic Beanstalk, you create an application, upload an application version in the form of an application source bundle (for example, a Java .war file) to Elastic Beanstalk, and then provide some information about the application AWS CloudFormation uses templates to deploy infrastructure as code. It is not a PaaS service like Elastic Beanstalk and is more focussed on infrastructure than applications and management of applications AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances, on-premises instances, serverless Lambda functions, or Amazon ECS services AWS OpsWorks is a configuration management service that provides managed instances of Chef and Puppet References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-elastic-beanstalk/ https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

---

Q73)

A development team are creating a Continuous Integration and Continuous Delivery (CI/CD) toolchain on the AWS cloud. The team currently use Jenkins X and Kubernetes on-premise and are looking to utilize the same services in the AWS cloud.

What AWS service can provide a managed container platform that is MOST similar to their current CI/CD toolchain?

- AWS CodePipeline
- Amazon ECS
- AWS Lambda
- ✅ Amazon EKS

Explanation:-Amazon EKS is AWS' managed Kubernetes offering, which enables you to focus on building applications, while letting AWS handle

managing Kubernetes and the underlying cloud infrastructure Amazon Elastic Container Service (ECS) does not use Kubernetes so it is not the most similar product AWS Lambda is a serverless service that executes code as functions AWS CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. It is not a container platform References: https://aws.amazon.com/eks/

Q74)

You are a Solutions Architect at Digital Cloud Training. One of your clients has requested some advice on how to implement security measures in their VPC. The client has recently been the victim of some hacking attempts. Fortunately, no data has been exposed at this point but the client wants to implement measures to mitigate further threats. The client has explained that the attacks always come from the same small block of IP addresses.

What would be a quick and easy measure to help prevent further attacks?

- Use CloudFront's DDoS prevention features
- Use a Security Group rule that denies connections from the block of IP addresses
- ✅ Use a Network ACL rule that denies connections from the block of IP addresses

Explanation:-With NACLs you can have permit and deny rules. Network ACLs contain a numbered list of rules that are evaluated in order from the lowest number until the explicit deny. Network ACLs have separate inbound and outbound rules and each rule can allow or deny traffic With Security Groups you can only assign permit rules, you cannot assign deny rules A bastion host is typically used for admin purposes, allowing access to a single endpoint in the AWS cloud for administration using SSH/RDP. From the bastion instance you then connect to other EC2 instances in your subnets. This is not used as a method of adding security to production systems and cannot stop traffic from hitting application ports CloudFront does have DDoS prevention features but we don't know that this is a DDoS style of attack and CloudFront can only help where the traffic is using the CloudFront service to access cached content References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- Create a Bastion Host restrict all connections to the Bastion Host only

---

Q75)

A company is investigating ways to analyze and process large amounts of data in the cloud faster, without needing to load or transform the data in a data warehouse.

The data resides in Amazon S3.

Which AWS services would allow the company to query the data in place? (choose 2)

- Amazon Kinesis Data Streams
- ✅ Amazon S3 Select

Explanation:-Amazon S3 Select is designed to help analyze and process data within an object in Amazon S3 buckets, faster and cheaper. It works by providing the ability to retrieve a subset of data from an object in Amazon S3 using simple SQL expressions Amazon Redshift Spectrum allows you to directly run SQL queries against exabytes of unstructured data in Amazon S3. No loading or transformation is required Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. It does not allow you to perform query-in-place operations on S3 Amazon Elasticsearch Service, is a fully managed service that makes it easy for you to deploy, secure, operate, and scale Elasticsearch to search, analyze, and visualize data in real-time Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/ https://aws.amazon.com/blogs/aws/s3-glacier-select/ https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-redshift-spectrum-is-now-available-in-four-additional-aws-regions-and-enhances-query-performance-in-all-available-aws-regions/

