

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 251 – Question 300**

[Visit PassLeader and Download Full Version AWS-Associate Exam Dumps](#)

QUESTION 251

You have a distributed application that periodically processes large volumes of data across multiple Amazon EC2 Instances. The application is designed to recover gracefully from Amazon EC2 instance failures. You are required to accomplish this task in the most cost-effective way. Which of the following will meet your requirements?

- A. Spot Instances
- B. Reserved instances
- C. Dedicated instances
- D. On-Demand instances

Answer: A

Explanation:

Using reserved instances is not the most cost-effective way.

<https://aws.amazon.com/blogs/aws/new-scheduled-reserved-instances/>

“Scheduled Reserved Instance model allows you to reserve instances for predefined blocks of time on a recurring basis for a one-year term, with prices that are generally 5 to 10% lower than the equivalent On-Demand rates.”

You can get spot instances with much lower prices:

<https://aws.amazon.com/ec2/spot/pricing/>

“Spot instances are also available to run for a predefined duration – in hourly increments up to six hours in length – at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak times for a total of up to 50% savings.”

QUESTION 252

Which of the following are true regarding AWS CloudTrail? Choose 3 answers.

- A. CloudTrail is enabled globally.
- B. CloudTrail is enabled by default.
- C. CloudTrail is enabled on a per-region basis.
- D. CloudTrail is enabled on a per-service basis.
- E. Logs can be delivered to a single Amazon S3 bucket for aggregation.
- F. CloudTrail is enabled for all available services within a region.
- G. Logs can only be processed and delivered to the region in which they are generated.

Answer: ACE

Explanation:

A: have a trail with the Apply trail to all regions option enabled.

C: have multiple single region trails.

E: Log files from all the regions can be delivered to a single S3 bucket.

Global service events are always delivered to trails that have the Apply trail to all regions option enabled. Events are delivered from a single region to the bucket for the trail. This setting cannot be changed.

If you have a single region trail, you should enable the Include global services option.

If you have multiple single region trails, you should enable the Include global services option in only one of the trails.

D: Incorrect. Once enabled it is applicable for all the supported services, service can't be selected.

QUESTION 253

You have a content management system running on an Amazon EC2 instance that is approaching 100% CPU utilization. Which option will reduce load on the Amazon EC2 instance?

- A. Create a load balancer, and register the Amazon EC2 instance with it
- B. Create a CloudFront distribution, and configure the Amazon EC2 instance as the origin
- C. Create an Auto Scaling group from the instance using the CreateAutoScalingGroup action
- D. Create a launch configuration from the instance using the CreateLaunchConfiguration action

Answer: C

Explanation:

You can create an ASG from instance ID

http://docs.aws.amazon.com/AutoScaling/latest/APIReference/API_CreateAutoScalingGroup.html

QUESTION 254

You have a load balancer configured for VPC, and all back-end Amazon EC2 instances are in service. However, your web browser times out when connecting to the load balancer's DNS name. Which options are probable causes of this behavior? Choose 2 answers.

- A. The load balancer was not configured to use a public subnet with an Internet gateway configured.
- B. The Amazon EC2 instances do not have a dynamically allocated private IP address.
- C. The security groups or network ACLs are not properly configured for web traffic.
- D. The load balancer is not configured in a private subnet with a NAT instance.
- E. The VPC does not have a VGW configured.

Answer: AC

Explanation:

There is no such thing as VGW. Hence E is not correct answer.

QUESTION 255

A company needs to deploy services to an AWS region which they have not previously used. The company currently has an AWS identity and Access Management (IAM) role for the Amazon EC2 instances, which permits the instance to have access to Amazon DynamoDB. The company wants their EC2 instances in the new region to have the same privileges. How should the company achieve this?

- A. Create a new IAM role and associated policies within the new region
- B. Assign the existing IAM role to the Amazon EC2 instances in the new region
- C. Copy the IAM role and associated policies to the new region and attach it to the instances
- D. Create an Amazon Machine Image (AMI) of the instance and copy it to the desired region using the AMI Copy feature

Answer: B

QUESTION 256

Which of the following notification endpoints or clients are supported by Amazon Simple Notification Service? Choose 2 answers.

- A. Email
- B. CloudFront distribution
- C. File Transfer Protocol
- D. Short Message Service
- E. Simple Network Management Protocol

Answer: AD

Explanation:

SNS Supported Endpoints

Email Notifications

Amazon SNS provides the ability to send Email notifications.