- ✅ Amazon RedShift Spectrum

Explanation:-Amazon S3 Select is designed to help analyze and process data within an object in Amazon S3 buckets, faster and cheaper. It works by providing the ability to retrieve a subset of data from an object in Amazon S3 using simple SQL expressions Amazon Redshift Spectrum allows you to directly run SQL queries against exabytes of unstructured data in Amazon S3. No loading or transformation is required Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. It does not allow you to perform query-in-place operations on S3 Amazon Elasticsearch Service, is a fully managed service that makes it easy for you to deploy, secure, operate, and scale Elasticsearch to search, analyze, and visualize data in real-time Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/ https://aws.amazon.com/blogs/aws/s3-glacier-select/ https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-redshift-spectrum-is-now-available-in-four-additional-aws-regions-and-enhances-query-performance-in-all-available-aws-regions/

- Amazon Elasticsearch

Q76)

One of you clients has asked for assistance with a performance issue they are experiencing. The client has a fleet of EC2 instances behind an Elastic Load Balancer (ELB) that are a mixture of c4.2xlarge instance types and c5.large instances. The load on the CPUs on the c5.large instances has been very high, often hitting 100% utilization, whereas the c4.2xlarge instances have been performing well.

The client has asked for advice on the most cost effective way to resolve the performance problems?

- Enable the weighted routing policy on the ELB and configure a higher weighting for the c4.2xlarge instances
- Add more c5.large instances to spread the load more evenly
- Add all of the instances into a Placement Group
- ✅ Change the configuration to use only c4.2xlarge instance types

Explanation:-The 2xlarge instance type provides more CPUs. The best answer is to use this instance type for all instances A placement group helps provide low-latency connectivity between instances and would not help here The weighted routing policy is a Route 53 feature that would not assist in this situation References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

Q77)

A company is deploying a new two-tier web application that uses EC2 web servers and a DynamoDB database backend. An Internet facing ELB distributes connections between the web servers.The Solutions Architect has created a security group for the web servers and needs to create a security group for the ELB.

What rules should be added? (choose 2)

- ✅ Add an Inbound rule that allows HTTP/HTTPS, and specify the source as 0.0.0.0/0

Explanation:-An inbound rule should be created for the relevant protocols (HTTP/HTTPS) and the source should be set to any address (0.0.0.0/0) The address 0.0.0.0/32 is incorrect as the 32 mask means an exact match is required (0.0.0.0) The outbound rule should forward the relevant protocols (HTTP/HTTPS) and the destination should be set to the web server security group Using the VPC CIDR would not be secure and you cannot specify an Internet Gateway in a security group (not that you'd want to anyway) FYI on the web server security group you'd want to add an Inbound rule allowing HTTP/HTTPS from the ELB security group References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Add an Outbound rule that allows HTTP/HTTPS, and specify the destination as VPC CIDR
- ✅ Add an Outbound rule that allows HTTP/HTTPS, and specify the destination as the web server security group

Explanation:-An inbound rule should be created for the relevant protocols (HTTP/HTTPS) and the source should be set to any address (0.0.0.0/0) The address 0.0.0.0/32 is incorrect as the 32 mask means an exact match is required (0.0.0.0) The outbound rule should forward the relevant protocols (HTTP/HTTPS) and the destination should be set to the web server security group Using the VPC CIDR would not be secure and you cannot specify an Internet Gateway in a security group (not that you'd want to anyway) FYI on the web server security group you'd want to add an Inbound rule allowing HTTP/HTTPS from the ELB security group References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

- Add an Outbound rule that allows ALL TCP, and specify the destination as the Internet Gateway

---

Q78)

You launched an EBS-backed EC2 instance into your VPC. A requirement has come up for some high-performance ephemeral storage and so you would like to add an instance-store backed volume.

How can you add the new instance store volume?