SMS Notifications

Amazon SNS provides the ability to send and receive Short Message Service (SMS) notifications to SMS-enabled mobile phones and smart phones.

<http://docs.aws.amazon.com/sns/latest/dg/welcome.html>

QUESTION 257

Which set of Amazon S3 features helps to prevent and recover from accidental data loss?

- A. Object lifecycle and service access logging
- B. Object versioning and Multi-factor authentication
- C. Access controls and server-side encryption
- D. Website hosting and Amazon S3 policies

Answer: B

Explanation:

Versioning-enabled buckets enable you to recover objects from accidental deletion or overwrite. In addition to that, they have made it a requirement that delete operations on versioned data can only be done using MFA (Multi factor authentication).

http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

QUESTION 258

A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers.

- A. Amazon Simple Email Service
- B. Amazon CloudWatch
- C. Amazon Simple Queue Service
- D. Amazon Route 53
- E. Amazon Simple Notification Service

Answer: BE

Explanation:

B: Amazon RDS provides metrics in real time for the operating system (OS) that your DB instance runs on. You can view the metrics for your DB instance using the console, or consume the Enhanced Monitoring JSON output from CloudWatch Logs in a monitoring system of your choice.

E: Use Amazon RDS DB events to monitor failovers. For example, you can be notified by text message or email when a DB instance fails over. Amazon RDS uses the Amazon Simple Notification Service (Amazon SNS) to provide notification when an Amazon RDS event occurs.

QUESTION 259

A company is preparing to give AWS Management Console access to developers. Company policy mandates identity federation and role-based access control. Roles are currently assigned using groups in the corporate Active Directory. What combination of the following will give developers access to the AWS console? Choose 2 answers.

- A. AWS Directory Service AD Connector
- B. AWS Directory Service Simple AD
- C. AWS Identity and Access Management groups
- D. AWS identity and Access Management roles
- E. AWS identity and Access Management users

Answer: AD

Explanation:

http://docs.aws.amazon.com/STS/latest/APIReference/API_AssumeRoleWithSAML.html

QUESTION 260

What is the durability of S3 RRS?

- A. 99.99%
- B. 99.95%
- C. 99.995%
- D. 99.999999999%

Answer: A

QUESTION 261

What does specifying the mapping /dev/sdc=none when launching an instance do?

- A. Prevents /dev/sdc from creating the instance.
- B. Prevents /dev/sdc from deleting the instance.
- C. Set the value of /dev/sdc to 'zero'.
- D. Prevents /dev/sdc from attaching to the instance.

Answer: D

QUESTION 262

You are deploying an application to track GPS coordinates of delivery trucks in the United States. Coordinates are transmitted from each delivery truck once every three seconds. You need to design an architecture that will enable real-time processing of these coordinates from multiple consumers. Which service should you use to implement data ingestion?

- A. Amazon Kinesis
- B. AWS Data Pipeline
- C. Amazon AppStream
- D. Amazon Simple Queue Service

Answer: A

Explanation:

<https://aws.amazon.com/streaming-data/>

QUESTION 263

A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

- A. SAML-based Identity Federation
- B. Cross-Account Access
- C. AWS Identity and Access Management roles
- D. Web Identity Federation

Answer: D

Explanation:

Web identity federation - You can let users sign in using a well-known third party identity provider such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC) 2.0 compatible provider. AWS STS web identity federation supports Login with Amazon, Facebook, Google, and any OpenID Connect (OIDC)-compatible identity provider.

QUESTION 264

You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (S3). Video uploads are taking longer than expected, resulting in poor application performance. Which method will help improve performance of your application?

- A. Enable enhanced networking
- B. Use Amazon S3 multipart upload
- C. Leveraging Amazon CloudFront, use the HTTP POST method to reduce latency
- D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

Answer: B

Explanation:

Using multipart upload provides the following advantages:

- Improved throughput - You can upload parts in parallel to improve throughput.
- Quick recovery from any network issues - Smaller part size minimizes the impact of restarting a failed upload due to a network error.
- Pause and resume object uploads - You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.
- Begin an upload before you know the final object size.
- You can upload an object as you are creating it.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.html>

QUESTION 265

A customer wants to track access to their Amazon Simple Storage Service (S3) buckets and also use this information for their internal security and access audits. Which of the following will meet the Customer requirement?