- You can use a block device mapping to specify additional instance store volumes when you launch your instance, or you can attach additional instance store volumes after your instance is running
- You must use an Elastic Network Adapter (ENA) to add instance store volumes. First, attach an ENA, and then attach the instance store volume
- ✅ You can specify the instance store volumes for your instance only when you launch an instance

Explanation:-You can specify the instance store volumes for your instance only when you launch an instance. You can't attach instance store volumes to an instance after you've launched it You can use a block device mapping to specify additional EBS volumes when you launch your instance, or you can attach additional EBS volumes after your instance is running An Elastic Network Adapter has nothing to do with adding instance store volumes References: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/add-instance-store-volumes.html

- You must shutdown the instance in order to be able to add the instance store volume

---

Q79)

You have just created a new Network ACL in your VPC. You have not yet created any rules.

Which of the statements below are correct regarding the default state of the Network ACL? (choose 2)

- There is a default inbound rule allowing traffic from the VPC CIDR block
- ✅ There is a default inbound rule denying all traffic

Explanation:-A VPC automatically comes with a default network ACL which allows all inbound/outbound traffic A custom NACL denies all traffic both inbound and outbound by default Network ACL's function at the subnet level and you can have permit and deny rules. Network ACLs have separate inbound and outbound rules and each rule can allow or deny traffic. Network ACLs are stateless so responses are subject to the rules for the direction of traffic. NACLs only apply to traffic that is ingress or egress to the subnet not to traffic within the subnet References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

- There is a default outbound rule allowing traffic to the Internet Gateway
- ✅ There is a default outbound rule denying all traffic

Explanation:-A VPC automatically comes with a default network ACL which allows all inbound/outbound traffic A custom NACL denies all traffic both inbound and outbound by default Network ACL's function at the subnet level and you can have permit and deny rules. Network ACLs have separate inbound and outbound rules and each rule can allow or deny traffic. Network ACLs are stateless so responses are subject to the rules for the direction of traffic. NACLs only apply to traffic that is ingress or egress to the subnet not to traffic within the subnet References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/

Q80)

A large quantity of data that is rarely accessed is being archived onto Amazon Glacier. Your CIO wants to understand the resilience of the service.

Which of the statements below is correct about Amazon Glacier storage? (choose 2)

- Data is replicated globally
- ✅ Data is resilient in the event of one entire Availability Zone destruction

Explanation:-Glacier is designed for durability of 99.999999999% of objects across multiple Availability Zones. Data is resilient in the event of one entire Availability Zone destruction. Glacier supports SSL for data in transit and encryption of data at rest. Glacier is extremely low cost and is ideal for long-term archival Data is not resilient to the failure of an entire region Data is not replicated globally There is no availability SLA with Glacier References: https://aws.amazon.com/s3/storage-classes/

- Data is resilient in the event of one entire region destruction
- ✅ Provides 99.999999999% durability of archives

Explanation:-Glacier is designed for durability of 99.999999999% of objects across multiple Availability Zones. Data is resilient in the event of one entire Availability Zone destruction. Glacier supports SSL for data in transit and encryption of data at rest. Glacier is extremely low cost and is ideal for long-term archival Data is not resilient to the failure of an entire region Data is not replicated globally There is no availability SLA with Glacier References: https://aws.amazon.com/s3/storage-classes/

Q81)

You are planning to launch a fleet of EC2 instances running Linux. As part of the launch you would like to install some application development frameworks and custom software onto the instances. The installation will be initiated using some scripts you have written.

What feature allows you to specify the scripts so you can install the software during the EC2 instance launch?

- Metadata
- AWS Config
- ✅ User data

Explanation:-When you launch an instance in Amazon EC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon EC2: shell scripts and cloud-init directives User data is data that is supplied by the user at instance launch in the form of a script and is limited to 16KB User data and meta data are not encrypted. Instance metadata is available at http://169.254.169.254/latest/meta-data. The Instance Metadata Query tool allows you to query the instance metadata without having to type out the full URI or category names The AWS Systems Manager run command is used to manage the configuration of existing instances by using remotely executed commands. User data is better for specifying scripts to run at startup. References: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/

- Run command

Q82) One of your clients has multiple VPCs that are peered with each other. The client would like to use a single Elastic Load Balancer (ELB) to route traffic to multiple EC2 instances in peered VPCs within the same region. Is this possible?