- A. Enable AWS CloudTrail to audit all Amazon S3 bucket access.
- B. Enable server access logging for all required Amazon S3 buckets.
- C. Enable the Requester Pays option to track access via AWS Billing.
- D. Enable Amazon S3 event notifications for Put and Post.

Answer: B

Explanation:

If it's just for internal audit, then Server access logging, I assume is sufficient.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html>

For external audits I would go for CloudTrail.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-logging.html>

QUESTION 266

A company is deploying a two-tier, highly available web application to AWS. Which service provides durable storage for static content while utilizing lower Overall CPU resources for the web tier?

- A. Amazon EBS volume
- B. Amazon S3
- C. Amazon EC2 instance store
- D. Amazon RDS instance

Answer: B

QUESTION 267

You are designing a web application that stores static assets in an Amazon Simple Storage Service (S3) bucket. You expect this bucket to immediately receive over 150 PUT requests per second. What should you do to ensure optimal performance?

- A. Use multi-part upload.
- B. Add a random prefix to the key names.
- C. Amazon S3 will automatically manage performance at this scale.
- D. Use a predictable naming scheme, such as sequential numbers or date time sequences, in the key names.

Answer: B

Explanation:

If you anticipate that your workload will consistently exceed 100 requests per second, you should avoid sequential key names. If you must use sequential numbers or date and time patterns in key names, add a random prefix to the key name. The randomness of the prefix more evenly distributes key names across multiple index partitions. Examples of introducing randomness are provided later in this topic.

QUESTION 268

When will you incur costs with an Elastic IP address (EIP)?

- A. When an EIP is allocated.
- B. When it is allocated and associated with a running instance.
- C. When it is allocated and associated with a stopped instance.
- D. Costs are incurred regardless of whether the EIP is associated with a running instance.

Answer: C

Explanation:

You are allowed one EIP to be attached to a running instance at no charge. otherwise, it will incur a small fee. in this case, the instance is stopped, and thus, the EIP will be billed at the normal rate.

<http://aws.amazon.com/ec2/pricing/>

QUESTION 269

A company has an AWS account that contains three VPCs (Dev, Test, and Prod) in the same region. Test is peered to both Prod and Dev. All VPCs have non-overlapping CIDR blocks. The company wants to push minor code releases from Dev to Prod to speed up time to market. Which of the following options helps the company accomplish this?

- A. Create a new peering connection Between Prod and Dev along with appropriate routes.

- B. Create a new entry to Prod in the Dev route table using the peering connection as the target.
- C. Attach a second gateway to Dev. Add a new entry in the Prod route table identifying the gateway as the target.
- D. The VPCs have non-overlapping CIDR blocks in the same account. The route tables contain local routes for all VPCs.

Answer: A

Explanation:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-pg.pdf#create-vpc-peering-connection>

QUESTION 270

Which of the following instance types are available as Amazon EBS-backed only? Choose 2 answers.

- A. General purpose T2
- B. General purpose M3
- C. Compute-optimized C4
- D. Compute-optimized C3
- E. Storage-optimized 12

Answer: AC

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html>

QUESTION 271

A customer is hosting their company website on a cluster of web servers that are behind a public-facing load balancer. The customer also uses Amazon Route 53 to manage their public DNS. How should the customer configure the DNS zone apex record to point to the load balancer?

- A. Create an A record pointing to the IP address of the load balancer.
- B. Create a CNAME record pointing to the load balancer DNS name.
- C. Create a CNAME record aliased to the load balancer DNS name.
- D. Create an A record aliased to the load balancer DNS name.

Answer: D

QUESTION 272

You try to connect via SSH to a newly created Amazon EC2 instance and get one of the following error messages: "Network error: Connection timed out" or "Error connecting to [instance], reason: -> Connection timed out: connect," You have confirmed that the network and security group rules are configured correctly and the instance is passing status checks. What steps should you take to identify the source of the behavior? Choose 2 answers.

- A. Verify that the private key file corresponds to the Amazon EC2 key pair assigned at launch.
- B. Verify that your IAM user policy has permission to launch Amazon EC2 instances.
- C. Verify that you are connecting with the appropriate user name for your AMI.
- D. Verify that the Amazon EC2 Instance was launched with the proper IAM role.
- E. Verify that your federation trust to AWS has been established.