- This is possible using the Classic Load Balancer (CLB) if using Instance IDs
- This is not possible with ELB, you would need to use Route 53
- No, the instances that an ELB routes traffic to must be in the same VPC
- ✅ This is possible using the Network Load Balancer (NLB) and Application Load Balancer (ALB) if using IP addresses as targets

Explanation:-With ALB and NLB IP addresses can be used to register: Instances in a peered VPC AWS resources that are addressable by IP address and port On-premises resources linked to AWS through Direct Connect or a VPN connection. References: https://aws.amazon.com/blogs/aws/new-application-load-balancing-via-ip-address-to-aws-on-premises-resources/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/

Q83)

One of the applications you manage receives a high traffic loads between 7:30am and 9:30am daily. The application uses an Auto Scaling Group (ASG) to maintain 3 EC2 instances most of the time but during the peak period requires 6 EC2 instances.

How can you configure ASG to perform a regular scale-out event at 7:30am and a scale-in event at 9:30am daily to account for the peak load?

- Use a Dynamic scaling policy
- Use a Simple scaling policy

✅ Use a Scheduled scaling policy

Explanation:-Simple – maintains a current number of instances, you can manually change the ASGs min/desired/max and attach/detach instances Scheduled – Used for predictable load changes, can be a single event or a recurring schedule Dynamic (event based) – scale in response to an event/alarm Step – configure multiple scaling steps in response to multiple alarms References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

⚪ Use a Step scaling policy

Q84)

An application is generating a large amount of clickstream events data that is being stored on S3. The business needs to understand customer behaviour and want to run complex analytics queries against the data.

Which AWS service can be used for this requirement?

✅ Amazon RedShift

Explanation:-Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and existing Business Intelligence (BI) tools RedShift is used for running complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution RDS is a relational database that is used for transactional workloads not analytics workloads Amazon Neptune is a new product that offers a fully-managed Graph database Amazon Kinesis Firehose processes streaming data, not data stored on S3. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/database/amazon-redshift/

⚪ Amazon Kinesis Firehose

⚪ Amazon RDS

⚪ Amazon Neptune

Q85)

You would like to create a highly available web application that serves static content using multiple On-Demand EC2 instances.

Which of the following AWS services will help you to achieve this? (choose 2)

✅ Elastic Load Balancer and Auto Scaling

Explanation:-None of the answer options present the full solution. However, you have been asked which services will help you to achieve the desired outcome. In this case we need high availability for on-demand EC2 instances. A single Auto Scaling Group will enable the on-demand instances to be launched into multiple availability zones with an elastic load balancer distributing incoming connections to the available EC2 instances. This provides high availability and elasticity Amazon S3 and CloudFront could be used to serve static content from an S3 bucket, however the question states that the web application runs on EC2 instances DynamoDB and ElastiCache are both database services, not web application services, and cannot help deliver high availability for EC2 instances Direct Connect is used for connecting on-premise data centers into AWS using a private network connection and does not help in this situation at all. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

✅ Multiple Availability Zones

Explanation:-None of the answer options present the full solution. However, you have been asked which services will help you to achieve the desired outcome. In this case we need high availability for on-demand EC2 instances. A single Auto Scaling Group will enable the on-demand instances to be launched into multiple availability zones with an elastic load balancer distributing incoming connections to the available EC2 instances. This provides high availability and elasticity Amazon S3 and CloudFront could be used to serve static content from an S3 bucket, however the question states that the web application runs on EC2 instances DynamoDB and ElastiCache are both database services, not web application services, and cannot help deliver high availability for EC2 instances Direct Connect is used for connecting on-premise data centers into AWS using a private network connection and does not help in this situation at all. References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-vpc/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elastic-load-balancing/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

⚪ DynamoDB and ElastiCache

⚪ Direct Connect

Q86)

You are a Solutions Architect at Digital Cloud Training. A client of yours is using API Gateway for accepting and processing a large number of API calls to AWS Lambda. The client's business is rapidly growing and he is therefore expecting a large increase in traffic to his API Gateway and AWS Lambda services.The client has asked for advice on ensuring the services can scale without any reduction in performance.

What advice would you give to the client? (choose 2)

✅ AWS Lambda scales concurrently executing functions up to your default limit

Explanation:-API Gateway can scale to any level of traffic received by an API. API Gateway scales up to the default throttling limit of 10,000 requests per second, and can burst past that up to 5,000 RPS. Throttling is used to protect back-end instances from traffic spikes Lambda uses continuous scaling scales out not up. Lambda scales concurrently executing functions up to your default limit (1000) API Gateway does not use provisioned throughput - this is something that is used to provision performance in DynamoDB API Gateway can scale past the default throttling limits (they are not fixed, you just have to apply to have them adjusted) References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-api-gateway/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-lambda/

⚪ API Gateway scales manually through the assignment of provisioned throughput

✅ API Gateway scales up to the default throttling limit, with some additional burst capacity available

Explanation:-API Gateway can scale to any level of traffic received by an API. API Gateway scales up to the default throttling limit of 10,000 requests per second, and can burst past that up to 5,000 RPS. Throttling is used to protect back-end instances from traffic spikes Lambda uses continuous scaling scales out not up. Lambda scales concurrently executing functions up to your default limit (1000) API Gateway does not use provisioned throughput - this is something that is used to provision performance in DynamoDB API Gateway can scale past the default throttling limits (they are not fixed, you just have to apply to have them adjusted) References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-api-gateway/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/aws-lambda/

⚪ API Gateway can only scale up to the fixed throttling limits

---

Q87)

An application you manage uses Auto Scaling and a fleet of EC2 instances. You recently noticed that Auto Scaling is scaling the number of instances up and down multiple times in the same hour. You need to implement a remediation to reduce the amount of scaling events. The remediation must be cost-effective and preserve elasticity

What design changes would you implement? (choose 2)

⚪ Modify the Auto Scaling group termination policy to terminate the oldest instance first

✅ Modify the Auto Scaling group cool-down timers

Explanation:-The cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that it doesn't launch or terminate additional instances before the previous scaling activity takes effect so this would help. After the Auto Scaling group dynamically scales using a simple scaling policy, it waits for the cooldown period to complete before resuming scaling activities The CloudWatch Alarm Evaluation Period is the number of the most recent data points to evaluate when determining alarm state. This would help as you can increase the number of datapoints required to trigger an alarm The order in which Auto Scaling terminates instances is not the issue here, the problem is that the workload is dynamic and Auto Scaling is constantly reacting to change and launching or terminating instances References: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html#alarm-evaluation https://digitalcloud.guru/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

✅ Modify the CloudWatch alarm period that triggers your Auto Scaling scale down policy

Explanation:-The cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that it doesn't launch or terminate additional

instances before the previous scaling activity takes effect so this would help. After the Auto Scaling group dynamically scales using a simple scaling policy, it waits for the cooldown period to complete before resuming scaling activities The CloudWatch Alarm Evaluation Period is the number of the most recent data points to evaluate when determining alarm state. This would help as you can increase the number of datapoints required to trigger an alarm The order in which Auto Scaling terminates instances is not the issue here, the problem is that the workload is dynamic and Auto Scaling is constantly reacting to change and launching or terminating instances References: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html#alarm-evaluation https://digitalcloud.guru/certification-training/aws-solutions-architect-associate/compute/aws-auto-scaling/

- Modify the Auto Scaling policy to use scheduled scaling actions

---

Q88)

An EBS-backed EC2 instance has been configured with some proprietary software that uses an embedded license. You need to move the EC2 instance to another Availability Zone (AZ) within the region.

How can this be accomplished? Choose the best answer.