Answer: AC

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html>

QUESTION 273

A customer is running a multi-tier web application farm in a virtual private cloud (VPC) that is not connected to their corporate network. They are connecting to the VPC over the Internet to manage all of their Amazon EC2 instances running in both the public and private subnets. They have only authorized the bastion-security-group with Microsoft Remote Desktop Protocol (RDP) access to the application instance security groups, but the company wants to further limit administrative access to all of the instances in the VPC. Which of the following Bastion deployment scenarios will meet this requirement?

- A. Deploy a Windows Bastion host on the corporate network that has RDP access to all instances in the VPC.
- B. Deploy a Windows Bastion host with an Elastic IP address in the public subnet and allow SSH access to the bastion from anywhere.
- C. Deploy a Windows Bastion host with an Elastic IP address in the private subnet, and restrict RDP access to the bastion from only the corporate public IP addresses.

- D. Deploy a Windows Bastion host with an auto-assigned Public IP address in the public subnet, and allow RDP access to the bastion from only the corporate public IP addresses.

Answer: D

QUESTION 274

A customer has a single 3-TB volume on-premises that is used to hold a large repository of images and print layout files. This repository is growing at 500 GB a year and must be presented as a single logical volume. The customer is becoming increasingly constrained with their local storage capacity and wants an off-site backup of this data, while maintaining low-latency access to their frequently accessed data. Which AWS Storage Gateway configuration meets the customer requirements?

- A. Gateway-Cached volumes with snapshots scheduled to Amazon S3
- B. Gateway-Stored volumes with snapshots scheduled to Amazon S3
- C. Gateway-Virtual Tape Library with snapshots to Amazon S3
- D. Gateway-Virtual Tape Library with snapshots to Amazon Glacier

Answer: A

Explanation:

<http://docs.aws.amazon.com/storagegateway/latest/userguide/storage-gateway-cached-concepts.html>

QUESTION 275

You are building an automated transcription service in which Amazon EC2 worker instances process an uploaded audio file and generate a text file. You must store both of these files in the same durable storage until the text file is retrieved. You do not know what the storage capacity requirements are. Which storage option is both cost-efficient and scalable?

- A. Multiple Amazon EBS volume with snapshots
- B. A single Amazon Glacier vault
- C. A single Amazon S3 bucket
- D. Multiple instance stores

Answer: C

QUESTION 276

You need to pass a custom script to new Amazon Linux instances created in your Auto Scaling group. Which feature allows you to accomplish this?

- A. User data
- B. EC2Config service
- C. IAM roles
- D. AWS Config

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html#user-data-shell-scripts>

Not B, because EC2Config is used for Windows instances:

http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/UsingConfig_WinAMI.html

QUESTION 277

Which of the following services natively encrypts data at rest within an AWS region? Choose 2 answers.

- A. AWS Storage Gateway
- B. Amazon DynamoDB
- C. Amazon CloudFront
- D. Amazon Glacier
- E. Amazon Simple Queue Service

Answer: AD

Explanation:

https://media.amazonwebservices.com/AWS_Securing_Data_at_Rest_with_Encryption.pdf (page 12)

QUESTION 278

A company is building software on AWS that requires access to various AWS services. Which configuration should be used to ensure that AWS credentials (i.e., Access Key ID/Secret Access Key combination) are not compromised?

- A. Enable Multi-Factor Authentication for your AWS root account.
- B. Assign an IAM role to the Amazon EC2 instance.
- C. Store the AWS Access Key ID/Secret Access Key combination in software comments.
- D. Assign an IAM user to the Amazon EC2 Instance.

Answer: B

Explanation:

Use roles for applications that run on Amazon EC2 instances.

Applications that run on an Amazon EC2 instance need credentials in order to access other AWS services. To provide credentials to the application in a secure way, use IAM roles. A role is an entity that has its own set of permissions, but that isn't a user or group. Roles also don't have their own permanent set of credentials the way IAM users do. In the case of Amazon EC2, IAM dynamically provides temporary credentials to the EC2 instance, and these credentials are automatically rotated for you.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-roles-with-ec2>

QUESTION 279

Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? Choose 2 answers.