- Take a snapshot of the instance. Create a new EC2 instance and perform a restore from the snapshot
- ✅ Create an image from the instance. Launch an instance from the AMI in the destination AZ

Explanation:-The easiest and recommended option is to create an AMI (image) from the instance and launch an instance from the AMI in the other AZ. AMIs are backed by snapshots which in turn are backed by S3 so the data is available from any AZ within the region You can take a snapshot, launch an instance in the destination AZ. Stop the instance, detach its root volume, create a volume from the snapshot you took and attach it to the instance. However, this is not the best option There痴 no way to move an EC2 instance from the management console You cannot perform a copy operation to move the instance References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ec2/ https://aws.amazon.com/premiumsupport/knowledge-center/move-ec2-instance/

- Use the AWS Management Console to select a different AZ for the existing instance
- Perform a copy operation to move the EC2 instance to the destination AZ

Q89)

Your manager has asked you to explain how Amazon ElastiCache may assist with the company痴 plans to improve the performance of database queries.

Which of the below statements is a valid description of the benefits of Amazon ElastiCache? (Choose 2)

- ElastiCache can form clusters using a mixture of Memcached and Redis caching engines, allowing you to take advantage of the best features of each caching engine
- ElastiCache is best suited for scenarios where the data base load type is OLTP
- ✅ The in-memory caching provided by ElastiCache can be used to significantly improve latency and throughput for many read-heavy application workloads or compute-intensive workloads

Explanation:-ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud The in-memory caching provided by ElastiCache can be used to significantly improve latency and throughput for many read-heavy application workloads or compute-intensive workloads ElastiCache is best for scenarios where the DB load is based on Online Analytics Processing (OLAP) transactions not Online Transaction Processing (OLTP) Elasticache EC2 nodes cannot be accessed from the Internet, nor can they be accessed by EC2 instances in other VPCs You cannot mix Memcached and Redis in a cluster References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-route-53/

- ✅ ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud

Explanation:-ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud The in-memory caching provided by ElastiCache can be used to significantly improve latency and throughput for many read-heavy application workloads or compute-intensive workloads ElastiCache is best for scenarios where the DB load is based on Online Analytics Processing (OLAP) transactions not Online Transaction Processing (OLTP) Elasticache EC2 nodes cannot be accessed from the Internet, nor can they be accessed by EC2 instances in other VPCs You cannot mix Memcached and Redis in a cluster References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/networking-and-content-delivery/amazon-route-53/

Q90)

You have been tasked with building an ECS cluster using the EC2 launch type and need to ensure container instances can connect to the cluster. A colleague informed you that you must ensure the ECS container agent is installed on your EC2 instances. You have selected to use the Amazon ECS-optimized AMI.

Which of the statements below are correct? (Choose 2)

- The Amazon ECS container agent is installed on the AWS managed infrastructure used for tasks using the EC2 launch type so you don't need to do anything
- The Amazon ECS container agent must be installed for all AMIs
- ✅ The Amazon ECS container agent is included in the Amazon ECS-optimized AMI

Explanation:-The ECS container agent allows container instances to connect to the cluster nd runs on each infrastructure resource on an ECS cluster. The ECS container agent is included in the Amazon ECS optimized AMI and can also be installed on any EC2 instance that supports the ECS specification (only supported on EC2 instances). It is available for Linux and Windows The ECS container agent does not need to be installed for all AMIs as it is included in the Amazon ECS optimized AMI With the EC2 launch type the container agent is not installed on AWS managed infrastructure - however this is true for the Fargate launch type You can install the EC2 container agent on Windows instances References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ecs/

- ✅ You can install the ECS container agent on any Amazon EC2 instance that supports the Amazon ECS specification

Explanation:-The ECS container agent allows container instances to connect to the cluster nd runs on each infrastructure resource on an ECS cluster. The ECS container agent is included in the Amazon ECS optimized AMI and can also be installed on any EC2 instance that supports the ECS specification (only supported on EC2 instances). It is available for Linux and Windows The ECS container agent does not need to be installed for all AMIs as it is included in the Amazon ECS optimized AMI With the EC2 launch type the container agent is not installed on AWS managed infrastructure - however this is true for the Fargate launch type You can install the EC2 container agent on Windows instances References: https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/amazon-ecs/