- A. Supported on all Amazon EBS volume types
- B. Snapshots are automatically encrypted
- C. Available to all instance types
- D. Existing volumes can be encrypted
- E. shared volumes can be encrypted

Answer: AB

Explanation:

This feature is supported on all Amazon EBS volume types (General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic). You can access encrypted Amazon EBS volumes the same way you access existing volumes; encryption and decryption are handled transparently and they require no additional action from you, your Amazon EC2 instance, or your application. Snapshots of encrypted Amazon EBS volumes are automatically encrypted, and volumes that are created from encrypted Amazon EBS snapshots are also automatically encrypted.

<http://docs.aws.amazon.com/kms/latest/developerguide/services-ebs.html>

QUESTION 280

A company is deploying a new two-tier web application in AWS. The company has limited staff and requires high availability, and the application requires complex queries and table joins. Which configuration provides the solution for the company's requirements?

- A. MySQL Installed on two Amazon EC2 Instances in a single Availability Zone
- B. Amazon RDS for MySQL with Multi-AZ
- C. Amazon ElastiCache
- D. Amazon DynamoDB

Answer: B

Explanation:

When is it appropriate to use DynamoDB instead of a relational database?

From our own experience designing and operating a highly available, highly scalable ecommerce platform, we have come to realize that relational databases should only be used when an application really needs the complex query, table join and transaction capabilities of a full-blown relational database. In all other cases, when such relational features are not needed, a NoSQL database service like DynamoDB offers a simpler, more available, more scalable and ultimately a lower cost solution.

QUESTION 281

A t2.medium EC2 instance type must be launched with what type of Amazon Machine Image (AMI)?

- A. An Instance store Hardware Virtual Machine AMI
- B. An Instance store Paravirtual AMI
- C. An Amazon EBS-backed Hardware Virtual Machine AMI
- D. An Amazon EBS-backed Paravirtual AMI

Answer: C

Explanation:

You must launch a T2 instance using an HVM AMI. For more information, see Linux AMI Virtualization Types. You must launch your T2 instances using an EBS volume as the root device. For more information, see Amazon EC2 Root Device Volume.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

QUESTION 282

You manually launch a NAT AMI in a public subnet. The network is properly configured. Security groups and network access control lists are properly configured. Instances in a private subnet can access the NAT. The NAT can access the Internet. However, private instances cannot access the Internet. What additional step is required to allow access from the private instances?

- A. Enable Source/Destination Check on the private Instances.
- B. Enable Source/Destination Check on the NAT instance.
- C. Disable Source/Destination Check on the private instances.
- D. Disable Source/Destination Check on the NAT instance.

Answer: D

Explanation:

Disabling Source/Destination Checks.

Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance. You can disable the SrcDestCheck attribute for a NAT instance that's either running or stopped using the console or the command line.

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

QUESTION 283

Which of the following approaches provides the lowest cost for Amazon Elastic Block Store snapshots while giving you the ability to fully restore data?

- A. Maintain two snapshots: the original snapshot and the latest incremental snapshot.
- B. Maintain a volume snapshot; subsequent snapshots will overwrite one another.
- C. Maintain a single snapshot the latest snapshot is both Incremental and complete.
- D. Maintain the most current snapshot, archive the original and incremental to Amazon Glacier.

Answer: C

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

QUESTION 284

An existing application stores sensitive information on a non-boot Amazon EBS data volume attached to an Amazon Elastic Compute Cloud instance. Which of the following approaches would protect the sensitive data on an Amazon EBS volume?

- A. Upload your customer keys to AWS CloudHSM. Associate the Amazon EBS volume with AWS CloudHSM. Re-mount the Amazon EBS volume.
- B. Create and mount a new, encrypted Amazon EBS volume. Move the data to the new volume. Delete the old Amazon EBS volume.
- C. Unmount the EBS volume. Toggle the encryption attribute to True. Re-mount the Amazon EBS volume.
- D. Snapshot the current Amazon EBS volume. Restore the snapshot to a new, encrypted Amazon EBS volume. Mount the Amazon EBS volume.

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

To migrate data between encrypted and unencrypted volumes:

1. Create your destination volume (encrypted or unencrypted, depending on your need) by following the procedures in Creating an Amazon EBS Volume.
2. Attach the destination volume to the instance that hosts the data to migrate. For more information, see Attaching an Amazon EBS Volume to an Instance.

3. Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Using. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.
4. Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.

QUESTION 285

A US-based company is expanding their web presence into Europe. The company wants to extend their AWS infrastructure from Northern Virginia (us-east-1) into the Dublin (eu-west-1) region. Which of the following options would enable an equivalent experience for users on both continents?

- A. Use a public-facing load balancer per region to load-balance web traffic, and enable HTTP health checks.
- B. Use a public-facing load balancer per region to load-balance web traffic, and enable sticky sessions.
- C. Use Amazon Route 53, and apply a geolocation routing policy to distribute traffic across both regions.
- D. Use Amazon Route 53, and apply a weighted routing policy to distribute traffic across both regions.

Answer: C

Explanation:

Geolocation routing lets you choose the resources that serve your traffic based on the geographic location of your users, meaning the location from which DNS queries originate. For example, you might want all queries from Africa to be routed to a web server with an IP address of 192.0.2.111.

Another possible use is for balancing load across endpoints in a predictable, easy-to-manage way, so that each user location is consistently routed to the same endpoint.

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-weighted>

QUESTION 286

Which of the following are use cases for Amazon DynamoDB? Choose 3 answers.

- A. Storing BLOB data.
- B. Managing web sessions.
- C. Storing JSON documents.
- D. Storing metadata for Amazon S3 objects.
- E. Running relational joins and complex updates.
- F. Storing large amounts of infrequently accessed data.

Answer: BCD

Explanation:

Ideal Usage Patterns.

Amazon DynamoDB is ideal for existing or new applications that need a flexible NoSQL database with low read and write latencies, and the ability to scale storage and throughput up or down as needed without code changes or downtime. Use cases require a highly available and scalable database because downtime or performance degradation has an immediate negative impact on an organization's business. for e.g. mobile apps, gaming, digital ad serving, live voting and audience interaction for live events, sensor networks, log ingestion, access control for web-based content, metadata storage for Amazon S3 objects, e-commerce shopping carts, and web session management.

QUESTION 287

A customer implemented AWS Storage Gateway with a gateway-cached volume at their main office. An event takes the link between the main and branch office offline. Which methods will enable the branch office to access their data? Choose 3 answers.

- A. Use a HTTPS GET to the Amazon S3 bucket where the files are located.
- B. Restore by implementing a lifecycle policy on the Amazon S3 bucket.
- C. Make an Amazon Glacier Restore API call to load the files into another Amazon S3 bucket within four to six hours.
- D. Launch a new AWS Storage Gateway instance AMI in Amazon EC2, and restore from a gateway snapshot.
- E. Create an Amazon EBS volume from a gateway snapshot, and mount it to an Amazon EC2 instance.
- F. Launch an AWS Storage Gateway virtual iSCSI device at the branch office, and restore from a gateway snapshot.

Answer: DEF

Explanation:

A is certainly not right, because files persisted by Storage Gateway to S3 are not visible, let alone be accessible. <https://forums.aws.amazon.com/thread.jspa?threadID=109748>

B is invalid option because you cannot apply Lifecycle Policies because AWS Storage Gateway does not give you that option. Cached Volumes are never stored to Glacier and hence "C" is not a valid.

QUESTION 288

A company has configured and peered two VPCs: VPC-1 and VPC-2. VPC-1 contains only private subnets, and VPC-2 contains only public subnets. The company uses a single AWS Direct Connect connection and private virtual interface to connect their on-premises network with VPC-1. Which two methods increase the fault tolerance of the connection to VPC-1? Choose 2 answers

- A. Establish a hardware VPN over the internet between VPC-2 and the on-premises network.
- B. Establish a hardware VPN over the internet between VPC-1 and the on-premises network.
- C. Establish a new AWS Direct Connect connection and private virtual interface in the same region as VPC-2.
- D. Establish a new AWS Direct Connect connection and private virtual interface in a different AWS region than VPC-1.
- E. Establish a new AWS Direct Connect connection and private virtual interface in the same AWS region as VPC-1.

Answer: BE

QUESTION 289

What is the minimum time Interval for the data that Amazon CloudWatch receives and aggregates?

- A. One second
- B. Five seconds
- C. One minute
- D. Three minutes
- E. Five minutes

Answer: C

Explanation:

Many metrics are received and aggregated at 1-minute intervals. Some are at 3-minute or 5-minute intervals.

QUESTION 290

Which of the following statements are true about Amazon Route 53 resource records? Choose 2 answers.

- A. An Alias record can map one DNS name to another Amazon Route 53 DNS name.
- B. A CNAME record can be created for your zone apex.
- C. An Amazon Route 53 CNAME record can point to any DNS record hosted anywhere.
- D. TTL can be set for an Alias record in Amazon Route 53.
- E. An Amazon Route 53 Alias record can point to any DNS record hosted anywhere.

Answer: AC

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

QUESTION 291

A 3-tier e-commerce web application is currently deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database fail-over capability, and uses several read slaves for scaling. Data on all servers and the distributed file system directory is backed up weekly to off-site tapes. Which AWS storage and database architecture meets the requirements of the application?

- A. Web servers store read-only data in S3, and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment and one or more Read Replicas. Backup web and app servers backed up weekly via MFS. Database backed up via DB snapshots.
- B. Web servers store read-only data in S3, and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment and one or more read replicas. Backup web servers, app servers, and database backed up weekly to Glacier using snapshots.
- C. Web servers store read-only data in S3 and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment. Backup web and app servers backed up weekly via AMFS. Database backed up via DB snapshots.
- D. Web servers store read-only data in an EC2 NFS server, mount to each web server at boot time. App servers share state using a combination of DynamoDB and IP multicast. Database uses RDS with multi-AZ deployment and one or more Read Replicas. Backup web and app servers backed up weekly via MFS. Database backed up via DB snapshots.

Answer: A

Explanation:

<https://d0.awsstatic.com/whitepapers/Storage/AWS%20Storage%20Services%20Whitepaper-v9.pdf>

Amazon Glacier doesn't suit all storage situations. Listed following are a few storage needs for which you should consider other AWS storage options instead of Amazon Glacier.

Data that must be updated very frequently might be better served by a storage solution with lower read/write latencies, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, or relational databases running on EC2.

QUESTION 292

Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database Which backup architecture will meet these requirements?

- A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore.
- B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
- C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore.
- D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

Answer: A

Explanation:

You need to use enterprise backup software to provide file level restore. See:

https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf

Page 18:

If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products. AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

QUESTION 293

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA. The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database. In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements?

- A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region.
- B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region.
- C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region.
- D. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region.
- E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process.

Answer: A

QUESTION 294

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.

- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

Answer: B

Explanation:

<https://aws.amazon.com/sqs/faqs/>

QUESTION 295

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named Score Data. When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the Game State S3 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

Explanation:

The requirements state "Users will log into the game using their existing social media account to streamline data capture." This is what Cognito is used for, i.e. Web Identity Federation. Amazon also recommends to "build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation."

QUESTION 296

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

- A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
- B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- C. Amazon ElastiCache to store the writes until the writes are committed to the database.
- D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput.

Answer: B

Explanation:

<https://aws.amazon.com/sqs/faqs/>

There is no limit on the number of messages that can be pushed onto SQS. The retention period of the SQS is 4 days by default and it can be changed to 14 days. This will make sure that no writes are missed.

QUESTION 297

You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 Instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The two EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes) for a total of 16,000 random IOPS on the instance. The EC2 Instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPS like the original four for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized: use an EBS-Optimized instance that provides larger throughput.

- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6.000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instant root volume to also be a 500GB 4.000 Provisioned IOPS volume.

Answer: E

QUESTION 298

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS. During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database. The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage. The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year improvements. To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

- A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
- B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
- C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
- D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Answer: B

Explanation:

The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months' worth of data. This rules out RDS as a possible DB solution which leaves you with RedShift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference.

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

QUESTION 299

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met. Provide the ability for real-time analytics of the inbound biometric data. Ensure processing of the biometric data is highly durable. Elastic and parallel. The results of the analytic processing should be persisted for data mining. Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data, analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EMR with Amazon Kinesis and save the results to DynamoDB.

Answer: B

QUESTION 300

You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project. What database implementation would better fit this scenario, keeping costs as low as possible?

- A. Use RDS Multi-AZ with two tables, one for "Active calls" and one for "Terminated calls". In this way the "Active calls_

table is always small and effective to access.

- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE" field that can be equal to "ACTIVE" or "TERMINATED". In this way the SQL query is optimized by the use of the Index.

Answer: B

Explanation:

<https://aws.amazon.com/dynamodb/faqs/>

Q: Can a global secondary index key be defined on non-unique attributes?

Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique.

Q: Are GSI key attributes required in all items of a DynamoDB table?

No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them, then that item will not be indexed by the corresponding GSI. In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.

[Visit PassLeader and Download Full Version AWS-Associate Exam Dumps](#)